Investment Policy 2026-2036

Revised Structured Edition

A clearer decision architecture with a deeper decade worldview, an explicit theory of value capture, conditional business-model priors, disfavored models, company research mapping, governance, and falsifiers.

Evidence cutoff	22 August 2026
Main structure	15 chapters and 6 audit annexes
Manuscript SHA-256	a993e207ce377bf00aa338850de02bad3a3636da55c8567bee12fa0278b39ef8

Policy boundary. This document authorizes no trade, order, execution, hedge, transfer, outbound message, account opening, target allocation, or portfolio change. Named companies are research examples at current admissible weight zero until every applicable gate and owner approval passes.

Contents

Part I - Policy in Brief 14

Chapter 1 - The Policy in Brief: Beliefs, Boundaries, and Current Readiness 15

What the policy believes . . . 15
The objective hierarchy . . . 16
How success is measured . . . 17
Durable policy, dated worldview, living research, and current state . . . 18
Current readiness and certification dashboard . . . 18
Current blockers at a glance . . . 20
Authority hierarchy . . . 20
Three conclusions that must remain separate . . . 21
The concise operating map . . . 21
1. Establish annual investor facts . . . 21
2. Protect survival and continuity . . . 21
3. Define the strategic baseline and risk budgets . . . 22
4. Look through the comparator before active ideas . . . 22
5. Admit active research only through non-compensating gates . . . 22
6. Keep speculative options inside the opportunity and loss budgets . . . 23
7. Freeze the decision record before approval . . . 23
8. Use narrow monitoring and explicit clocks . . . 23
9. Audit governance and active-program viability annually . . . 23
10. Certify only when every joint blocker is complete . . . 24
How to read the rest of the book . . . 24
What the reader should know before continuing . . . 24

Part II - Durable Portfolio Policy 26

Chapter 2 - The Investment Constitution 27

The constitution begins with an order of purposes . . . 27
Material claims must carry epistemic labels . . . 28
Risk is governed through three separate budgets . . . 28
Measurement starts with liabilities, not the return league table . . . 29
The comparator bears risk and sets the burden of proof . . . 30
Shortfall Parity prevents safety from becoming silent failure . . . 30
Evidence governs tolerable loss; valuation governs admission . . . 31
Recourse leverage is prohibited by default . . . 31
Every exposure must be viewed on a look-through basis . . . 32
Broad strategic rules may move faster than active judgments . . . 32
Tax is friction, not an investment thesis . . . 33
Persistent breaches require amendment, not excuses . . . 33
Governance must be viable enough to endure . . . 34

Research never authorizes implementation . . . 35
Why the principles must coexist . . . 35

Chapter 3 - Investor Facts, Funded State, and Liquidity 37

Funded status is the selector, not a return forecast . . . 37
Hypothetical worked example - calculating FS and its sensitivity . . . 37
The four funded states . . . 38
State 1 - FS below 1.00 . . . 38
State 2 - FS from 1.00 to 1.24 . . . 39
State 3 - FS from 1.25 to 1.74 . . . 39
State 4 - FS at or above 1.75 . . . 39
What each table column controls . . . 39
The second axis - capacity beyond the FS decimal . . . 40
Hypothetical worked example - the second axis . . . 40
Hysteresis and state transitions . . . 41
Hypothetical worked example - promotion and immediate demotion . . . 41
Speculative options are nested and charged at total loss . . . 41
Hypothetical worked example - market value versus loss charge . . . 42
The Investor Liquidity Floor . . . 42
Hypothetical worked example - ILF target and hard floor . . . 43
Hypothetical worked example - the D24 election . . . 43
Rebalancing architecture - fast for broad sleeves, slow for active risk . . . 44
Hypothetical worked example - the two rebalancing thresholds . . . 44
Decision clocks and fail-closed responses . . . 44
Immediate freeze and a 0-5-trading-day resolution window . . . 44
Thirty-day full re-underwrite and no add . . . 44
Valuation-only deterioration . . . 45
Hypothetical worked example - different evidence, different clocks . . . 45
Entry pacing is an information rule, not calendar theater . . . 45
Policy Exception Events and the freeze on discretion . . . 45
Hypothetical worked example - the rolling exception rule . . . 46
How the pieces fit together . . . 46

Chapter 4 - The Three Risk Budgets and the Grind Test 47

Why a grind test is necessary . . . 47
The four default-admissible buckets . . . 48
Equity . . . 48
Matched nominal duration . . . 48
Verified duration-matched inflation-linked bonds . . . 48
ILF cash . . . 49
How to read the table . . . 49
Clearly labeled hypothetical application examples . . . 50
Hypothetical example 1 - applying the inherited mix . . . 50
Hypothetical example 2 - applying the boundary mix . . . 50

Hypothetical example 3 - a mislabeled ILB bucket . . . 50
Hypothetical example 4 - an omitted withdrawal reverses the pass . . . 50
Hypothetical example 5 - the liability estimate changes . . . 51
What the test rules out . . . 51
What the test cannot prove . . . 51

Chapter 5 - Comparator, Look-Through Limits, and Certification 53

The comparator is a real position, not a neutral backdrop . . . 53
What proxy scoring is trying to measure . . . 53
How the top-20-plus-buckets method works . . . 54
What the inherited proxy found - and what it did not prove . . . 55
The forced-choice rule . . . 56
Why a temporary diagnostic ceiling is tightly bounded . . . 56
What must be studied before either permanent branch is credible . . . 57
Certification is a blocker, not a blessing . . . 57
Limitations and failure modes . . . 58

Part III - Decade Investment Worldview 59

Chapter 6 - The Decade Investment Worldview: Why These Forces May Shape 2026-2036 60

Status, governing boundary, and evidence date . . . 60
Method: separating evidence, mechanism, cycle, forecast, and implication . . . 60
The value-capture layer: creation is not ownership . . . 61
Integrated value-capture matrix . . . 63
Twelve forces and one counterforce . . . 64
Force 1 - AI and software diffusion: capability must become governed work . . . 64
Starting evidence and epistemic limit . . . 64
Why the force may persist . . . 64
Capital cycle and value distribution . . . 65
Strongest counterthesis, indicators, and model effects . . . 65
Force 2 - Compute, data centers, and semiconductors: a capital cycle with different clocks . . . 66
Starting evidence and layered economics . . . 66
Persistence and the capital-cycle response . . . 66
Who captures and who bears the loss . . . 67
Countertheses, falsifiers, and model effects . . . 67
Force 3 - Electricity, grids, and power: global abundance can coexist with local scarcity . . . 67
Starting evidence and the correct unit of analysis . . . 67
Structural mechanism, supply response, and possible equilibria . . . 68
Value capture: the rent can move faster than the electrons . . . 69
Countertheses, falsifiers, and model effects . . . 69
Force 4 - Electrification and industrial automation: diffusion through the last meter . . . 70
Starting evidence and persistence . . . 70
Capital cycle and distribution . . . 70
Countertheses, falsifiers, and model effects . . . 71

Force 5 - Aging, health, and labor: slow cohort arithmetic meets a hard payer gate . . . 71
Starting evidence and geographic divergence . . . 71
Why labor scarcity may persist without becoming provider profit . . . 71
Health demand, innovation, and the contested surplus . . . 72
Capital cycles and value distribution . . . 73
Countertheses, falsifiers, and model effects . . . 73
Force 6 - Fiscal activism, defense, and the state: funded demand with a monopsony . . . 74
Starting evidence and persistence . . . 74
Capital cycle and procurement chain . . . 74
Distribution of surplus and the per-share test . . . 75
Countertheses, falsifiers, and model effects . . . 75
Force 7 - Fragmentation and supply-chain resilience: selective duplication within continued interdependence . . . 76
Starting evidence and a more precise thesis . . . 76
Resilience is not autarky . . . 76
Capital response and distribution . . . 76
Counterthesis, falsifiers, and model effects . . . 77
Force 8 - Inflation, real rates, debt, and financial repression: a wider regime distribution . . . 77
Starting evidence without a one-way forecast . . . 77
Structural mechanisms and alternative regimes . . . 78
Distribution: inflation is a bargaining process . . . 78
Countertheses, falsifiers, and model effects . . . 79
Force 9 - Natural resources and the harvest phase: depletion is not a shortcut to owner return . . . 79
Starting evidence and the scarcity distinction . . . 79
Persistence, capital cycle, and verified harvest . . . 80
Value distribution . . . 80
Countertheses, falsifiers, and model effects . . . 81
Force 10 - Emerging-market formalization and digital platforms: access can outrun owner capture . . . 81
Starting evidence and a decade mechanism . . . 81
Platform and financial capital cycles . . . 82
Distribution and the eight bridges to each share . . . 82
Countertheses, falsifiers, and model effects . . . 82
Force 11 - Climate adaptation and physical resilience: necessary expenditure is not automatically collectible . . . 83
Starting evidence and persistence . . . 83
From avoided loss to paid cash . . . 83
Distribution . . . 84
Countertheses, falsifiers, and model effects . . . 84
Force 12 - Urbanization, infrastructure, and the ordinary location rent . . . 85
Starting evidence and persistence . . . 85
Funding, capital cycle, and distribution . . . 85
Countertheses, falsifiers, and model effects . . . 86
Counterforce - The ordinary decade and the ordinary compounder . . . 86
Force-to-business-model matrix . . . 87
Cross-force interactions: the decade is a system, not twelve sleeves . . . 88

AI, labor, health, and the productivity split . . . 88
Compute, power, land, and finance . . . 89
State procurement, debt, and crowding . . . 89
Fragmentation, resources, and the double capital cycle . . . 89
Urbanization, formalization, and land capture . . . 90
Inflation, replacement cost, and apparent profit . . . 90
Coherent decade regimes and how value moves . . . 90
Productive diffusion . . . 91
Inflation and repression grind . . . 91
Fragmented blocs . . . 91
Crash and sequence . . . 91
Adaptation and physical constraint . . . 92
Plateau and muddle . . . 92
Where capture may strengthen, weaken, or remain radically uncertain . . . 92
Monitoring architecture: from narrative to observable rates . . . 93
Evidence discipline, source conflicts, and model uncertainty . . . 94
Practical research boundary . . . 94
Precommitted triggers for revision . . . 95
What would change this worldview . . . 95

Chapter 7 - Favored and Disfavored Conditional Business-Model Priors 96

Status, purpose, and the three conclusions that must never be collapsed . . . 96
The ranking at a glance . . . 96
How the ranking is formed . . . 97
Strongest conditional priors . . . 98
Installed-base aftermarket with qualification and high-consequence service . . . 98
Ratings and select mandated assurance - strong economics with binary legitimacy tails . . . 99
Exchanges, clearing, market infrastructure, and payment rails . . . 100
Reliability distribution and low-ticket mission-critical inputs . . . 100
Record-to-action software and lawful proprietary data . . . 101
Favorable but valuation-sensitive priors . . . 102
Semiconductor, equipment, power-control, and industrial bottlenecks . . . 102
Disciplined insurance float, deposit franchises, and other balance-sheet businesses . . . 103
Selective brands with cohort desire and controlled distribution . . . 104
Funded health innovation . . . 105
Constructive regulated franchises and funded state procurement . . . 105
High-quality serial acquirers - favorable only after the treadmill is charged as reinvestment . . . 106
Emerging and unproven priors . . . 107
Currently unsupported or without a qualifying example . . . 107
Natural-resource and infrastructure harvest . . . 108
Audit, asset management, broad regulated resilience, and generic EM growth . . . 108
Recurring negative priors . . . 109
Cross-class value-capture map for 2026-2036 . . . 110
Required evidence and class-level falsification . . . 111

Positive examples and counterexamples are paired research designs . . . 112
Capital-cycle and valuation traps common to the favored list . . . 112
What not to infer from the ranks . . . 113
How these priors should change research, not authorize ownership . . . 114

Chapter 8 - From Worldview to Investment Themes and Business-Model Winners and Losers 115

The master crosswalk . . . 115
Interaction map: the forces are not independent themes . . . 117
How to analyze the distribution of surplus . . . 118
AI and software diffusion . . . 119
Compute, data centers and semiconductor capital cycles . . . 120
Electricity demand, grids and power constraints . . . 121
Electrification and industrial automation . . . 121
Aging, healthcare innovation and labor scarcity . . . 122
Fiscal activism, defense and state procurement . . . 123
Deglobalization, supply-chain resilience and economic security . . . 124
Inflation, real rates, debt and financial repression . . . 125
Natural-resource depletion and harvest cycles . . . 126
Emerging-market formalization and digital platforms . . . 126
Urbanization and infrastructure as a cross-cutting force . . . 127
Climate adaptation and physical resilience . . . 128
Verification, standards and market infrastructure as cross-cutting models . . . 128
The plateau/muddle counter-scenario and ordinary compounders . . . 129
Where value capture may strengthen, weaken or remain radically uncertain . . . 129
Research protocol for using the crosswalk . . . 130
The bridge is a research hypothesis, not expected-return proof . . . 130

Part IV - Security Research and Operating System 131

Chapter 9 - From Business-Quality Prior to Security-Return Conclusion 132

The five-pillar profile, not an additive quality score . . . 132
Pillar A - Demand and lock-in . . . 133
Pillar B - Industry structure, cycle, and payer . . . 134
Pillar C - Capital economics . . . 135
Pillar D - Governance and per-share capture . . . 136
Pillar E - Valuation and falsifiability . . . 137
Evidence coverage and the standard of proof . . . 137
Class overlays: asking the question appropriate to the model . . . 138
From an admitted thesis to an allowed loss . . . 139
Reading the implied maximum-weight table . . . 141
How the gates interact in practice . . . 141
Boundaries and failure modes of the scorecard itself . . . 142

Chapter 10 - Research Priority, Classification, and the Eight-Underwrite Capacity Rule 143

Classification vocabulary and explicit legacy mapping . . . 143

Why 42 notes cannot become 42 live underwrites . . . 146
Expected information value, not expected share-price return . . . 147
The eight selection dimensions . . . 148
Expected information value . . . 148
Model-prior strength . . . 149
Valuation distance . . . 149
Portfolio and comparator uniqueness . . . 149
Evidence freshness . . . 150
Counter-canon and non-US balance . . . 150
Complexity . . . 150
Probability of changing a decision . . . 150
The queue construction procedure . . . 151
A practical selection ledger . . . 152
Wmax must be removed from queue logic . . . 153
Current-state consequence: this chapter does not appoint eight names . . . 153
Refresh, displacement, expiry, and sunset . . . 154
Capacity governance and failure modes . . . 154
Research priority is not portfolio priority . . . 155

Chapter 11 - The Canonical Decision Workflow and a Complete Worked Example 156

The canonical sequence . . . 156
Case file: the entirely fictional Northbridge household . . . 157
The liability register . . . 158
Funded status and the mandatory liability sensitivities . . . 158
The Investor Liquidity Floor . . . 159
Strategic architecture and the grind test . . . 160
Comparator look-through and the forced choice . . . 161
Selecting a research prior without creating a buy list . . . 163
Full company underwrite: fictional Alder Route Systems plc . . . 163
Evidence Coverage . . . 164
The five non-compensating pillars . . . 164
Reverse DCF, conservative return, and falsifiers . . . 165
Allowed-loss arithmetic and portfolio budgets . . . 165
The derating-loss budget . . . 166
Independent challenge and separate fictional owner approval . . . 166
Decision journal at the point of the fictional approval . . . 167
Monitoring the case without converting noise into action . . . 168
Closure, expiry, and research sunset . . . 169
Control ledger: where the case would fail closed . . . 169
What the arithmetic does and does not say . . . 170
The governing boundary after a fully worked case . . . 171
What the example demonstrates - and what it cannot establish . . . 171

Chapter 12 - Monitoring, Falsification, Sell Discipline, and Research Sunset 172

Monitoring is a hierarchy, not a stream of alerts . . . 172
Evidence has an expiry date . . . 173
Falsification separates an investment thesis from a story . . . 174
Price volatility is not thesis breakage . . . 175
The exit hierarchy and its decision clocks . . . 175
The valuation-drawdown budget is set before purchase . . . 177
Research sunset is an attention rule . . . 177
Failure modes the discipline is meant to prevent . . . 178

Part V - Governance, Risks, and Certification **180**

Chapter 13 - Risks, Unresolved Disputes, and Portfolio-Level Falsifiers **181**

Uncertified liability closure . . . 181
A comparator born in breach . . . 182
The active program may be uneconomic . . . 183
Concentration, quality crowding, and the danger of hidden sameness . . . 184
Procyclical diversifier evidence and the unresolved insurance question . . . 184
Narrative and capital-expenditure risk . . . 185
Political, legal-title, custody, and currency risk . . . 186
Obsolescence and dishonest economic lives . . . 186
Other unresolved allocation disputes . . . 187
How risks combine into correlated failure . . . 187
Company-level and portfolio-level falsification . . . 188
What evidence would change the overall view . . . 189

Chapter 14 - Governance, Decision Rights, Maintenance, and Certification Readiness **190**

Governance rests on an authority hierarchy . . . 190
Functional roles and decision rights . . . 191
Incapacity continuity is a survival control . . . 193
The maintenance calendar . . . 194
Monthly cycle . . . 194
Quarterly cycle . . . 195
Annual cycle . . . 196
Event-driven clocks . . . 197
Shelf lives and the 12-, 24-, 36-month, and five-year horizons . . . 197
Amendment, compliant design, approval, exception, and execution flow . . . 199
Ordinary evidence and decision flow . . . 199
Formal amendment flow . . . 199
Compliant design flow . . . 200
Emergency exception flow . . . 200
Approval and execution separation . . . 200
Current certification-readiness and blocker register . . . 200
Readiness outputs and records . . . 202
Failure states and governed defaults . . . 203

Annual certification-readiness statement . . . 204
The policy is maintainable only if responsibility is explicit . . . 204

Chapter 15 - Certification, Ownership, and the Policy Boundary 205

Final as a document, not certified as an investment plan . . . 205
What the edition establishes . . . 205
What remains uncertified . . . 206
The controlling document hierarchy . . . 207
The evidence cutoff and the meaning of a source snapshot . . . 208
Ownership, maintenance, challenge, and approval . . . 208
Prohibited uses of the policy . . . 209
How final amendments and future amendments should be read . . . 210
Boundary failure modes . . . 210

Annex A - Full Company Research Map 212

A map of questions, not a portfolio . . . 212
How to read the common fields . . . 213
Core Research Candidates - quality plausible, valuation and full evidence still open . . . 213
1. S&P; Global - SPGI, NYSE; United States . . . 214
2. CME Group - CME, Nasdaq; United States . . . 214
3. Visa - V, NYSE; United States . . . 214
4. Fastenal - FAST, Nasdaq; United States . . . 215
5. Berkshire Hathaway - BRK.B, NYSE; United States . . . 215
6. Chubb - CB, NYSE; Switzerland . . . 216
7. DBS Group - D05, SGX; Singapore . . . 216
8. Taiwan Semiconductor Manufacturing - TSM/2330, NYSE/TWSE; Taiwan . . . 217
9. ASML - ASML, Euronext Amsterdam/Nasdaq; Netherlands . . . 217
10. Schneider Electric - SU, Euronext Paris; France . . . 218
11. RELX - REL, LSE/NYSE; United Kingdom . . . 218
Watchlist/Price Sensitive - model plausible, gate or expectation unresolved . . . 218
12. GE Aerospace - GE, NYSE; United States . . . 219
13. Microsoft - MSFT, Nasdaq; United States . . . 219
14. Alphabet - GOOGL, Nasdaq; United States . . . 219
15. NVIDIA - NVDA, Nasdaq; United States . . . 220
16. Eaton - ETN, NYSE; Ireland . . . 220
17. London Stock Exchange Group - LSEG, LSE; United Kingdom . . . 221
18. Moody's - MCO, NYSE; United States . . . 221
19. Constellation Software - CSU, TSX; Canada . . . 222
20. Eli Lilly - LLY, NYSE; United States . . . 222
21. Novo Nordisk - NOVO-B, Nasdaq Copenhagen; Denmark . . . 222
22. MercadoLibre - MELI, Nasdaq; United States legal domicile, Latin America economic exposure . . . 223
23. Ferrari - RACE, NYSE/Euronext Milan; Netherlands/Italy . . . 223
24. Rheinmetall - RHM, Xetra; Germany . . . 224
Speculative Options - research optionality only, with total loss assumed . . . 224

25. CoreWeave - CRWV, Nasdaq; United States . . . 224
26. Samsara - IOT, NYSE; United States . . . 225
27. Symbotic - SYM, Nasdaq; United States . . . 225
28. Nu Holdings - NU, NYSE; Cayman Islands/Brazil . . . 225
29. Sea Limited - SE, NYSE; Singapore/Southeast Asia . . . 226
30. BYD - 1211, HKEX / 002594, Shenzhen; China . . . 226
Avoid/Counterexamples - appealing narratives that fail the gates . . . 227
31. Palantir Technologies - PLTR, Nasdaq; United States . . . 227
32. Intel - INTC, Nasdaq; United States . . . 227
33. Boeing - BA, NYSE; United States . . . 228
34. Coinbase - COIN, Nasdaq; United States . . . 228
35. PG&E; - PCG, NYSE; United States . . . 228
36. Plug Power - PLUG, Nasdaq; United States . . . 229
37. Nike - NKE, NYSE; United States . . . 229
38. Teladoc Health - TDOC, NYSE; United States . . . 230
39. Flagstar Financial - FLG, NYSE; United States . . . 230
40. Bausch Health - BHC, NYSE/TSX; Canada . . . 230
41. First Quantum Minerals - FM, TSX; Canada/Zambia-Panama exposure . . . 231
42. Adani Enterprises - ADANIENT, NSE; India . . . 231
Reading across business-model classes . . . 232
How the map should be used . . . 234
How the map must not be used . . . 234

Annex B - Detailed Operating Checklist 236

1. Annual facts . . . 236
2. Survival . . . 237
3. Strategic design . . . 238
4. Look-through . . . 239
5. Active admission . . . 240
6. Options . . . 240
7. Decision record . . . 241
8. Monthly/quarterly . . . 242
9. Annual governance . . . 242
10. Certification blockers and stop rules . . . 243

Annex C - Debate Adjudication and Hard Questions 245

C.1 Status and purpose . . . 245
C.2 Controlling hierarchy . . . 245
C.3 Five-stage debate sequence . . . 245
C.4 Round 4 adjudication . . . 246
C.5 Final Round 5 amendments . . . 247
C.5.1 Controlling D-1 reopening rule . . . 248
C.6 Rejected alternatives, minority views, and unresolved disputes . . . 248
C.7 The twelve final hard questions - closure map . . . 249

C.8 Practical boundary . . . 249

Annex D - Bias, Prohibition, and Implementation Self-Audit 250

1. No buy call, target weight, entry instruction, order, or portfolio mutation . . . 250
2. No narrative shortcut to expected return . . . 251
3. No return to the old additive Business-Model Score . . . 251
4. No short-history challenger as a Core Research Candidate . . . 252
5. Do not ignore exposure already embedded in the comparator . . . 252
6. No inflation of weak, conflicting, or missing evidence . . . 253
7. Harvest status is neither normal sizing nor renewable by inertia . . . 253
8. State budgets are not collectible revenue without the Payer Quality Gate . . . 254
9. No personalized tax, account, domicile, currency, or custody conclusion . . . 254
10. No silent resolution of disputes or professional-review gates . . . 255

Reading the audit result correctly . . . 255
What the final bias actually is . . . 256
Hidden consensus: when a skeptical process still assumes activity . . . 256
Recency bias: a current map is not a decade-long base rate . . . 257
Survivorship bias: visible winners are not a validated sample . . . 258
Narrative fallacy: causal coherence is not cash capture . . . 259
Concentration: the comparator can violate the risk the active sleeve is forbidden to add . . . 260
Obsolescence: correct evidence may arrive after the loss . . . 260
Political and regulatory dependence: law can dominate business quality . . . 261
Implementation complexity: the institutional process fails the one-person test . . . 262
The final one-person verdict and the approval boundary . . . 263

Annex E - Glossary, Formula Registry, and Parameter Register 264

E.1 Purpose, scope, and interpretive rules . . . 264
E.2 Glossary . . . 265
E.2.1 Portfolio, liability, and governance terms . . . 265
E.2.2 Research, evidence, and security terms . . . 267
E.2.3 Factor, scenario, provenance, and notation terms . . . 269
E.3 Formula registry . . . 270
E.4 Parameter register . . . 272
E.4.1 Document, certification, and standing boundaries . . . 272
E.4.2 Funded-state bands and portfolio architecture . . . 272
E.4.3 Liquidity, rebalancing, emergency, and exception parameters . . . 274
E.4.4 Scorecard, evidence, loss, and sizing parameters . . . 276
E.4.5 Implied best-case maximum-weight matrix . . . 277
E.4.6 Common-factor caps and comparator-governance parameters . . . 278
E.4.7 Monitoring, evidence-decay, and research-attention parameters . . . 280
E.4.8 Active-program economics and hours . . . 281
E.4.9 Valuation, tax, and exit parameters . . . 282
E.4.10 Custody, jurisdiction, illiquidity, FX, and professional-review parameters . . . 282
E.4.11 Grind-test parameters and diagnostic boundary . . . 283

E.5 Research-falsifier threshold register . . . 283
E.6 Superseded hierarchy and residual contradictions . . . 286
E.7 Amendment and breach authority matrix . . . 287
E.8 Completeness checklist and practical boundary . . . 287

Annex F - Traceability, Open Disputes, Provenance, and Final Amendments 288

Appendix A - Full debate traceability . . . 288
A1. Policy articles P1-P52 . . . 288
A2. Amendment ledger F1-F4 . . . 290
Appendix B - Open-dispute register . . . 292
Appendix C - Fable provenance receipts . . . 292
Appendix D - Final Fable amendments F5-1 through F5-14 . . . 293
Final unresolved minority views and adjudicating observables . . . 293

Part I - Policy in Brief

Chapter 1 - The Policy in Brief: Beliefs, Boundaries, and Current Readiness

This book is a governance policy for investing from 2026 to 2036. It is not a portfolio, a recommendation, or a certificate that any investor or security is ready for implementation. Its central belief is simple: investment success is not the highest reported return. Success is preserving legal ownership and decision continuity, meeting essential liabilities through adverse paths, and then compounding purchasing power after tax, cost, and error. Optional goals come only after those obligations are protected.

The policy begins from uncertainty rather than confidence. Investor facts are incomplete. The actual liability schedule, funded status, market-drawdown capacity, tax basis, domicile and custody structure, named comparator, factor look-through, withdrawal path, and current company evidence and valuations have not been verified. The rules in this book explain how those facts must be assembled and how later decisions would be governed. They do not supply the missing facts and do not authorize a transaction.

That boundary is categorical. Nothing in this book authorizes a trade, order, hedge, transfer, account opening, target allocation, entry price, execution instruction, external message, or portfolio change. A scenario is not a forecast instruction. A favored business-model prior is not a buy list. A company classification is a place in a research queue, not an ownership decision. A diagnostic ceiling is a loss-control boundary, not a target weight. Every named company has a current admissible weight of zero until a complete and current evidence and valuation review, verified investor facts, portfolio look-through, relevant tax and legal review, all policy gates, and separate implementation approval have been completed.

The evidence cutoff for this edition is 22 August 2026. A source snapshot means only that a disclosure or data source existed by that date. It does not independently verify a management claim, establish that a budget will become collectible cash, complete a valuation, or prove that a statement remains current after the cutoff. Company classifications expire after 12 months without refresh and then revert to Unclassified Research Note. The dated worldview should be reconsidered as evidence changes. The durable constitution should change only by explicit amendment.

What the policy believes

The policy believes that the investor's obligations, legal rights, and ability to continue the plan should govern the portfolio before any security story does. Assets are useful only insofar as their rights, liquidity, currency, timing, and returns serve real objectives. A holding that performs well in isolation can still fail the plan if it creates forced-sale risk, duplicates an existing factor, cannot be accessed, or is held at a price that offers inadequate owner return. Conversely, a simple strategic asset can be valuable even when it is not exciting because it matches a liability, preserves liquidity, or supplies a low-cost baseline that active work must beat.

The policy also believes that risk is plural. Permanent impairment, mark-to-market drawdown, and decade-long real grind are different failures. Permanent impairment is destruction or non-recovery of the economic claim through insolvency, dilution, fraud, confiscation, lost title, broken economics, or an unrecoverable purchase price. Drawdown is a decline in quoted value that may or may not reflect permanent damage, but that can still cause forced sale or behavioral failure. Real grind is the slow erosion caused by weak after-tax returns, inflation, withdrawals, currency mismatch, and liability revaluation. Passing one test does not compensate for failing another.

The policy believes that knowledge and price have different jobs. Evidence determines how much permanent loss may be tolerated after all other gates pass. Valuation determines whether a security is admissible and whether its conservative expected return clears the relevant comparator. Better evidence cannot waive a hard factor cap, cure illiquidity, reduce an anti-gaming severity floor, or make a poor price attractive. A lower price cannot repair defective title, governance, or business economics. Confidence and narrative quality grant no weight directly.

The policy believes that the low-cost implementable comparator is a real portfolio choice, not a neutral background. It must be named, scored, and looked through to its underlying issuers, factors, sovereigns, currencies, and counterparties. Active complexity bears the burden of proving value after tax, fees, research hours, data, implementation friction, and error. The admissible active allocation can therefore be zero. A simple comparator plus protected liquidity and verified nominal and real liability matching is a complete governed state, not a failure to find ideas.

The policy believes that thematic importance is not security return. Its central value-capture chain is:

social value -> customer benefit -> bargaining split among labor, rent, capital, the state, and entrepreneurship -> company cash -> diluted per-share owner cash -> price -> portfolio role

Every arrow is a gate, not a rhetorical transition. A useful technology may create enormous social value yet give most of the benefit to customers through lower prices. Productivity gains may flow to wages, taxes, interest, land or power rents, new entrants, or a bottleneck supplier rather than to the focal company. Accounting profit may be consumed by maintenance capital, working capital, debt service, acquisitions, stock-based compensation, minority claims, or state leakage before cash reaches the parent and each diluted share. Even durable per-share cash can offer a poor security return if the purchase price already assumes too much. Even an attractively priced security can have no portfolio role if liabilities, liquidity, currency, tax, factor headroom, or loss capacity reject it.

This chain is therefore a central belief rather than a detail reserved for company research. For each decade force, the policy asks who owns the scarce input, who has bargaining power, how regulation and financing shape the split, how new supply or substitution could reverse early rents, and which indicators show where value is actually accumulating. A social need can grow while suppliers overbuild. A technology can transform society while customers, workers, the state, or new entrants capture most of the value. A government budget can rise without becoming an appropriated, contracted, collectible, profitable cash flow. A business can be excellent while its security is priced for a result it cannot deliver. The decade worldview therefore generates conditional research hypotheses and falsifiers. It never supplies expected-return proof by itself.

The policy believes that simplicity is a risk control only when it preserves the necessary controls. An elaborate system that one person cannot maintain produces stale classifications, selective monitoring, and false confidence. The recurring operating system is consequently bounded: verify liabilities and funded status; protect the Investor Liquidity Floor; name and look through one comparator; match nominal and real liabilities; enforce issuer and factor limits; maintain one journal and one exception log; run no more than eight live full underwrites, and fewer if the hours or economics do not work; and test program viability annually. Missing work narrows active scope. It never lowers the evidence standard.

Finally, the policy believes that written rules must constrain behavior when conditions are uncomfortable. Broad strategic sleeves may rebalance through approved mechanical bands. Single-name decisions require current evidence because a lower price can mean either better value or a broken claim. Persistent breaches must become an explicit amendment or a compliant design change, not a renewable exception. Emergency authority is narrow, risk-monotone, recorded within 24 hours, independently reviewed within seven days, and expires within 30 days. Two material Policy Exception Events in 12 months freeze new discretion pending review.

The objective hierarchy

The policy orders objectives rather than treating them as an undifferentiated wish list. A lower objective cannot consume resources required by a higher one.

1. **Survival and legal ownership.** The investor must retain enforceable title, usable access, and the ability to withstand plausible adverse paths without ruin or forced sale. Custody resilience, beneficiary access, and incapacity arrangements are part of investment survival.

2. **Near-term essential liabilities and continuity.** Essential spending, known taxes, fixed commitments, and dependable decision authority must remain funded even when markets are impaired. The Investor Liquidity Floor exists for this purpose and is not opportunistic buying power.
3. **Liability closure across time.** The portfolio must be tested against the timing, currency, inflation character, probability, and flexibility of liabilities, including withdrawals. A starting balance-sheet surplus is not enough if sequence risk or real grind later breaks Coverage.
4. **Tax-aware real compounding.** After the higher claims are protected, the portfolio should increase purchasing power after tax, fees, transaction costs, and the inflation experienced by the liabilities. Nominal return is an incomplete measure.
5. **Optional aspiration.** Additional legacy, philanthropy, discretionary consumption, thematic expression, prestige, and intellectual interest are legitimate only to the extent that they do not endanger the higher objectives.

This hierarchy does not authorize permanent over-caution. Shortfall Parity requires a modeled liability shortfall caused by insufficient productive exposure to receive a review as serious as an excess-risk breach. The remedies are not automatically more market risk. Contributions, spending flexibility, timing, and objective changes are considered first. Risk may increase only within verified financial capacity and stable trait tolerance. Need for return is never treated as capacity to bear loss.

How success is measured

The policy reports success in the investor's terms before reporting the portfolio's market return. No single number is sufficient. The measures below form a hierarchy of evidence about whether the plan is working.

Measure	What it asks	Required interpretation
Legal title and access	Can the investor or authorized successor enforce and use the claim when needed?	A statement balance is insufficient if title, segregation, transfer, beneficiary, incapacity, or sanctions access is uncertain.
Investor Liquidity Floor	Are near essential liabilities, known taxes, fixed commitments, and the applicable income interruption protected?	Report target, hard floor, currency, duration, custody integrity, access, and real drag. Never buy risk from the hard floor.
Funded status, FS = A* / PVL*	Do investable assets after known taxes and transaction costs cover probability-weighted present liabilities plus safety margin?	Report base PVL* and +/-10% and +/-20%. If +/-10% changes the band, use the lower state.
Stressed Coverage	Do assets continue to meet liabilities along crash, inflation, income-loss, currency, and withdrawal paths?	A starting FS above 1.00 is not certification. The path, not only the endpoint, matters.
Permanent-loss budget	How much irreversible active error can the plan survive?	Measure expected loss exposure, not conviction. Evidence, liquidity, severity, and factor headroom constrain it.
Mark-to-market drawdown budget	Can the investor finance and behaviorally endure the quoted loss path?	Separate business impairment from derating, but budget both before purchase.
Real-grind budget	Can the plan withstand a decade of weak real after-tax outcomes?	Include liability inflation, withdrawals, currency mismatch, and basis risk. Do not hide spending outside the stress.
Real after-tax progress	Is purchasing power growing after taxes, fees, transaction friction, and relevant inflation?	Report this before celebrating nominal return.
Comparator-relative value	Does active complexity improve the result relative to a named implementable alternative?	Compare after tax, cost, hours, error, and implementation. Separate comparator-design value from company-selection value.
Governance viability	Can the actual operator maintain the required controls at a justified cost?	Print the hours and costs. Below the prospective 2x benefit-to-cost threshold, active capital remains in paper mode.

Return still matters. Shortfall Parity and the comparator burden prevent the policy from congratulating itself for safety while objectives become unreachable or a simple alternative performs the same job better. But return is judged after title, liability, liquidity, tax, cost, and path have been incorporated. Outperforming a market index while essential liabilities become less secure is not success. Lagging a rising index while a deliberately matched liability is funded is not automatically failure. The appropriate comparison must reflect both the objective and the opportunity cost.

Durable policy, dated worldview, living research, and current state

Readers should not give every statement in the book the same shelf life or authority. The revised edition separates five layers so that a changing forecast cannot silently rewrite the constitution and a stale company note cannot masquerade as current evidence.

Policy layer	Typical content	Status and shelf life	What can change it	What it cannot do
Durable constitution	Purpose hierarchy, epistemic labels, three risk budgets, liability-first measurement, comparator burden, evidence-to-loss, look-through, leverage default, tax discipline, amendment rules, minimum governance, research non-authorization	Durable until explicit amendment	A documented policy amendment with rationale, effects, challenge, and effective date	It cannot manufacture investor facts, evidence, valuation, or implementation authority.
Calibrated and provisional rules	FS bands, sleeve ranges, error budgets, issuer base-loss allowances, severity floors, factor caps, liquidity multipliers, rebalancing and review clocks, viability thresholds	Controlling policy boundaries, but dated hypotheses awaiting investor facts and ongoing validation	Verified operating evidence or formal amendment; never convenience or a quiet exception	A ceiling is not a target. A range does not compel exposure. A numerical pass is not transaction approval.
Decade worldview	2026-2036 scenarios, structural forces, capital-cycle mechanisms, favored and disfavored conditional priors, counterthemes, indicators and falsifiers	Dated analytical layer as of 22 August 2026	New primary evidence, observed milestones, failed causal links, regime change, or scheduled worldview review	It cannot prove expected return, select a security, or waive security and portfolio gates.
Living company research	Research classifications, five-pillar assessments, EC, class batteries, reverse DCFs, severities, milestones, counterthemes, and evidence pointers	Provisional and time-sensitive; classifications expire after 12 months without refresh	Current filings, independent corroboration, valuation, legal and tax work, portfolio look-through, and adversarial review	It cannot confer ownership approval. Every named company is currently at admissible weight zero.
Audit, provenance, and current readiness	Amendment ledger, dispute register, decision journal, exception log, evidence cutoff, certification blockers, ownership and approval record	Records what governed, what was known, what remains open, and who had authority at a point in time	Dated audit evidence, explicit resolution, or a new certified review	Provenance is not investment evidence. Documentary finality is not portfolio certification.

The labels used throughout the book reinforce these layers. A **fact** is observed and dated but may still contain measurement error. An **inference** is a conclusion drawn from facts. An **assumption** is a provisional input. A **forecast** concerns an uncertain future. A **classification** organizes research under stated criteria. **Unresolved** marks a live question, missing input, or dispute. A narrative never becomes evidence through repetition, and a model output is only as reliable as its inputs and structure.

The distinction between durable and dated content is especially important for the decade worldview. The policy may expect AI diffusion, compute investment, power constraints, aging, fiscal activism, fragmentation, climate adaptation, or other forces to matter. Those are conditional hypotheses about mechanisms. Capital supply can respond, customer economics can disappoint, regulation can redirect rents, and high returns can attract competition. Even if a force persists, security holders may not capture the value. The ordinary-compounder and low-cost-comparator counter-scenario remains live precisely because dramatic themes can disappoint.

Current readiness and certification dashboard

The text is complete as a debate-integrated policy edition. The plan is not certified for implementation. That distinction should remain visible on every reading. The dashboard below records the current state rather than the desired state.

Domain	Current status	What is established	What is missing or blocking	Consequence now
Policy text and durable constitution	**Document complete; not an investment certificate**	The governance rules, adopted numerical boundaries, non-authorization, evidence cutoff, and unresolved disputes are recorded.	Investor-specific application and approval are outside the text.	The book may guide analysis but authorizes no transaction.

Domain	Current status	What is established	What is missing or blocking	Consequence now
Investor facts and liability closure	**Uncertified - blocking**	The required definitions and sensitivity method exist.	Reconciled A*; dated and probability-weighted PVL*; safety margin; currencies; nominal/real character; taxes; commitments; income; withdrawal path; flexibility; succession and authority.	No current FS band, strategic design, risk budget, or personal conclusion may be asserted.
Funded status and trait capacity	**Uncertified - blocking**	Four conditional FS bands, +/-10% and +/-20% sensitivity, hysteresis, and the second-axis rule are defined.	Actual FS, two-confirmation history for any promotion, trait tolerance, MDB, and income-market correlation.	Use no permissive band by assumption; affected discretion remains closed.
ILF, title, custody, and continuity	**Uncertified - blocking**	Target and hard-floor formulas, Tier A/B/C structure, access principles, and annual transfer test are specified.	Verified D36 or valid annual D24 election; Tax12; commitments; H; title; beneficiary and incapacity access; Tier A second custodian and successful transfer test; professional work for any Tier B/C structure.	Survival is not certified; ILF assets and custody changes are not authorized by this book.
Withdrawal-path and real-grind closure	**Open - blocking**	The diagnostic stress shows one illustrative mix failing immediately and another boundary mix passing only with genuine real-liability matching.	Actual withdrawal sequence, liability index and duration, implementable inflation-linked match, basis, title, tax, and market-integrity evidence.	No diagnostic mix is a recommendation; liability closure remains open.
Actual comparator and strategic portfolio	**Uncertified - blocking**	The comparator burden, top-20-plus-buckets protocol, factor methodology, and four-quarter forced-choice rule exist.	One named implementable comparator; archived top 20; remainder buckets; two independent first-year scoring passes; sensitivity ranges; tax and cost; executable capped/blended study; tracking-error and abandonment analysis.	The diagnostic proxy is not ratified and cannot be treated as the approved portfolio.
Factor and issuer headroom	**Uncertified - blocking for active admission**	Original risk caps govern while comparator feasibility work is incomplete. The diagnostic proxy indicated AI and joint AI-plus-power breaches.	Reproducible actual-comparator scoring, current holdings, direct and embedded exposures, joint stresses, and currency gap report.	Active additions with material AI/power sensitivity remain at zero; any hard factor breach makes Fmult = 0.
Company evidence and valuation	**Uncertified - blocking for every name**	A research map, provisional classifications, five non-compensating pillars, class overlays, severity floors, and falsifiers exist.	Current extracted Tier 1 evidence, independent corroboration, EC audits, resolved evidence conflicts, reverse DCF, tax and cost, loss severity, portfolio role, factor headroom, and approval.	Every named company remains non-buy with current admissible weight zero.
Active-program economics and operator capacity	**Uncertified; paper mode unless proved**	Maximum eight live full underwrites, about 40 recurring governance hours, 15-25 hours per live underwrite, and a prospective 2x fully loaded cost threshold are defined.	A prospective hours budget, valued opportunity cost, evidenced premium and hit rate, data and implementation costs, and separated attribution.	Active capital remains zero below the threshold; reduce the live set prospectively, possibly to zero.
Tax, legal, domicile, and account fit	**Uncertified - blocking where relevant**	Tax is treated as friction, custody tiers are defined, and professional gates are preserved.	Actual basis, statutory rate, wrappers, domicile, citizenship, estate situs, sanctions, title, transfer, reporting, and account facts.	No personalized tax, legal, currency, custody, account, or cross-border conclusion may be drawn.
Implementation authority	**Absent from this document - blocking**	The policy requires separate implementation approval after all applicable gates pass.	Named owner or approver, mandate, account, execution channel, and current approval record. The book intentionally does not invent them.	No order, hedge, transfer, message, account opening, or portfolio mutation is authorized.

The four joint certification blockers that must be printed prominently are: liability closure; the withdrawal-path and real-grind budget; actual-comparator scoring; and company EC audits. Other missing facts - actual FS, trait tolerance, MDB, tax basis, domicile and custody, valuation, portfolio role, and implementation approval - remain independently binding. Completing one blocker does not cure another. A current company valuation cannot repair an unmeasured liability schedule. Verified liabilities cannot create factor headroom. Approval cannot turn unsupported evidence into fact.

Current blockers at a glance

Blocker	Minimum completion evidence	Fail-closed state
Liability closure	Reconciled assets after known frictions; dated liability cash flows; probabilities and safety margin; currency and nominal/real mapping; taxes, commitments, income, withdrawals, and stressed Coverage	No certified FS, sleeve design, or discretionary risk budget
Withdrawal-path and grind budget	Multi-period asset-liability paths including actual withdrawals, inflation, currency, income loss, matching basis, and implementable default-admissible assets	No claim that the plan survives the decade; diagnostic mixes remain illustrative
Actual comparator scoring	Named implementable comparator; archived top 20; sector buckets; reproducible FEW/FSL and joint stresses; second scoring pass; capped/blended alternative; tax, cost, tracking error, and abandonment point	Proxy remains unratified; original caps govern; affected active factor additions remain zero
Company EC audits and current valuation	Extracted and verified required evidence; conflicts resolved; all A-E pillars at least 2; EC at least 60; class batteries; reverse DCF; severity and liquidity; portfolio role; tax/cost and factor headroom	Every company remains research-only or inadmissible, with current admissible weight zero
Title, liquidity, and continuity	ILF target and hard floor; appropriate currency and duration; enforceable title; beneficiary/incapacity access; Tier A second custodian and annual transfer test	Survival uncertified; do not fund risk from the ILF or assume custody diversification
Operator viability	Prospective expected program contribution at least 2x fully loaded cost; explicit annual hours; maintainable comparator and live-underwrite scope	Paper-journal mode with zero active capital, or comparator plus ILF plus verified matching
Separate approval authority	Identified authorized decision maker, mandate, current approval, and execution route after all gates pass	No transaction or external instruction

Authority hierarchy

When formulations differ, readers should use the following hierarchy. This protects final adopted rules from being displaced by earlier debate language or explanatory shorthand.

1. **The controlling concise policy and its final adopted amendments.** The concise policy is the governing source. Later final amendments control over earlier language they explicitly correct. In particular, the 0.50 floor applies only to the Evidence Coverage multiplier; liquidity and factor penalties retain full force; a hard factor breach produces zero. Anti-gaming severity floors apply. The concentration question may be reconsidered only under the final five-year, at-least-20-decision audited reopening conditions, not an earlier superseded formulation.
2. **The durable rules in this revised structured edition.** This edition reorganizes and explains the controlling policy. It may improve clarity and navigation, but it does not invent authority, resolve an open dispute silently, or relax a numerical boundary.
3. **Current verified investor facts and approved implementation records.** Facts apply the policy. They can determine FS, show that a gate passes or fails, or support a formal amendment. They cannot silently rewrite the policy. A separate approval record is required for any implementation.
4. **Current research files, worksheets, professional reviews, and company evidence.** These supply inputs for comparator, security, tax, legal, and custody decisions. A research file can support or defeat admission. It does not itself confer execution authority.
5. **The expanded explanatory chapters.** These define terms, show interactions, and provide hypothetical examples. They are subordinate to the concise controlling rules if a conflict remains.
6. **Debate traceability, minority views, and provenance annexes.** These explain why the rule exists, preserve rejected or unresolved views, and record model or run history. They are audit material, not an alternative menu of governing rules and not investment evidence.

Authority also separates roles. The policy owns the rules. Evidence owners maintain the factual record. The operator maintains only the assigned governance process. An independent challenger tests specified assumptions and decisions. Tax and legal professionals establish relevant facts and constraints. None of those functions automatically holds trading authority. The actual investor or other authorized approver and the execution channel must be named outside this document. Until then, authority is absent rather than implied.

Three conclusions that must remain separate

The book repeatedly distinguishes three questions that are often collapsed into one.

Business-quality conclusion: Does the business appear to have durable demand, a rational structure and funded payer, attractive capital economics, and credible governance and per-share capture? This is assessed through Pillars A through D and any class-specific battery. A favorable answer is a research result, not a security-return conclusion.

Security-return conclusion: At the current price, do conservative after-tax owner returns clear the named comparator after appropriate cost and uncertainty? This is Pillar E, supported by reverse DCF, valuation sensitivity, and explicit falsifiers. A strong business can fail this question.

Portfolio-role conclusion: Even if business and price pass, does the security fit the actual investor's liabilities, liquidity, currency, tax, loss budget, drawdown budget, factor headroom, tracking error, and operator capacity? A security can pass underwriting and still have no admissible role because the comparator already embeds the exposure or the investor lacks capacity.

Only after all three conclusions are current and favorable, all remaining gates pass, and separate approval is supplied could implementation even be considered. This book completes none of those chains for a named company.

The concise operating map

The operating sequence is intentionally front to back. It prevents an attractive security from becoming the starting point and the investor's objectives from becoming an afterthought.

1. Establish annual investor facts

Reconcile A*, the assets actually investable after known tax and transaction friction. Build PVL*, the dated, probability-weighted present value of liabilities plus explicit safety margin. Record currencies, nominal or real character, taxes, fixed commitments, dependable income, income-market correlation, spending flexibility, decision makers, incapacity arrangements, succession, domicile, and custody facts.

Calculate FS at base PVL* and at +/-10% and +/-20%. If the +/-10% range changes the band, use the lower state. Promotion requires two consecutive annual confirmations. Death, incapacity, income loss, new leverage, litigation, custody impairment, or newly inflexible spending can cause immediate demotion. Missing or ambiguous facts never grant the more permissive state.

2. Protect survival and continuity

Calculate the ILF target:

ILF target = D36 + Tax12 + Commit + H x max(0, essential spending beyond D36 - dependable income covering that later interval)

The hard floor is six months of uncovered essential spending plus twelve months of fixed commitments plus known taxes. H covers only essential spending beyond D36 so the same spending is not counted twice. D24 may replace D36 only by annual documented election when FS is at least 1.75 and spending flexibility is genuinely high. ILF

assets require liability currency, short duration, title and custody integrity, and practical access. Their real drag is reported as insurance cost. It is not repaired by taking risk inside the floor.

Verify beneficiary and incapacity access, Tier A same-jurisdiction second custody, and an annual transfer test. Tier B plain instruments in a second legal system require the specified tax, reporting, title, transfer, sanctions, and estate-situs review. Tier C entities, trusts, and complex wrappers default to zero until the full professional gauntlet passes.

3. Define the strategic baseline and risk budgets

Name one low-cost implementable comparator. Separate nominal duration matched to nominal liabilities from inflation-linked duration matched to verified real liabilities. Treat matched duration separately from directional rates exposure and report its basis risk. Record strategic and opportunity sleeve ranges, the permanent-loss error budget, the mark-to-market drawdown budget, the speculative-loss ceiling, the illiquid ceiling, and active tracking-error limit. These are ceilings and ranges, not targets. Unused opportunity capacity remains strategic; it is not a requirement to find active ideas.

At equal FS, shorten MDB and EB by one state when essential liabilities are near or inflexible or when income is highly correlated with market losses. A high FS does not enlarge the security-level severity or factor budget. The lowest applicable constraint binds.

4. Look through the comparator before active ideas

Archive and score the comparator's top 20 constituents individually. Bucket the remainder through defensible sector proxies. Aggregate direct and embedded issuers, sovereigns, currencies, factors, financing chains, and counterparties. Calculate Factor Equivalent Weight, Factor Stress Loss, and joint stresses. Refresh annually and on material events; repeat independently after six months in year one. Missing look-through is unknown and constraining, never zero exposure.

If an embedded hard-cap breach persists for four consecutive quarters, make a written choice: adopt an implementable capped or blended comparator or raise the cap through formal amendment. There is no renewable-exception path. A raised cap accepts and discloses risk; it does not diversify it. Until actual-comparator work is complete, original caps govern and affected active additions remain zero.

5. Admit active research only through non-compensating gates

For a company, assess Pillar A demand and lock-in; B structure, cycle, and payer quality; C capital economics; D governance and per-share capture; and E valuation and falsifiability. Every pillar must score at least 2. EC must be at least 60, with verified fields weighted toward filings, contracts, and regulators. Conflicting material evidence scores zero until resolved. Principal financial businesses use their class batteries; other classes use their specified overlays.

Only after admission gates pass calculate:

AllowedLoss = BaseIssuerLoss(FS) x max(0.50, ECmult) x Lmult x Fmult

and

MaxWeight = AllowedLoss / S

The floor applies only to ECmult. Liquidity penalties retain full force. A hard factor breach sets Fmult and therefore allowed loss to zero. Bear-case permanent-impairment severity S is at least 40% for ordinary single equity, at least 60% for a levered financial, cyclical, or roll-up, and 80-100% for binary-license, pre-cash-flow, or flagged emerging-market exposure. A lower severity requires written evidence and independent challenge. The result is a diagnostic ceiling, not a recommendation or authority.

Budget active derating loss at purchase using the current and conservative valuation cases. Aggregate derating loss should fit within 50% of MDB after strategic-sleeve stress; up to 75% requires written amendment and independent challenge. Tax, cost, and the comparator replacement return remain part of admission and hold decisions.

6. Keep speculative options inside the opportunity and loss budgets

Speculative options are assumed to lose 100% of their market value. They are normally 0.25-0.75% per name and remain inside the applicable speculative sleeve, EB, opportunity sleeve, and all other gates. There is no averaging down based on price alone. A two-stage harvest candidate expires after 24 months unless preregistered exit and shareholder-cash evidence arrives; the status does not roll because the story remains appealing.

7. Freeze the decision record before approval

Record evidence timestamps and tiers, fact and inference labels, the causal thesis, strongest countertheme, measurable falsifiers, tax lots, loss severity, factor role, action clock, any event-anchored pacing, and the words "non-buy until approved." If entry is ever authorized elsewhere, tranches must attach to filings or preregistered milestones and end by day 180. Where all gates and the conservative case clear and no information event is pending, full deployment is the statistical prior; behavioral pacing must record its expected opportunity cost. Neither route is authorized here.

8. Use narrow monitoring and explicit clocks

Monthly monitoring covers ILF, custody and title alerts, and hard-gate events. Quarterly work reconciles holdings and look-through, issuer and factor headroom, evidence capable of breaking a pillar, tax lots, and the exception log. Do not fully rescore every company because a price moved.

Title, fraud, de-licensing, and other hard-gate threats freeze affected discretion immediately and enter a 0-5-trading-day resolution window. Pillar 0, a hard-cap issue, two missed milestones, or EC below 60 requires a full re-underwrite within 30 days and no add. Valuation-only deterioration causes no mechanical trim, but a failed valuation admission forbids an add and requires the standardized after-tax comparison. Strategic sleeves may rebalance when more than 5 percentage points or more than 20% relative outside the approved band, subject to the ILF, taxes, and higher gates. A broken theme is not mechanically refilled.

9. Audit governance and active-program viability annually

Expected active-program contribution is:

portfolio value x active sleeve x evidenced premium x evidenced hit rate

It must be at least 2 times fully loaded annual cost before active capital is used. Cost includes research hours at the operator's opportunity cost, data, tax, and implementation. The recurring system is estimated at about 40 governance hours plus 15-25 hours per live full underwrite. Eight live underwrites imply about 160-240 total hours before one-off comparator and custody work. Eight is an absolute operating maximum, not a quota. Reduce the live set before work begins if the time is unavailable.

Below the prospective 2x threshold, use paper-journal mode with zero active capital. Two consecutive negative deployed viability years suspend new active work until re-justified. Attribute value from comparator and liability-matching design separately from value from company selection. Audit ILF drag, exceptions, tax deferrals and challengers, diversifier start-date sensitivity, map expiry, counter-canon and non-US screens, and amendments.

10. Certify only when every joint blocker is complete

Certification requires the withdrawal-path and grind budget, liability closure, actual-comparator scoring, and company EC audits, plus all other applicable investor, valuation, legal, tax, custody, portfolio, governance, and authority conditions. A complete process packet may be handed to a separate approver. It is not self-executing. Failure of maintenance or reproducibility contracts the system to the named comparator plus ILF plus verified nominal and real liability matching. This is a governance default, not a transaction instruction.

How to read the rest of the book

Part II, Durable Portfolio Policy, should be read as the constitution and portfolio architecture. It explains objectives, funded status, liquidity, comparator design, the three risk budgets, strategic sleeves, look-through, loss arithmetic, taxes, leverage, custody, and amendment rules. These chapters define the durable sequence. Their examples illustrate mechanisms but do not fill missing investor facts.

Part III, Decade Investment Worldview, is the dated analytical layer. Read it as a set of causal hypotheses, countertheses, capital-cycle questions, value-capture chains, indicators, milestones, and falsifiers. It distinguishes structural forces from cyclical conditions and forecasts. Do not read its favored or disfavored priors as forecasts of security returns or portfolio instructions. The S6 plateau or muddle-through scenario is a first-class alternative, not a footnote.

Part IV, Security Research and Operating System, explains how a business-quality prior becomes, or fails to become, a security-return conclusion and then a portfolio-role conclusion. It contains the five-pillar scorecard, class overlays, evidence coverage, reverse DCF, severity, factor headroom, the company research queue, decision records, monitoring, exits, and research sunset. Research classifications remain provisional and dated. The live operating scope is no more than eight full underwrites.

Part V, Governance, Risks and Certification, explains roles, cadence, exceptions, amendments, unresolved disputes, bias controls, operator capacity, readiness, evidence cutoff, and approval boundaries. Read its certification language literally. The policy can be final as a document while every implementation remains unauthorized.

The annexes preserve detailed company maps, operating checklists, adjudication, bias review, traceability, and provenance. They support audit and deeper work. They do not control over the final concise policy or provide an alternate route to action. Debate history is available there for readers who need rationale; it is not required to understand the main policy path.

For any material statement, ask four questions. What layer is this - durable rule, provisional calibration, dated worldview, living research, or audit record? What epistemic label applies - fact, inference, assumption, forecast, classification, or unresolved? Which higher-level constraint governs it? What evidence and authority would be required before it could affect a portfolio? Those questions prevent the book's breadth from being mistaken for implementation readiness.

What the reader should know before continuing

The policy puts survival, legal title, essential liabilities, and continuity before return; tax-aware real compounding before optional aspiration; and verified evidence before conviction. It measures success through funded status, stressed Coverage, liquidity, real after-tax progress, three separate risk budgets, comparator-relative value, and maintainable governance. It expects the comparator to bear scrutiny and active complexity to earn its full cost. It treats themes as questions, company classifications as research states, and ceilings as boundaries.

The durable constitution is controlling until amended. The numerical architecture is calibrated and provisional. The 2026-2036 worldview is dated as of 22 August 2026. Company research is living and expires without refresh. Provenance explains the record but is not investment evidence.

Current readiness is fail-closed. Liability closure, the withdrawal-path grind budget, actual-comparator scoring, company EC audits, actual FS, trait tolerance, MDB, tax basis, domicile and custody, factor headroom, current valuation, operator viability, and implementation authority remain incomplete or unverified. Every named company has current admissible weight zero. The diagnostic comparator is not ratified. Material AI/power-sensitive active additions remain zero while original factor caps govern. Active capital remains in paper mode unless the prospective 2x fully loaded cost threshold is proved.

The practical reading rule is therefore: learn the hierarchy first, then verify the investor, protect survival, design and look through the strategic baseline, and only then examine active research. At every stage, missing evidence narrows scope; it never grants permission. No transaction is authorized by this book.

Part II - Durable Portfolio Policy

Chapter 2 - The Investment Constitution

An investment constitution is not a market forecast, a list of favored securities, or a plan for maximizing the next reported return. It is the durable order in which decisions must be made when evidence is incomplete, objectives compete, and market conditions change. The policy horizon is 2026-2036, but the principles in this chapter are meant to outlive any particular view of inflation, technology, interest rates, or market leadership. They govern how facts are classified, how risks are separated, how liabilities are protected, how active judgment earns permission to matter, and how exceptions are prevented from becoming a second, unwritten policy.

The constitution also has a strict boundary. It is a policy and research map, not a portfolio or recommendation. Nothing in this chapter authorizes a trade, order, hedge, transfer, external message, account opening, or portfolio change. A hypothetical example explains a rule; it is not a forecast, target allocation, entry price, or live instruction. A theme can identify a causal question, and a company classification can organize diligence, but neither creates an investable conclusion. Under the governing policy, every named company's current admissible weight remains zero until verified investor facts, a complete current evidence and valuation review, portfolio look-through, any relevant tax and legal review, and every applicable gate have all passed through a separate approval process.

The constitution begins with an order of purposes

The first durable principle is a purpose hierarchy. Survival, legal title, near-term essential liabilities, and incapacity continuity precede nominal return. Tax-aware real compounding follows. Optional aspiration comes last.

Each term matters. Survival means that plausible adverse paths must not destroy the investor's ability to meet essential needs or continue holding assets without forced sale. Legal title means more than seeing a balance on a screen: the owner must have an enforceable claim, appropriate custody, and practical access under the relevant legal system. Near-term essential liabilities are cash needs that cannot simply be postponed because markets are down. Incapacity continuity means the plan must still function if the usual decision maker is ill, absent, or unable to act. Nominal return is the percentage gain measured in currency units before considering inflation and, depending on the presentation, taxes and costs. Real compounding is growth in purchasing power, and tax-aware real compounding recognizes that fees, taxes, withholding, transaction friction, and the investor's own liability inflation determine how much economic progress actually remains.

The ordering prevents a lower goal from consuming the resources required by a higher one. An attractive expected return cannot justify placing essential spending at the mercy of a forced sale. A tax benefit cannot justify defective title. A thematic aspiration cannot make incapacity arrangements unnecessary. Optional aspiration includes goals that are genuinely valuable but not essential, such as additional legacy, philanthropy, prestige, thematic expression, or intellectual excitement. Calling them last does not call them worthless. It means they are funded from capacity that remains after the higher claims are protected.

This hierarchy is not the same as maximizing cash or minimizing all price movement. That would quietly substitute comfort for the actual goal. The constitution therefore pairs survival with Shortfall Parity, discussed below. Survival comes first because ruin and forced sale can make recovery impossible; it does not receive an unlimited license to hold so little productive risk that modeled liabilities fail anyway.

A clearly labeled hypothetical illustrates the distinction. Suppose an investor has an essential payment due next year and a separate aspiration to fund a discretionary project in eight years. A volatile asset might offer a plausible high return over eight years, but it is unsuitable funding for next year's unavoidable payment if a drawdown could force liquidation. The same asset might still be researchable for genuinely long-duration surplus capital. The hierarchy changes the role of the asset; it does not issue a verdict on the asset itself.

The most common failure mode is hierarchy inversion. It occurs when nominal performance, tax avoidance, benchmark envy, or enthusiasm for a story is allowed to endanger title, liquidity, or essential funding. A quieter inversion occurs when the investor treats being able to operate the plan as an afterthought. An elaborate process

that fails during incapacity, or that only one person understands, is a survival defect even if its securities are individually sound.

Material claims must carry epistemic labels

Every material assertion must be tagged as a fact, inference, assumption, forecast, classification, or unresolved. This is the constitution's epistemic grammar: a discipline for stating what kind of claim is being made before using it in a decision.

A fact is an observed or logically necessary proposition, dated and sourced when material. It can still contain measurement error. An inference is a conclusion drawn from facts through an explicit chain of reasoning; it should not be promoted to fact merely because the reasoning is persuasive. An assumption is a provisional input adopted because analysis cannot proceed without it. It needs review when better information arrives. A forecast is a claim about an uncertain future and should be expressed through ranges, scenarios, or probabilities rather than as certainty. A classification places an item in a research category under stated criteria; it is an organizational result, not an investment approval. Unresolved identifies a live dispute, missing fact, or question for which the record does not support closure.

The rule applies to material assertions rather than requiring decorative tags on every sentence. Its purpose is to stop category errors. A management statement can be a fact about what management said while the substance remains unverified. An official institution's projection is still a forecast. A model output can be a fact about what the model produced while its inputs and interpretation remain assumptions. A company can be classified as a research candidate without being admissible for ownership.

A narrative never becomes evidence by repetition. Repeating that a market is large, a management team is excellent, or a technology is inevitable does not establish shareholder capture, valuation, legal title, or survivability. Repetition can increase familiarity and subjective confidence without increasing the quality of the underlying evidence. The labeling rule makes that separation visible.

Hypothetical example: "Demand rose 12% last year" may be a fact if it comes from verified data. "The increase shows permanent pricing power" is an inference. "Demand will rise 10% annually for five years" is a forecast. "The company is a Core Research Candidate" is a classification. "Whether new entrants will erase the margin" may remain unresolved. Writing all five as variations of "the business is strong" would destroy information the policy needs.

The failure modes are predictable: forecasts smuggled in as facts, assumptions left undated until they harden into custom, classifications read as recommendations, and uncertainty hidden behind the authority of a source. The remedy is not false precision. It is an auditable record in which contrary evidence can attack the right link.

Risk is governed through three separate budgets

Permanent impairment loss, mark-to-market drawdown, and decade-long real grind are distinct budgets. Passing one does not pass another.

Permanent impairment is a durable loss in the value of the owner's claim, caused by such mechanisms as destroyed earning power, insolvency, dilution, fraud, loss of title, confiscation, or a purchase price that cannot be recovered through plausible owner cash within the relevant horizon. The permanent-loss budget, also called an error budget in later operating sections, asks how much such irreversible loss the total plan can survive. It is not the same as the largest temporary price decline seen on a chart.

Mark-to-market drawdown is the peak-to-trough decline in quoted value. A sound claim can suffer a large drawdown because discount rates, liquidity, or sentiment change while its operating economics remain intact. That decline matters because the owner may be unable to tolerate it, liabilities may force sale, collateral terms may turn it into a permanent loss, or the original purchase may have been too exposed to valuation derating. But a

drawdown alone is not proof that the thesis failed. Treating every price decline as permanent impairment can force sales of intact claims; treating every decline as temporary can protect broken theses from re-underwriting.

The grind budget addresses a different failure: years of weak real after-tax returns interacting with withdrawals, inflation, currency mismatch, and liability revaluation. A portfolio can avoid dramatic crashes and still fail slowly if nominal returns do not keep pace with the investor's actual cost basket. A ten-year average return does not reveal whether withdrawals arrived early in the bad years or whether liabilities inflated faster than the reported benchmark.

Hypothetical example: a 4% portfolio position falls 40% solely because its valuation multiple contracts while verified cash economics remain intact. The immediate mark-to-market effect is 1.6% of the total portfolio. That result consumes drawdown capacity and may reveal that the purchase was oversized for valuation risk. It does not automatically establish a 1.6% permanent impairment, and it says nothing by itself about whether a decade of withdrawals remains funded. Conversely, a bond held for a dated liability may show interim price volatility but still match the payment if title, credit, currency, and maturity remain sound. A long real-return grind can nevertheless threaten other liabilities even without a default.

The budgets coexist because each answers a different question: Can the claim be permanently damaged? Can the owner endure and finance the path? Can the full liability schedule survive years of poor real outcomes? A position or portfolio must pass all applicable tests. No strong showing in one budget offsets failure in another.

Measurement starts with liabilities, not the return league table

The constitution requires funded status, Coverage, real after-tax returns, currency mismatch, and liquidity to be reported before portfolio return. The central funded-status formula is:

FS = A* / PVL*

A* means investable assets after known taxes and transaction costs. It excludes value that is inaccessible or cannot honestly be used for the modeled obligations. PVL* means the probability-weighted present value of liabilities plus an explicit safety margin. Present value translates payments at different dates into a common current measure; probability weighting recognizes that some liabilities are contingent; the safety margin recognizes model and estimation error. Both sides therefore require verified investor facts rather than a generic asset allocation.

Coverage is the corresponding ability of stressed assets to meet stressed liabilities along a scenario path. Funded status is a starting balance-sheet ratio; Coverage tests what happens when asset returns, inflation, income, withdrawals, currency, and liability values change. Real after-tax return measures purchasing-power progress after relevant tax and friction. Currency mismatch asks whether assets and obligations respond to different currencies. Liquidity asks whether cash can be obtained when needed without relying on an assumed sale at an unstressed mark.

Hypothetical example, not a personal calculation: if eligible assets after known tax and exit costs are 1.50 million and probability-weighted liabilities plus safety margin are 1.10 million, then FS = 1.50 / 1.10, or about 1.36. That means 1.36 units of eligible assets per unit of modeled liability at the starting measurement. It does not certify the plan. A currency shock, underestimated liability, inaccessible account, income loss, or adverse withdrawal sequence can reduce Coverage even though the initial ratio exceeds 1.

This ordering changes what counts as success. A portfolio may beat a market index and still fail if liabilities grow faster, tax leakage is large, or assets cannot be accessed on time. It may lag a rising benchmark while successfully funding essential objectives. Neither conclusion should be used to ignore opportunity cost, which is why comparator and Shortfall Parity rules remain necessary.

The failure mode is return-first reporting: celebrating nominal or benchmark-relative gains before asking whether the owner is more securely funded. Another is false precision in A* or PVL*, such as counting illiquid reported

values at par, omitting tax and transaction costs, or treating uncertain liabilities as zero. When inputs remain unverified, the conclusion must remain conditional.

The comparator bears risk and sets the burden of proof

The policy requires a named, low-cost, implementable comparator. It is an active position with look-through risks, not a neutral default exempt from analysis. Any active complexity must beat it after tax, cost, hours, and error.

"Named" means the alternative must be concrete enough to implement and measure. "Low-cost" recognizes that fees, turnover, tax, and operational friction compound with greater certainty than most claims of excess return. "Implementable" rules out an academic index or historical strategy the investor cannot lawfully, practically, or tax-efficiently hold. The comparator supplies the opportunity-cost baseline: if the active program is not better than the simple alternative on the investor's actual terms, complexity has not earned a role.

Calling the comparator an active position is crucial. A broad index may contain concentrated issuer, country, currency, duration, sector, or factor exposure. A liability-matched bond sleeve contains credit, custody, currency, basis, and inflation risks. These risks do not disappear because the holdings are passive or popular. The comparator therefore enters the same look-through process as active holdings.

The principle holds two dangers in tension. Treating the comparator as neutral can hide large embedded bets and allow them to bypass limits. Constantly redesigning the comparator can become an elaborate valuation-timing strategy whose tracking error, taxes, and abandonment risk destroy the simplicity it was meant to provide. The durable rule does not settle every comparator-design dispute. It demands that the actual comparator be named, inspected, and used as the burden-of-proof baseline.

The phrase "after tax, cost, hours, and error" makes the comparison fully loaded. Hours have an opportunity cost. Data, legal review, custody complexity, and implementation consume resources. Error includes model error, execution error, missed monitoring, behavioral mistakes, and the risk that a complicated process will not be maintained. Zero active exposure remains a legitimate outcome if the active program cannot establish net value.

Shortfall Parity prevents safety from becoming silent failure

Excess caution that creates a modeled liability shortfall is governed as seriously as excess risk. This is Shortfall Parity. It corrects the natural asymmetry of governance systems: breaches caused by action are visible, while years of underinvestment can look calm even as real goals become unreachable.

Parity means a modeled shortfall receives a documented review comparable in seriousness to a risk-limit breach. It does not mean the remedies are identical, and it never licenses risk beyond verified capacity or trait tolerance. Capacity is the financial ability to bear loss while meeting obligations. Trait tolerance is the investor's stable, calm-state ability to adhere through losses and time underwater, rather than the fear or exuberance of the current market state.

The remedy sequence matters because Shortfall Parity can otherwise become a disguised instruction to buy risk after prices have risen and projected foregone returns look painful. A shortfall should first force explicit consideration of contributions, spending or timing flexibility, and objectives. Only then may admissible risk within existing capacity be considered. The constitution does not permit return anxiety to enlarge the admissible region.

Hypothetical example: a plan using conservative assumptions shows that essential and chosen liabilities cannot both be funded. Shortfall Parity does not say "buy more equities." It says the failure may not be ignored. Higher contributions, delayed discretionary spending, reduced aspirations, or a different admissible risk mix must be compared openly. If the investor lacks capacity for additional risk, the honest conclusion may be an objective change rather than a more aggressive portfolio.

The principle coexists with the purpose hierarchy because each checks the other's characteristic error. Survival prevents reach-for-return behavior from endangering essential claims. Shortfall Parity prevents survival language from ratifying an unbounded reserve, chronic cash drag, or refusal to hold productive risk. Neither can overrule title, liquidity, incapacity continuity, capacity, or trait tolerance.

Evidence governs tolerable loss; valuation governs admission

Evidence controls tolerable loss. Valuation controls admission and the return hurdle. Neither story quality nor confidence directly grants weight.

Evidence includes the reliability, completeness, relevance, and independence of what is known about the claim. Better evidence can justify using more of a pre-existing loss budget because uncertainty has been reduced. It cannot make an adverse outcome harmless, erase a severity floor, waive liquidity and factor constraints, or turn a failed gate into a pass. Weak evidence therefore leads to smaller tolerable loss or research-only status, not merely a lower adjective in the prose.

The core quantity is loss exposure, not narrative conviction. At the simplest level, plausible permanent loss is position weight multiplied by bear-case permanent-impairment severity. This distinction allows the policy to think coherently about bounded optionality: a low-evidence proposition with genuinely capped downside can be evaluated as a small total-loss risk without pretending that its evidence is strong. That does not authorize such a position. It merely avoids the error of treating notional exposure, market value, and maximum tolerable loss as interchangeable.

Valuation has a different job. It asks whether the rights and expected owner cash available at the current price offer a sufficient real after-tax return relative to the comparator, given uncertainty and capital needs. A wonderful business can fail the valuation gate. A large price decline can improve prospective return without repairing governance, title, or business economics. Because valuation governs admission and the hurdle, it must not be used as story-based permission to inflate tolerable loss.

Hypothetical example: two claims have the same market weight. One has a plausible 40% permanent impairment and strong verified evidence; the other has an assumed total-loss outcome and weak evidence. Their loss exposures are not equal. Conversely, if the strongly evidenced claim is priced so aggressively that its conservative expected return does not clear the comparator, the evidence does not authorize admission. Evidence answers "how much loss can be tolerated if all gates pass"; valuation answers "does the claim clear the price-based gate at all."

The common failure modes are conviction sizing, double-counting the same uncertainty in multiple places, and allowing a spreadsheet's apparent precision to dominate weak source evidence. A fluent thesis is not a risk budget. Confidence is not a position weight. The constitution requires uncertainty to show up in admissible loss, gates, diversification, and review rules.

Recourse leverage is prohibited by default

The base rule is no recourse leverage by default. Recourse leverage is borrowing or a contingent structure that can reach the investor's other assets, create a collateral call, force liquidation, or impose a payment deadline that converts a temporary price path into permanent loss. Debt changes the nature of risk because time and counterparty rights can become more important than the eventual correctness of the thesis.

The default is not defeated by labels. Mortgages remain balance-sheet liabilities even when they are not treated as borrowing undertaken to increase market exposure. Their service, currency, maturity, and effect on funded status and liquidity still matter. Paid options and collateralized structures retain analysis of notional exposure, path, expiry, basis, and repeat-premium cost. Paying the premium up front may cap one contract's direct loss, but it does not make the economic exposure, concentration, expiry mismatch, or a recurring series of premiums disappear.

Collateralization can reduce some counterparty or recourse channels while leaving basis risk, liquidity requirements, and path dependence.

Notional is the reference exposure on which gains and losses respond. Path risk means the sequence of prices or cash demands matters, not only the endpoint. Expiry risk is the possibility that a thesis is directionally right but realizes after the instrument ends. Basis risk is the gap between the instrument used and the exposure it is intended to hedge or express. Repeat-premium analysis asks what happens if supposedly small defined losses recur year after year.

Hypothetical example: a one-year paid option has a maximum direct premium loss of 0.3% of portfolio value. Describing it only as a 0.3% position would omit its notional sensitivity, the chance it expires before the anticipated event, the mismatch between its payoff and the actual liability, and the cost of repurchasing it annually. The example is not permission to use the option. It shows why defined loss is necessary but not sufficient analysis.

The failure mode is leverage by euphemism: treating a mortgage as economically irrelevant, a securities-backed line as liquidity rather than forced-sale risk, or a derivative premium as the whole exposure. The constitutional default remains prohibition unless a separately governed exception elsewhere proves it is bounded, non-forcing, compatible with the liquidity floor, and superior to an unlevered alternative. This chapter itself grants no exception.

Every exposure must be viewed on a look-through basis

The policy aggregates issuers, sovereigns, factors, currencies, financing chains, counterparties, and the comparator. Look-through means tracing economic dependence through wrappers, funds, derivatives, corporate structures, and account labels until the underlying common risks are visible.

Issuer aggregation prevents the same company from appearing diversified merely because it is held through several funds or share classes. Sovereign aggregation captures law, sanctions, taxation, capital controls, state payer exposure, and legal-system dependence. Factor aggregation identifies positions that may respond to the same underlying driver, such as duration, credit, artificial-intelligence capital spending, power constraints, or one commodity cycle. Currency aggregation follows revenue, financing, listing, and liability currencies rather than relying on the trading currency alone. Financing-chain analysis asks who supplies capital, collateral, guarantees, leases, and refinancing. Counterparty analysis includes custodians, brokers, banks, clearing systems, and contractual payers. The comparator must be included because passive exposure is still exposure.

Hypothetical example: an investor owns a global index fund, a technology fund, and two individual companies. The account statement shows four holdings. Look-through analysis may show that the same large issuer appears in both funds and directly, while all four depend on the same currency and long-duration growth factor. The number of line items increased; economic diversification may not have.

This principle interacts with legal title and incapacity continuity. Multiple accounts at one legal custodian are not necessarily custody diversification. A globally diversified revenue base does not necessarily diversify the law governing the security or the account. A hedge with a different ticker may add the same counterparty or financing-chain risk it was intended to reduce.

The failure mode is label-based diversification: counting products instead of economic claims. Another is missing data treated as zero exposure. Unknown look-through must remain unknown and constraining; it cannot be credited as diversification. Look-through is therefore both measurement and humility made operational.

Broad strategic rules may move faster than active judgments

"Fast broad, slow active" separates precommitted portfolio maintenance from discretionary security judgment. Broad strategic sleeves rebalance mechanically. Single names move only after evidence review, except for enumerated solvency or integrity events.

A broad strategic sleeve is the approved, diversified policy exposure intended to be held across regimes. Mechanical rebalancing follows prewritten bands and does not require a fresh macro forecast each time. Its purpose is to keep the approved architecture from drifting with price. "Fast" refers to executing an already governed rule when its observable condition is met, subject to higher constraints such as liquidity, taxes, title, and Coverage. It does not mean impulsive trading.

Active positions are slower because a falling price can have two very different meanings. It may improve expected return while the owner claim remains intact, or it may be the market's response to deteriorating business economics, financing, governance, or legal rights. Price changes can alter valuation; they do not by themselves repair or destroy the thesis. A single name therefore requires fresh evidence review before an add, reduction, or thesis-dependent response. A broken theme is never mechanically refilled merely to restore its former weight.

Enumerated solvency and integrity events are the narrow exception to slow active review. Examples in the governing framework include title or custody failure, evidenced fraud, de-licensing, sanctions or capital controls affecting access, a legal compliance deadline, or a threat under an already approved financing covenant. These events can require prompt, minimum-scope protection because waiting may destroy access or legal compliance. Even then, the emergency route is risk-monotone and does not authorize buying a dislocation.

The central tension is behavior under stress. If fear could suspend every mechanical rebalance, the policy would become procyclical. If rebalancing automatically restored every collapsed security or theme, the rule would average down into permanent impairment. Broad rules can be fast because diversification and bands were approved in advance. Active changes are slow because the evidence specific to the claim must be renewed.

Tax is friction, not an investment thesis

Tax matters because it compounds negatively and can materially change the value of one action relative to another. It is friction in the economic comparison, not proof that the underlying asset remains sound. The required discipline is to re-underwrite a material holding as if it were unowned.

"As if unowned" removes cost basis, pride, familiarity, and the disposition effect from the first question. At today's price, given today's evidence and the available comparator, would the claim still deserve to be held? Only after that tax-blind economic assessment should the actual basis, statutory rate, deferral horizon, loss offsets, wrappers, estate rules, and replacement return enter the implementation comparison.

Deferral can support hold-not-add. This is a deliberately narrow status. A holding may remain economically preferable to an immediate taxable sale even though the investor would not add new capital at the current price. Deferral can never rescue a failed business case indefinitely. It cannot restore lost title, cure fraud, repair a failed hard gate, or convert inadequate expected return into a permanent thesis.

Hypothetical example: an investor would choose zero new exposure to a holding at today's price, but selling would trigger a material tax and the verified thesis remains intact. A standardized hold-versus-redeploy comparison may conclude that retaining the existing position temporarily has greater after-tax value. That conclusion does not authorize an add, and it must be reviewed rather than repeated forever. If the business case fails, the tax cost is a cost of action, not a reason to pretend the claim is sound.

The failure modes are tax tail wagging the portfolio, tax estimates based on generic rates rather than actual facts, and annual repetition of "not yet" until deferral becomes an undocumented permanent exception. Tax-aware compounding and thesis integrity coexist by separating economic underwriting from implementation friction.

Persistent breaches require amendment, not excuses

The constitution is stable but not frozen. Persistent breaches must become explicit amendments or compliant design changes. They may not be carried indefinitely as temporary exceptions.

An amendment says openly that the governing rule is changing, why it is changing, what evidence supports the change, what contrary case remains, and how the effects on liabilities, liquidity, tax, currency, concentration, behavior, and implementation will be reviewed. A compliant design change instead alters the portfolio or process, through separately authorized action, so the existing rule is once again met. The constitution prefers either route to the quiet normalization of breach.

Emergency exceptions are narrower. They must be risk-monotone, recorded within 24 hours, independently reviewed within seven days, and expire within 30 days. Risk-monotone means the emergency action is limited to reducing or not increasing the relevant risk, preserving liquidity, or restoring compliance. It cannot be used to raise gross exposure, add illiquidity, or buy an apparent bargain under the label of emergency. The written record within 24 hours prevents memory from being rewritten. Independent review within seven days tests necessity, scope, and conflicts. Automatic expiry within 30 days forces any continuing arrangement into the normal amendment process.

The numbers are not suggestions. "Within 24 hours" is the deadline for the record, "within seven days" is the deadline for independent review, and "within 30 days" is the maximum life of the emergency exception. A recurring or structural problem cannot be renewed forever through a chain of 30-day labels.

Hypothetical example: a custodian announces an access restriction affecting an account. A minimum-scope transfer or freeze may qualify if it preserves title or liquidity and does not increase risk. The action must be documented within 24 hours, challenged within seven days, and the exceptional authority ends within 30 days. Using the same event to add a speculative security would fail the risk-monotone requirement.

The main failure mode is exception laundering: treating repeated breach as harmless because each instance has a note. Another is amending limits merely because prices moved or recent performance is uncomfortable. The constitution changes for documented objectives, constraints, durable evidence, structure, capability, or implementation failure - not to avoid admitting an error or to chase a forecast.

Governance must be viable enough to endure

Minimum viable governance is mandatory. It consists of the Investor Liquidity Floor check, the look-through dashboard, the decision journal, the annual funded-status review, the exception log, and active-program viability. Everything else must earn its hours.

The Investor Liquidity Floor, or ILF, is the protected liquidity assigned to near essential needs, known taxes and commitments, and relevant income interruption. The ILF check confirms that this floor remains funded, accessible, appropriately timed and denominated, and not treated as opportunistic buying power. The look-through dashboard aggregates the underlying risks described above rather than counting account labels. The decision journal preserves the evidence type, thesis, valuation, alternatives, downside, disconfirmers, tax effects, and decision logic available at the time. The annual funded-status review refreshes A*, PVL*, Coverage, liabilities, currencies, income correlation, capacity, trait tolerance, succession, and policy fit. The exception log records observable deviations and prevents them from vanishing into memory. Active-program viability asks whether active work produces enough evidenced decision value to justify research hours, data, tax, implementation, and error costs relative to the comparator.

"Everything else must earn its hours" is itself a control. Additional dashboards, meetings, models, roles, data feeds, jurisdictions, and instruments may be useful, but their benefit must exceed the monitoring burden and new ways they can fail. The mandatory minimum may not be discarded in the name of simplicity. Equally, optional scaffolding may not grow until the operator cannot maintain the core controls.

A common failure mode is governance theater: exhaustive paperwork that does not change a decision and crowds out the few checks that prevent ruin. The opposite failure is informal memory, where theses, exceptions, and past assumptions are reconstructed after outcomes. Minimum viable governance aims at a small, durable audit trail whose decision value can itself be reviewed.

Research never authorizes implementation

The final durable principle is a standing separation between research and authority. Themes identify causal questions. Company classifications organize diligence. Only a current full review can create an investable conclusion, and this document performs none.

A theme such as artificial intelligence, electrification, healthcare, or state-directed capital can direct attention to needs, payers, bottlenecks, capital cycles, and possible beneficiaries. It cannot establish which security captures the value, whether the legal claim is sound, whether the current price offers an adequate return, whether the portfolio already owns the same factor through the comparator, or whether the investor has capacity for the loss. Social importance and technological success do not imply shareholder return.

A company classification is similarly limited. "Core Research Candidate," "Watchlist/Price Sensitive," "Speculative Option," or "Avoid/Counterexample" describes a research state under a dated framework. It is not a buy list, target weight, or instruction. A favorable classification can expire, evidence can conflict, valuation can fail, and the actual investor's liabilities, tax, custody, currency, and factor headroom can make an otherwise attractive claim inadmissible.

A current full review means more than rereading an old memo. It requires current evidence, valuation, disconfirmation, legal and economic rights, permanent-loss severity, liquidity, factor headroom, tax and cost where relevant, and portfolio role, followed by the separate approval process. Research output may recommend further work or abstention. It may not execute an order, send an external instruction, mutate a portfolio, or represent a diagnostic ceiling as a target.

This boundary protects all the preceding principles. Without it, an engaging narrative could bypass epistemic labels; a company map could bypass the comparator; a sizing calculation could be mistaken for authorization; or an emergency discussion could become an execution channel. The separation is categorical: this chapter explains governance and authorizes nothing.

Why the principles must coexist

The constitution is a system of checks, not fourteen independent slogans. The purpose hierarchy gives survival priority, while Shortfall Parity makes chronic under-risking visible. Liability-first measurement defines success in the investor's terms, while the comparator preserves opportunity-cost discipline. The comparator is simple enough to be the burden-of-proof baseline, while look-through prevents simplicity from masquerading as neutrality.

Evidence-to-loss and valuation admission separate two questions that narratives often blend: how much damage can be tolerated, and whether the current price clears the hurdle at all. The three budgets then prevent permanent loss, temporary drawdown, and slow real shortfall from being compressed into one volatility number. No-recourse-leverage protects the ability to hold, while fast broad rebalancing preserves disciplined action inside already approved capacity. Slow active review prevents the same mechanical rule from averaging down into a broken issuer.

Tax-aware measurement recognizes genuine friction, while tax-blind re-underwriting prevents friction from becoming thesis. Policy stability protects calm-period commitments, while the amendment rule permits honest change and the narrow emergency exception preserves title, liquidity, and compliance during discontinuity. Minimum viable governance ensures the system can actually be maintained, and research non-authorization ensures that analysis does not quietly seize decision or execution authority.

These pairings also identify what evidence would change a conclusion. Verified liabilities can change funded status and the role of risk. Better source evidence can change tolerable loss but not waive valuation. A different price can change the admission hurdle but not business quality or title. A persistent comparator breach can require an explicit design change or amendment, not denial that the exposure exists. A genuine life, health, legal, custody, or income change can alter capacity and continuity. Repeated process failure can justify simplifying governance.

None of these observations is a trade instruction; each identifies the proper constitutional channel for reconsideration.

Chapter 3 - Investor Facts, Funded State, and Liquidity

This chapter turns the policy's liability-first principles into an operating system. Its first question is not how much the portfolio earned, whether a theme is attractive, or how confident the investor feels. It is whether investable assets, after known frictions, are sufficient for the present value of the liabilities the assets must support. That relationship determines how much room exists for strategic market exposure, active decisions, illiquidity, tracking error, and mistakes.

The system is deliberately conditional. The policy does not yet contain verified investor assets, liabilities, taxes, income, domicile, custody arrangements, spending flexibility, or behavioral capacity. Therefore none of the four funded-state bands is a current personal classification, and none of the percentages in this chapter is a target allocation or authority to transact. They are dated policy boundaries that become usable only after the underlying facts are verified. Where a required fact is missing or ambiguous, the architecture fails closed: use the lower funded state, freeze the affected discretion, or leave the relevant exposure at zero.

Funded status is the selector, not a return forecast

The central ratio is:

FS = A* / PVL*

FS means funded status. A* means investable assets after known taxes and transaction costs. The asterisk matters. Gross account value is not the numerator if part of that value would be lost to a known tax charge, exit cost, settlement cost, legal restriction, or other known implementation friction. Assets that cannot actually be accessed for the liabilities in question also should not be treated as freely investable merely because they appear on a statement.

PVL* means the probability-weighted present value of liabilities plus an explicit safety margin. A liability is any future claim the portfolio is expected to fund, not only formal debt. Essential living costs, taxes, health and education costs, insurance gaps, property obligations, fixed commitments, and contingent obligations may all belong in the liability schedule. "Present value" converts future payments into a current amount using assumptions about timing, inflation, currency, and discount rates. "Probability-weighted" recognizes that some claims are contingent rather than certain. The explicit safety margin recognizes that a neat central estimate is not enough when omitted expenses, estimation error, or adverse timing could matter.

An FS of 1.00 means adjusted investable assets equal the stressed present value of liabilities. An FS below 1.00 indicates a modeled shortfall. An FS above 1.00 indicates a modeled surplus, but the amount of surplus matters. The ratio does not say that assets will earn a particular return, that liabilities are perfectly known, or that the investor can tolerate the path required to earn that return. It is a solvency selector, not a forecast.

The ratio also prevents a common category error. A person may say that a 30% portfolio decline is tolerable, but a percentage decline has no useful meaning until it is compared with the liability denominator. A large decline can be survivable for a deeply overfunded investor and ruinous for an investor barely above FS = 1.00. Conversely, an investor below 1.00 is not authorized to take more risk simply because higher returns would be convenient. The remedy sequence remains contributions, spending or timing changes, objective changes, and only then admissible risk within verified capacity. Shortfall parity treats excessive caution seriously when it creates a modeled shortfall, but it never converts a need for return into capacity to bear loss.

Hypothetical worked example - calculating FS and its sensitivity

This example is illustrative only. It is not investor data, a forecast, an allocation, or an instruction.

Assume gross financial assets are 1,500,000. Known taxes associated with making the assets available are 100,000, and known transaction or exit costs are 50,000. Then:

A* = 1,500,000 - 100,000 - 50,000 = 1,350,000

Assume base PVL*, including its explicit safety margin, is 1,000,000. Base funded status is:

FS = 1,350,000 / 1,000,000 = 1.35

On the base estimate, 1.35 falls in the 1.25-1.74 band. The policy does not stop at the base estimate. It requires PVL* to be tested at base, -10%, +10%, -20%, and +20%. Holding A* constant for this diagnostic produces:

PVL* case	PVL*	FS	Apparent band
-20%	800,000	1.6875	1.25-1.74
-10%	900,000	1.50	1.25-1.74
Base	1,000,000	1.35	1.25-1.74
+10%	1,100,000	1.227	1.00-1.24
+20%	1,200,000	1.125	1.00-1.24

The +10% liability case changes the band. The rule therefore assigns the lower state, 1.00-1.24, even though the base calculation is 1.35. This is not a pessimistic forecast. It is an ambiguity rule designed to prevent uncertain liability estimates from granting risk capacity. The -20% and +20% cases still have to be reported because they show the breadth of the model uncertainty, even though the binding band test is whether a +/-10% change moves the classification.

The four funded states

The policy uses four states. Each state defines a range for the strategic sleeve, opportunity sleeve, active permanent-loss error budget, illiquid ceiling, and active tracking-error ceiling.

FS	Strategic sleeve	Opportunity sleeve	Error budget (EB)	Illiquid ceiling	Active tracking-error ceiling
<1.00	95-100%	0-5%	0-3%	0-5%	0-2%
1.00-1.24	90-95%	5-10%	3-5%	0-8%	0-3%
1.25-1.74	80-90%	10-20%	5-8%	5-12%	0-4%
>=1.75	75-90%	10-25%	8-10%; hard 12%	10-15%; hard 20%	0-5%

These are ranges, not targets. The upper end is never automatically earned. Missing evidence, weak behavioral capacity, inflexible liabilities, correlated income, a failed liquidity test, or an unavailable implementation can require the lower end or zero. The strategic and opportunity ranges are paired parts of one portfolio design; their maxima cannot simply be added. A coherent selection must account for all investable capital and preserve the Investor Liquidity Floor described later in this chapter.

State 1 - FS below 1.00

FS <1.00 is the shortfall state. Adjusted assets are less than the stressed present value of liabilities. The strategic sleeve occupies 95-100%, while the opportunity sleeve is limited to 0-5%. The active permanent-loss error budget is 0-3%, illiquid assets are capped at 0-5%, and active tracking error is capped at 0-2%.

The zeroes in these ranges are meaningful. An investor in shortfall may have no admissible opportunity sleeve, no active loss budget, no illiquid allocation, and no active tracking error. The upper bounds do not constitute permission to gamble out of the shortfall. They only define the outer boundary within which a verified capacity and shortfall remedy process might operate. If near liabilities, title, liquidity, or behavioral tolerance do not support risk, the admissible figure remains zero.

The policy favors strategic assets here because complexity and permanent impairment can turn a modeled shortfall into an irreversible one. Strategic does not mean riskless. It means the preapproved, low-cost architecture

built around the named comparator, liability matching, liquidity, and only validated diversifiers. The purpose is to keep the portfolio governed while the shortfall is addressed through real-world levers such as contributions, timing, spending, and objectives.

State 2 - FS from 1.00 to 1.24

FS 1.00-1.24 is funded but marginal. The strategic sleeve is 90-95%, the opportunity sleeve 5-10%, EB 3-5%, the illiquid ceiling 0-8%, and active tracking error 0-3%.

At the bottom of this band there is essentially no cushion between assets and modeled liabilities. At the top there is some surplus, but liability error, taxes, income loss, or a custody problem can consume it quickly. The architecture therefore permits somewhat more active and illiquid capacity than the shortfall state while keeping most capital in the strategic sleeve. The words "funded" and "surplus" should not be confused: 1.01 and 1.24 both fall in this state, yet their margins of safety differ. The selected point within each range must reflect the actual facts rather than the label alone.

State 3 - FS from 1.25 to 1.74

FS 1.25-1.74 is a funded-with-surplus state. The strategic sleeve is 80-90%, the opportunity sleeve 10-20%, EB 5-8%, the illiquid ceiling 5-12%, and active tracking error 0-4%.

The larger surplus allows a wider opportunity sleeve and greater capacity for mistakes, illiquidity, and deviation from the comparator. It does not relax the security-level admission gates. A company that lacks sufficient evidence, fails a business-model pillar, breaches a factor cap, or does not clear the valuation hurdle remains inadmissible regardless of the portfolio's funded status. Funded status supplies a portfolio-level budget; it does not turn an unqualified asset into a qualified one.

The 5-12% illiquid band is also not an endorsement of private assets. The policy's private and illiquid asset dispute remains open, and any specific vehicle still has to pass manager, cash-flow, governance, valuation, access, and liquid-alternative tests. The lower bound in a state table does not compel ownership when no admissible implementation exists.

State 4 - FS at or above 1.75

FS >=1.75 is the strongly funded state. The strategic sleeve remains the majority at 75-90%. The opportunity sleeve may be 10-25%. EB is normally 8-10% with a hard cap of 12%. The illiquid ceiling is normally 10-15% with a hard cap of 20%. Active tracking error is capped at 0-5%.

This state has the broadest ranges, but it is not a license for concentration or complexity. The EB hard cap of 12% and the illiquid hard cap of 20% are ceilings above the normal ranges, not routine settings. Moving beyond the normal range requires the relevant evidence and governance; the hard cap may not be exceeded. More importantly, high funded status does not enlarge issuer loss budgets merely because spending is flexible. Security-level loss arithmetic, severity floors, liquidity penalties, and factor hard caps continue to bind. A hard factor breach still means no active addition.

What each table column controls

The strategic sleeve is the durable policy allocation. It contains the named, implementable comparator; high-quality nominal duration matched to nominal liabilities; inflation-linked duration matched to verified real liabilities when the index, law, and market integrity are credible; the ILF; and only diversifiers that have passed their validation tests. A strategic position is not neutral merely because it is broad. The comparator has issuer, factor, country, currency, and valuation risks that must be looked through and governed.

The opportunity sleeve is the bounded home for fully underwritten active names and any speculative options. It is not a return requirement. If no active idea clears every gate, unused opportunity capacity remains in the strategic sleeve. It is not held as extra "dry powder" outside the ILF, and it is not filled merely to reach the top of the range.

The error budget, or EB, is the maximum portfolio budget for active permanent loss, subject to tighter limits from actual capacity, trait tolerance, currency loss capacity, and stressed liability coverage. Permanent loss means destruction or non-recovery of the investor's economic claim, such as business impairment, fraud, unrecoverable custody loss, or an option assumed to expire worthless. A temporary price decline with the claim intact is different and belongs to the mark-to-market drawdown budget. EB is neither an expected loss nor an amount that should be spent. It is an outer loss capacity that protects the liability plan when research is wrong.

The illiquid ceiling limits assets that cannot be converted into reliable spending power on the required timetable. Illiquidity includes more than a stated lock-up. Gates, delayed redemption, uncertain secondary markets, capital calls, legal title problems, or a vehicle whose reported value cannot be realized can all make an asset operationally illiquid. Because the percentages are ceilings, the actual admissible illiquid amount can be zero in every state.

Active tracking error measures how much the active design is allowed to deviate from the named comparator. It is a relative-risk ceiling, not a promise of outperformance and not a permanent budget that must be used. A comparator modification or active sleeve can look sensible in isolation yet produce an intolerable pattern of relative losses. The ceiling forces that path risk to be acknowledged. If the actual comparator is missing or cannot be scored, the control is not satisfied by substituting an informal benchmark.

The second axis - capacity beyond the FS decimal

Two investors with the same funded status can have very different ability to bear loss. The policy therefore applies a second axis: the duration and flexibility of essential liabilities, and the correlation of labor or business income with portfolio losses.

Essential liabilities are "near" when cash is required soon enough that market recovery cannot safely be assumed. They are "inflexible" when the amount or timing cannot realistically be reduced, delayed, or substituted. A required medical cost next year is different from optional travel ten years from now, even if their present values happen to be equal. Near and inflexible liabilities reduce the ability to wait through a drawdown.

Income-market correlation asks whether dependable income is likely to weaken at the same time as the portfolio. A business owner whose customers, enterprise value, and personal income all fall in an equity recession has less true capacity than a worker with diversified, stable income, even if both report the same current FS. Correlated income can also create a double claim on liquidity: the portfolio falls while withdrawals rise.

At equal FS, the rule is to shorten the mark-to-market drawdown budget (MDB) and EB by one funded state when essential liabilities are inflexible or near, or when labor or business income is highly correlated with portfolio losses. MDB is the maximum peak-to-trough market-value loss that the investor can endure without violating liabilities, liquidity, or behavior. It is separate from permanent-loss EB. "One state" means the next more conservative funded-state allowance governs those budgets; it does not falsify the reported FS number. At the lowest state, there is no lower band to borrow from, so the most restrictive treatment already applies. This chapter does not invent a personal MDB percentage: that requires verified liability and tolerance facts.

Hypothetical worked example - the second axis

Assume two investors each have a verified base FS of 1.50 and remain in the 1.25-1.74 band after the required liability sensitivities. Investor A has distant, substantially flexible spending and stable income that has historically been weakly related to market losses. Investor B has large essential payments within two years and owns a cyclical business whose income is likely to fall during an equity downturn.

The table's EB band for their reported funded state is 5-8%. Investor B must shorten EB one state, so the more conservative 3-5% band governs EB. MDB must also be shortened one state once its separately verified calibration exists. Investor A does not earn the top of 5-8%; A simply avoids this particular one-state reduction. Other gates can still require less. The purpose of the example is to show why 1.50 is not a complete description of risk capacity.

Flexibility operates asymmetrically. At FS >=1.75, high spending flexibility may support the documented annual election to replace D36 with D24 in the ILF calculation. It does not enlarge issuer loss budgets. The policy permits less near-term liquidity insurance in a narrowly defined strongly funded, flexible state; it does not translate lifestyle flexibility into permission to assume that a company cannot suffer permanent impairment.

Hysteresis and state transitions

A useful policy must resist both false precision and rapid state flipping. Present values move when discount rates, inflation assumptions, spending classifications, currencies, or probabilities change. Market values also move daily. Without persistence rules, a portfolio could become more aggressive after one favorable mark and defensive after one adverse mark, buying and selling around noisy thresholds.

The sensitivity rule is the first defense. Calculate FS at base PVL* and at PVL* +/-10% and +/-20%. If the +/-10% range changes the band, use the lower state. This treats close calls as uncertain rather than as invitations to choose the more permissive label.

The promotion rule is the second defense. Promotion to a higher funded state requires two consecutive annual confirmations. A one-year crossing does not enlarge the applicable policy ranges. If an investor moves from 1.23 to 1.28 at one annual review, the higher state is provisional. Only a second consecutive annual review confirming the higher state can promote the architecture, and the liability sensitivity and second-axis tests still apply.

Demotion is intentionally faster. Death, incapacity, income loss, leverage, litigation, custody impairment, or inflexible spending can demote the state immediately. These are not merely market movements. They are capacity events that can reduce accessible assets, increase liabilities, weaken decision authority, or make near-term cash needs less flexible. Waiting for a second annual review would give adverse facts time to damage the portfolio while the policy continued to assume old capacity.

This asymmetric transition rule is hysteresis: good news must persist before discretion expands, while verified deterioration tightens the system without delay. Hysteresis is not permanent pessimism. It is a control against threshold gaming and procyclical risk taking. It also distinguishes a genuine life event from ordinary market noise. Daily price changes should update monitoring, but they do not automatically promote the policy state.

Hypothetical worked example - promotion and immediate demotion

Suppose an investor is governed by the 1.00-1.24 state. At the 2027 annual review, verified FS is 1.30 and remains in 1.25-1.74 under the +/-10% liability test. The portfolio is not yet promoted. At the 2028 annual review, FS is again verified in that band and all other capacity conditions remain sound. Promotion can then occur.

Now suppose that three months later the investor loses the dependable income used in the liability model. The policy does not wait until 2029. PVL*, the income coverage assumptions, and the second axis are recalculated immediately, and the applicable budgets can demote at once. If the lost income was highly relevant to near essential spending, the ILF and discretionary-risk rules tighten before any new active decision is considered.

Speculative options are nested and charged at total loss

Speculative options sit inside the opportunity sleeve. They are never a third allocation added on top of the strategic and opportunity ranges. The control is based on speculative loss, not merely quoted market value. Across the four funded-state bands, speculative loss is capped respectively at:

- • 0-2% when FS <1.00;
- • 2-3% when FS is 1.00-1.24;
- • 3-5% when FS is 1.25-1.74; and
- • a preferred 5% with a hard cap of 7% when FS >=1.75.

Each speculative option normally has market value of 0.25-0.75% of the portfolio and is assumed to be a total loss. "Normally" does not create a route around the sleeve loss cap. The assumption prevents probability stories, recovery hopes, or mark-to-model values from shrinking the charge. A pre-cash-flow company, paid option, or research position may have upside, but the portfolio must remain sound if its value is zero.

Hypothetical worked example - market value versus loss charge

Assume an investor is validly governed by the 1.25-1.74 state. Six speculative options each have a market value of 0.50% of the total portfolio. Their combined market value is 3.00%. Because each is assumed to lose 100%, the speculative-loss charge is also 3.00%. That sits at the bottom of the 3-5% sleeve-loss range, subject to every other gate.

If an analyst instead assigns only a 30% chance of failure, the policy charge does not fall to 0.90%. Total loss remains the governing assumption. Nor does the existence of a 5% outer cap require adding more speculative positions. Zero or 3% can be the correct result if no further idea qualifies. Any such example remains diagnostic only; this policy authorizes no position.

The Investor Liquidity Floor

The Investor Liquidity Floor (ILF) is the survival reserve that prevents near liabilities from becoming forced sales of risky assets. Its target formula is:

ILF target = D36 + Tax12 + Commit + H x max(0, essential spending beyond D36 - dependable income covering that later interval)

The variables separate claims that are often blurred together.

D36 is the amount required for dated liabilities within the next 36 months. It captures obligations by their time and currency rather than using a generic cash percentage. Where the permitted election applies, D24 replaces this term with dated liabilities within 24 months.

Tax12 is known tax due within the next 12 months. A tax bill is not discretionary capital and cannot be left invested in risk merely because payment is months away.

Commit is the amount needed for fixed or already undertaken commitments. The policy's hard floor separately makes clear that twelve months of fixed commitments must be protected. Known legal, medical, contractual, or similar obligations belong in the verified schedule rather than in an informal emergency estimate.

H is the applicable income-interruption horizon, expressed in units consistent with the spending input. It scales only the uncovered essential spending in the later interval. The final policy does not certify a universal numerical value for H; it must be derived from verified income stability, concentration, market correlation, and the liability schedule. Most importantly, H covers only essential spending beyond what D36 already funds. This de-overlap rule prevents counting the same first 36 months once in D36 and again in an income-interruption buffer.

The expression inside max(0, ...) is later essential spending minus dependable income covering that same later interval. The max function means the term cannot be negative. If dependable income fully covers or exceeds that spending, this component contributes zero; excess income does not reduce dated liabilities, taxes, or commitments elsewhere in the formula. If dependable income covers only part of spending, H multiplies the uncovered gap.

The target is not the same as the hard floor. The hard floor is:

six months of uncovered essential spending + twelve months of fixed commitments + known taxes

The target is the planned reserve. The hard floor is the emergency minimum that cannot be used to buy risk. Falling below the target calls for restoration; falling below the hard floor is a more severe survival breach. Neither cash drag nor a compelling investment story authorizes reaching into the hard floor.

ILF assets must use the home-liability currency, short duration, and high title and custody integrity. "Home-liability currency" means the currency in which essential liabilities must actually be paid, not necessarily the currency printed on the investor's preferred account. Short duration limits the risk that the reserve itself falls materially when cash is needed. High title and custody integrity means the investor's legal claim, access, segregation, transferability, and incapacity arrangements matter as much as nominal safety.

The ILF's real return can be negative during inflation. That drag is reported as an insurance cost, not hidden and not "repaired" by taking credit, duration, complexity, or equity risk inside the emergency reserve. The ILF's job is to preserve payment capacity and reduce forced-sale risk, not to compete with the long-term expected return of equities.

Hypothetical worked example - ILF target and hard floor

This is a hypothetical calculation only.

Assume D36 = 180,000, Tax12 = 20,000, and Commit = 40,000. Beyond the period already covered by D36, annual essential spending is 60,000 and dependable income covering that later interval is 45,000. Assume the documented H for this example is two years. The uncovered later spending is:

max(0, 60,000 - 45,000) = 15,000

The horizon component is:

2 x 15,000 = 30,000

The target is therefore:

ILF target = 180,000 + 20,000 + 40,000 + 30,000 = 270,000

Now assume six months of uncovered essential spending is 30,000, twelve months of fixed commitments is 24,000, and known taxes are 20,000. The hard floor is:

ILF hard floor = 30,000 + 24,000 + 20,000 = 74,000

The difference between 270,000 and 74,000 is not spare risk capital. It is the difference between the target reserve and the emergency minimum. The target remains the planning objective, while the hard floor identifies the amount below which discretionary risk cannot be funded.

Hypothetical worked example - the D24 election

Assume the same investor later has verified FS = 1.80, remains at or above 1.75 under the required sensitivity test, and receives a documented high-flexibility finding at the annual review. If dated liabilities are 180,000 over 36 months but 120,000 over 24 months, the investor may elect D24 = 120,000 in place of D36 = 180,000 for that annual ILF calculation.

The election reduces the dated-liability term by 60,000 for that year. It does not happen automatically, it must be documented annually, and losing either the FS >=1.75 condition or the high-flexibility finding removes the basis for it. It also does not increase the loss budget for any issuer. The rule changes liquidity matching, not the severity of company risk.

Rebalancing architecture - fast for broad sleeves, slow for active risk

The policy uses different clocks for different kinds of decisions. Strategic sleeves are broad, preapproved components, so they can rebalance mechanically when they leave their bands. Active names and themes depend on current evidence and valuation, so they cannot be refilled simply because their prices fell. This is the practical meaning of "fast broad, slow active."

Strategic sleeves trigger action when they are more than 5 percentage points or more than 20% relative from the approved band. The two tests catch different problems. A large allocation can move materially in percentage-point terms without a large relative change, while a small allocation can move 20% relative with only a few percentage points of movement. The action remains subject to the emergency floor and taxes. A soft deviation does not override the ILF, and tax cost must be considered through the policy's governed comparison.

Hypothetical worked example - the two rebalancing thresholds

Suppose a strategic component has an approved upper band boundary of 50% and reaches 56%. It is 6 percentage points above the boundary, so the greater-than-5-percentage-point trigger is met.

Now suppose a smaller component has an approved reference boundary of 10% and reaches 12.1%. The change is only 2.1 percentage points, but it is 21% relative to 10%, so the greater-than-20%-relative trigger is met. These examples only illustrate the trigger calculations. Any action still has to preserve the ILF, respect taxes and costs, and follow the separate implementation authority.

Strategic rebalancing does not mechanically refill a broken theme. If a broad equity sleeve is below its approved band after a sector collapse, the broad sleeve may be restored under its standing rule. The failed sector, issuer, or theme does not receive an automatic top-up. Its business pillars, evidence, valuation, and factor headroom must be re-underwritten.

Decision clocks and fail-closed responses

The operating architecture assigns clocks according to the type of failure. Faster clocks govern legal existence, title, fraud, and licensing because delay can make the investor's claim unrecoverable. Slower but still bounded clocks govern analytical deterioration because a full review takes time. In every case, the stated clock is a governance deadline, not trade authorization from this chapter.

Immediate freeze and a 0-5-trading-day resolution window

A hard gate, title problem, fraud concern, or de-licensing event freezes the affected discretion immediately and must be resolved within 0-5 trading days. "Freeze" means no add and no discretionary increase while the claim is uncertain. The response is risk-monotone: it should reduce or contain the affected risk rather than use an emergency as an excuse to buy a dislocation. Market integrity, legal constraints, and the separate authorization process still matter, but they do not turn an unresolved title or fraud problem into a normal 30-day research item.

Thirty-day full re-underwrite and no add

A Pillar 0, a hard-cap problem, two missed milestones, or Evidence Coverage (EC) below 60 requires a full re-underwrite within 30 days and forbids adding during that period. The pillars are the policy's non-compensating tests of demand, structure and payer quality, capital economics, governance and per-share capture, and valuation and falsifiability. A zero means a foundational part of the case has failed. EC <60 means fewer than 60% of applicable required evidence fields are verified under the policy's coverage method. Two missed milestones prevent a ten-year narrative from indefinitely excusing failure to produce near-term evidence.

The 30-day rule does not assume the position will recover, and it does not assume sale is always correct. It requires the investment case to be rebuilt before any add. If the hard gate or cap remains failed, the later policy's reduction or research-only rules govern. Uncertainty never grants a temporary right to increase exposure.

Valuation-only deterioration

Valuation deterioration is treated differently from deterioration in the owner claim. A price change alone does not force a mechanical trim. If Pillar E, the valuation and falsifiability pillar, falls below admission, adding is forbidden and the position is re-underwritten. An intact thesis may remain only if a standardized after-tax comparison with the named comparator clears.

This separation prevents two opposite errors. It prevents selling a sound claim merely because its quoted price fell, and it prevents using an intact business story to hold an asset whose forward return no longer compensates for price, tax, cost, and risk. "No mechanical trim" is not "hold regardless." It redirects the decision to a standardized comparison rather than a price-only reaction.

Hypothetical worked example - different evidence, different clocks

Assume a holding receives an allegation that legal title to a key asset is invalid. The affected decision freezes immediately and enters the 0-5-trading-day resolution process.

Assume instead that a routine evidence audit lowers EC from 64 to 58. There is no automatic purchase or averaging down. A full re-underwrite is due within 30 days, and no add is permitted.

Finally, assume the business evidence remains intact but a higher price causes Pillar E to fall below admission. The policy does not mechanically trim merely because valuation worsened. It forbids an add, requires re-underwriting, and allows an intact thesis to remain only if the standardized after-tax comparator test clears. The same word, "deterioration," therefore does not produce the same action when the underlying failure mode is different.

Entry pacing is an information rule, not calendar theater

If an entry is ever authorized through a separate process, pacing must be anchored to filings or preregistered milestones and must finish within 180 days. A tranche is not justified merely because a month has passed. The purpose of pacing is to let a known information event test the thesis, not to disguise indecision as risk control.

When all gates pass, the conservative valuation case clears, and no information event is pending, full deployment is the statistical default. Behavioral pacing may still be recorded, but its expected opportunity cost must be recognized. This keeps the lump-sum prior for fully qualified decisions while allowing event-based caution when evidence is genuinely scheduled to arrive.

The rule contains two fail-closed boundaries. An uncompleted schedule does not roll forever after day 180, and a pending information event cannot be bypassed by calendar tranches. Most importantly, this chapter authorizes neither entry method. It states how a separately authorized entry would have to be governed.

Policy Exception Events and the freeze on discretion

A policy is not reliable if repeated deviations merely generate explanations. Policy Exception Events make behavioral and process failures observable. Material examples include using the ILF to fund risk, taking an unjournaled panic action, missing a mandatory review without a written deferral, or deviating from an approved rule without an amendment. The exception log records what happened; it is not a device for normalizing the deviation.

Two material Policy Exception Events in any rolling 12 months freeze new discretion pending review. The threshold is two, not "several," and the lookback is 12 months. Mechanical actions that remain valid under the policy can continue where the other gates allow, but new discretionary risk cannot be introduced until governance has reviewed why the written system and actual behavior diverged.

The exception rule is a safeguard against self-assessed calm. People are poor judges of whether emotion affected them in the moment. Account records, missed deadlines, journal entries, and ILF balances are more observable. The freeze does not punish volatility; it responds to failures of the operating process.

Hypothetical worked example - the rolling exception rule

Assume an investor uses part of the emergency ILF floor to fund a risk purchase in March. That is logged as a material Policy Exception Event. In October, the investor misses a mandatory review and has no written deferral. That is a second material event within 12 months. New discretion freezes pending review, even if markets have risen and both actions appear profitable after the fact. Outcome does not erase process failure.

If the first event falls outside the rolling 12-month window before another material event occurs, the numerical trigger is not met by those two events together. The event history should still remain in the log, and any separate hard gate continues to apply. The rolling rule prevents a permanent freeze for stale events while still treating clustered deviations as evidence that discretion is not currently controlled.

How the pieces fit together

The architecture works in a strict order. First, verify adjusted assets and the liability schedule. Second, calculate FS under base and liability sensitivity cases and select the lower state when +/-10% creates ambiguity. Third, apply the second axis for liability timing, spending flexibility, and income correlation. Fourth, build and protect the ILF. Fifth, set coherent strategic and opportunity ranges, EB, MDB, illiquidity, speculative-loss, and tracking-error limits without treating any ceiling as a target. Only after those portfolio controls exist can later chapters test a particular active idea.

The order matters because every later permission depends on earlier survival facts. A strong company cannot repair an uncertified liability denominator. A high FS cannot repair weak title. A low market price cannot repair a Pillar 0. A broad comparator cannot be treated as riskless merely because it is passive. A speculative option cannot be charged at less than total loss. And a repeated exception cannot be converted into normal policy through silence.

There is also a deliberate asymmetry between expansion and contraction. More discretion requires verified facts, two annual confirmations for promotion, qualifying evidence, and compliance with every gate. Less discretion can follow immediately from death, incapacity, income loss, leverage, litigation, custody impairment, inflexible spending, a hard gate, or repeated exception events. That asymmetry is the policy's central defense against optimism becoming permanent capital loss.

Chapter 4 - The Three Risk Budgets and the Grind Test

The grind test asks a narrow but essential question: after a prolonged real-value stress, do the assets that the policy permits by default still cover the investor's stressed real liabilities? It is not a return forecast, an optimization exercise, or a proposed allocation. It is a solvency diagnostic designed to prevent an apparently diversified portfolio from being called safe merely because its components have different names.

The test exists because the policy separates three kinds of loss. Permanent impairment is the loss of an economic claim itself. A mark-to-market drawdown is a fall in quoted value that may or may not become permanent. A real grind is different from both: inflation, weak real returns, repeated rallies and declines, and ongoing liabilities can erode purchasing power for years even when no single asset suffers a conventional default. Passing a permanent-loss budget does not pass a drawdown budget, and passing either one does not pass the grind test.

This distinction matters most when the assets commonly described as safe are safe only in nominal terms. Cash can preserve a number on an account statement while losing purchasing power. High-quality nominal bonds can pay exactly as promised while their real value falls. Reducing equity exposure can soften an abrupt crash but still leave the portfolio exposed to a decade in which nominal claims are steadily impaired by inflation. The grind test therefore measures assets against real liabilities, not against their starting nominal value and not against an equity index.

The section's numbers are assumptions, not forecasts. Starting investable assets, written A*, are assumed to be 1.50 million. Stressed real liabilities, written PVL*, are assumed to be 1.10 million. The symbol A* means investable assets after known taxes and transaction costs under the broader policy definition. PVL* means the probability-weighted present value of liabilities plus an explicit safety margin. In this chapter's diagnostic, the asterisk is a reminder that neither number is a raw brokerage balance or an informal spending guess. The letter m means million.

The test then applies four real stress results: equity -45%; matched nominal duration -25%; verified duration-matched inflation-linked bonds, or ILBs, 0% real; and Investor Liquidity Floor, or ILF, cash -10% real. No withdrawal is hidden in the arithmetic. These stress values are dated policy hypotheses as of the policy evidence date, 22 August 2026. They are deliberately not presented as probability estimates or promises about 2026-2036.

Why a grind test is necessary

An abrupt crash is visible. A decade-long grind is more deceptive. Nominal balances may rise, coupons may be paid, and occasional market recoveries may create the impression that the plan is working. Yet the liability basket can reprice upward faster than the assets that are supposed to fund it. A household can remain nominally solvent and become progressively less able to meet its real obligations.

The debate that produced this section began with the observation that a diversified core could survive ordinary crash mechanics and still deliver a long negative real result. The policy therefore refused to use discipline, rebalancing, or patience as substitutes for liability arithmetic. "Hold and endure" is a behavior, not evidence that essential liabilities remain funded. If coverage falls below one, the plan has not closed merely because the investor can tolerate the experience.

An earlier diagnostic appeared to pass only because it gave a permanent real diversifier a positive 15% result. That was inconsistent with the policy's own admissibility rules: trend and commodity exposures were conditional, not default assets, and their favorable crisis behavior had not been established after full-cycle carry, costs, taxes, and implementation limits. The final grind test removes that favorable assumption. ILBs receive no speculative 15% gain. Their result is set to 0% real, and only verified liability-matched ILBs are admitted. This change is the central honesty control in the chapter: the test cannot be rescued by an asset whose eligibility or payoff is assumed for the purpose of making the table pass.

The test also resolves an important part of the policy's dispute over tail protection. Lower gross risk remains the first comparison, or null hypothesis, for an abrupt crash and sequence event. Against a real grind, however, simply owning less equity is not enough when cash and nominal duration are also losing purchasing power. For verified real liabilities, the structural instrument is verified inflation-linked matching. Trend, paid tail protection, and commodity or other real diversifiers remain conditional. They require full-cycle net carry evidence and an actual improvement in coverage; they are not silently inserted into this default test.

The four default-admissible buckets

"Default-admissible" does not mean universally suitable, riskless, or authorized for purchase. It means that the asset class has a permitted structural role in this diagnostic without first winning the separate case required of a conditional diversifier. Actual implementation still depends on investor facts, legal title, taxes, currency, duration, instrument design, and approval outside this policy. The four buckets are equity, matched nominal duration, verified duration-matched ILBs, and ILF cash.

Equity

The equity bucket represents productive ownership claims and receives a -45% real stress. In the two diagnostic mixes, the equity contribution to portfolio loss is its weight multiplied by 45%. A 45% equity weight therefore contributes 45% x 45% = 20.25% of starting portfolio assets to the loss. A 40% equity weight contributes 40% x 45% = 18.00%.

The test does not claim that every equity security will fall exactly 45%, nor does it distinguish broad equity from individual active holdings. It also does not treat business quality, thematic appeal, or nominal revenue growth as a real hedge. Those questions belong to other policy gates. Here equity is a stressed aggregate bucket. The purpose is to prevent optimistic security selection from weakening a liability test.

Matched nominal duration

Matched nominal duration means high-quality nominal claims whose currency and timing are aligned with fixed nominal liabilities. The word "matched" is essential. A generic bond allocation is not automatically a liability match, and a liability match is not a claim that bonds will provide generic crash alpha. In the grind scenario, matched nominal duration receives a -25% real stress because an instrument can honor its nominal cash flows while those cash flows lose real purchasing power.

At a 25% portfolio weight, this bucket contributes 25% x 25% = 6.25% to the total portfolio loss. At a 20% weight, it contributes 20% x 25% = 5.00%. Its role remains legitimate for nominal liabilities. The negative real stress simply stops that role from being confused with protection against an inflation-linked liability.

Verified duration-matched inflation-linked bonds

ILBs are bonds whose principal or payments are linked to an inflation index. In this policy they are structurally admissible only as matching assets for verified real liabilities, and only when the relevant index, law, and market integrity are credible. The match must also survive currency, duration, basis, title, and tax review. "Basis" is the difference between the inflation experienced by the liability and the inflation index used by the bond. A security linked to the wrong index or currency may be inflation-linked in name and still fail to match the actual obligation.

The grind table gives verified duration-matched ILBs a 0% real result. That means their assumed real value is unchanged over the measurement used in the table. It does not mean the securities have no risk, no interim price movement, or a guaranteed real return. It also does not apply automatically to any fund with "inflation-linked" in its name. The policy's 0% treatment is defensible only for the verified matching function described in the section. If the investor may be forced to sell an unmatched or long-duration vehicle before the relevant liability date, the simple endpoint assumption cannot establish survival through the path.

Most importantly, ILBs get no speculative +15% gain. The test credits matching, not alpha. At either a 20% or 30% weight, their direct loss contribution in the printed arithmetic is zero. The larger allocation helps only because it replaces buckets that carry negative real stresses. That is a very different claim from saying ILBs will appreciate.

ILF cash

The ILF is the Investor Liquidity Floor: the reserve intended to meet near-term essential spending, known taxes, fixed commitments, and other covered liquidity needs without forcing the sale of distressed risk assets. Under the broader policy, it belongs in the home-liability currency, at short duration, with high title and custody integrity. It is survival insurance, not discretionary dry powder and not a yield-seeking sleeve.

In the grind test, the 10% ILF allocation loses 10% in real terms. The contribution to total portfolio loss is therefore 10% x 10% = 1.00%. The nominal cash may still be present; the assumed loss is in purchasing power. This treatment makes the trade-off explicit: liquidity can protect near-term execution while imposing real drag. The policy accepts and reports that drag rather than reaching for risk inside the emergency floor.

The 10% shown in both rows is diagnostic, not an ILF target. Actual ILF size is determined from the investor's essential spending, taxes, commitments, dependable income, and flexibility under the policy's separate liquidity rules. Nothing in this table overrides that calculation.

How to read the table

The table is a one-period stressed balance-sheet snapshot. For each bucket, multiply its portfolio weight by the magnitude of its assumed real loss. Add those weighted losses. Subtract the sum from one. Multiply the remainder by the 1.50 million of starting investable assets. Finally, divide end assets by the 1.10 million of stressed real liabilities.

The coverage formula is:

Coverage = stressed end assets / stressed real liabilities

Coverage of 1 means assets exactly equal liabilities. Coverage above 1 means the diagnostic has a positive cushion. Coverage below 1 means an immediate shortfall. A pass is therefore not "positive return" and a fail is not "lost money." The pass/fail boundary is whether stressed assets cover stressed real liabilities.

Diagnostic mix	Weighted real loss	End assets	Coverage	Admissibility result
45% equity / 25% nominal / 20% ILB / 10% ILF	20.25% + 6.25% + 0% + 1.00% = 27.50%	1.5 x (1 - .2025 - .0625 - 0 - .01) = 1.0875m	0.989	**FAIL by 12,500 immediately**; any spending extends the failure.
40% equity / 20% nominal / 30% ILB / 10% ILF	18.00% + 5.00% + 0% + 1.00% = 24.00%	1.5 x (1 - .18 - .05 - 0 - .01) = 1.1400m	1.036	**PASS by 40,000**, only if roughly 30% is genuinely real/index-matchable and basis/title integrity holds.

The first row starts with the inherited illustrative mix. Equity reduces the portfolio by 303,750, nominal duration by 93,750, ILBs by zero under the matching assumption, and ILF cash by 15,000. Total loss is 412,500, leaving 1,087,500. Against liabilities of 1,100,000, the shortfall is 12,500. The unrounded coverage is about 0.9886, printed as 0.989. This is an immediate failure before adding any spending that was not already captured in PVL*.

The second row is the boundary diagnostic. Equity reduces the portfolio by 270,000, nominal duration by 75,000, ILBs by zero, and ILF cash by 15,000. Total loss is 360,000, leaving 1,140,000. Against the same 1,100,000 liability value, the cushion is 40,000. Coverage is about 1.036. The word "boundary" matters: this is a narrow conditional pass, not evidence that the mix is optimal or safe in every relevant path.

The difference between the rows is not mysterious diversification. Five percentage points move out of equity, five move out of nominal duration, and ten move into verified ILBs, while ILF remains at ten. Under the stipulated stresses, that reduces loss by 52,500: 33,750 from the lower equity weight and 18,750 from the lower nominal

weight. The result moves from 12,500 below liabilities to 40,000 above them. The arithmetic also shows how fragile the classification is. If the supposed ILB allocation is not a genuine real-liability match, the zero-loss input is unavailable and the printed pass is not valid.

Clearly labeled hypothetical application examples

Hypothetical example 1 - applying the inherited mix

Assume, solely for illustration, that an investor's verified inputs exactly equal the table's A*=1.50m and PVL*=1.10m, and that every bucket qualifies for its stated treatment. The portfolio is 45% equity, 25% matched nominal duration, 20% verified duration-matched ILBs, and 10% ILF cash.

The weighted losses are 20.25%, 6.25%, 0%, and 1%, for a total of 27.5%. End assets are 72.5% of 1.50 million, or 1,087,500. Coverage is 1,087,500 divided by 1,100,000, or 0.989 after rounding. The correct label is **FAIL by 12,500 immediately**. The investor may still own liquid, high-quality assets and may still be comfortable with market volatility. Neither fact changes the shortfall.

This example is not a recommendation, a current portfolio assessment, or a forecast.

Hypothetical example 2 - applying the boundary mix

Now assume the same starting assets and stressed liabilities, but use 40% equity, 20% matched nominal duration, 30% verified duration-matched ILBs, and 10% ILF cash. The total weighted real loss is 24%, leaving 1,140,000. Coverage is 1.036 and the cushion is 40,000.

The correct label is **PASS by 40,000**, subject to the full condition. Roughly 30% of the portfolio must genuinely match real liabilities, and the index basis and legal title must remain sound. If the liabilities are in a different currency, follow a different inflation basket, arrive on a different duration schedule, or cannot legally be met from the instrument, the table has not demonstrated a pass.

This example is a boundary diagnostic only. It is not a recommendation or authorization to adopt a 40/20/30/10 allocation.

Hypothetical example 3 - a mislabeled ILB bucket

Suppose the second row is presented as 30% ILBs, but diligence shows that only 20% is genuinely matched. The other 10% is an unmatched nominal exposure. For a deliberately simple illustration, classify that 10% under the table's -25% nominal stress. It adds a 2.5% portfolio loss, or 37,500, to the printed result. End assets fall from 1,140,000 to 1,102,500, leaving only 2,500 of coverage above 1.10 million.

The lesson is not that 2,500 is a valid new policy cushion. The lesson is that labels do not earn admissibility. Until currency, index, duration, basis, title, and tax are verified, the 0% real ILB treatment cannot be assumed. With a different mismatch, the result could be worse. This is hypothetical arithmetic, not a new stress rule.

Hypothetical example 4 - an omitted withdrawal reverses the pass

Assume the boundary mix leaves 1,140,000, but a 50,000 real withdrawal was omitted from PVL* and must be funded during the stressed period. A simplified endpoint subtraction leaves 1,090,000 against 1,100,000 of liabilities. Coverage becomes about 0.991 and the plan fails by 10,000.

This simplified example does not model timing, returns before and after withdrawal, taxes, or ILF use. It shows why the chapter cannot infer a grind duration from the printed endpoint. The sequence of withdrawals is the grind's defining mechanism. Since no withdrawal path was supplied to the final policy, the duration of accepted slow failure is unbounded once withdrawals occur. The withdrawal-path grind budget remains a blocking open item for certification.

Hypothetical example 5 - the liability estimate changes

Assume the second row's stressed assets remain 1,140,000, but a verified liability review raises PVL* from 1,100,000 to 1,150,000. Coverage becomes about 0.991 and the result fails by 10,000. No asset price changed. The classification changed because the denominator was corrected.

This example illustrates liability-first measurement. The table cannot compensate for an incomplete liability schedule. A portfolio can appear to pass only because liabilities, taxes, or currency effects were omitted. This is hypothetical and does not alter the section's official 1.10 million assumption.

What the test rules out

First, it rules out an unearned positive diversifier return. A conditional real asset, commodity program, managed-trend strategy, or tail hedge cannot be inserted with a favorable payoff merely to close the table. The earlier +15% real-diversifier assumption is explicitly removed. Trend, tail-premium, and commodity exposures remain outside the default set unless full-cycle net carry and actual coverage improvement justify them under their separate evidence tests.

Second, it rules out treating lower gross risk as a universal answer. Lower gross risk remains the null for crash risk. It is not the grind instrument because cash and nominal duration also receive negative real results in this stress. For verified real liabilities, verified inflation-linked matching is the structural response.

Third, it rules out calling nominal safety real safety. A matched nominal bond can be exactly right for a fixed nominal liability and still lose 25% in real terms in this diagnostic. ILF cash can protect spending access and still lose 10% of purchasing power. The table permits both statements to be true at once.

Fourth, it rules out certifying the inherited 45/25/20/10 mix. That mix fails by 12,500 immediately, at coverage of 0.989. Any spending omitted from the liability measure extends the failure. The policy does not round 0.989 up to solvency and does not excuse a small shortfall because the inputs are illustrative.

Fifth, it rules out presenting the 40/20/30/10 boundary as a portfolio recommendation. Its 40,000 cushion exists only under the stipulated real stresses and only if roughly 30% of assets is genuinely real/index-matchable with sound basis and title integrity. It does not authorize a trade, target allocation, instrument, account, or implementation path.

Finally, it rules out hiding unresolved liabilities behind portfolio analysis. Certification requires a verified liability cash-flow schedule, currencies, index basis, title, and tax treatment. The liability currency, index, and duration schedule are not optional refinements; without them, the match cannot be certified.

What the test cannot prove

The table cannot prove how many years the portfolio will survive. It is an immediate stressed endpoint, not a yearly cash-flow simulation. A genuine grind budget would project assets and remaining liabilities through time, including the timing of withdrawals, inflation, income impairment, taxes, foreign exchange, and returns along the path. Those data were not supplied. The final policy therefore states that the duration of accepted slow failure is unbounded once withdrawals occur. This is not a prediction of infinite failure; it is a statement that no finite survival duration can be established from the available inputs.

The test cannot prove that 0% is the right result for every inflation-linked security. The assumption belongs to verified duration-matched ILBs serving a real-liability function. An ILB fund, a long-duration linker, an unmatched inflation index, or a position exposed to forced interim sale can have material mark-to-market risk. The debate specifically identified vehicle and interim-path risk, but the final section did not adopt a separate numeric interim yield-shock threshold. The prudent conclusion is therefore limited: the endpoint arithmetic is valid only after the stated matching and integrity conditions are met; it does not certify path liquidity.

The test cannot prove legal access, title, convertibility, or tax efficiency. A claim can be economically matched and practically unavailable. That is why basis and title integrity appear in the pass condition, and why the policy's key-risk observable includes the liability cash-flow schedule, currencies, index basis, title, and tax.

The test also cannot prove that equities will fall 45%, nominal duration 25%, or ILF cash 10%, nor that these assets will move at the same time in exactly this way. The stresses are assumptions used to expose a boundary. A pass under one severe-but-plausible path does not prove robustness to all paths, while a fail identifies a concrete insufficiency under the stated path.

It cannot determine the correct allocation. The two rows are diagnostics selected to display a failure and a conditional boundary pass. They do not optimize expected return, minimize risk, account for the investor's behavior, or replace the funded-state, liquidity, comparator, tax, custody, and implementation rules elsewhere in the policy.

Most importantly, the grind test cannot certify liability closure. Liability closure remains uncertified. The withdrawal-path grind budget, the liability schedule, and the actual matching implementation remain blockers. The table is valuable precisely because it states what follows from the known assumptions and refuses to manufacture what the missing facts cannot support.

Chapter 5 - Comparator, Look-Through Limits, and Certification

The comparator is a real position, not a neutral backdrop

The comparator is the investable baseline against which the policy judges active complexity. It is not merely an index printed in a performance report. It owns securities, carries issuer and factor concentrations, incurs fees and taxes, and can experience permanent or long-lasting losses. Because most capital may sit in the strategic sleeve, the design and look-through risk of the comparator can matter more than any single active research decision. Calling an exposure "passive" does not make it harmless or exempt it from the portfolio's risk limits.

The investor's actual named comparator is missing. That missing fact is decisive. The policy therefore uses, for diagnosis only and not for ratification, a draft proxy consisting of 60% MSCI ACWI and 40% liability-matched high-grade fixed income. The equity reference is supported only at the level of holdings availability by the MSCI ACWI factsheet and the iShares ACWI holdings file. Those sources can show that constituent information is available. They do not establish the analyst's sensitivity scores, prove the factor model, identify the investor's actual implementable fund, or settle tax, tracking-error, custody, and liability-matching questions.

The phrase "for diagnosis only" is therefore substantive. The 60%/40% proxy can reveal that the proposed governance system may be born with embedded risk. It cannot serve as a portfolio recommendation, a target allocation, or a substitute for verified investor facts. It does not authorize the purchase of MSCI ACWI exposure, bonds, or any other instrument. It also cannot ratify a higher cap merely because its preliminary output exceeds the existing one.

This distinction prevents a circular conclusion. Without it, an analyst could choose a concentrated proxy, observe that the proxy breaches the policy, and then use the breach to raise the policy cap. The comparator would manufacture the permission needed to accommodate itself. The final rule blocks that move: the diagnostic output may identify a problem, but only a reproducible actual-comparator study can support a formal resolution.

What proxy scoring is trying to measure

Ordinary sector labels are too coarse for this purpose. A software company, semiconductor producer, electrical-equipment supplier, and bond can all respond to a common AI-capex, power, or real-rate shock even though they sit in different sectors. Conversely, two companies in the same sector can have very different sensitivities. The proxy method therefore asks a causal question: how much of the total portfolio behaves as though it were exposed to a specified common factor, and how much could the portfolio lose if that common factor breaks?

The policy uses two related measures. For factor f, factor-equivalent weight is:

$$FEW_f = \Sigma(w_i \times x_{if})$$

Here, i identifies a holding or a proxy bucket. w_i is that holding's weight in the total portfolio, not merely its weight inside an equity fund. x_{if} is its estimated exposure to factor f, expressed from 0 to 1. The underlying scoring convention is 0 for none, 0.25 for indirect, 0.50 for material, 0.75 for high, and 1.0 for dominant exposure. Σ means that the weighted exposures are added across all relevant holdings and buckets.

Factor stress loss is:

$$FSL_f = \Sigma(w_i \times \text{loss}_{if} \text{ under the common scenario})$$

loss_{if} is the estimated fractional loss for holding i under the common stress scenario for factor f. Where the worksheet separates sensitivity from the stressed repricing assumption, the same calculation may be written as

$FSL_f = \Sigma(w_i \times x_{if} \times s_{if})$, with s_{if} denoting the loss rate applied to the exposed portion. FEW describes equivalent exposure; FSL translates that exposure into portfolio loss under a defined shock. Both are diagnostics, and neither is a forecast of what will happen.

The five mandatory factor families are AI-capex sustainability, state-payer exposure, China/export-control exposure, real rates/refinancing/quality duration, and power availability/cost. The inherited risk framework also tests two joint shocks: AI-capex plus power, and state payer plus rates. Joint tests matter because separate factor totals can miss a single economic chain. A data-center buildout, for example, can support chip demand and electrical-equipment orders at the same time; a break in that chain may affect both exposures together.

The original common-factor framework supplies the risk reference points that remain in force while the actual-comparator work is incomplete. AI-capex has a preferred FEW cap of 15% of the total portfolio or 25% of risk assets and a 5% hard FSL cap. State-payer exposure has a preferred FEW cap of 20% of the total portfolio, with one sovereign at 15%, and a 5% hard FSL cap. China/export-control exposure has a preferred FEW cap of 15%, a 10% limit for one irreplaceable node, and a 5% hard FSL cap. Directional rates/refinancing/quality-duration exposure has a preferred net FEW cap of 25% and a 6% hard FSL cap. Power availability/cost has a preferred FEW cap of 20% and a 4% hard FSL cap. The joint AI-capex plus power shock and the joint state-payer plus rates shock each have an 8% portfolio-loss cap.

These are portfolio look-through limits. Comparator exposures consume headroom before an active addition is considered; the active sleeve does not receive a separate factor allowance. In the loss-budget arithmetic elsewhere in the policy, a proposed addition that would breach a hard factor cap has $Fmult = 0$, which sets allowed loss to zero. No epistemic floor rescues that hard-cap state.

One important correction applies to rates. Liability-matched duration is excluded from directional rates FEW. If a high-grade bond position genuinely offsets the duration and currency of a verified liability, treating the asset's standalone rate sensitivity as a directional bet would manufacture a false breach. The matched position must instead be reported separately, together with basis risk: the risk that the instrument, index, law, currency, tax treatment, cash-flow timing, or market behavior does not match the liability as assumed. Exclusion from directional rates FEW is not permission to ignore that basis risk.

How the top-20-plus-buckets method works

The full constituent-by-constituent ideal is not operationally realistic for one person. It can also create false precision: hundreds or thousands of tiny scores would consume hours without necessarily improving the decision. The accepted proxy method therefore scores the comparator's top 20 constituents individually and groups the remainder into sector proxy buckets.

For each top-20 constituent, the worksheet must preserve the date, portfolio weight, factor score, stress-loss assumption, supporting evidence, and scorer's rationale. The analyst should distinguish the channels that create exposure, such as revenue, gross profit, backlog, capex, asset replacement, supplier dependence, customer financing, legal restrictions, or custody. A diversified revenue label does not cancel a shared capex or discount-rate dependency. If multiple channels point to different sensitivities, the inherited method uses the highest channel unless a cash-flow model supports a lower score.

The remainder is then grouped into documented sector buckets. Each bucket needs its total weight, composition source, factor scores, stress-loss assumptions, and limitations. "Other technology" or "industrials" is not enough if those buckets conceal materially different AI, power, or export-control channels. At the same time, the method does not require invented security-level precision for immaterial positions. The purpose of bucketing is a bounded, reproducible approximation, not an excuse to place unexamined exposure in an "other" category.

The scoring is refreshed annually and on a material event. In year one, it is repeated independently after six months. The final Section 7 also requires two independent scoring passes before ratification. Independence means

the second pass is not a cosmetic review of the first person's totals. It should be capable of assigning different sensitivities from the archived evidence and then reconciling the differences explicitly.

The worksheet must also report ranges under ±0.25 sensitivity perturbations. This is not a claim that every error is exactly 0.25. It is a standardized robustness check around the scoring grid. A conclusion that changes from "within cap" to "in breach" when a plausible sensitivity moves by one scoring step is not robust enough to ratify a cap. The uncertainty belongs in the decision record rather than being hidden behind a single decimal.

Hypothetical example - FEW calculation only, not a forecast or live portfolio. Suppose one top constituent is 4% of total portfolio and receives an AI score of 0.75. Its AI FEW contribution is 4% × 0.75 = 3.0%. Suppose a sector bucket is 12% of total portfolio and receives a score of 0.50. Its contribution is 12% × 0.50 = 6.0%. Together those two lines contribute 9.0% of AI FEW before any other constituents are counted. The example shows why portfolio weights must be measured at the total-portfolio level and why apparently different holdings can accumulate into one factor exposure. It does not say that either score is correct for any real holding.

Hypothetical example - FSL and perturbation only, not a forecast or instruction. If the same 4% constituent has AI sensitivity 0.75 and the exposed portion is stressed by 60%, its FSL contribution is 4% × 0.75 × 60% = 1.8% of total portfolio. Moving the sensitivity by ±0.25 changes that line's FSL by 4% × 0.25 × 60% = 0.6 percentage points. A few such differences across the top 20 can move a total across a 5% cap. This is why an unarchived set of analyst sensitivities may be plausible at an order-of-magnitude level while still being too unstable for ratification.

What the inherited proxy found - and what it did not prove

The inherited diagnostic proxy produced AI FSL of 8.2% and joint AI+power FSL of 9.0%. These exceed the original 5% AI cap and 8% joint cap. The gaps are therefore 3.2 percentage points and 1.0 percentage point, respectively. The policy describes the figures as plausible order-of-magnitude assumptions, not reproducible measurements.

Hypothetical currency translation - illustration only. If total portfolio assets were 1,000,000 currency units, an 8.2% AI FSL would represent 82,000 of modeled stress loss, compared with 50,000 under a 5% cap. The diagnostic gap would be 32,000. A 9.0% joint AI+power FSL would represent 90,000, compared with 80,000 under the 8% cap, for a 10,000 gap. The actual policy requires the risk-versus-feasibility gap to be reported in currency, but this illustration is not based on verified investor assets and does not authorize a position or cap change.

Several missing elements prevent those outputs from carrying more weight. The actual named comparator is unknown. The top-20 worksheet has not been archived. The sector-bucket assumptions are not published. The second independent scoring pass has not occurred. The ±0.25 perturbation ranges are absent. The investor's tax basis, wrappers, implementation costs, and executable capped or blended alternatives are unknown. As a result, the proxy can support a conservative interim restriction, but it cannot support ratification of a new feasibility cap.

This is an important epistemic distinction. "Plausible" means the number is not obviously impossible given the broad composition and stress logic. "Reproducible" means another scorer can reconstruct the number from archived weights, rules, evidence, and assumptions. "Robust enough for a cap decision" means plausible disagreements and sensitivity perturbations do not reverse the decision. The proxy presently reaches only the first level.

The policy therefore does not use 8.2% or 9.0% as forecasts, accepted target risks, or newly authorized caps. Nor does it infer that the AI buildout must fail, that power-sensitive businesses must underperform, or that a capped comparator must win. The figures identify a governance conflict: a supposedly neutral default may already carry more of the named factor risk than the original limits contemplated.

The forced-choice rule

The debate exposed a failure mode in which an embedded breach could survive indefinitely as a "documented exception." An active position would be prohibited, yet the same exposure inside the comparator would remain untouched because changing the comparator required an expected-return proof that could not honestly be supplied. The result would be paperwork without resolution.

The durable remedy is a forced choice after four consecutive quarters of an embedded issuer or factor hard-cap breach measured on the actual comparator. The policy must then choose in writing between two branches. One branch adopts an implementable capped or blended comparator, with its tracking-error budget, costs, tax consequences, and abandonment point specified. The other branch raises the cap through a formal amendment after the embedded risk has been re-underwritten. A renewable exception is not a third branch.

Section 7 does not execute either permanent branch on the incomplete proxy. Instead, it retains the original risk caps and sets every active addition with material AI/power sensitivity to zero until the actual-comparator study is complete. This is the present forced-choice state. It is deliberately conservative about new active exposure while refusing to pretend that the proxy is sufficient to redesign the strategic allocation.

"Active addition" is narrower than a trade instruction and must not be read as one. The rule says the governance calculation has no headroom for a new AI/power-sensitive active addition while the original caps are exceeded. It does not direct a sale, rebalance, hedge, or comparator change. Any implementation action would require the separate approval process, verified facts, tax and legal review where relevant, and all other policy gates.

The rule also makes the standing anti-AI-active position visible. That stance may later prove too strict, too loose, or directionally right for the wrong reason. It remains a dated risk control, not a forecast resolution of AI overbuild, accelerator lives, defensive capex, or power scarcity. Disclosure is essential because a prohibition on active additions is itself a consequential portfolio-design position even when no order is placed.

Why a temporary diagnostic ceiling is tightly bounded

The policy recognizes that a diagnostic or feasibility ceiling may sometimes be discussed while the formal work is underway. It imposes three safeguards. First, a temporary diagnostic ceiling may never exceed twice the original cap. For the original 5% AI FSL cap, twice the cap is 10%; for the original 8% joint AI+power cap, twice the cap is 16%. These are absolute outer bounds on a temporary diagnostic construct, not current caps and not permissions to take that much risk.

Second, the ceiling does not float with index concentration. A rule such as "current comparator FSL plus 1 percentage point" would rise automatically as the embedded risk rises. If concentration worsened, the control would loosen, which reverses the function of a cap. The final rule therefore says the ceiling floats neither with index concentration nor with analyst convenience.

Third, a temporary diagnostic ceiling auto-reverts after 12 months. It cannot be renewed simply by changing its label or updating the same incomplete worksheet. The actual-comparator study must either support a formal amendment or support an implementable capped/blended design. Otherwise the original caps govern and the forced choice runs again with the capped-comparator study mandatory.

Hypothetical example - temporary ceiling logic only. Assume an incomplete diagnostic study estimates a factor FSL at 6.4% against an original 5% cap. Writing 7.4% because the estimate is 6.4% and adding 1 percentage point would create a floating cap. If the estimate later rose to 8.1%, the same formula would raise the cap to 9.1%. The final policy forbids that ratchet. Any temporary ceiling would need a fixed documented number, could not exceed 10% for that 5% original cap, and would lapse after 12 months without proper ratification. In the current Section 7 state, however, the original 5% cap continues to govern; this example grants no authority to adopt 7.4%, 9.1%, or 10%.

What must be studied before either permanent branch is credible

Ratification requires an evidence package, not a conclusion in search of support. The package begins with the actual named comparator: the investable allocation, vehicles or instruments, rebalancing rule, costs, tax treatment, and liability-matching role that apply to this investor. A placeholder cannot survive as the comparator after the policy's 12-month limit.

The package must contain archived top-20 weights and sector buckets for the remainder. It must show two independent scoring passes, including the second pass completed after six months in year one, and ±0.25 sensitivity perturbation ranges. It must calculate FEW, FSL, and the joint stresses from visible assumptions. It must report liability-matched duration separately from directional rates FEW and disclose the remaining basis risk.

It must also present a capped or blended alternative rather than assuming none is executable. "Capped" generally means that a rules-based design limits a constituent or factor concentration; "blended" means the strategic comparator combines exposures under a stated rule. Those descriptions are explanatory only. This chapter selects neither design and supplies no blend percentage. The study must compare after-tax and after-cost implementation, expected tracking error, turnover, tax realization, operational availability, and the possibility that a modification is abandoned at the worst time.

Tracking error is the divergence between the modified design and the uncapped comparator. It is not merely a statistical nuisance. A capped design can lag an increasingly concentrated benchmark for years, creating behavioral and governance pressure to abandon the rule just before it might help. The study therefore needs a tracking-error budget and a predeclared abandonment point. An abandonment rule is a trigger for review under the policy, not an automatic trade instruction.

The historical comparison must include Japan-1989, US-1992-2025 and US-2013-2025, Nortel/TSE, and Nokia/HEX. These exhibits point in both directions. Japan-1989 and the Nortel and Nokia episodes warn that cap-weighted comparators can embed extreme concentration and that structural capping may protect against an unwind. The US periods warn that capping, equal weighting, or valuation tilting can lag an uncapped winner-dominated benchmark for a long time. The policy does not allow one exhibit to erase the others. Their role is to force the study to confront concentration loss, long tracking-error pain, tax and turnover costs, and the risk of procyclical abandonment.

This balanced exhibit set preserves the open D-2 dispute. The procedure for deciding is resolved, but the actual capped-versus-uncapped design remains open. The study is not permitted to smuggle in a market forecast as a certainty. Raising a cap accepts more embedded risk; it does not remove the risk. Adopting a cap exchanges some embedded concentration for tracking error, implementation cost, and possible regret; it does not create a free hedge.

Certification is a blocker, not a blessing

Actual-comparator scoring is a certification blocker. In practical terms, the policy cannot be represented as fully operationally certified while the actual comparator, its look-through, and its resolution of any persistent breach remain missing. The proxy cap decision is unratified. The original caps govern. The embedded risk indicated by the proxy must be disclosed rather than excused.

Certification here does not mean that a scoring worksheet certifies an investment. It does not validate a company, create a buy list, establish a target weight, or approve an order. Completing the actual-comparator study would remove one blocker only. Liability closure, actual funded status, trait tolerance, market drawdown budget, tax basis, domicile and custody structure, the withdrawal-path grind budget, company evidence-coverage audits, and every current valuation remain separate unverified items under the full policy.

The distinction matters because a well-formatted worksheet can create institutional theater. A file can contain 20 rows, two signatures, and decimal outputs while still using stale weights, management-only evidence, broad

buckets, inconsistent stress losses, or coordinated scorers. Certification requires reproducibility and resolution of the identified governance question, not merely document completion.

The policy also recognizes an executability limit. The comparator's top-20-plus-buckets work is refreshed annually and on material events, rather than rescored in full every quarter. If the actual comparator cannot be scored reproducibly within 16 operator hours per year, the full policy treats that as a portfolio-level falsifier for new active work. The default then remains the named comparator plus the Investor Liquidity Floor and verified nominal/real liability matching, not an improvised replacement. This is still a governance default, not an authorization to transact.

Limitations and failure modes

The first limitation is proxy mismatch. The investor may not use the assumed 60% MSCI ACWI and 40% liability-matched high-grade fixed-income structure. The actual equity vehicle may track a different index, use sampling, hold cash, lend securities, or have different tax and custody consequences. The actual liabilities may not make the assumed fixed-income sleeve genuinely matched. A precise analysis of the wrong comparator is still wrong.

The second is stale composition. Cap-weighted indexes change continuously, and a material event can alter both weights and sensitivities before the annual review. This is why material-event refreshes supplement the annual cadence. It is also why a cap cannot float automatically with the latest concentration: measurement must update, but the control cannot chase the risk upward.

The third is sensitivity subjectivity. Scores of 0, 0.25, 0.50, 0.75, and 1.0 are disciplined labels, not natural constants. A company may be exposed through customer budgets, suppliers, financing, valuation duration, or regulation in ways that filings do not isolate. The second pass and ±0.25 ranges reveal this uncertainty but do not eliminate it.

The fourth is bucket error. Individually scoring the top 20 captures the largest weights but may understate a factor spread across many medium and small constituents. Sector averages can also obscure business-model differences. The bucket assumptions must therefore be archived, challenged, and interpreted as approximations.

The fifth is stress-model error. FSL depends on the selected scenario and the assumed transmission from a factor shock to each holding. Joint shocks can overlap, so adding two standalone FSLs may double-count losses; treating them as independent may undercount a common chain. The policy therefore requires explicit joint AI+power and state+rates tests rather than relying only on separate totals. Even those tests cannot enumerate every path.

The sixth is matching error. Excluding liability-matched bonds from directional rates FEW is correct only to the extent the match is real. Currency, duration, index, legal, tax, reinvestment, and market-liquidity differences can leave basis risk. Labeling an asset "matched" without a verified liability schedule would hide rather than remove exposure.

The seventh is governance laundering. A proxy breach can be used to argue for a higher cap; a capped alternative can be presented without its tax and tracking-error burden; or a temporary diagnostic ceiling can quietly become permanent. The absolute 2× ceiling, non-floating rule, 12-month auto-reversion, mandatory alternative study, and certification block exist to stop those transformations.

The eighth is behavioral abandonment. An uncapped comparator can feel intolerable near a concentration peak, while a capped comparator can feel intolerable during a prolonged winner-led market. The historical exhibits show why either design can become hardest to maintain at precisely the wrong time. A tracking-error budget and abandonment point must be written before that pressure arrives.

The ninth is false actionability. FEW and FSL can determine that an active addition has no factor headroom, but they do not by themselves determine whether an existing exposure should be sold, how quickly anything should change, or what instrument should replace it. Those questions belong to separate approval, tax, liquidity, legal, and implementation processes. This chapter authorizes none of them.

Part III - Decade Investment Worldview

Chapter 6 - The Decade Investment Worldview: Why These Forces May Shape 2026-2036

Status, governing boundary, and evidence date

This chapter is a dated worldview, not durable portfolio law. Its evidence cutoff is 22 August 2026. It identifies forces that may shape the decade, the mechanisms that could make them persistent, the capital cycles they may summon, and the parties that may capture or lose the resulting value. It does not forecast a single path, recommend a security, set a target allocation, authorize an order, or change any existing policy gate. Every company remains at a current admissible weight of zero until the separate evidence, valuation, governance, portfolio, tax, legal, liquidity, factor, and owner-approval requirements have been completed. Nothing here cures the unresolved comparator, liability, AI/power factor, or implementation questions in the controlling policy.

The purpose of a worldview is narrower and more useful than prediction. It should tell the researcher what causal link to test, what evidence would make a belief stronger, what evidence would falsify it, and where an apparently powerful social force can fail to reach a diluted shareholder. The chapter therefore uses phrases such as "potentially helped," "potentially harmed," and "favored research prior" only to describe operating hypotheses. They are not expected-return conclusions. A business can benefit from a force and still be a poor security because the benefit is competed away, captured by another claimant, consumed by reinvestment, trapped at a subsidiary, diluted, taxed, financed expensively, or already more than reflected in price. A security can be attractive in isolation and still have no portfolio role because its factor is already embedded in the named low-cost comparator or because liability and liquidity constraints dominate.

Method: separating evidence, mechanism, cycle, forecast, and implication

Every force in this chapter is examined through five distinct layers.

Observed starting evidence is a dated fact or estimate available by the cutoff. It may show adoption, spending, age structure, physical exposure, debt, trade policy, capacity, utilization, or price. It is not allowed to become a decade conclusion merely because it is vivid. A queue is not an energized load. A budget is not an appropriation. An appropriation is not a profitable contract. A population is not a funded customer base. A clinical endpoint is not a collectible payer saving. A backlog is not cash.

Structural mechanism is the causal process that could persist for much of 2026-2036. Cohorts already born will age. Software replication remains cheap while organizational change remains slow. Transmission and specialist training take years. Public debt refinances gradually. Installed cities, coastlines, grids, supply chains, and care systems cannot be rebuilt in one cycle. Structural does not mean inevitable, monotonic, or profitable; it means the causal clock is longer than a normal inventory or credit cycle.

Cyclical condition is a reversible market state: shortage pricing, duplicate orders, unusually cheap or expensive finance, inventory accumulation, emergency procurement, speculative connection requests, commodity spikes, project backlogs, venture funding, or temporary margin. Cyclical scarcity can be real and still disappear before long-lived capital earns its cost. Conversely, a severe bust can set up a later harvest if capacity truly exits and survivors become disciplined. The chapter therefore describes scarcity, buildout, competitive flood, and consolidation or harvest as states, not permanent labels.

Dated forecast is a scenario estimate such as power demand in 2030, urban investment through 2050, debt in 2029, or an age-related spending path. Forecasts provide scale and lead-time tests. They do not receive the status of policy parameters, and their error bands are not averaged away. Persistent revisions are themselves evidence.

Portfolio implication is deliberately the weakest output. It is a research direction, a stress test, a business-model prior, a correlated-factor warning, or an abstention case. It is never a transaction instruction. The governing chain remains:

> *Human or social need -> funded payer -> adopted solution -> delivered volume and quality -> division of surplus -> operating-company cash after full reinvestment -> parent-company cash after all prior claims -> diluted per-share owner cash -> valuation -> portfolio role.*

Each arrow requires evidence. The chapter repeatedly distinguishes five ledgers that are often collapsed in thematic investing:

1. **Social value:** customer time, health, security, reliability, resilience, quality, productivity, and avoided loss, net of external costs.
2. **Operating profit:** revenue less the labor, purchased inputs, power, rent, taxes, depreciation, and other costs recognized at the operating entity.
3. **Parent cash:** cash actually available to the listed or investable parent after debt service, leases, working capital, maintenance and replacement at economic lives, minority interests, legal claims, subsidiary restrictions, acquisitions, and taxes.
4. **Diluted per-share cash:** parent cash divided by the fully diluted share count after stock compensation, convertibles, options, acquisition shares, recapitalizations, and recurring issuance.
5. **Investment and portfolio fit:** the return implied by the price paid, after tax, cost, liquidity, factor overlap, opportunity cost, and the liability plan.

A successful worldview should survive a socially successful but financially disappointing decade. AI may improve service while model rents collapse. Adaptation may lower losses while contractors earn poor returns. Defense may improve security while the state captures technical surplus. Formalization may broaden access while public rails compress private tolls. Health innovation may improve function while payers and patients retain most of the gain. This possibility is not a footnote; it is the central discipline.

The value-capture layer: creation is not ownership

The decade cannot be reduced to "capital versus labor." The possible claimants include workers with different skills and mobility; landowners and holders of power, water, mineral, permit, and connection rights; creditors and lessors; common equity; tax authorities and public payers; entrepreneurial founders; incumbent operating companies; entrants; platforms; qualified bottleneck suppliers; customers; and communities bearing external costs. Several may gain together if productivity enlarges the pool. Several may report growing activity while real welfare or per-share cash falls.

A social-value ledger can be written as:

incremental social value = customer surplus + labor income + land/resource rents + supplier surplus + entrepreneurial profit + incumbent profit + creditor income + taxes and public value - uncompensated external costs.

The operating-company bridge is narrower:

customer receipts - purchased inputs - labor - land, power and resource claims - cash tax - maintenance and replacement capital - working capital - interest and lease claims = operating-company residual cash.

The parent bridge then deducts minority interests, trapped cash, pension and legal claims, recurring acquisitions required to sustain the franchise, affiliate support, and capital controls. The share bridge adjusts for the fully diluted denominator. A company can therefore show higher revenue, EBITDA, and aggregate free cash while each continuing owner's economic claim stagnates. Buybacks are valuable only net of issuance, at a defensible price, and without weakening the balance sheet.

Six recurring determinants of bargaining power organize the analysis.

- • **Scarcity:** Is the needed input truly scarce at the time and location of demand, or merely expensive today?
- • **Substitutability:** Can the customer change supplier, technology, geography, chemistry, task design, or consumption without large cost?
- • **Mobility and coordination:** Can labor move or organize, customers aggregate procurement or multi-home, and entrants finance a response?
- • **Contract and law:** Who bears inflation, utilization, cancellation, liability, asset-life, and currency risk?
- • **Capital response:** How rapidly do high returns attract capacity, migration, automation, open standards, vertical integration, or subsidized entry?
- • **Political legitimacy:** Can a visible rent tied to an essential service, public subsidy, health need, power constraint, defense requirement, or natural resource survive tax, regulation, compulsory access, or renegotiation?

The most defensible capture architecture is usually narrow rather than grand: a qualified, failure-critical input that is a small share of customer value; a funded payer; slow qualification; favorable risk allocation; moderate capital need; transparent parent cash; and enough customer surplus to preserve demand and political legitimacy. The least defensible combines easy imitation, concentrated or unfunded payers, capital flooding in, short economic lives, long fixed-price projects, high leverage, opaque group structures, and visible rents dependent on public support.

Integrated value-capture matrix

Claimant	Likely 2026-2036 bargaining direction	Where power may strengthen	How the gain appears	Main leakage or reversal	Best indicators
Labor	Bifurcated rather than uniformly weak or strong	Licensed, physical, accountable, relationship, exception-handling, integration, and scarce technical roles	Real wages, benefits, schedule control, safer work, shorter training, or more autonomy	Automation, migration, training, offshoring, employer monopsony, fiscal payer caps, housing cost	Real compensation versus productivity; hours; vacancies; quits; entry hiring; labor share; rent-adjusted wages
Land, power, water, permits, and other rent claims	Stronger at constrained productive nodes; uncertain for generic assets	Firm power, interconnection, fiber, transport, water, zoning, legal title, mineral and concession rights	Rent, lease premium, tariff, royalty, connection payment, land appreciation	New infrastructure, rezoning, substitution, relocation, land-value tax, hazard repricing, stranded projects	Firm-MW rent; connection cost; industrial/housing rent; congestion; water price; property tax; utilization
Creditors, lessors, and senior capital	Larger prior cash claim, with wider default and repression tails	Seniority, durable collateral, floating/reset yield, covenants, guarantees, stable funding	Interest, fees, lease yield, residual guarantees, equity kickers	Inflation, funding-cost rise, default, restructuring, obsolete collateral, directed holdings, capital controls	Effective yield after loss; interest/lease coverage; maturity wall; covenant headroom; collateral vintage; recoveries
Common equity	Radically uncertain in aggregate	Self-funding, low replacement capital, real pricing power, transparent cash, disciplined allocation	Residual cash growth and buybacks or dividends funded from operations	Wages, input rents, tax, interest, capex, working capital, dilution, minority claims, parent leakage	Real mature-cohort ROIC; economic depreciation; parent cash; fully diluted cash/share; net issuance
State and taxpayers	Likely stronger as payer, regulator, tax sovereign, licensor, and insurer of last resort	Essential services, procurement, resource ownership, concessions, permits, rescue dependency, public rails	Tax, royalties, strategic capacity, security, resilience, healthier labor, public-service savings	Subsidy races, poor execution, contingent liabilities, corruption, weak collection, political turnover	Net fiscal benefit after subsidy/infrastructure; effective state take; appropriation-to-cash; interest/revenue; clawbacks
Entrepreneurial founders	More experiments, but highly unequal realized capture	Narrow insight, speed, proprietary lawful data, direct distribution, low capital need	Founder equity, new profit pools, option value, acquisition proceeds	Easy copying, preference stacks, repeated funding, platform dependence, regulation, incumbent bundling	Formation to employer conversion; cohort retention; acquisition cost; funding terms; founder dilution; cash break-even
Incumbent firms	Stronger where complements matter; weaker where legacy units are attacked	Installed base, system of record/action, trust, identity, distribution, compliance, balance sheet	Retention, price, service attach, lower implementation cost, defended franchise cash	Cannibalization, inertia, open interfaces, antitrust, obsolete assets, customer self-supply	Normalized retention; switching; attach; organic price/volume; incremental cash; replacement economics
New entrants	Stronger in narrow applications; weaker in capital-heavy nodes	Cheap intelligence, modular infrastructure, underserved workflow, open standards	New products, price competition, customer surplus, selective entrepreneurial rent	Distribution and compute tolls, trust, licenses, financing cost, incumbent copying, customer acquisition	Entry and survival; independent distribution; contribution after platform/compute; capital required to scale
Platforms	Strong if they control execution and trust; weaker if only discovery	Identity, payment, dispute resolution, reputation, workflow, transaction data, low multi-homing	Take rate, software, advertising, float, credit, logistics spread	Public rails, portability, agents, direct channels, regulation, supplier exit, privacy law	Take rate net of incentives/fraud; multi-homing; merchant margin; direct traffic; switching; counterparty savings
Qualified bottleneck suppliers	Strong early potential with a shortening clock	Long qualification, failure criticality, low share of system cost, scarce capacity, service base	Price/mix, backlog, deposits, recurring service and consumables	Capacity entry, standardization, substitution, customer integration, export controls, inventory correction	Lead time; cancellations; utilization; new capacity; customer concentration; incremental cash through normalization
Customers and consumers	Potentially largest cumulative beneficiary, often delayed	Procurement scale, portability, competition, open standards, abundant supply, ownership of context	Lower quality-adjusted price, time saved, reliability, access, new capability	Scarcity, personalized pricing, tariffs, lock-in, land/power constraints, rationing	Total cost of use; quality; wait time; switching; customer ROI; savings retained; access distribution
Communities and non-contracting parties	Highly variable; often residual risk bearers	Political voice, property rights, regulatory standing, local tax authority	Employment, tax base, infrastructure, reliability, environmental restoration	Bills, displacement, congestion, pollution, water use, hazard, stranded public assets	Local bills/tax receipts; emissions/water; outage; housing burden; health; compensation and remediation

The matrix describes a moving distribution. Early scarcity commonly rewards bottleneck suppliers, labor, land, and finance before the operating company proves a return. Buildout moves cash toward equipment vendors and

construction. Competitive flood transfers value toward customers and senior claimants. Regulation can socialize downside and cap upside. A later harvest may improve survivor cash only after capacity is removed, weak finance is closed, and management stops rebuilding the surplus. The first recipient is not necessarily the decade owner.

Twelve forces and one counterforce

The forces below are not independent sectors. AI, semiconductors, grids, automation, defense, localization, adaptation, resources, healthcare, and urban infrastructure can share the same skilled labor, power, sovereign payer, financing cost, East Asian fabrication nodes, or commodity inputs. Emerging-market platforms and balance-sheet franchises can share public payment rails, currency, regulation, and credit. The chapter therefore treats interactions and common-factor failure as part of each force.

Force 1 - AI and software diffusion: capability must become governed work

Starting evidence and epistemic limit

AI use had clearly moved beyond a specialist niche by the cutoff, but depth remained limited. The US Census Bureau reported that 18% of firms used AI in at least one function in November 2025-January 2026, or 32% when weighted by employment; most adopters used it in only a few functions and tasks, and 66% reported augmentation only. A Federal Reserve synthesis put work-related generative-AI use around 41% of workers in November 2025 but daily use near 12%. Firm adoption, worker exposure, task breadth, and use frequency are different denominators rather than contradictory facts (US Census Bureau, *The Microstructure of AI Diffusion*, May 2026; Federal Reserve, *Monitoring AI Adoption in the US Economy*, April 2026).

Narrow field experiments show real productivity in selected settings. A customer-support study found issues resolved per hour increased about 14%, with much larger gains for novice and lower-skilled workers; software-development field experiments found higher completed-task output but heterogeneous effects. These studies establish that some workflows have value. They do not establish general equilibrium, permanent employment effects, vendor capture, or economy-wide productivity (Brynjolfsson, Li, and Raymond, *Generative AI at Work*, NBER and QJE; Cui et al., *The Effects of Generative AI on High-Skilled Work*, August 2025). Falling capability-adjusted price is also observable: Stanford's 2025 AI Index estimated that the price of a fixed benchmark-level capability fell by orders of magnitude from 2022 to 2024. Benchmark progress remains narrower than generalized real-world performance, as NIST cautions (Stanford HAI, *AI Index 2025*; NIST, *Expanding the AI Evaluation Toolbox*, February 2026).

Why the force may persist

The structural case does not require uninterrupted frontier progress. Falling cost expands the feasible task set. Software can be replicated cheaply. AI can diffuse tacit practice, compress training time, and let scarce experts supervise more work. Incumbent distribution can place capability inside office suites, clouds, industrial controls, health workflows, devices, and systems of record. Aging and labor scarcity can improve automation payback. Defensive spending can continue when an incumbent must protect search, advertising, cloud, or software profit even if direct AI revenue is hard to isolate.

Persistence is also created by friction. Deep adoption requires permissioned data, process maps, integration, evaluation, training, incentive redesign, legal review, security, accountability, and exception handling. Those organizational complements move much more slowly than model releases. A decade-long diffusion curve can therefore coexist with a shorter, more violent infrastructure cycle. The constraint stack includes technical reliability, data rights and quality, workflow redesign, liability, verification, total economics, skills, budgets, trust, and worker cooperation. NIST identifies confabulation as a natural model risk in consequential settings; the European AI framework adds documentation, logging, oversight, robustness, and transparency obligations. Regulation can slow use while creating demand for testing, identity, audit trails, and monitoring (NIST, *Generative AI Profile for the AI

Risk Management Framework*, July 2024; European Commission, *AI Act regulatory framework*, updated August 2026).

Capital cycle and value distribution

The software layer can experience abundance before organizations become productive. Model performance may converge; open-weight alternatives, smaller models, caching, distillation, and customer routing can reduce price. The customer can then capture the gain even while usage rises. Foundation developers may be financing an expensive public good if each generation commoditizes the prior one before fixed cost is recovered. Durable vendor capture is more plausible where differentiated capability is combined with distribution, proprietary lawful context, workflow authority, identity, billing, reliable tools, and enough cash to finance the next generation.

Labor's result is likely to be polarized. Workers whose output is codifiable and easy to monitor may lose bargaining power or entry pathways. Workers who own relationships, licenses, exception judgment, physical execution, or legal accountability can become more valuable because AI expands the scale of their decisions. The novice-productivity result could compress within-role wage dispersion while increasing employer margin. Alternatively, training pathways could weaken if apprenticeship tasks disappear. The correct indicators are entry hiring, promotion, hours, real wages, task mix, errors, and output - not a count of exposed occupations.

Incumbents can capture through systems of action and distribution, yet seat-based software is vulnerable if a customer needs fewer employees or agents substitute for user licenses. Entrants gain a lower cost of invention but still pay for distribution, trust, liability, and compute. Platforms can copy generic features or bundle them at zero incremental price. Customers gain when they multi-home, own their context, demand outcome pricing, or receive AI as a table-stakes feature. The state gains through tax, procurement, safety obligations, competition policy, and possibly public compute or data access. Creditors and infrastructure providers can take a growing prior claim if the build is financed with debt and leases.

The full owner bridge must subtract inference and human review, recurring model development, sales, security, stock compensation, acquisitions, legal exposure, and the share denominator. Reported recurring revenue is not durable if retention is held by bundling and the economic billing unit is shrinking.

Strongest counterthesis, indicators, and model effects

The favorable productivity thesis fails if use remains shallow; error and review cost absorb savings; firms do not redesign production; and aggregate or sector productivity does not improve after reasonable lags. The producer-capture thesis can fail in a socially successful abundance scenario: capability becomes cheap, customers retain savings, and model or software margins normalize. A third path is concentrated capture, where a small number of workflow or distribution incumbents retain value while many foundation and application entrants destroy capital.

Constructive evidence includes recurring governed production across more functions; verified customer savings after integration and error; stable retention and price after seat normalization; compute cost per outcome falling faster than vendor price; and diluted cash per share growing after recurring development. Re-underwrite if daily intensity and task breadth remain flat for two years, customers repeatedly cancel after pilots, AI gross profit fails to cover incremental compute and development, or productivity appears only in vendor-defined activity.

Potentially helped models include trusted systems of record that become systems of action, verification and standards, cybersecurity and identity, domain workflows with proprietary lawful context, and integrators with measurable customer outcomes. Potentially harmed models include generic wrappers, undifferentiated information tolls, labor-heavy outsourcing priced by hours, and seat software unable to migrate to outcome economics. None is a security conclusion.

Force 2 - Compute, data centers, and semiconductors: a capital cycle with different clocks

Starting evidence and layered economics

Physical investment had moved ahead of mature proof of customer return. The IEA estimated capital expenditure by five large technology companies exceeded $400 billion in 2025 and could rise another 75% in 2026; AI-factory capacity had more than tripled in 18 months. The same report warned that financing needs had outgrown balance sheets alone, increasing sensitivity to capital-market sentiment and expected AI returns (IEA, *Key Questions on Energy and AI*, April 2026). Federal Reserve work described the observable economy as a buildout phase rather than broad labor displacement and warned that imported equipment can make gross investment overstate domestic contribution (Federal Reserve, *The AI Buildout and the Economy*, July 2026).

"AI infrastructure" hides at least six economics: frontier-model development; internal hyperscaler capacity used across cloud and incumbent franchises; merchant or neocloud rental capacity; accelerators, memory, networking, and advanced packaging; foundries and semiconductor equipment; and buildings, powered land, cooling, and connectivity. Each has a different owner, customer, asset life, pricing mechanism, financing structure, and residual risk.

Frontier developers face a tournament-versus-commoditization question. Hyperscalers can earn across infrastructure rental, APIs, software, advertising, search defense, and internal productivity, making attribution difficult. Merchant capacity carries high utilization, rental-price, customer-concentration, collateral, lease, and refinancing risk. Accelerators can capture a system rent through performance, software, networking, and supply, but annual cadence and customer silicon shorten the clock. Foundries and unique tools have long qualification, high fixed cost, and few substitutes, while state-backed duplication and simultaneous fab additions can later depress utilization. Data-center shells are valuable only as part of a bundle that includes firm power, permits, fiber, water or suitable cooling, tenant credit, and re-leasing protection.

Persistence and the capital-cycle response

The structural mechanism is continued digital demand, more complex models and workloads, and high qualification barriers in selected manufacturing layers. The cyclical mechanism is emergency reservation, shortage pricing, duplicate orders, non-cancelable commitments, optimistic asset lives, and financing against peak rents. The cycle can be written:

> *Scarcity -> customer forecasts -> supply reservations -> fab, packaging, memory, network, building, and power expansion -> debt, leases, vendor finance, and subsidy -> simultaneous capacity -> falling rental and token prices -> utilization dispersion -> cancellation, impairment, covenant stress, consolidation, and possible harvest.*

Different clocks prevent one synchronized top. Packaging can be tight while generic shells are abundant. Memory can tighten after accelerator supply improves. Powered sites can remain constrained after merchant compute prices fall. Foundry and equipment demand can rise because each geopolitical bloc duplicates capacity even while global utilization falls. Long qualification lets rents persist longer than in a simple commodity, while efficiency, custom silicon, or monetization disappointment can accelerate the reversal.

Economic depreciation is the central accounting test. Physical operation does not establish economic competitiveness. Evidence should include vintage rental price, utilization, energy efficiency, resale value, retirement, supported software, maintenance, and cash margin. A server operating after six years does not prove the accelerator inside it earned a return for six years. Contractual commitments, finance leases, residual guarantees, customer advances, and off-balance-sheet partnerships must be consolidated into the owner-cash bridge.

Who captures and who bears the loss

Bottleneck chip, packaging, memory, networking, equipment, and cooling suppliers can take early rent through price, deposits, and non-cancelable orders. Their own upstream constraint may be skilled labor, electrical steel, export approval, or customer financing. Hyperscalers can retain value if compute supports several revenue pools or defends a larger franchise at an economic cost below the rents protected. Merchant operators can show rapid revenue while lenders, landlords, chip suppliers, and anchor customers own stronger contractual claims. Landowners and utilities capture where power and connection are scarce. States subsidize strategic capacity, condition output, regulate exports, and later tax visible profit. Customers may be the durable winner if compute abundance lowers price more rapidly than volume expands.

Operating profit is especially poor evidence here. Parent cash must be measured after economic replacement, power and rent, leases, interest, purchase commitments, minority infrastructure vehicles, stock compensation, taxes, and the capital required for the next architecture. Diluted per-share cash can fall even if installed compute, reported depreciation, and revenue rise. Equity receives the residual after a tournament in which creditors often hold a fixed claim on assets whose economic life can shorten suddenly.

Countertheses, falsifiers, and model effects

The fiber-like counterthesis is socially successful abundance: capability and usage soar, but rental and hardware prices collapse, economic lives shorten, and builders lose. The scarcity counterthesis is the reverse: reasoning, video, agents, science, and industrial control absorb efficiency gains; memory, packaging, transformer, turbine, connection, and skilled-labor constraints persist. A geopolitical counterthesis has nominal orders rising while duplicate subscale capacity destroys owner returns.

Track rental price and utilization by vintage; contracted revenue concentration; cancellations; lead times; inventory and provisions; accelerator resale and retirement; cloud gross profit after economic depreciation; capex and contractual commitments versus incremental customer cash; fab yield and utilization; project starts versus energization; and interest plus lease coverage. A builder-rent thesis needs re-underwriting if rental prices fall sharply while utilization is below economic break-even, inventory charges and cancellations persist, or capex remains above operating-cash improvement. A commoditization thesis weakens if lead times, prices, utilization, and customer full-cost returns remain firm after qualified supply enters.

Potentially helped models include unique semiconductor equipment and advanced packaging during genuine qualification scarcity, recurring equipment service, power-control and cooling bottlenecks, and disciplined survivors in a confirmed harvest. Potentially harmed are leveraged merchant compute, generic data-center real estate without firm power, builders relying on long accounting lives, and subsidized duplicated capacity without commercial utilization. Leading accelerators, hyperscalers, foundries, and large data-center platforms remain radically uncertain because market power and reinvestment burden are both real.

Force 3 - Electricity, grids, and power: global abundance can coexist with local scarcity

Starting evidence and the correct unit of analysis

Data centers are material to electricity planning without dominating the global system. The IEA estimated data-center consumption rose to about 485 terawatt-hours in 2025, with AI-focused facilities growing faster, and projected roughly 950 TWh in 2030, near 3% of world electricity use. In advanced economies and particular US regions, the share of incremental demand is much larger. A single campus can resemble a power-intensive industrial complex, while substations, transmission, dispatchable capacity, water, and permits are local rather than globally fungible (IEA, *Key Questions on Energy and AI*, April 2026; IEA, *Energy and AI executive summary*, April 2025).

The economically relevant unit is therefore the node, substation, utility territory, balancing area, and hour - not global TWh. National generation abundance does not release a constrained interconnection. Average energy price does not describe the value of firm peak capacity. A nominal land allocation does not create a financed substation. A queue does not prove a customer will energize or use the requested megawatts.

The IEA reported more than 2,500 GW of generation, storage, and large-load projects in connection queues and estimated grid planning and construction can take 5-15 years, compared with 1-3 years for data centers. It projected annual grid investment would need to rise roughly 50% from about $400 billion by 2030 under its demand scenarios. Component prices and lead times had increased materially (IEA, *Electricity 2026: Grids*). Yet NERC's 2025 reliability work applied large realization discounts to requested loads in ERCOT: studied sites used about half requested capacity on average, delays were common, and only a portion energized. PJM reduced near-term forecasts after better vetting while retaining substantial longer-run growth (NERC, *2025 Long-Term Reliability Assessment*; PJM, *2026 Long-Term Load Forecast*). The queue is both constraint evidence and option-taking.

Structural mechanism, supply response, and possible equilibria

Electrification broadens persistence beyond AI. Industrial motors, electric vehicles, cooling, heat pumps, automation, housing, and new manufacturing add load. Replacement and resilience investment is needed even where energy demand is modest. Transmission, distribution, transformer factories, turbines, permitting, and skilled labor respond more slowly than many loads. Once installed, the network must be maintained and adapted for decades.

The dated demand forecasts remain wide. Efficiency per compute task can improve rapidly while heavier workloads and more volume offset it. Electric-vehicle paths differ by region. Industrial projects can relocate. Economic weakness can delay load. The structural thesis is therefore not a universal shortage. It is repeated local mismatches between fast load, slow networks, uncertain generation, and contract design.

Five equilibria can resolve the mismatch:

1. **Build-through:** credible customers sign long commitments, fund upgrades and generation, and support utilization.
2. **Flex-through:** non-firm service, storage, geographic routing, time shifting, dynamic ratings, reconductoring, topology optimization, and demand response add hosting capacity without every forecast mile of new grid.
3. **Price-through:** high energy and capacity prices ration low-value loads and reward flexibility or relocation.
4. **Delay-through:** speculative projects lose queue positions, customers cancel, and equipment backlogs unwind.
5. **Socialization-through:** utilities build for forecast demand that arrives late, and residential or small-business customers bear fixed cost until political intervention.

The marginal megawatt can come from renewable generation plus storage and network support, gas with fuel and turbine constraints, existing nuclear, new nuclear or geothermal with longer clocks, onsite generation, efficiency, or flexible load. No technology is a universal answer. The all-in firmed and connected cost, time to service, emissions and water obligations, and risk allocation determine customer value.

The grid opportunity is not synonymous with more regulated rate base. IEA analysis suggests operational and network measures could release substantial advanced-stage queued capacity, including non-firm connections and grid-enhancing technologies. These are modeled possibilities, not revenue. They imply that control, sensing, efficient conductors, storage, and contract design can weaken pure scarcity rents while improving system value.

Value capture: the rent can move faster than the electrons

Powered land, interconnection rights, water, permits, fiber, and generation access can capture first. A site premium may rise even if the downstream data-center operator earns an ordinary margin. Grid and equipment suppliers benefit from qualification, limited plants, failure severity, and long lead times, but capacity, standardization, and customer inventory respond. Skilled electricians, lineworkers, engineers, technicians, and construction labor can absorb part of project economics through wages and schedule delay.

Utilities can report rate-base growth without generating per-share value. Capture requires timely recovery, an allowed return above actual financing cost, competent construction, affordable bills, and financing that does not require persistent dilution. Merchant generators can receive scarcity prices while fuel, hedging, environmental obligations, market design, and windfall policy claim the gain. Public and cooperative systems can transfer more of it directly to customers.

Large flexible customers can capture through dedicated supply, interruption contracts, location choice, or self-build. Households and small businesses can lose if expansion costs are socialized. FERC's June 2026 large-load initiative emphasized faster integration, transparency, reliability, consumer protection, and cost-recovery agreements that protect other customers if a large load does not appear. That contract design can move value from site owners and tenants toward utilities and existing ratepayers (FERC, *Targeted Action to Speed Large Load Integration*, June 2026). Virginia's legislative research agency illustrated the affordability issue: current data-center tariffs could cover allocated present service while future generation and transmission still raised household bills in modeled scenarios (Virginia JLARC, *Data Centers in Virginia*, 2024).

The state owns permits, public land, regulation, taxation, and sometimes the utility. It can appropriate infrastructure-created land value or transfer it to private owners. Communities receive jobs and tax base but may bear bills, water use, emissions, noise, and stranded assets. Political legitimacy therefore becomes part of the rent duration.

Countertheses, falsifiers, and model effects

The scarcity thesis fails if requested loads remain phantom; realized-to-requested ratios stay low; forecast revisions continue downward; lead times normalize; regional prices fall; and new capacity remains financed. The construction-only thesis fails if non-firm contracts, reconductoring, storage, controls, and demand response connect large loads cheaply. The utility-capture thesis fails if rate-base growth produces disallowance, arrears, declining affordability, below-cost earned returns, or recurring equity issuance. A scarcity-abundance oscillation is also plausible: local shortages persist while global equipment supply overbuilds.

Track requested, vetted, contracted, financed, under-construction, connected, energized, peak-used, billed, and collected megawatts separately. Track hourly congestion, curtailment, firmed cost, queue withdrawals, customer deposits, transformer and turbine lead times, factory additions, large-load tariff terms, allowed versus earned returns, bill arrears, water use, and cash after replacement capital. The decisive test is not a growing queue but collected revenue and per-share cash from a delivered system.

Potentially helped models include qualified electrical equipment and recurring service, grid control and efficiency, reliability distribution, funded flexible generation or storage, and operators with strict cost allocation and balance-sheet discipline. Potentially harmed models include speculative powered-land claims, fixed-price construction, utilities with suppressed tariffs or catastrophic liability, low-value power-intensive users at constrained nodes, and equipment capacity built on duplicate orders. Regulated owners remain radically uncertain because essential demand and political appropriation are inseparable.

Force 4 - Electrification and industrial automation: diffusion through the last meter

Starting evidence and persistence

Electrification is broader than vehicles and AI. The IEA expected electricity use to grow materially faster than total energy demand through 2030, supported by industry, vehicles, cooling, heat pumps, and data centers (IEA, *Electricity 2026: Demand*). More than 20 million electric cars were sold in 2025, about one quarter of global sales, but regional shares ranged from very high in China to below 10% in the United States. Chinese output exceeded domestic demand and many models competed for a concentrated sales pool. Unit adoption, social benefit, and manufacturer profit already diverged (IEA, *Global EV Outlook 2026*).

Industrial-robot installations exceeded 500,000 for a fourth year in 2024 and were more than twice the level a decade earlier. Adoption remains concentrated in large firms and standardized industries because integration has high fixed cost. Observational Census evidence associates advanced-technology adopters with higher productivity, but productive firms may also be more able to adopt (International Federation of Robotics, *World Robotics 2025*; US Census Bureau, *Automation and the Workforce*, 2022).

The persistence mechanism is economic: automation is adopted when the risk-adjusted cost per acceptable unit beats labor and alternative capital over realistic utilization. Payback improves with wages, vacancies, dangerous work, quality cost, training difficulty, 24-hour throughput, and stable products. It weakens with volatile volumes, short runs, variable physical objects, safety constraints, scarce integrators, financing cost, maintenance, and unwillingness to redesign the process. AI can expand perception, planning, natural-language programming, simulation, predictive maintenance, and adaptive control. Physical reliability, grippers, sensors, certification, cycle time, integration, and service remain the last-meter constraint.

Capital cycle and distribution

Electrification manufacturing demonstrates how volume growth can invite destructive supply. Battery, vehicle, inverter, charging, and component plants respond to subsidies, regional policy, learning curves, and high growth expectations. Simultaneous capacity, chemistry changes, tariffs, retooling, warranty, and price wars can move value to customers while manufacturers consume capital. Grid and charging constraints can slow downstream use after plants are built.

Automation has a different cycle. Demonstrations attract venture and corporate budgets; integrators and components become scarce; pilot backlogs rise. Many pilots fail at uptime, exception handling, or customer process redesign. Project vendors can recognize milestones while working capital, guarantees, and installation delays consume cash. Successful systems create an installed base, shifting economics toward spares, software upgrades, consumables, maintenance, and integration knowledge. Open interfaces, right-to-repair, third-party service, or fleet retirement can then erode that base.

Labor can capture more of automation's productivity where scarce workers remain necessary for exceptions, safety, customer relationships, and physical work. Employers retain savings where tasks standardize and product markets are not competitive. Customers retain them where competition forces lower prices. Integrators capture when deployment skill is scarcer than hardware. Land and power claims absorb gains where automation enables more output at a constrained site. Creditors gain when equipment is financed, but collateral can be highly specific. The state may subsidize manufacturing, regulate safety, tax displacement gains, or support training.

The owner-cash test is site-vintage rather than order-book based: installation cash, uptime, throughput, quality, labor minutes, maintenance, warranty, service attach, and customer payback. Revenue growth is not evidence if cumulative project cash remains negative or the vendor must finance customers. For vehicle and battery manufacturing, per-share cash must include retooling, warranty, inventory, price concessions, dealer or customer incentives, tariffs, environmental obligations, and dilution.

Countertheses, falsifiers, and model effects

One counterthesis is rapid service-productivity improvement: AI and robotics remove enough labor minutes that customers gain capacity and lower cost, weakening wage and pure-scarcity stories. The opposite is integrator bottleneck: physical variability and liability keep adoption narrow, shifting value to field labor and installed-base service rather than general-purpose robot builders. A third is manufacturing abundance: electrification continues socially, but batteries, vehicles, and standard equipment earn poor returns.

Track customer payback by site and vintage; pilot-to-production conversion; uptime and contracted cycle time; labor minutes and quality per unit; integrator hours; warranty; service attach; customer concentration; backlog-to-cash; working capital; vehicle or battery utilization; inventory; price and mix; chemistry substitution; and economic returns after subsidies. A rapid-autonomy thesis fails when deployments repeatedly miss throughput and uptime or require unchanged staffing. A provider-capture thesis fails when customer savings exceed vendor price gains and competing systems reduce renewal or service economics.

Potentially helped models include installed-base automation service, qualified safety and controls, machine vision and sensing tied to outcomes, reliability distribution, and capital-light integrators with repeatable deployments. Potentially harmed are bespoke fixed-price integrators, legacy proprietary equipment without service or interoperability advantage, commodity battery or vehicle manufacturing, and labor-outsourcing models priced by hours. Automation manufacturers remain cycle-sensitive; growth in robots does not settle the return on their builders.

Force 5 - Aging, health, and labor: slow cohort arithmetic meets a hard payer gate

Starting evidence and geographic divergence

Demography is one of the more persistent forces because much of the 2026-2036 cohort structure is already known. The United Nations estimated world population at 8.2 billion in 2024, but one quarter of people lived in countries or areas whose population had already peaked, while many countries in South Asia and Africa were still growing rapidly. More than half of countries and areas had fertility below replacement. Low births today affect the distant labor force; the next decade is shaped mainly by smaller cohorts already born and large cohorts already approaching retirement (UN Population Division, *World Population Prospects 2024*).

The OECD old-age dependency ratio rose from 19% in 1980 to 31% in 2023 and is projected to 52% in 2060; its working-age population is projected to fall 8% from 2023 to 2060, with much larger contractions in some countries. These are scenarios beyond the chapter horizon, not 2036 forecasts, but they show the scale of participation and productivity response required (OECD, *Employment Outlook 2025*, demographic change chapter). Near-horizon age bands are less uncertain than distant population totals: a 55-year-old in 2026 is 65 in 2036 if alive.

There is no single global aging trade. Advanced East Asia faces compressed aging and limited labor time. Europe combines aging with social insurance and administered prices. Migration-supported economies can replenish some labor but face housing, licensing, and political constraints. Young countries have abundant potential workers but may lack education, power, formal jobs, and purchasing power. Rapid middle-income transitions can age before care systems and savings are complete. Chronological age alone is insufficient; healthy life expectancy, actual employment, hours, disability, tax contributions, and care dependency are the useful economic variables.

Why labor scarcity may persist without becoming provider profit

Retirements, care demand, licensing lead times, geography, and job quality can create shortages through soft economic cycles. The response can be higher wages, later work, participation, migration, task shifting, capital deepening, automation, lost output, waiting lists, burnout, or lower quality. The same country can have unemployment and severe shortages because skills and locations do not match.

Scarce nurses, caregivers, technicians, electricians, engineers, and construction workers gain bargaining power only against a funded and flexible payer. In care, the state or insurer is often a monopsony and tariffs adjust slowly. Scarcity can therefore produce vacancies and rationing rather than provider margin or wage. Informal care is part of the economy even when absent from corporate revenue: family members pay through time and reduced labor participation. Formalization of care can raise measured spending while merely replacing unpaid labor.

Participation and migration can materially change the outcome. OECD simulations show that later exit and greater participation among older workers, women, and underrepresented groups could offset much demographic drag. Migration moves faster than births but not frictionlessly; credentials, language, occupation, geography, housing, politics, and source-country effects determine whether it closes a shortage. Better health is productive capital because it reduces absence and disability, extends work, and frees caregivers. The value may accrue to the patient, employer, state, or family rather than the treatment provider.

Health demand, innovation, and the contested surplus

Aging changes the mix toward multimorbidity, medication management, diagnostics, rehabilitation, mobility, hearing and vision, home modification, community care, caregiver support, and end-of-life services. Many systems will try to move suitable work away from expensive acute facilities. That can help home and community models with density and labor retention while harming marginal facilities built on simple age curves. Demand is funded differently by tax systems, social insurance, employers, commercial plans, households, and informal caregivers.

Biomedical innovation remains active across targeted therapies, immune and inflammatory disease, cell and gene approaches, diagnostics, devices, and digital monitoring. A differentiated treatment with hard outcomes, enforceable intellectual property, difficult manufacturing, and a clearly identified high-risk group can capture temporarily. A diagnostic without a treatment pathway may not. A one-time therapy may face manufacturing, specialist-center, durability, and payment-timing constraints. A digital tool can scale cheaply but remain easy to copy and expensive to integrate.

Intellectual-property rent decays through competing mechanisms, negotiated prices, rebates, litigation, generic or biosimilar entry, and the replacement pipeline. Medicare's first price-negotiation cycle illustrates state bargaining: CMS estimated meaningful savings across selected high-spend drugs under negotiated 2026 prices, though counterfactuals remain contestable (CMS, *Negotiated Prices for Initial Price Applicability Year 2026*). FDA biosimilar approvals and proposed simplification point toward further transfer from reference products to payers and patients (FDA, *Accelerating Biosimilar Development*, 2025).

Obesity and metabolic medicines show the complete chain. WHO estimated more than 890 million adults lived with obesity in 2022. Trials have shown large weight reductions for incretin medicines, and SELECT found a 20% relative reduction in major cardiovascular events for a defined high-risk population, supporting an FDA cardiovascular indication (WHO, *Obesity and overweight*; New England Journal of Medicine, *SELECT cardiovascular outcomes*; FDA, cardiovascular indication). Clinical value does not settle broad access, adherence, long-run safety, net price, or budget impact. Withdrawal studies show substantial weight regain, supporting chronic use while magnifying payer cost. WHO's 2025 guideline remained conditional because of long-term, readiness, cost, and equity uncertainty, and expected supply to reach only a minority of potentially benefiting people by 2030 (WHO, global GLP-1 guideline announcement).

The value can divide among patients, employers, payers, clinicians, innovators, capacity suppliers, states, and later entrants. An illustrative CBO Medicare analysis found medical savings much smaller than direct drug cost over its horizon, showing why avoided events do not automatically finance broad premium-price access (CBO, budget effect of Medicare anti-obesity-drug coverage). Prevention can reduce procedure, dialysis, sleep-apnea, hospital, or chronic-care demand, moving value away from incumbents. Longer life can create other care. The payer's time horizon and beneficiary churn decide how much saving is collectible.

Capital cycles and value distribution

Visible age curves attract senior housing, clinics, facilities, home-care networks, diagnostics, manufacturing, software, and therapeutic investment. A software product can launch quickly; a specialist, biologic plant, or licensed facility takes years. Facility economics are levered to construction, property or lease claims, labor, and reimbursement. Occupancy can be high while agency labor and rent consume cash. Home care looks asset-light but can be working-capital and labor intensive. Acquisitions can assemble revenue without proving organic cohort return.

Biomedicine has long lags from discovery to trial, approval, payer access, and real-world outcome. Capital can bid up scientists, trial sites, manufacturing, and acquisition prices before successful supply expands. A platform can produce candidates while consuming more capital per approval than expected. A mature franchise can overpay to replace a patent cliff. No terminal value should bridge exclusivity by assumption.

Labor captures through scarcity wages and conditions where outside options exist. Landlords capture around hospitals, facilities, and productive cities. Debt and property owners receive senior claims. Providers retain savings only where reimbursement rewards efficiency rather than activity. Payers and the state can negotiate or reset tariffs. Patients receive much social value that never appears as provider revenue. Innovators retain rent only while evidence, access, differentiation, and exclusivity survive. Each share receives the residue after R&D, manufacturing, sales, rebate, tax, acquisitions, debt, stock compensation, and replacement pipeline.

Countertheses, falsifiers, and model effects

Healthy longevity can postpone disability, extend careers, and reduce age-specific utilization. Migration and participation can close worker gaps. Automation can produce a service-productivity surprise. Fiscal constraint can ration demand despite need. Innovation can add chronic cost without enough outcome. Competitive supply can destroy early therapeutic or facility rents. The modal path may be gradual aging, incremental innovation, ordinary margins, and no dramatic shortage.

Track healthy life expectancy, disability, participation at 60-74, effective retirement, migration by occupation, license processing, vacancies, real wages and productivity, agency labor, waiting lists, covered services, net revenue and labor minutes per quality-adjusted episode, payer denials, total cost of care, trial endpoints, real-world adherence and safety, indication-level net price, exclusivity, manufacturing utilization, facility starts, occupancy, rent and interest coverage, and diluted cash per share. A provider thesis fails if reimbursement lags labor and property or if demand rises while cash conversion deteriorates. An innovation thesis fails if hard outcomes do not reproduce, access remains negligible, net price falls faster than volume, or the replacement pipeline consumes the franchise.

Potentially helped models include funded health innovation with hard outcomes, care workflows that remove verified scarce labor minutes, accountable primary or risk-bearing models able to retain savings, home/community delivery with density and retention, diagnostics tied to action, and specialized capacity during real shortage. Potentially harmed are undifferentiated labor-heavy providers under fixed tariffs, facilities built on age curves alone, digital health without payer savings, mature IP without replacement, and market-size stories that ignore affordability. Broad metabolic treatment, risk-bearing care, disease-modifying therapies, and clinical AI remain radically uncertain.

Force 6 - Fiscal activism, defense, and the state: funded demand with a monopsony

Starting evidence and persistence

The state is becoming more important as customer, financier, regulator, insurer, and residual claimant. Aging, defense, strategic autonomy, climate, energy, and debt service compete for the same tax base. Fiscal activism includes grants, tax preferences, guarantees, subsidized credit, procurement, local-content rules, and state-owned balance sheets used to change what and where the economy produces. Its persistence comes from security and distributional objectives that private markets do not fully price, and from constituencies created by supported projects. Its constraint is narrowing fiscal space.

Industrial policy can improve targeted capability while imposing cost and spillovers. IMF work found moderate average gains in targeted comparative advantage but uneven results, negative effects in some untargeted sectors, and significant fiscal cost, with better outcomes when interventions address clear upstream distortions and institutions constrain capture (IMF, *Industrial Policies - Handle with Care*). A subsidy can bring capacity forward, solve coordination, or finance learning. It can also induce simultaneous supply and lower later returns. The analytical error is capitalizing the subsidy twice - as higher demand and as undiminished terminal margin.

Defense provides the clearest state-procurement case. SIPRI estimated world military spending reached $2.887 trillion in 2025, the eleventh consecutive annual increase and 41% above 2016 in real terms. NATO's 2035 commitment, European financing and procurement initiatives, depleted inventories, long platform lives, and qualification can support persistent demand (SIPRI, *Trends in World Military Expenditure 2025*; NATO, defence investment and the 5% commitment). Yet budgets are not delivered output. GAO's 2025 assessment found large cost growth and schedule delay across major US programs, showing backlog can mean scarcity, commitment, poor execution, or all three (US GAO, *Weapon Systems Annual Assessment 2025*).

Capital cycle and procurement chain

Fiscal activism passes through policy scarcity, subsidized buildout, competitive flood, and rationalization or permanent support. Grants lower private risk; guarantees protect lenders; tariffs and local content reserve demand. Early surplus goes to land, construction, equipment, labor, and finance. When several jurisdictions support the same capacity, global supply can rise while each bloc perceives local shortage. If weak projects are kept open for employment or security, the harvest is delayed and taxpayers fund persistent excess capacity.

Defense has an explicit chain: threat assessment; capability target; budget authority; appropriation; program; contract; qualification; production; acceptance; payment; installed-base support; parent cash; and share. Contract type decides the split. Cost-plus can stabilize operating profit while caping upside and rewarding a larger cost base. Fixed-price development can create large losses. Customer advances can make early cash look strong and reverse during delivery. The state can fund development, control exports and data, audit cost, block mergers, and remain the only buyer. Technical scarcity is compatible with monopsony.

Capacity expansion attracts entrants, factories, shifts, working capital, and private finance. Input inflation can cause nominal budgets to buy less real equipment. Standardization and collaborative procurement can deepen runs; national preferences can fragment them. A ceasefire or political change can slow procurement before specialized assets earn their cost. The NATO 2029 review is a milestone, not proof that the 2035 ratio becomes contracts.

Distribution of surplus and the per-share test

The state receives security, resilience, employment, tax, and strategic option value. Those benefits are partly nonfinancial and should not be inferred from supplier profit. Taxpayers bear debt, tax, inflation, or displaced services. Labor captures through cleared engineering, welding, software, skilled manufacturing, compliance, and construction shortages. Land and infrastructure owners capture at yards, depots, ranges, ports, and power-connected sites. Qualified component suppliers across multiple programs may have better economics than prime integrators carrying fixed-price risk. Lenders receive protected claims and guarantees. New entrants can win innovation awards but face long sales cycles, qualification, and a single customer.

Prime contractors retain value only if escalation, execution, volume, and negotiated margin exceed cost and working-capital burden. A large backlog is not enough. The research bridge should include funded appropriation, cancellation rights, program-vintage cost-to-complete, advances, acceptance, receivables, pension and environmental claims, fixed-price exposure, acquisitions, restricted cash, and diluted per-share cash. State support can lower funding cost while obligations, government warrants, local investment, tax, or pricing appropriate the upside.

Entrepreneurial effects are two-sided. Contestable milestone grants, first-customer procurement, and shared infrastructure can solve coordination failures. Bespoke incumbent support, complex eligibility, collateral requirements, and protected relationships crowd out young firms. A policy labeled pro-investment can socialize favored finance while raising the hurdle elsewhere.

Countertheses, falsifiers, and model effects

Fiscal constraint may dominate activism as interest, pensions, and care consume budgets. A peace dividend or political turnover can reverse defense. State demand can succeed socially and still destroy shareholder value through monopsony, taxes, labor and land inflation, and duplicate capacity. Conversely, standardized multiyear procurement, customer advances, disciplined contracts, and recurring installed-base service can support real cash.

Track announced, authorized, appropriated, awarded, advanced, delivered, accepted, billed, and collected amounts separately. Monitor real equipment budgets, program quantities, cost and schedule variance, contract mix, cancellations, collaborative procurement, input wages, plant utilization, grants paid versus announced, commercial utilization after subsidy, net fiscal benefit, procurement concentration, and diluted cash per share. Re-underwrite if targets do not become contracts, cost outruns budget, backlog cash conversion weakens, supported plants require recurring subsidy, or state take rises while incremental return falls below funding cost.

Potentially helped models include diversified qualified components, legitimate testing and certification, maintenance and spares attached to a net installed base, reliability logistics, production software that shortens cycle time, and dual-use products with private as well as public payers. Potentially harmed are one-program fixed-price developers, facilities built on aspirations without funded orders, tax-credit capacity without unsubsidized economics, roll-ups capitalizing backlog, and businesses carrying slow public receivables. National champions remain radically uncertain because protection and extraction are two sides of the same relationship.

Force 7 - Fragmentation and supply-chain resilience: selective duplication within continued interdependence

Starting evidence and a more precise thesis

"Deglobalization" is too blunt if it means international exchange is ending. WTO evidence showed sharply higher tariff and policy activity in 2025-2026 and a decline in the share of merchandise trade conducted at most-favored-nation rates, yet most trade remained under the multilateral framework. Goods and commercial-services trade reached about $34.9 trillion in 2025, with services at a record share; trade continued to reroute and grow even amid disruption (WTO, *Global Trade Outlook and Statistics*, March 2026; WTO, *World Trade Statistics 2025*).

The structural hypothesis is **selective fragmentation plus redundant interdependence**. Strategic technologies, minerals, data, capital, infrastructure, defense inputs, pharmaceuticals, and semiconductors face more export controls, investment screening, sanctions, local-content rules, stockpiles, and security procurement. Ordinary goods and services continue to trade. Firms optimize a portfolio of cost, legal eligibility, recovery time, inventory, route, and geopolitical exposure rather than lowest landed cost alone.

The persistence mechanism is insurance. Policy protects against a low-frequency disruption, so it can remain even when near-term cost is visible. Bureaucracies, control lists, screening systems, standards, and political constituencies survive individual shocks. Concentrated refining, fabrication, shipping, cloud, and data nodes make vulnerability salient. Rules can also adjust rather than only ratchet: licensing can replace bans, allied sourcing can replace domestic sourcing, and political cost can narrow controls.

Resilience is not autarky

The OECD's 2025 resilience review estimated that full relocalization could materially reduce trade and global GDP and increase volatility in many economies. It distinguishes agility, adaptability, and alignment from self-sufficiency (OECD, *Supply Chain Resilience Review 2025*). Reshoring protects against some foreign interruptions while concentrating domestic disaster, labor, and energy risk. Friend-shoring can create a new bloc concentration. Diversification adds qualified sources and routes. Visibility and modular design reduce recovery time without replicating every asset. Stockpiling buys time but ties capital, risks obsolescence, and amplifies bullwhip ordering.

The efficient response depends on failure cost, product life, qualification time, inventory decay, substitutes, and who pays. A small critical component with a two-year qualification can justify expensive redundancy; a common input with many sources may not. Resilience should be measured by recovery time and cash loss under defined failures, not domestic-content share.

Capital response and distribution

Tariffs, subsidies, minimum prices, offtake, and public equity can finance duplicate mines, refineries, fabs, battery plants, logistics routes, inventories, payment systems, and factories. Early benefit flows to construction, equipment, skilled labor, land, power, logistics, compliance, and qualified incumbents. Simultaneous projects later reduce price and utilization. Strategic foreign incumbents may keep producing below entrants' full cost, prolonging the flood. Domestic share can rise while the protected plant earns poor returns.

The state receives strategic option value and tax base but pays subsidy, procurement premium, administration, retaliation, and contingent liability. Customers pay higher price during duplication and may later receive lower price or reduced disruption loss. Labor gains local demand but can lose purchasing power to tariffs and suffer closure later. Land, grid, and logistics owners often capture first. Lenders finance inventories and duplicate capital while bearing policy reversal and obsolete collateral. Platforms for customs, testing, trade finance, and traceability collect a complexity toll until customers insource or standards commoditize it.

Incumbent suppliers gain qualification but invite price caps, compulsory access, public competitors, or windfall taxes. Entrants receive support but begin at high replacement cost and often lack scale. Entrepreneurs outside preferred sectors face higher rates and scarce labor; entrepreneurs that enable substitution, verification, modularity, or recovery can gain. Minority owners receive the residual after localization capex, working capital, tariffs on inputs, compliance, tax, financing, currency, related parties, and dilution. National importance is not shareholder bargaining power.

Counterthesis, falsifiers, and model effects

The strongest reversal is cheaper resilience through diversified global sourcing, interoperable standards, faster customs, shared stockpiles, reciprocal investment, and narrow controls. Trade may adapt rather than retreat. Political coalitions may reverse broad barriers when consumer prices, retaliation, and fiscal cost become visible. Technology can substitute away from a controlled input.

Track trade under multilateral terms, tariff-covered flows, control lists, screening, local-content conditions, supplier concentration, cross-bloc investment, route diversity, inventory days, total landed cost, recovery time, grants and capacity by jurisdiction, mature utilization, export-market loss, and commercial cash after subsidy. The fragmentation thesis weakens if restrictions narrow, international diversification reduces concentration, recovery improves without permanent support, and localized plants compete after protection ends. It strengthens operationally if legal eligibility becomes a durable procurement factor and multiyear funded capacity operates efficiently. Shareholder capture still requires a separate cash test.

Potentially helped models include multi-jurisdiction logistics, customs and trade-finance rails, testing and certification, traceability, modular and substitute designs, maintenance and reliability distribution, and scarce processors with diversified feedstock and customers. Potentially harmed are single-source importers, low-margin assemblers exposed to input and output tariffs, exporters concentrated in one contested market, working-capital-heavy stockpilers, and subscale protected plants. Domestic champions, strategic miners, and regulated infrastructure remain uncertain because state support and state appropriation rise together.

Force 8 - Inflation, real rates, debt, and financial repression: a wider regime distribution

Starting evidence without a one-way forecast

The strongest defensible claim is higher uncertainty and supply sensitivity, not permanently high inflation or a permanently high neutral real rate. Inflation had partly normalized before a new 2026 energy disruption. The World Bank's June 2026 forecast raised global inflation to 4.0% for 2026 while lowering growth, explicitly conditional on the shock (World Bank, *Global Economic Prospects*, June 2026). The ECB's 2025 strategy review argued that geopolitical fragmentation, climate and nature risks, digitalization, and demography can make supply disturbances more frequent, while acknowledging two-sided effects on equilibrium real rates (ECB, monetary-policy strategy overview, 2025). Productivity, saving, safe-asset demand, aging, risk premium, and investment can push real rates in opposite directions.

Debt service is a clearer claim on resources. The IMF placed global gross government debt just under 94% of GDP in 2025 and projected 100% by 2029 under then-current policy; interest payments had risen from about 2% to nearly 3% of GDP in four years. Social needs, defense, strategic autonomy, and aging were persistent pressures (IMF, *Fiscal Monitor*, April 2026). No universal debt threshold causes crisis: currency, maturity, investor base, growth, primary balance, and credibility matter. Interest nevertheless transfers tax revenue or displaces services before common equity receives anything.

Sovereign markets and finance are coupled through banks, nonbanks, repo, derivatives, collateral, pension and insurance portfolios, and central-bank operations. The BIS describes a fiscal-financial-stability nexus in which higher yields can increase fiscal risk, cause mark-to-market loss, reduce intermediary risk capacity, and

simultaneously pay more income to bondholders (BIS, *Annual Economic Report 2026*, Chapters I-II). The net demand and inflation effect is state-dependent.

Structural mechanisms and alternative regimes

Refinancing is slow but cumulative. Long maturities delay the budget effect of higher yields while placing duration risk on holders; short maturities transmit cash cost quickly. Structural spending claims overlap. Governments often support nominal household income after supply shocks, preserving demand while real supply remains constrained. Fragmentation duplicates capital, which can be inflationary during construction and deflationary in harvest. Prices, wages, tariffs, taxes, and coupons reset at different speeds, so the initial winner can lose when contracts catch up.

At least five coherent regimes should remain open:

- **Productive normalization:** AI, automation, energy, health, and supply redesign raise capacity; inflation anchors; mildly positive real rates coexist with better productivity; debt stabilizes through growth and selective reform.
- **Fiscal repair with expensive capital:** taxes or spending restraint stabilize inflation, but real yields and term premia remain above 2010s troughs because sovereign supply and uncertainty stay high.
- **Inflation-repression grind:** primary adjustment is insufficient; repeated inflation, directed saving, negative real funding, price control, and taxation allocate the loss.
- **Fiscal dominance or sovereign repricing:** term premium, currency risk, or spreads rise abruptly; central banks face price-stability and market-functioning conflict; adjustment may involve austerity, inflation, control, or restructuring.
- **Debt-deflation and theme bust:** strategic capex fails to produce enough demand or productivity; excess capacity, credit losses, and unemployment lower inflation while real debt burdens rise.

The chapter does not assign probabilities. Regimes can occur in sequence or across different jurisdictions. A productive US technology cycle can coexist with repression elsewhere; a commodity inflation shock can precede a capacity bust; fiscal repair can preserve selected defense or climate programs while crowding other activity.

Distribution: inflation is a bargaining process

Labor captures when skills are scarce, contracts index, unions coordinate, and productivity supports wage gains. Labor loses when wages reset slowly, employment weakens, or migration and automation expand effective supply. A nominal raise below household inflation is not capture. Land, power, and resource claims can reprice quickly when new investment needs scarce sites and connections. Lenders gain higher coupons only if inflation, funding, default, and regulation do not overwhelm them. Floating-rate credit transfers immediate stress to the borrower and can turn spread into loss.

The state gains through bracket creep, nominal-income and consumption taxes, royalties, seigniorage, regulated yields, and captive demand. It also pays indexed benefits, wages, subsidies, and higher refinancing cost. Surprise inflation's fiscal benefit diminishes as debt and obligations reset. Incumbents can reprice faster than wages or suppliers early; later wage catch-up, entry, substitution, price caps, and windfall tax erode the rent. Customers win when supply and productivity lower quality-adjusted prices, and lose when bottlenecks and tariffs pass through.

Financial repression includes policies that lower public funding cost or direct savings: administered rates, reserve and liquidity preferences for sovereign debt, balance-sheet support, directed lending, pension or insurance allocation, capital controls, dividend restrictions, subsidized credit, or prudential treatment. Some have stability or development rationales. Their incidence can favor the treasury, protected borrowers, state enterprises, selected sectors, and intermediaries while burdening depositors, pension beneficiaries, insurers, unpreferred entrepreneurs, importers, and currency holders. IMF research released in 2026 found historical captive government-bond demand associated with high debt, restricted mobility, and crowding out of private credit; it supports a regime possibility, not a forecast (IMF Working Paper, *The Coming Great Repression?*, July 2026).

Entrepreneurship can benefit from directed strategic finance and first-customer procurement but lose when captive banks have less capacity for unpreferred small firms, collateral dominates intangible value, subsidized incumbents bid up talent and sites, or capital controls reduce exit. A pro-investment label is not enough; contestability, sunset, milestone discipline, and broad private credit matter.

For common equity the important bridge is real, not nominal: realized price and volume; wage and input reset; inventory and receivables; replacement-cost capital; cash tax; effective interest; leases and pensions; parent access; net shares. Nominal EBITDA can rise while real diluted cash falls. For creditors the relevant return is interest actually received after funding, default, inflation, and control.

Countertheses, falsifiers, and model effects

The 2020s shocks may prove cyclical as energy expands, trade adapts, participation improves, AI raises productivity, and credible policy anchors expectations. Repression may be limited by independent institutions, mobile capital, and transparent pension rules. Fiscal repair may dominate activism. The opposite tail is sovereign repricing, where supposedly safe balance sheets become correlated and neither high coupons nor regulated returns protect equity.

Track inflation dispersion and expectation tails; real forwards and neutral-rate ranges; primary balance, interest-to-revenue, maturity and auctions; captive sovereign shares; real deposit and pension return; directed credit; private-credit breadth; corporate effective interest; debt-service ratios; wage and unit-profit contribution; replacement cost; state take; and real diluted cash per share. The persistence thesis weakens after a full five-year period of anchored expectations, low dispersion, fiscal repair, falling interest burden, market-priced savings, and broad private investment. A repression thesis strengthens when captive holdings and controls rise while real saver returns and private credit fall.

Potentially helped models include short-cycle necessary distributors, businesses with short repricing and low refinancing need, sound balance-sheet franchises with disciplined funding and credit, and ordinary companies with real pricing supported by customer value. Potentially harmed are long-duration fixed-price obligations, leveraged assets, captive finance, nominal-growth stories with high replacement capital, and visible scarcity rents exposed to tax. Resource owners and regulated infrastructure remain uncertain because a correct inflation view does not determine state take, rate recovery, or valuation.

Force 9 - Natural resources and the harvest phase: depletion is not a shortcut to owner return

Starting evidence and the scarcity distinction

Global material throughput remains large and rising. UNEP reports extraction of biomass, fossil fuels, metals, and non-metallic minerals tripled over five decades and could rise further under its baseline. Electrification, grids, cities, defense, data centers, housing, and water systems require physical inputs (UNEP, *Global Resources Outlook 2024*). Yet "resources" is not one force: geology, biological regeneration, substitution, recycling, inventories, processing, law, and market structure differ.

Reserves are dynamic economic and legal categories, not fixed geological inventories. USGS defines reserves as the economically and legally extractable portion of identified resources at a date; price, technology, exploration, recovery, permit, and law change the boundary (USGS, *Mineral Commodity Summaries 2026*). Dividing reserves by production and declaring an exhaustion date is unsound.

The stronger mineral thesis is constrained response and concentrated processing. The IEA identified potential mid-2030s gaps for selected materials if projects fail to arrive, falling copper grades and long lead times, but also recorded lithium's sharp price collapse after supply expanded and a decline in critical-mineral investment in 2025. The leading refiner dominates many energy-related minerals, creating vulnerability without literal exhaustion (IEA,

Global Critical Minerals Outlook 2025; IEA, *Global Critical Minerals Outlook 2026*). A concentrated low-cost incumbent can both create a security concern and destroy the return on higher-cost diversification.

Biological resources have stock-flow constraints. Agriculture uses much global freshwater and faces degraded soil and water systems. Wild capture fisheries have broadly plateaued while aquaculture supplies growth, transferring constraint toward feed, disease, genetics, water, and energy. Forest area can improve in aggregate while primary forest, biodiversity, or local quality deteriorates (FAO, *State of the World's Land and Water Resources 2025*; FAO, *State of World Fisheries and Aquaculture 2026*; FAO, *Global Forest Resources Assessment 2025*).

Persistence, capital cycle, and verified harvest

Selected resource constraints can persist because demand, long project clocks, concentration, by-product supply, qualification, ecological regeneration, community consent, power, water, and permits adjust at different speeds. Lower grade can increase rock moved, energy, tailings, capital, and remediation per unit. Structural does not mean prices rise steadily. Construction, China, inventory, credit, substitution, scrap, chemistry, recovery, and state-supported supply are cyclical or technological counterforces.

The capital cycle has four states:

1. **Scarcity:** inventories fall, utilization and realized price rise, and producers debottleneck.
2. **Buildout:** exploration, project approval, equipment, financing, subsidy, and entry accelerate; input inflation claims part of the rent.
3. **Competitive flood:** supply arrives after demand slows; inventories, discounts, impairments, covenant stress, and state support rise.
4. **Consolidation or harvest:** uneconomic capacity actually closes, weak firms cannot refinance, survivor capex is disciplined, utilization recovers, balance sheets survive, and cash is returned after full sustaining and closure cost.

Harvest is often declared too early. A lower share price, accounting impairment, merger, or dividend does not remove capacity. Public loans, employment mandates, customer prepayments, covenant waivers, or local policy can keep marginal supply alive. Capex below accounting depreciation can omit replacement cost, stripping, reserve development, tailings, water, remediation, and closure. A dividend funded by debt or asset sale is not harvest.

Verified harvest needs net capacity exit in the relevant delivered market, closed financing for weak operators, utilization and price after inventory normalization, and per-share cash after sustaining capital, reserve replacement where necessary, reclamation, royalties, taxes, interest, working capital, pensions, community obligations, dilution, and related parties. Renewable harvest also requires productive stock to stabilize: biomass, catch per effort, forest increment versus removal, aquifer recharge, soil quality, and regeneration. Profit from liquidating natural capital is not a franchise.

Value distribution

The scarce input may be the deposit, land, water right, permit, engineer, processor, port, power connection, shipping route, customer qualification, or finance rather than the mine. The state often owns subsoil resources and collects royalties, tax, rent tax, export duty, production share, or state-company dividends. Fiscal terms can become more progressive when prices and profits are visible. Land, water, and concession owners take access claims. Communities can receive employment, infrastructure, royalty, or equity and can delay projects where consent and environmental cost are unresolved.

Labor and specialist suppliers capture during buildout through wages, equipment, drilling, haulage, explosives, maintenance, and engineering. Financiers and offtakers take secured debt, streams, royalties, prepayments, hedges, and restrictive contracts ahead of common equity. Processors, ports, pipelines, assay and certification may own a steadier bottleneck. Customers benefit from competitive supply, substitution, efficiency, and recycling,

and policy may deliberately transfer value downstream by supporting uneconomic domestic capacity. Controlling owners can capture through private contractors, management agreements, asset transfers, or group structures. Minority owners receive the final residue.

Countertheses, falsifiers, and model effects

The principal counterthesis is abundance through technology, substitution, recycling, efficiency, and policy-supported supply. Higher physical demand can coexist with low producer returns, as the battery-material cycle showed. Demand can disappoint before long-lead supply arrives. Scarcity can also occur while state, labor, land, suppliers, or customers capture it.

Track normalized inventories, utilization, lead time and realized price; project sanctions and financing; grade and recovery; substitution and recycling; net closures versus starts and restarts; state support; sustaining and closure cost; reserve and biological stock; fiscal terms; labor and supplier inflation; debt and dilution; parent cash. A constraint thesis fails when additions, chemistry, recycling, or lower intensity close the delivered gap. A harvest thesis fails when state or cheap capital preserves capacity, dividends rely on borrowing, or 24 months pass without net exit and owner cash.

Potentially helped models include genuinely low-cost long-life assets with clear title and conservative finance, diversified scarce processors, mandatory testing and traceability, installed-base service, efficient recycling with feedstock rights, and water/material efficiency with short customer payback. Potentially harmed are high-cost single assets, peak-price projects, unpriced remediation, fixed-price processors with floating inputs, weak biological-stock managers, and equipment suppliers relying on every announced project. Producers remain radically uncertain by cost curve, jurisdiction, state take, cycle, and valuation.

Force 10 - Emerging-market formalization and digital platforms: access can outrun owner capture

Starting evidence and a decade mechanism

The World Bank's Global Findex 2025 reported that 79% of adults globally owned an account and 86% owned a mobile phone, with access gaps by income and gender. At the same time, the ILO estimates about 2 billion workers and 80% of enterprises remain informal. Digital access is real; comprehensive formality is not (World Bank, *Global Findex 2025*; ILO, transition to formality). An account can be dormant. A wallet user can cash out immediately. A merchant can be visible to a platform but not the tax or labor system. Registration can occur without filing, protection, enforceable contract, or formal finance.

Formalization may persist because identity, instant payment, invoices, transaction records, payroll, procurement, cloud tools, logistics, e-commerce, and government transfers reinforce one another. A merchant with a reliable transaction history becomes easier to underwrite. A worker with recognized identity and contributions can access protection. Formal buyers prefer traceable suppliers. Urban density reduces some service costs. The state gains an observable tax base and a more efficient transfer rail.

The counterforce is immediate compliance cost against uncertain benefit. Firms resist when tax, reporting, labor, and licensing cost exceed access to finance, procurement, or customers. Workers can lose benefits without receiving reliable protection. Arbitrary enforcement, surveillance, scams, cyber loss, outages, and weak dispute resolution reduce trust. Formalization is not a binary transition but a series of enforceable relationships.

Population and national growth are also insufficient. Frontier markets house a large and growing population but account for a much smaller share of output, with slowing per-capita growth and frequent sovereign default in recent decades. A young population creates opportunity and financing need, not convergence or hard-currency owner return (World Bank, *Frontier Market Economies 2026*).

Platform and financial capital cycles

Platforms often start with subsidized growth. Equity funds consumer discounts, free delivery, merchant incentives, driver guarantees, warehouses, and unsecured credit. Gross merchandise value, users, payment volume, and originations can rise while contribution after incentive, fraud, loss, logistics, and central cost remains negative. Funding tightens, promotion falls, credit cohorts season, and weaker firms exit. Consolidation can improve density and pricing while inviting fee caps, interoperability, labor rules, antitrust, or public alternatives.

Fast payment systems in more than 100 jurisdictions increasingly behave as interoperable public or regulated rails. BIS evidence associates broader adoption with public ownership, nonbank participation, broad use cases, and cross-border links. Socially, this lowers transaction cost and expands participation. Economically, it can commoditize the payment toll and move private value toward merchant software, deposits, fraud control, credit, logistics, cross-border service, and workflow (BIS, *Fast payments: design and adoption*).

Digital credit is an especially dangerous apparent harvest. Early cohorts are immature, nominal growth masks drift, fraud evolves, and foreign funding appears cheap before depreciation. Full-cycle evidence requires seasoned loss, collections, funding currency, deposit behavior, regulatory capital, customer acquisition, and parent cash. Logistics and warehouse assets can grow faster than route density. A closed wallet can lose its toll when public rails, portability, or open banking expand.

Distribution and the eight bridges to each share

Consumers receive lower search, payment, and delivery cost, access, and selection. Merchants gain reach, records, procurement, and credit while paying commissions, advertising, fulfillment, software, tax, and finance. Workers gain demand, payment certainty, and reputation while platforms may control pay, ranking, and deactivation. The state gains visibility, tax, and benefit delivery and can impose local data, ownership, credit, interoperability, or fee rules. Banks, telecoms, and infrastructure capture deposits, float, lending spread, identity, connectivity, and settlement.

Platforms capture take rate, advertising, logistics, software, data, and credit only if density lowers cost or improves outcomes and users cannot easily multi-home. Land and infrastructure owners capture warehouses, retail, towers, fiber, transport, housing, and power. Foreign-currency lenders can be protected while local equity absorbs depreciation. Controlling shareholders can capture through private suppliers, guarantees, brand or data ownership, intercompany loans, and asset transfers. Minority protection, beneficial ownership, and related-party disclosure are therefore causal, not merely governance decoration (G20/OECD, *Principles of Corporate Governance 2023*).

The GDP-to-share bridge has at least eight stages: real national output to nominal sector revenue; sector revenue to addressable platform transactions; transaction value to net take; net revenue to contribution after incentives and fraud; contribution to operating cash after central cost; operating cash to free cash after logistics, software, credit, and regulatory capital; operating-company cash to parent cash after tax, controls, currency, and minorities; and parent cash to each diluted share. Population and GDP help only the first bridge.

Countertheses, falsifiers, and model effects

Formalization can stall because compliance is costly and benefits unreliable. Public identity and payment rails can commoditize private tolls. Digital access can accelerate fraud and unsecured credit faster than supervision. Data and licenses can reinforce state or incumbent control rather than broad entry. Currency and funding stress can erase local nominal growth, and cash can remain trapped.

Track active rather than opened accounts; digital merchant use; savings and payroll; invoices and filing; procurement and worker protection; contribution after incentive and fraud; route density; multi-homing; take rate; credit loss by seasoned cohort; funding currency; regulatory capital; public-rail adoption; related parties; guarantees; parent cash; diluted shares; and hard-currency per-share value. A platform thesis fails when subsidy returns, take rate falls without cost offset, public rails erase the toll, credit seasons poorly, or reported profit cannot

be upstreamed. A formalization thesis fails when access rises without productivity, trust, compliance, or worker benefit.

Potentially helped models include merchant workflow above low-cost rails, identity and fraud infrastructure, tax/accounting/compliance systems, disciplined logistics density, transparent deposit franchises, marketplaces with organic cohorts, and platforms with simple audited ownership and parent cash. Potentially harmed are closed payment tolls, subsidized marketplaces, foreign-currency-funded unsecured lenders, regulatory-arbitrage labor models, dormant-wallet stories, asset-heavy logistics without density, and opaque conglomerates. EM platforms remain a governance-and-capture research question, never a geographic allocation instruction.

Force 11 - Climate adaptation and physical resilience: necessary expenditure is not automatically collectible

Starting evidence and persistence

Physical risk is observable as well as modeled. The World Meteorological Organization reported that 2015-2025 were the hottest eleven years in the instrumental record and that ocean heat reached a record in its series, alongside costly heat, rainfall, cyclone, fire, drought, and flood events (WMO, *State of the Global Climate 2025*). Event loss depends on hazard, exposure, vulnerability, and asset values; it should not be attributed to climate alone. IPCC assesses that adaptation opportunities become more constrained and less effective as warming increases and distinguishes soft limits of finance, governance, and technology from hard physical limits (IPCC AR6 WGII, Technical Summary).

Persistence has three sources. Atmospheric and ocean conditions, sea level, and installed exposure change slowly. Long-lived assets were designed to historical thresholds that heat, rainfall, salinity, fire, wind, and drought can breach nonlinearly. Urbanization and electrification add assets faster than many jurisdictions can plan, fund, and maintain them. Even rapid mitigation would not remove the need to adapt existing cities, grids, factories, farms, health systems, and coastlines during 2026-2036.

Adaptation includes short-cycle operating measures - forecasts, maintenance, vegetation, cooling schedules, water-loss control, crop choice, emergency response, insurance terms - and modular equipment such as pumps, sensors, switchgear, storage, fire protection, filtration, and efficient cooling. Sea defenses, drainage, reservoirs, transmission, relocation, code turnover, and ecosystem restoration have longer clocks. Better land use, maintenance, efficiency, nature-based systems, and retreat can avoid some headline construction.

The funding gap is the defining evidence. UNEP estimated developing-country needs in 2035 at roughly $310-$365 billion annually versus about $26 billion of international public adaptation finance in 2023. That 12-14 times gap shows large social need and weak paid demand simultaneously (UNEP, *Adaptation Gap Report 2025*). Water finance and national plans have grown, but an approval is finance, not a sustainable tariff, and a plan is not collectible revenue.

From avoided loss to paid cash

Adaptation often creates avoided loss dispersed among households, employers, landowners, hospitals, insurers, cities, and future residents. The World Bank has estimated large social benefit per dollar of resilient infrastructure in lower-income countries; that ratio is not a private return (World Bank, *Lifelines: The Resilient Infrastructure Opportunity*). A collectible model must answer four questions: who receives the benefit; who can sign and pay; what makes payment enforceable; and whether cash exceeds maintenance, replacement, working capital, warranty, claims, finance, tax, and dilution.

The evidence hierarchy runs from hazard estimate, plan, target, and unfunded mandate; through appropriation, reviewable rate allowance, insurance premium approval, or cancellable order; to funded indexed contract, approved rate with demonstrated collection, or recurring private demand with measured short payback. Cheap

finance cannot solve missing funding. A guarantee changes timing and risk allocation but leaves taxpayers with the residual. A public-private partnership does not create affordability.

Disasters and public targets attract engineering firms, developers, equipment plants, and private funds. Scarce labor, cable, transformer, cement, and permitting raise cost. Contractors bid fixed-price work on incomplete designs and later take charges. Utilities finance rate base before recovery. Insurers reprice, withdraw, or litigate. When fiscal conditions tighten, projects can be canceled after suppliers have built capacity. Adaptation is a set of overlapping capital cycles rather than one smooth growth market.

Distribution

Households receive safer shelter, water, power, health, and lower interruption, but pay through taxes, tariffs, premiums, rent, or exclusion from protected areas. Affordability politics can suppress provider return. Public resilience and transport can capitalize into land; new hazard maps and insurance terms can destroy exposed land value. Land-value capture and property tax determine whether public investment returns to the state or remains a private windfall.

Engineering, line, inspection, maintenance, and construction labor capture scarcity while heat itself lowers safe work hours. The state is planner, insurer of last resort, regulator, borrower, and payer; it can preserve tax base or shift loss to owners through disallowance, price caps, liability, compulsory investment, or tax. Financiers earn senior interest and fees; refinancing can absorb an allowed return. Insurers can benefit from risk selection and price adequacy but face correlated loss, stale models, litigation, and rate suppression. Infrastructure owners earn only where availability payment or regulated recovery exceeds full finance and replacement cost. Equipment, data, inspection, and recurring service often have better capture than bespoke fixed-price construction.

An adaptation investment can reduce a company's loss without creating revenue. A contractor can grow revenue without earning a return. A utility can grow rate base without parent cash. A property can be protected while its landowner captures most of the benefit. These ledgers must remain separate.

Countertheses, falsifiers, and model effects

Adaptation spending may remain modest because fiscal capacity and household affordability bind harder than physical risk. Efficient adaptation may rely more on land use, maintenance, forecast, distributed systems, and retreat than megaprojects. Maladaptation can encourage more exposure or shift water and flood risk. Loss trends can be misread if higher asset values and population explain much of the nominal increase.

Track location-adjusted loss and outage, heat and water stress, insurance price and terms, exposure concentration, enacted funding, rate and premium recovery, paid orders, working capital, fixed-price charges, service attachment, measured avoided loss, construction claims, financing and dilution, and owner cash after maintenance. A collectible thesis fails if plans stay unfunded, rates are suppressed, arrears grow, project cash is negative, or measured resilience does not improve. A regulated-owner thesis fails if nominal asset growth depends on repeated issuance and earned return remains below cost.

Potentially helped models include mandatory testing and certification, inspection and maintenance, water-loss and energy-efficiency systems with short payback, grid control, installed-base spares, reliability distribution, embedded risk data, and disciplined insurers or lenders able to reprice and limit exposure. Potentially harmed are exposed real estate with long debt, utilities with uncompensated catastrophic liability, undiversified insurers, one-water-source agriculture, fixed-price contractors, and municipalities dependent on volatile transfers. Regulated resilience remains a payer-quality test, not a safety label.

Force 12 - Urbanization, infrastructure, and the ordinary location rent

Starting evidence and persistence

Urbanization is a distinct force even though it interacts with emerging-market formalization and adaptation. UN DESA's 2025 harmonized revision estimates continued absolute growth and concentration in cities and towns, heavily in developing regions (UN DESA, *World Urbanization Prospects 2025*). The physical requirement includes housing, roads, drainage, water, sanitation, electricity, transport, logistics, schools, health, waste, broadband, cooling, and hazard protection. World Bank estimates of resilient and low-carbon urban investment in lower- and middle-income economies are very large, but they are modeled requirements rather than contracts (World Bank, *Banking on Cities*).

Agglomeration creates value when density improves matching, specialization, customer access, fixed-cost sharing, and knowledge. Congestion destroys it through unreliable services, pollution, crime, informality, long commutes, and housing cost. Infrastructure is a system. A metro without accessible housing can push workers away from jobs. Housing without water or transport creates sprawl. Generation without transmission sits in a queue. Fiber without power or affordable devices remains underused. The weakest complementary link can determine the whole return.

Path dependence makes the force structural. Street grids, rights of way, density, and hazard exposure become expensive to change after construction. Fast-growing regions still have much infrastructure to place, while mature cities must renew and adapt. Young cities need housing, schools, jobs, and transport; aging cities need access, care, and maintenance. Population is a volume indicator, not payer quality.

Funding, capital cycle, and distribution

Infrastructure starts with funding, not finance. Funding is the recurring claim that pays - taxes, tariffs, user charges, availability payments, land sales, development levies, freight, passenger fees, or private customer savings. Debt, equity, concession, development-bank credit, and public-private structures merely allocate timing and risk. When social return lacks a collectible source, cheap finance postpones the gap.

Projects move from plan to land, feasibility, funding, procurement, construction, utilization, maintenance, and renewal. Pipelines overstate demand. Inflation consumes budgets. Fixed-price contracts create claims. Arrears consume working capital. Low utilization depresses revenue. Maintenance is deferred. High return and subsidy attract equipment and construction capacity, which can overshoot after a program peaks. Toll roads, ports, parks, warehouses, towers, and data centers can be duplicated because every sponsor models the same growth. A credible harvest requires stable utilization, cost recovery, maintenance, debt service, and cash distribution after concession and currency obligations.

Labor receives higher wages from productivity, but housing and commuting can absorb the gain. Landowners capture access and zoning rents; public transport, drainage, and safety can capitalize into land. The state can reclaim uplift through property tax, leasehold, development charge, land sale, and special assessment, but weak cadastres and political resistance often leave it private. Infrastructure operators receive tolls or regulated returns where volume and funding are balanced. Contractors, labor, materials, and equipment capture early. Financiers receive interest and guarantees. Consumers and firms gain reliable services and productivity. Platforms use density for delivery, mobility, property, payments, and advertising, while paying higher land, driver, warehouse, and regulatory costs.

UN-Habitat's 2026 African cities work emphasizes land-value capture to recycle infrastructure-induced appreciation into services. Successful public capture can improve funding and simultaneously reduce the private windfall assumed by a landowner or developer (UN-Habitat, *State of African Cities 2026*, executive summary). Socially successful urbanization need not maximize incumbent property profit.

Minority owners receive the residue after land, build cost, interest, tax, concession, maintenance, currency, and dilution. Asset replacement value is not owner value if tariff cannot rise or cash cannot upstream. Speculative land banks can report appreciation without permission, connection, or paying users.

Countertheses, falsifiers, and model effects

Urbanization without productivity is possible: people move because rural options are worse, while employment remains informal and congestion grows. Fiscal exhaustion can crowd out maintenance. Remote work, distributed energy, wireless systems, and modular infrastructure can reduce centralized megaproject demand. Overbuild can precede users. Land can capture nearly all productivity and leave workers and operators worse off.

Track output and formal employment per worker; travel and service reliability; housing burden; formal jobs; funding collection; property tax and land-value capture; project utilization; customer savings; complementary links; claims and delay; arrears; maintenance; debt and currency match; and after-maintenance parent cash. The productive thesis fails when population rises while productivity, access, formal work, and service reliability stagnate. The collectible thesis fails when unidentified finance, unaffordable tariff, low utilization, foreign-currency debt, or replacement consumes operations.

Potentially helped models include small-ticket mission-critical distribution, maintenance, inspection, metering, leakage detection, interoperable payment and identity rails, high-failure-cost equipment, and operators with indexed enforceable funding and conservative currency structure. Potentially harmed are fixed-price megaproject contractors, speculative land banks, foreign-currency-funded local tariffs, aggressive volume concessions, high-loss utilities, and single-city platforms exposed to regulation. Landowners are conditionally helped only where title, access, supply, tax, and hazard all pass.

Counterforce - The ordinary decade and the ordinary compounder

The most important counterforce is not that technology, climate, aging, security, resources, urbanization, or formalization are false. It is that they become socially important and financially ordinary. AI diffuses but capability becomes cheap. Compute expands but builders overpay. Grid investment grows but regulated and equipment returns normalize. Electrification continues amid manufacturing price wars. Aging raises care demand but wages and payers absorb it. Defense budgets rise but states exercise monopsony. Supply chains diversify without wholesale autarky. Resource supply, recycling, and substitution prevent a durable harvest. Adaptation proceeds through maintenance, standards, public works, and avoided loss rather than premium-margin projects. Emerging-market payment rails expand while tolls commoditize. Urban value goes to customers, land, labor, and the state.

This is a failure of exceptional expectations, not of the underlying forces. A broad low-cost comparator can capture the diffuse productivity and customer surplus without requiring the investor to identify the narrow claimant or time the capital cycle. An ordinary compounder with recurring demand, modest capital need, conservative finance, transparent governance, real incremental returns, and disciplined per-share allocation may outperform a celebrated theme owner. A strong brand, distributor, service network, balance-sheet franchise, or mundane workflow can adapt to several forces without depending on one forecast.

"Ordinary" is not a quality waiver. A boring business can overpay for acquisitions, face disruption, lose customer relevance, or be purchased at an expectation-heavy multiple. The ordinary-compounder prior must still prove demand, structure and payer, capital economics, governance and per-share capture, valuation, and falsifiability. Its advantage is lower dependence on a single bottleneck and lower research error, not immunity.

The ordinary-decade counterforce also disciplines activity. Complex research has costs in hours, data, tax, spreads, abandoned comparator exposure, and decision error. If a theme produces no journaled decision or explicit abstention threshold within the policy's 36-month horizon, it should move to low-frequency monitoring. If

active work cannot demonstrate value beyond the comparator after fully loaded cost, abstention is a positive policy outcome, not analytical defeat.

The counterforce weakens only if productive diffusion, durable market structure, and diluted per-share cash exceed what starting prices implied across a full capital cycle, and if active identification improves the portfolio after tax, cost, factor overlap, and error. Strong social adoption or high revenue alone cannot falsify it.

Force-to-business-model matrix

The matrix below is a crosswalk, not a score and not a recommendation. "Helped" means the force can improve operating demand or bargaining power if the stated capture condition holds. "Harmed" means the force can weaken the model or move value to another claimant. "Ambiguous" means both channels are material. The final column identifies the bridge most likely to break between social value and each share.

Force	Models potentially helped, conditional on proof	Models potentially harmed	Primary capture condition	Capital-cycle reversal	Per-share failure to watch
AI and software diffusion	Systems of record becoming systems of action; proprietary lawful workflow; verification; cybersecurity; identity; outcome-based domain tools	Generic wrappers; routine information tolls; hourly outsourcing; seat software with shrinking users	Recurring governed use leaves customer surplus while vendor retains outcome price after compute and review	Low-cost models, incumbent bundling, open standards, customer multi-homing	AI revenue grows while gross profit after compute, development, stock pay, and dilution does not
Compute, data centers, semiconductors	Qualified semicap, advanced packaging, memory/network bottlenecks, recurring equipment service, disciplined harvest survivors	Leveraged merchant compute; generic shells; assets under-depreciated economically; duplicated subsidized capacity	Qualification and system value persist after new supply, with conservative asset lives	Simultaneous capacity, architecture shift, custom silicon, rental-price decline, refinancing	Capex, leases, commitments, interest, and replacement grow faster than parent cash/share
Power and grids	Grid control, efficiency, qualified electrical equipment, service, flexible load and funded reliability	Speculative queue/land; fixed-price build; rate base without recovery; low-value constrained loads	Delivered and collected hosting capacity with upgrade and stranded cost assigned	Flexible connections, reconductoring, generation, storage, cancellation, equipment entry	Allowed return or backlog grows while earned cash return and diluted cash/share trail funding cost
Electrification and automation	Installed-base service; qualified controls; sensing; repeatable integration; reliability distribution	Commodity EV/battery plants; bespoke project integrators; legacy closed equipment; hour-based outsourcing	Site-vintage labor, quality, uptime, and energy saving exceed all-in cost	Manufacturing flood, chemistry shift, open interfaces, customer process failure	Revenue and backlog rise while project cash, warranty, retooling, and shares consume value
Aging, health, and labor	Funded hard-outcome innovation; labor-saving care workflow; dense home/community delivery; diagnostics tied to action	Fixed-tariff labor-heavy care; facilities built on age curves; digital health without payer saving; mature IP without replacement	Payer access and collectibility; measurable outcome inside payer horizon; scarce labor minutes removed	Healthy aging, participation, migration, automation, price negotiation, substitute entry	Clinical or volume success fails to become parent cash after rebates, R&D, acquisition, leases, and dilution
Fiscal activism and defense	Diversified qualified components; installed-base maintenance; legitimate testing; dual-use private/state payer; short-cycle necessary supply	One-program fixed-price development; aspiration-backed plants; subsidy economics; backlog roll-ups	Funded appropriation, favorable risk allocation, acceptance and collection, commercial economics after support	Fiscal repair, peace dividend, procurement reform, input capacity, permanent subsidized flood	Budget and backlog rise while overruns, working capital, tax, guarantees, debt, and shares absorb cash
Fragmentation and resilience	Multi-jurisdiction logistics; customs/trade finance; traceability; verification; modularity; diversified processing	Single-source trade; low-margin tariff-exposed assembly; stockpilers; protected subscale plants	Measured recovery improvement at acceptable total landed cost and utilization	Trade accommodation, global diversification, substitution, retaliation, excess local capacity	Domestic revenue rises while export loss, duplication, compliance, currency, and state take reduce cash/share
Inflation, debt, repression	Short-cycle repricers; low-refinancing necessary services; disciplined balance-sheet franchises; real customer-value pricing	Long fixed-price duration; leveraged assets; captive finance; nominal growth with high replacement cost	Real price and volume survive wage/input reset; financing and legal integrity hold	Productivity, energy abundance, fiscal repair, demand bust, repression of visible rents	Nominal earnings rise while replacement capital, interest, tax, working capital, and dilution lower real cash/share

Force	Models potentially helped, conditional on proof	Models potentially harmed	Primary capture condition	Capital-cycle reversal	Per-share failure to watch
Natural resources and harvest	Low-cost long-life clear-title assets; scarce diversified processors; monitoring; recycling with feedstock; efficiency	High-cost single assets; peak-price projects; weak closure funding; fixed-price processors; unverified biological harvest	Delivered scarcity or confirmed net exit, stable fiscal terms, full sustaining/closure cash	Substitution, recycling, new supply, state support, demand disappointment, tax reset	Commodity profit rises while royalties, capex, remediation, debt, related parties, and shares take the surplus
EM formalization and platforms	Merchant workflow above public rails; identity/fraud; tax/accounting; disciplined logistics; transparent deposits; organic marketplace cohorts	Closed payment tolls; subsidized marketplaces; FX-funded credit; labor arbitrage; opaque groups	Active use, counterparty surplus, full-cycle credit, simple audited parent waterfall	Interoperable public rails, fee caps, multi-homing, fraud, currency stress, regulation	Local volume/profit grows while trapped cash, FX, credit, related parties, preferences, or dilution erase per-share growth
Climate adaptation	Inspection, maintenance, efficiency, water-loss control, risk data, installed service, disciplined underwriting	Exposed property; liability-heavy utilities; stale-model insurers; fixed-price contracting; weak municipalities	Enforceable funded payer and measurable avoided loss with maintenance included	Affordability, efficient non-capex adaptation, fiscal delay, maladaptation, standardization	Plans and asset base grow while disallowance, claims, financing, maintenance, and issuance consume parent cash
Urbanization and infrastructure	Small-ticket maintenance, metering, leakage control, interoperable rails, funded operators, high-failure-cost equipment	Speculative land banks; FX-funded tariffs; aggressive concessions; megaproject contractors; high-loss utilities	Productive agglomeration plus sustainable funding, utilization, maintenance, and currency match	Congestion, remote substitution, overbuild, land capture, fiscal exhaustion	Asset value and EBITDA rise while land, debt, concession, arrears, replacement, and currency block cash/share
Ordinary compounder/comparator	Recurring modest-capital demand; adaptive service; disciplined balance sheet and capital return; low-cost broad exposure	Narrative-dependent roll-ups; thematic pure plays priced for exceptional duration	Adequate return without heroic force forecast, after fully loaded research cost	Theme finally produces durable exceptional per-share cash at an undemanding price	Quality label or low volatility conceals valuation duration, factor concentration, or poor incremental allocation

Several patterns repeat across rows. Verification, installed-base service, reliability distribution, funded labor-saving workflow, disciplined balance sheets, and small-ticket failure-critical products appear more often than large builders. This is not because they are guaranteed winners. It is because their path from customer benefit to collected cash can be shorter, capital need can be smaller, and their fee can leave enough surplus to preserve legitimacy. Conversely, generic capacity, fixed-price megaprojects, unfunded public need, and acquisition- or debt-led growth appear repeatedly on the harmed side because activity can expand while prior claimants absorb the gain.

The matrix also exposes correlated exposure. Semicap, grid equipment, automation, and data-center service can all be one AI-capex and power-load factor. Defense, adaptation, local manufacturing, and regulated grids can be one sovereign-payer and scarce-labor factor. Resources, electrification, and data centers can share power, China, shipping, and state-policy exposure. EM platforms, market infrastructure, and balance-sheet franchises can share currency, regulation, cyber, and liquidity. Diversification by company name or sector label can therefore be illusory.

Cross-force interactions: the decade is a system, not twelve sleeves

AI, labor, health, and the productivity split

AI and aging pull labor in opposite directions at the task level. AI can weaken routine cognitive work while aging raises demand for physical care, maintenance, field service, engineering, and accountable judgment. The interaction could widen bargaining differences within the same occupation. A clinician whose documentation is automated may treat more patients, yet the saving can go to the payer through a tariff reset, to the hospital through capacity, to the vendor through an outcome fee, or to the clinician through lower workload. A caregiver may remain physically scarce even as scheduling and monitoring improve.

If AI mainly raises novice productivity, firms can broaden the pool of acceptable labor, compress wage dispersion, and retain margin. If experts remain the liability-bearing bottleneck, their compensation can rise. If entry-level

apprenticeship tasks disappear, future expert supply can shrink, creating a later scarcity that is not visible in current headcount. Monitoring must therefore follow hiring, training, promotion, error, and productivity cohorts rather than one aggregate exposure measure.

Health improvements feed back into labor supply. Better metabolic, cardiovascular, hearing, vision, mental, and chronic-disease outcomes can reduce absence, delay disability, and free caregivers. That can ease wage scarcity and improve tax revenue. The intervention provider may not receive the productivity value if the employer, later public payer, or family gains on a different clock. Contracts that share savings are attempts to bridge this split but can fail on attribution, data, turnover, and lag.

Compute, power, land, and finance

AI appears digital at the user interface and physical in its bottlenecks. Model demand requires chips, packaging, memory, networking, cooling, buildings, transformers, turbines, transmission, water, land, and skilled labor. Local power scarcity can transfer an apparent model rent to a landlord, utility, generator, equipment supplier, financier, or municipality. A software business can save customer labor while raising electricity and capital intensity elsewhere.

Higher rates magnify the asset-life problem. When merchant compute and data centers are financed with debt, leases, and residual guarantees, a small reduction in useful life moves a larger share of revenue to senior claimants. Rate cuts can reopen financing before excess capacity exits, delaying harvest. Subsidies can preserve strategically desired but uneconomic fabs or power projects. Thus cheap capital can accelerate both capability abundance and equity loss.

Power policy can also decide which social group pays. Customer-funded upgrades protect ratepayers and reduce tenant or land rent. Socialized expansion initially supports site economics and utility growth but can provoke political repricing. Flexible service can preserve more system surplus by connecting high-value loads sooner without forcing all users to fund firm capacity. The value of an equipment supplier depends on which equilibrium occurs.

State procurement, debt, and crowding

Defense, grids, semiconductor localization, climate adaptation, health, pensions, and interest all demand fiscal resources and the same engineers, construction labor, equipment, land, and credit. When buildouts overlap, nominal appropriations can rise while real delivered quantity disappoints. Higher wages and equipment prices may represent legitimate scarcity capture, but they reduce supplier margin where contract escalation is weak and increase public debt where budgets are maintained.

Monetary tightening intended to contain inflation raises sovereign interest and project hurdles. Governments can respond with guarantees and directed credit, increasing the state's allocation role. Banks and pensions may hold more sovereign and preferred-sector claims, crowding unrelated entrepreneurship. The state can therefore be simultaneously a larger customer and a harsher tax or price claimant. A business that depends on state demand should be underwritten against the same state's debt service and political priorities, not against the program in isolation.

Defense and adaptation can share the same contract risk: a socially essential need, concentrated buyer, specialized capital, and a large gap between target and paid delivery. State monopsony can protect continuity while limiting margin. A supplier's best economics may reside in a small cross-program component or recurring service rather than the headline platform.

Fragmentation, resources, and the double capital cycle

Economic-security policy can prolong scarcity by limiting suppliers and then destroy it by financing redundant capacity. A critical mineral can be geologically ample but concentrated in processing. Subsidized diversification can reduce geopolitical risk and create a global supply flood. Fabs can be duplicated within blocs while each individual facility lacks scale. Stockpiling can amplify orders during buildout and produce destocking later.

The state may rationally accept poor financial return for strategic capacity. That public value must be recorded separately from shareholder value. A domestic processor can benefit from a floor price while downstream customers pay more and exporters face retaliation. A foreign low-cost producer may expand into weakness to deter entry. The eventual survivors may have attractive economics only after support is withdrawn or rationalized, a condition that can be politically delayed for years.

Climate adds another physical interaction. Heat, drought, fire, flood, and water stress can interrupt mines, refineries, ports, fabs, farms, and grids, creating temporary commodity and equipment prices. Those prices attract capacity that may arrive after the disruption. Adaptation can increase near-term demand for metals, cement, pumps, cooling, and controls while efficiency and design reduce lifetime material intensity. A resource thesis that counts adaptation demand but omits recycling, substitution, and state-supported supply is one-sided.

Urbanization, formalization, and land capture

Identity, payments, invoices, land records, property tax, broadband, transport, and power can jointly raise urban productivity and municipal collection. Density improves delivery and matching, helping platforms and merchants. It also raises housing, warehouse, driver, labor, and regulatory costs. If housing supply is inelastic, land absorbs wage and productivity gains, deters migration, and makes industrial policy less feasible. A nominally high wage or platform density can coexist with poor worker welfare and weak operator margin.

Formalization can improve infrastructure funding by broadening tax and billing records. It can also impose compliance without delivering service, causing firms and workers to retreat into cash. Public payment rails can lower cost and reduce leakage, benefiting the state and customers while commoditizing private tolls. Land-value capture can finance transport and resilience while reducing the return expected by incumbent developers. A socially coherent city can therefore redistribute from private land and platform rents toward public services and new entrants.

Inflation, replacement cost, and apparent profit

Every capital-intensive force is vulnerable to nominal-accounting illusion. A utility, miner, data center, factory, hospital, vehicle plant, or infrastructure concession can report higher revenue and EBITDA because prices rise. Replacement equipment, labor, working capital, interest, tax, and insurance may rise faster. Accounting depreciation based on historical cost understates the cash required to reproduce the asset. Contracted nominal revenue can be fixed while inputs reset. A regulator can allow a nominal return that trails actual financing and replacement cost.

Real diluted cash per share is the common denominator. It must be measured through a normal cycle and against actual economic lives. The parent must have legal access. The share count must include recurring compensation and financing. Valuation must then compare that cash with what the price already assumes. No thematic force bypasses this bridge.

Coherent decade regimes and how value moves

The controlling policy retains six scenarios rather than a point forecast. This chapter expands their distribution logic.

Productive diffusion

AI, automation, health, energy investment, and supply redesign raise productive capacity. Inflation anchors and real wages can rise with output. Compute and power capacity become more abundant after a costly buildout. Customers capture lower quality-adjusted prices; skilled labor captures where judgment remains scarce; disciplined incumbents and entrants retain part of measurable savings; the state receives a broader tax base. Early infrastructure builders can still disappoint if depreciation and finance were aggressive. Models helped are workflow, verification, efficiency, installed service, and ordinary adopters. Models harmed are pure scarcity claims and underutilized buildout.

This regime is supported by widening productivity beyond technology sectors, stable or rising labor participation, falling customer cost, healthy credit, utilization of new assets, and real cash growth after depreciation. It is weakened by shallow adoption, idle capacity, weak wage pass-through, and recurring subsidies.

Inflation and repression grind

Repeated supply shocks, incomplete fiscal adjustment, and directed capital keep nominal growth high and real distribution contested. Land, power, resource, and scarce labor can reprice. The state and creditors hold larger prior claims, though fixed-rate savers and weak lenders lose in real or default terms. Visible profit invites tax, price control, and compulsory investment. Strong balance sheets, short repricing, and low replacement needs matter more; long fixed contracts, leverage, and nominal accounting are harmed.

This regime is supported by repeated pass-through, negative real regulated savings, captive sovereign holdings, widening state take, weak private credit, and replacement cost exceeding accounting depreciation. It is falsified by productivity-supported disinflation, credible fiscal repair, positive market savings, and broad entrepreneurship.

Fragmented blocs

Controls, subsidies, defense, and duplicate supply chains raise orders for testing, qualified equipment, power, logistics, and strategic processing. Land, labor, construction, and finance capture during buildout. States obtain security value and exercise monopsony. Customers pay a resilience premium. Later, duplicated supply can lower price and strand late capital. Trusted bottlenecks and compliance may retain value; undifferentiated localized capacity and export-dependent models are harmed.

Support requires durable funded procurement, measured recovery improvement, commercially utilized capacity, and cash that survives after tariff and subsidy. Retaliation, subscale operation, continuing support, and customer-price damage weaken the shareholder thesis even if the security objective remains.

Crash and sequence

Valuation, collateral, and credit unwind together. The physical buildout can become a source of forced selling as leases, debt, and commitments remain fixed while demand and asset values fall. Creditors gain seniority but can lose principal. Customers later receive cheap capacity. Strong balance sheets and liquid quality may gain after fresh underwriting; the Investor Liquidity Floor and matched liabilities protect the household from forced sales. No thematic "winner" is purchased mechanically.

Sequence separates temporary market loss from household impairment. If liabilities force sales before recovery, duration has become permanent. A broad comparator can be concentrated in the same AI or country factor, so low cost is not no risk. The regime does not authorize tactical buying; it requires liability and access first, then evidence.

Adaptation and physical constraint

Power, water, heat, hazard, minerals, insurance, and land become gating inputs. Rent moves toward connected sites, qualified labor, equipment, financiers, and the state. Customers bear higher tariffs, premiums, delay, or rationing. Efficiency, control, maintenance, risk selection, and funded resilience may retain value. Unfunded plans, liability-heavy owners, and fixed-price projects are harmed.

This regime is supported when location-adjusted losses, congestion, insurance stress, and paid adaptation persist and interventions measurably lower loss. It is weakened if forecasting, land use, flexibility, supply, and maintenance solve much of the constraint at ordinary margins.

Plateau and muddle

AI becomes useful but ordinary, inflation and term premium fluctuate without a permanent new regime, health innovation advances incrementally, power and infrastructure grow near demand, and geopolitical fragmentation remains selective. No bottleneck earns a decade rent. Customers and broad adopters receive much of the gain. Ordinary compounders, disciplined capital return, and the low-cost comparator can prevail. Research theater and overtrading become the main avoidable loss.

This scenario should be treated as a base-rate challenge, not a residual bucket. It is falsified only by sustained exceptional per-share cash and active attribution beyond the comparator after full costs, not by headlines or market-size growth.

Where capture may strengthen, weaken, or remain radically uncertain

Capture may strengthen, subject to valuation and evidence, in narrow qualified complements: regulated or customer-required verification; identity and cybersecurity tied to action; installed-base service where alternatives remain limited; reliability distribution with density and working-capital discipline; semicap and power-control bottlenecks before supply response; labor-saving workflows with funded payers; and balance-sheet capacity that finances necessary activity without taking uncontrolled credit, duration, or liquidity risk. These models are favored research priors because the fee can be small relative to avoided failure and because cash evidence may arrive inside 36 months. Qualification, contract, customer surplus, and capital need still decide the result.

Capture may weaken in undifferentiated digital features, seat models exposed to fewer users, generic merchant capacity financed against peak rents, speculative land or queue options, fixed-price megaprojects, undifferentiated localized factories, long-duration contracts with weak indexation, labor-heavy care with fixed reimbursement, mature IP without replacement, closed payment tolls facing public rails, and acquisition-led roll-ups requiring repeated debt or shares. Many sit in expanding markets. The weak prior comes from contestability, financing, payer, and per-share leakage rather than a prediction that their end demand falls.

Capture remains radically uncertain for hyperscalers, frontier-model developers, leading accelerators, foundries, large regulated utilities, general-purpose robotics, national champions, strategic resource producers, obesity platforms, risk-bearing healthcare, large infrastructure concessions, insurers facing changing physical risk, and dominant EM platforms. Each can control a genuine bottleneck or network and face a powerful capital-cycle, state, liability, or dilution response. Aggregate profit or technical leadership cannot resolve the conflict.

For labor, the prior is bifurcated. For land and connected-site rents, it is stronger but political and local. For creditors, it is contractually stronger with larger tail risk. For the state, it is stronger in both claim and responsibility. For entrepreneurs, experimentation becomes cheaper while scaling remains concentrated. For incumbents, distribution and installed base matter until technology lowers switching or regulation opens the layer. For entrants, narrow application opportunities expand while capital-heavy bottlenecks remain hard. For platforms, execution and trust are more defensible than discovery. For customers, cumulative capture may be largest after supply response. For common equity in aggregate, there is no strengthening prior independent of price and per-share evidence.

Monitoring architecture: from narrative to observable rates

Monitoring should prefer rates, cohorts, cash bridges, and transitions to thematic totals. Separate announced, requested, funded, delivered, utilized, billed, collected, and upstreamed quantities. Use medians and vintages rather than only aggregate growth. Measure price and volume separately. Record forecast revisions. Compare economic with accounting depreciation. Treat management attribution of "AI," "defense," "resilience," or "strategic" revenue as a hypothesis until reconciled to cash.

Research question	Leading indicators	Constructive 36-month evidence	Falsifier or mandatory re-underwrite
Is AI diffusion productive?	Daily use, task breadth, governed workflows, error-adjusted output, customer implementation cost	Recurring use broadens and verified savings exceed compute, review, error, and integration	Adoption remains shallow for two years or repeated customers fail cash payback
Does software retain value?	Seat-normalized retention, real price, compute cost/outcome, switching, multi-model use	Customers save while vendor retention, gross profit, and diluted cash/share hold	AI is bundled or bypassed; gross profit after compute/development remains inadequate
Does compute earn its capital?	Vintage utilization, rental price, resale, retirement, capex and commitments, customer ROI	Incremental gross profit and defended-franchise cash cover economic depreciation and funding	Prices/utilization fall, asset lives shorten, or capex stays above cash improvement through six quarters
Is power scarcity real and collectible?	Vetted MW, deposits, interconnection, hourly congestion, lead time, tariff terms	Requests become energized peak load and billed cash with stranded cost assigned	Requested-to-realized remains low; cancellations and equipment capacity rise; cost is socialized politically
Does grid investment reach equity?	Allowed/earned return, recovery lag, bills/arrears, debt and issuance, project cash	Delivered reliability and after-finance cash grow without repeated dilution	Rate base grows while earned return stays below cost or disallowance/liability rises
Does automation remove scarce work?	Pilot conversion, uptime, labor minutes, quality, warranty, service attach	Repeated sites meet throughput and payback with lower labor and equal quality	Pilots need unchanged staffing, miss uptime, or cumulative project cash remains negative
Is aging creating funded demand?	Coverage, tariffs, waiting, disability, migration, vacancies, care hours	Funded community and care capacity expands while outcomes and cash conversion improve	Need rises but eligibility, collection, staffing, or access deteriorates
Does health innovation create net value?	Hard endpoints, safety, access, adherence, net price, total cost, exclusivity	Defined cohorts reproduce outcomes at acceptable budget impact and parent cash grows	Surrogates fail, safety worsens, persistence collapses, or price-volume cash falls
Is fiscal activism productive?	Awarded/paid support, private co-investment, capacity, learning, commercial demand	Unsubsidized unit cost and utilization improve; net public benefit is measurable	Recurring subsidy, low utilization, no learning, or crowding and protection increase
Is defense demand collectible?	Appropriations, contract awards, quantities, contract type, cost/schedule, advances	Delivery and cash conversion improve after capacity investment	Targets fail to become contracts, costs outrun budget, or cancellations and working capital rise
Is fragmentation deepening efficiently?	Tariff-covered trade, control lists, screening, source diversity, recovery time	Resilience improves with commercially used diversified supply	Restrictions narrow or duplication raises cost without lowering disruption loss
Is the inflation regime structural?	Dispersion, expectations, wage/price reset, real forwards, neutral-rate range	Repeated supply shocks pass through and contract behavior changes beyond one cycle	Five years of low dispersion, anchored expectations, and non-inflationary shock absorption
Is debt constraining allocation?	Interest/revenue, primary gap, maturity, auctions, captive holdings, private credit	Interest displaces spending or directed finance and captive claims rise	Credible adjustment lowers interest burden while private credit and capital mobility broaden
Is a resource constraint durable?	Inventories, utilization, grade, projects, substitution, recycling, export policy	Delivered tightness persists after destocking and financeable supply stays limited	New supply, chemistry, recycling, or weaker intensity closes the gap
Is harvest confirmed?	Net exits, funding windows, state support, survivor capex, utilization, owner cash	Capacity is removed; weak firms cannot refinance; survivors return cash after full sustaining cost	Capacity is preserved, dividends use debt, or 24 months pass without confirmation
Is EM formalization substantive?	Active use, payroll, invoices, filing, procurement, worker protection	Use moves beyond cash-out and firms gain productivity, finance, and enforceability	Access/registration grows without use, trust, productivity, or reciprocal benefit
Does a platform create rather than tax value?	Multi-homing, take rate, incentives, counterparty margin, fraud, direct traffic	Platform economics improve with merchant, worker, and customer surplus	Take rises while exit/evasion/regulation grows or public rails erase the toll
Is adaptation funded and effective?	Location-adjusted loss, enacted funding, rates/premiums, paid orders, service attach	Collected cash and measured loss reduction appear after full maintenance	Plans remain unfunded, prices are suppressed, fixed-price losses and arrears grow
Is urban buildout productive and paid?	Formal jobs, output/worker, travel, housing burden, collection, utilization	Complementary assets improve access and funding covers operation and debt	Population rises with congestion, informality, low use, arrears, and land capture
Does equity receive the residual?	Economic maintenance, working capital, interest, tax, parent access, diluted shares	Real parent cash per fully diluted share grows through a normal cycle	EBITDA or aggregate profit rises while real cash/share and mature ROIC stagnate

Cadence should match the clock. Demographic cohort and long infrastructure forecasts need annual or material-revision review. Adoption, payer, price, capacity, utilization, credit, and cash need quarterly or semiannual attention. Safety, legal title, reimbursement, license, export control, sanctions, catastrophe, and covenant events require event-driven review. Daily narrative monitoring adds little to a decade view and can increase decision error.

Every live thesis should have at least one leading indicator, a 36-month causal milestone, and a hard falsifier. The milestone tests the link actually underwritten, not completion of the decade. A general-purpose technology need not transform national productivity in three years, but a workflow should demonstrate customer saving and retention. A grid thesis should produce delivered projects. A state-payer thesis should produce appropriated and collected cash. A platform should show organic contribution and parent cash. Two missed primary milestones require a full re-underwrite. Continued social need cannot excuse payer, financing, legal, or per-share failure.

Evidence discipline, source conflicts, and model uncertainty

The sources broadly agree on direction in several areas while disagreeing on magnitude or realization. Data-center electricity projections have wide bands because server shipments, efficiency, utilization, task intensity, and load definitions are uncertain. Grid queues show congestion and speculation at once. Demographic forecasts are firmer for near-age cohorts than for migration, fertility, health, or labor behavior. Adaptation-needs estimates show scale but not paid demand. Industrial-policy and defense budgets show state intent but not program economics. Critical-mineral gaps depend on announced projects, chemistry, recycling, and policy. These are not conflicts to average into false precision.

Where evidence is company-specific, it is used to illustrate a mechanism rather than establish a security conclusion. Large supply commitments illustrate inventory and architecture risk; large contractual capacity and lease commitments illustrate financing and residual risk; long equipment backlogs illustrate qualification and capital response. No named-company observation creates a positive weight or bypasses current valuation and portfolio review.

The policy should preserve unresolved disputes rather than bury them in prose. AI overbuild remains open by layer. Accelerator economic lives remain open. Defensive capex versus empire-building remains open. AI's effect on assurance rents remains open. State-funded capex can create public value while shareholder return remains open. The ordinary-decade countercase remains open. A worldview becomes more credible when it states what it does not know.

Practical research boundary

The chapter changes the research map, not the portfolio. It supports conditional priors toward failure-critical verification, trusted installed-base service, reliability distribution, funded labor-saving workflow, disciplined balance sheets, selective qualified bottlenecks, and models with transparent per-share cash. It supports skepticism toward generic capacity, unfunded public need, fixed-price long-cycle work, subsidy-dependent unit economics, opaque group structures, short economic lives financed with long claims, and market-size narratives that skip the payer and share bridge.

It does not establish expected return. It does not select a company. It does not authorize a weight, entry, exit, hedge, or message. It does not override the Investor Liquidity Floor, verified liability matching, the comparator burden, the five non-compensating business-model pillars, Evidence Coverage, conservative valuation, loss budgets, issuer and factor limits, custody and legal review, tax analysis, or explicit owner approval. The current admissible weight of every company remains zero. In particular, no positive discussion of AI, semiconductors, power, or automation changes the controlling bar on active AI/power additions while the unresolved comparator condition persists.

The most useful output is a disciplined question: after customers, labor, land and power, suppliers, creditors, the state, and controlling owners exercised their claims, did a funded and defensible operating advantage produce

parent cash and grow real cash per fully diluted share through the capital cycle, at a price and portfolio role that beat the low-cost comparator after all costs? If the answer is not evidenced, the correct conclusion is research-only or abstention.

Precommitted triggers for revision

The worldview should be changed by evidence, not by price action or narrative fatigue. It would require material revision if several of the following occur together: AI use remains shallow and customer payback absent despite continued capability gains; compute and power utilization collapse after announced capacity; grid flexibility releases hosting capacity much faster than construction; healthy longevity, migration, participation, and automation neutralize worker scarcity; fiscal repair lowers debt-service pressure without repression; defense targets fail to become contracts or a durable peace dividend emerges; global diversification restores resilience without broad duplication; inflation dispersion and real rates return sustainably to pre-2022 patterns; recycling, substitution, and new supply remove repeated resource tightness; formalization stalls or public rails erase private platform economics; adaptation funding remains persistently uncollectible or risk stabilizes after exposure adjustment; or ordinary compounders and the comparator capture the decade's productivity while thematic owners fail to produce superior per-share cash.

It would also require revision in the opposite direction if productive diffusion broadens across sectors; customers document strong returns after full implementation; bottleneck lead times and prices remain firm after funded entry; grids convert requested loads into paid utilization with equitable cost allocation; labor-saving tools improve outcomes and wages without displacement stress; state procurement becomes standardized, funded, and cash-converting; localized capacity earns competitive returns without continuing support; real rates and inflation volatility stay high across a full cycle; verified resource capacity exits create disciplined harvest; EM platforms show audited hard-currency parent cash with controlled dilution; and adaptation measurably lowers losses under enforceable funding. Even then, the business, security, price, and portfolio conclusions remain separate.

What would change this worldview

These are precommitted revision conditions, not rhetorical caveats. When the monitored evidence changes the causal mechanism, capital cycle, claimant split, or per-share capture path, the dated worldview and its research priorities must be revised without changing the durable constitution by implication.

Chapter 7 - Favored and Disfavored Conditional Business-Model Priors

Status, purpose, and the three conclusions that must never be collapsed

This chapter ranks business-model classes for research over 2026-2036. It does not rank securities by expected return, specify a portfolio allocation, or authorize any transaction. The classifications are dated conditional priors: starting hypotheses about where durable owner economics may be more or less likely to exist, given the evidence available through 22 August 2026. They must be updated when the causal mechanism, capital cycle, payer, market structure, or distribution of economic surplus changes.

Three different conclusions must be written separately in every use of these priors.

1. A **business-quality prior** asks whether a class is structurally capable of producing recurring demand, bargaining power, attractive incremental cash returns, and reliable per-share capture. It helps decide where scarce research time should begin. It is not a conclusion about a particular company.
2. A **security-return conclusion** asks whether a specific legal security, at its current price and after dilution, debt, tax, cost, economic depreciation, and scenario risk, offers a conservative return that clears the named low-cost comparator. A favorable business-quality prior does not answer this question. Starting valuation can turn a wonderful business into a poor security, and a verified harvest can sometimes make an historically mediocre industry researchable.
3. A **portfolio-role conclusion** asks whether the security improves the investor's funded plan after the Investor Liquidity Floor, liability matching, legal title, tax, custody, liquidity, issuer limits, common-factor limits, and the opportunity cost of active research are considered. A security that clears its own return hurdle can still have no portfolio role because the comparator already contains the same factor, the risk budget is exhausted, or the active program is uneconomic.

The controlling policy has not certified the investor facts, actual comparator, factor look-through, tax and legal position, company Evidence Coverage audits, or current valuations required for implementation. Every company named in this chapter therefore retains a current admissible weight of **zero**. A positive example means only that the company is a useful place to test a model. A counterexample means only that the current research record exposes a failure mechanism. Neither classification predicts the next share-price move, and neither is an instruction to buy, sell, short, hold, hedge, open an account, or communicate an order.

The chapter also separates operating success from owner capture. Social demand can create revenue without attractive operating profit. Operating profit can accrue to employees, landlords, power providers, lenders, the state, suppliers, or a controlling shareholder before it reaches the parent. Parent cash can then be consumed by acquisitions, stock compensation, debt service, minority interests, trapped subsidiaries, taxes, or maintenance investment. Only the residual cash and reinvested value attributable to each diluted share are relevant to the security owner.

The ranking at a glance

The ranks below are intentionally broad. They describe the burden and direction of research, not a score and not an automatic ladder. A class can move between ranks as its capacity cycle and price change. Within a class, one company can pass while another fails because payer quality, governance, capital structure, or legal rights differ.

Rank	Model classes	Why the prior starts here	Main way the prior can be wrong	Security and portfolio boundary
Strongest conditional	Installed-base aftermarket; ratings and select mandated assurance; exchanges, clearing and payment rails; reliability distribution; genuine record-to-action data/workflow systems	Necessity, qualification, density, license, workflow control, or network liquidity can make the customer's avoided loss much larger than the fee while incremental capital remains moderate	Regulation, open access, AI substitution, fleet retirement, customer distress, labor scarcity, fee compression, cyber or liability event, and premium starting valuation	Research priority only. Each security still needs all five pillars, current reverse DCF, factor headroom, and portfolio approval
Favorable but valuation-sensitive	Semiconductor and industrial bottlenecks; power control and automation; disciplined insurance float and deposit franchises; selective brands; funded health innovation; high-quality serial acquirers; constructive regulated franchises	Scarce qualification, trusted balance sheets, clinical differentiation, controlled supply, or repeatable decentralized reinvestment can create attractive economics	Capital flooding, economic-life error, reserve or credit loss, payer capture, regulation, acquisition treadmill, state direction, dilution, and long-duration valuation	Business quality may be high while security return is poor; correlated factor exposure can make portfolio role zero
Emerging/unproven	Short-history challengers; merchant AI capacity and neoclouds; young vertical software and automation platforms; early EM digital platforms; new therapeutic or diagnostic platforms	New distribution, technical cost advantage, data loops, or formalization may create a future moat	Short evidence history, customer and funding concentration, leases and debt, stock compensation, project accounting, credit cycles, legal title, and rapid imitation	Research-only or speculative-option framework; assumed-total-loss discipline; no promotion by narrative or growth rate
Currently unsupported/no qualifier	Full natural-resource or infrastructure harvest; public audit pure-play; asset-management franchise; fully qualified defense/state-procurement compounder; broad regulated resilience; generic EM growth	A plausible class or cycle thesis exists, but the present record lacks a candidate that passes the required class battery and valuation work	Search bias may hide a genuine candidate; waiting for perfect evidence may miss part of a rerating	Absence is a research finding, not proof that the class is uninvestable. It authorizes neither abstention forever nor early entry
Recurring negative priors	Financed capacity floods; leveraged merchant compute; acquisition treadmills; undifferentiated seat software; backlog-without-cash; capital-intensive turnarounds; legacy brands without cohort demand; commodity scarcity with weak title; high-growth platforms with weak per-share capture	The value pool is structurally exposed to entry, lenders, suppliers, labor, state capture, dilution, or economic depreciation	A genuine financing closure, consolidation, governance repair, or product transition can reverse the prior	The default is reject or keep research-only until the named falsifier is met; low price or exciting demand cannot waive a failed pillar

These ranks do not claim that favored classes will outperform disfavored classes. The strongest classes overlap heavily with the 2010-2025 consensus-quality factor. That overlap is itself a risk. If real rates stay high, the quality premium derates, AI weakens information tolls, or ordinary low-cost market exposure captures the decade's gains more cheaply, a collection of excellent businesses can underperform. Conversely, a class with a poor long-run history can become attractive after financing closes, capacity exits, survivor balance sheets heal, and price stops assuming permanent scarcity. Capital-cycle state and valuation can reverse the security conclusion without changing the class's long-run operating character.

How the ranking is formed

The ranking applies the five non-compensating pillars rather than an additive quality score. Pillar A tests demand and lock-in. Pillar B tests structure, cycle, and payer. Pillar C tests capital economics after true maintenance and economic depreciation. Pillar D tests governance and per-share capture. Pillar E tests valuation and falsifiability. Ordinary admission requires every pillar to reach at least 2; a zero is a hard failure. Strength in recurrence cannot compensate for a financed capacity flood, and a scarce technical asset cannot compensate for weak title or a price that requires an implausible terminal value.

The prior also asks a distributional question before admiring aggregate profit: who owns the scarce input and who can credibly walk away? A license holder can have bargaining power over customers, but a regulator can change the fee. An equipment maker can control qualification, but a concentrated customer can sponsor a substitute. A software vendor can own a workflow, but a platform can bundle the feature and employees can capture scarce engineering rents. A utility can receive an allowed return, but households and the state can use affordability policy to deny realized cash. A resource owner can hold ore in the ground, but the sovereign owns taxation and title

enforcement. The durable model is not merely the firm with high current margins; it is the firm whose counterparties retain enough surplus to sustain the relationship while the firm still converts an acceptable portion into diluted per-share cash.

Diagnostic question	Evidence that strengthens a class prior	Evidence that weakens it	Typical claimant on the value pool
What is scarce?	Accreditation, trust, qualified installed base, dense liquidity, proprietary lawful data, low-cost funding, clinical differentiation, or a genuinely constrained bottleneck	A marketing label, temporary queue, commodity hardware, public data, or subsidy-created demand	Bottleneck supplier if scarcity is durable; customers or entrants if it is reproducible
Who can walk away?	Numerous funded customers, fragmented suppliers, high switching cost, and a fee small relative to avoided loss	One state payer, one hyperscaler, one distributor, one key employee group, or refinancing dependence	Concentrated payer, labor, lender, platform, or state when the firm cannot walk away
Who finances growth?	Customer advances, retained cash, low-cost durable float or deposits, and modular projects with kill points	Expensive equity, short debt, leases, supplier finance, or acquisition consideration needed to sustain reported growth	Debt and lessors claim cash first; new equity dilutes existing owners
Who keeps productivity gains?	Outcome pricing, scarce qualification, proven price retention, or a closed-loop workflow that improves results	Seat contraction, procurement rebids, regulated fee caps, wage escalation, or rapid entry	Customers, labor, platforms, and the state can absorb gains even as industry output rises
Does operating profit reach each share?	Cash conversion, stable share count, disciplined reinvestment, clean upstreaming, conservative leverage, and transparent related parties	Stock compensation, roll-up spend, trapped cash, minority leakage, capital calls, windfall taxes, or refinancing	Management, sellers, creditors, state, or controllers can intercept the residual
What would entry do?	Qualification delays, license scarcity, network depth, or local density cannot be bought quickly; incumbent investment earns above cost	Financing and subsidies create capacity faster than demand; switching and data portability improve	Customers benefit from competition; early scarcity rents disappear

Strongest conditional priors

Installed-base aftermarket with qualification and high-consequence service

The strongest operating prior begins with an installed base that is measured net of retirements, not with cumulative shipments. The model earns money through parts, repair, inspections, software updates, shop visits, consumables, uptime guarantees, and long-term service agreements attached to equipment already embedded in a customer's operation. The strongest versions combine proprietary qualification, safety or regulatory records, rapid local response, and a fee that is small relative to the cost of downtime or failure. That combination can make recurring service cash less cyclical than new-equipment revenue and allow moderate incremental capital to support a long runway.

The bargaining-power map is favorable only under specific conditions. The original equipment maker or specialist service provider holds power when it controls certified parts, design knowledge, performance data, or the liability for returning an asset safely to service. Customers hold power when fleets are concentrated, service contracts are rebid, or independent repair becomes credible. Skilled technicians and engineers can capture a meaningful share through wages when labor is scarce. Suppliers can capture value through proprietary subcomponents. Lessors and lenders matter when service providers finance inventories or customer equipment. The state can redirect value through safety rules, right-to-repair law, export controls, or defense procurement. The strongest class prior therefore belongs to service cash after labor, inventory, warranty, and sustaining capital, not to gross aftermarket revenue.

Durability can persist because qualification and service history accumulate. A provider that repeatedly restores uptime gains incident data, customer trust, parts logistics, and installed workflow that a low-price entrant cannot instantly reproduce. Yet the same installed base can become a wasting asset. A fleet peak followed by retirements, a transition to equipment with lower service content, open architecture, additive manufacturing, predictive maintenance, customer financial distress, or a manufacturer that discounts service to win new equipment can reverse the economics. Aging may raise service events for a time while shortening terminal duration.

The required evidence is a beginning-to-ending installed-base bridge, retirements, attach rate, service events per unit, real price and volume, parts and labor mix, contract renewal and cancellation, independent alternatives,

service free cash flow, inventory turns, warranty claims, and customer return on the total relationship. The falsifier is not one weak equipment quarter. It is sustained divergence between the installed base and service cash, accelerating retirement, falling attach or real price, loss of qualification, or incremental service capital returns falling below cost.

GE Aerospace is the positive research example used by the existing map because an installed engine base can support long-duration service demand; it remains Watchlist/Price Sensitive because supply constraints, contract economics, economic lives, and valuation remain open. Boeing is the counterexample: duopoly position, backlog, defense exposure, and an installed fleet do not compensate for safety, certification, fixed-price program, leverage, and governance failures. Boeing's own reports and SEC materials are evidence inputs, not a cure for the failed gates. Both names remain at current admissible weight zero.

The security-return trap is to capitalize a service annuity at a premium multiple just as the underlying fleet matures or customers renegotiate. The portfolio trap is to count an engine aftermarket, a semiconductor equipment service base, and an industrial-maintenance distributor as diversified when each depends on the same airline, fabrication, or industrial capital cycle. The business-quality prior is strongest conditional; the security and portfolio conclusions remain entirely open.

Ratings and select mandated assurance - strong economics with binary legitimacy tails

Mandated assurance is not one homogeneous class. Ratings has the strongest prior; testing, inspection, and certification, or TIC, is conditional but favorable; public audit has the weakest prior and belongs later in the unsupported or negative discussion. A ratings or selected assurance business earns fees because law, regulation, contract, insurance, or risk policy requires a trusted third party. The fee is often small relative to the transaction or avoided loss. Accreditation, accumulated methodologies, issuer and investor workflow, reputation, data, and regulatory recognition can limit entry. When incremental delivery is information-based rather than labor-based, additional volume can carry attractive margins.

The apparent toll is politically granted, so bargaining power is shared. The assurance provider can price when customers need its recognized credential and switching would delay a financing or product launch. Issuers can resist fees when concentrated or able to shop opinions. Regulators and courts possess the strongest latent power because they can open entry, rotate providers, cap fees, impose liability, or withdraw recognition. Skilled analysts, auditors, and inspectors can capture the productivity gain through wages, especially where human accountability cannot be automated. Customers may keep AI-driven savings if regulation requires more audit effort while procurement compresses the fee. Plaintiffs and insurers can claim the value pool after a failure.

This is why operating margin is an incomplete measure. The owner needs claims and penalties, staff cost, technology investment, acquisitions, liability insurance, and the possibility of a discontinuous franchise loss deducted before per-share cash is judged. Single-firm exposure also carries a binary severity that is not visible in an asset-light balance sheet. A business can report excellent margins until a conflict or inspection failure destroys legitimacy quickly.

Evidence should cover mandated revenue, organic price and volume, accreditation retention, regulator inspection or methodology findings, conflict controls, claim and penalty history, client and sector concentration, fee relative to customer economics, employee productivity and turnover, liability limits, and outcomes rather than management descriptions of quality. An unresolved conflict between filings and regulator or customer evidence scores zero. The class falsifier is a trend of worsening quality, independence breaches, liability relative to fees, forced rotation, material loss of recognized status, or price failing to cover required effort.

S&P Global is a positive research example for the ratings, indices, data, and workflow branch. Its map status is Core Research Candidate, not a completed underwrite; its current filing and operating evidence begins with company disclosures such as the Q2 2026 release. The counterexample is conceptual and historical rather than a current listed audit pure-play: a mandate did not protect Arthur Andersen from a rapid loss of legitimacy. That history explains why the class prior must be split and why no generic assurance label earns a pass.

The valuation trap is duration. A legally recurring fee stream can still deliver poor returns if the price assumes perpetual volume, margin, and regulatory privilege. The capital-cycle trap is subtler: technology can lower evidence-collection cost and attract software entrants while incumbents compete away the savings, or consolidation can raise acquisition prices for fragmented TIC firms. Ratings remains strongest conditional, TIC is favorable with labor and roll-up reservations, and audit has no automatic favorable prior.

Exchanges, clearing, market infrastructure, and payment rails

These models earn through transaction fees, clearing, market data, technology, connectivity, collateral services, value-added risk tools, and, in some payment systems, cross-border and authentication services. The best networks become more useful as liquidity, acceptance, trust, uptime, and workflow participation deepen. High fixed costs and low marginal costs can convert incremental activity into cash without proportionate new capital. Risk management, payment, and settlement needs recur even when the specific volume mix changes.

The value split is not automatically favorable to the platform. A venue has bargaining power when its liquidity minimizes total execution cost or its acceptance network raises conversion and reduces fraud. Large banks, brokers, merchants, and asset managers can multi-home, internalize flow, demand incentives, or lobby for fee changes. Public instant-payment rails can make basic transfer cheaper. Regulators can cap interchange, require access, mandate consolidated data, or change collateral rules. Customers often capture technological savings through tighter spreads or rewards. Clearing members and collateral providers bear or charge for tail risk. Cybersecurity and scarce engineering labor absorb part of the surplus. The owner captures durable value only if net take rate and service cash persist after incentives, losses, regulatory capital, and reinvestment.

The required evidence is product-level open interest and volume quality, liquidity and spread, take rate net of incentives, recurring data and technology mix, multi-homing, customer concentration, clearing resources and loss history, uptime, fraud and credit loss, capital requirements, and the regulatory share of economics. Volume generated by exceptional volatility should be normalized. A large user count or payment volume is not the same as pricing power.

CME Group is the map's positive exchange-and-clearing research example; its 2026 investor commentary is an evidence starting point. Visa is the positive global payment-rail example; current research begins with its fiscal 2026 results. London Stock Exchange Group and Moody's remain valuation-sensitive research examples for data, indices, workflow, and ratings. Coinbase is the counterexample to equating volatile crypto activity with durable infrastructure economics: transaction sensitivity, regulation, custody, asset-price exposure, and customer switching must be separated from the attractive word "exchange." Every name remains zero-weight.

The strongest falsifiers are simultaneous core-product share loss and material fee compression, public-rail substitution without offsetting service economics, a clearing or settlement impairment, regulatory destruction of the take rate, or a cyber event that causes persistent user exit. The security-return trap is extrapolating crisis volume or passive-asset growth into perpetuity. The portfolio trap is hidden commonality: an exchange, payments network, financial-data vendor, and bank can all depend on the same market liquidity, regulation, and financial-asset duration.

Reliability distribution and low-ticket mission-critical inputs

Reliability distribution earns a spread and service fee by placing the correct small part, consumable, tool, or safety item where and when a customer needs it. The customer is buying avoided downtime, fewer procurement touches, lower inventory complexity, and dependable local response rather than merely a physical item. Density can improve fill rates and route economics, while vending, on-site inventory, technical advice, and procurement integration can raise retention.

The distributor holds bargaining power when product cost is small relative to failure cost, local availability matters, and its branch or route density is difficult to reproduce. Manufacturers can take surplus through exclusive products or direct sales. Large customers can consolidate procurement and demand price transparency. Marketplaces can

commoditize search and basic fulfillment. Warehouse, driver, and sales labor can absorb productivity gains, particularly where service is local and urgent. Working-capital lenders and suppliers can capture cash when inventory and receivables expand. Customers retain value when the distributor's efficiency appears as lower stock, fewer outages, or lower purchasing labor. The shareholder receives the residual only if density raises gross profit and service cash faster than inventory, receivables, labor, and acquisition cost.

Evidence must show same-site or same-branch organic growth, price and units separately, gross profit per route, branch, or employee, customer retention, service levels and stockouts, inventory turns and obsolescence, receivables, supplier and customer concentration, digital order economics, and acquisition versus organic contribution. A large revenue base and thin margin do not establish a moat. Nor does an acquisition roll-up demonstrate organic density.

Fastenal is the positive research example because on-site service and replenishment can embed the distributor in customer operations. Its Q2 2026 release supplies dated operating evidence that still requires full-cycle normalization and cohort work. The present map did not source a fully developed reliability-distribution counterexample; that absence is research incompleteness, not proof the class lacks failures. A conceptual counterexample is an acquisitive distributor whose reported growth rises while inventory turns, same-branch volume, cash conversion, and leverage deteriorate. The missing named negative example must remain visible in future search rather than being converted into false confidence.

The falsifier is density investment failing to produce share and service gains, persistent cash conversion below the required full-cycle level, direct-channel substitution, or working-capital needs driving incremental returns below cost. The valuation trap is treating operational excellence as an infinite-duration bond. The portfolio trap is counting several distributors as diversified when their end customers share one industrial or construction cycle.

Record-to-action software and lawful proprietary data

The strongest version of vertical software is not a subscription with high historical retention. It is a system of record that controls consequential actions, approvals, integrations, audit trails, and liability-sensitive decisions, while lawful, unique, fresh feedback improves the next outcome. It earns money through subscription, usage, transaction, or outcome pricing and may expand through additional modules. The customer accepts the fee because migration risk and the cost of error are high relative to price.

This class is bifurcating rather than uniformly favored. The vendor has bargaining power if it owns workflow control, permissioned data, integration depth, and verified outcome improvement against the best foundation-model-plus-customer-data substitute. Foundation-model providers and cloud platforms gain power when they can bundle equivalent features. Customers gain power when data are portable and agents reduce paid seats. Scarce engineers, data licensors, and compute providers can absorb gross-margin gains. Regulators can raise assurance cost or create a mandate that many vendors can satisfy. The employee or customer can capture most of the productivity surplus even while the software remains useful.

The owner-cash bridge must therefore deduct compute, support, implementation, stock compensation, acquisition spend, and the cost of maintaining model and data quality. Reported free cash flow before recurring stock compensation can overstate diluted per-share capture. Revenue retention can also flatter economics when automatic renewals conceal declining usage or when price offsets seat loss.

Required evidence includes gross and net retention normalized for price, seats, usage, and acquisitions; module attach; critical workflow events; migration cost and defector evidence; data ownership, exclusivity, freshness, exportability, and deletion rights; outcome lift against current substitutes; AI compute and support cost; price per outcome; security incidents; and diluted owner cash. The bear-case falsifier is successful migration from seats to outcome or transaction pricing, stable real price and gross retention after AI deployment, proven unique data lift, and growing owner earnings after compute and dilution. The bull-case falsifier is seat contraction, bundled substitution, outcome parity using customer-owned data, or weakening per-share cash despite reported growth.

RELX is the positive map example for embedded legal, scientific, and risk workflow, with AI acting as both enhancer and substitute; current evidence begins with its 2025 results release published in 2026. Microsoft and Alphabet are valuation-sensitive examples rather than proof that scale platforms capture all AI surplus. Samsara is an emerging research option because connected-operations workflow and telemetry may develop a system-of-record moat, but short history, competition, device economics, and stock compensation prevent a mature prior. Palantir is the admired counterexample in the current map: an attractive data and workflow story cannot compensate for valuation and per-share-capture concerns. All remain zero-weight.

The capital-cycle trap is the ease with which venture or public capital funds near-substitutes and customer acquisition. The valuation trap is capitalizing past seat-based retention across an AI-driven pricing transition. The portfolio trap is owning several apparently vertical applications whose true scarce supplier is the same foundation model, cloud platform, or quality-duration factor.

Favorable but valuation-sensitive priors

Semiconductor, equipment, power-control, and industrial bottlenecks

The favorable prior applies to scarce, qualified enabling nodes, not to every beneficiary of compute, electrification, automation, or data-center growth. Examples include leading process manufacturing, lithography and other difficult equipment, advanced packaging, power management, switching, cooling, drives, controls, and specialized components whose failure can prevent a much larger system from operating. These firms earn through equipment sales, content per system, software, spares, service, and sometimes process or intellectual-property royalties.

The bottleneck supplier has bargaining power while qualification is slow, technical substitutes are scarce, lead times matter, and the component is a small share of customer economics. Large hyperscalers, semiconductor manufacturers, automakers, or industrial customers can regain power by dual-sourcing, sponsoring capacity, redesigning systems, or insourcing. Skilled engineering labor captures value through wages. Land, grid connection, and power owners can take an increasing share during data-center buildout. Suppliers of scarce materials and tools may capture upstream rent. Governments influence export permission, subsidies, local capacity, and state take. Debt and equity providers finance long buildouts and can capture much of a late entrant's value if utilization disappoints. Customers eventually benefit when new capacity and learning lower system cost.

This is a classic capital-cycle problem. High margins, backlogs, and long lead times invite entrants, customer prepayments, state subsidies, and incumbent expansion. The same observable can be both evidence of scarcity and a warning that supply is coming. A strong class prior requires evidence that the scarce process cannot be reproduced on the forecast timetable and that incremental capacity earns above its cost after price normalization. It does not permit a terminal value that freezes peak utilization, price, or market share.

TSMC and ASML are positive research examples for process, yield-learning, trust, and lithography scarcity; their current research sources include TSMC's Q2 2026 results and ASML's Q2 2026 filing exhibit. Schneider Electric is a positive power-management and automation example, with dated operating evidence on its financial-results page. Eaton is another valuation-sensitive power-control example. NVIDIA is an admired business whose current map fails the valuation pillar, demonstrating why extraordinary operating quality is not security-return proof. Intel is the principal counterexample to declaring a semiconductor harvest before net capacity exit, financing closure, balance-sheet survival, disciplined capex, and owner cash have been demonstrated.

Evidence must include customer concentration, qualification time, node or product share, utilization by vintage, price and volume, bookings and cancellations, service revenue per installed unit, economic asset lives, inventory and impairment, R&D and maintenance capital, customer return on the equipment, capacity additions across competitors, export policy, and free cash after capex. A two-year depreciation-life extension that raises reported margin is not proof of better economics. Option or reserve capacity may be credited only as an explicitly designated option at its observable market cost, never through a convenient longer asset life.

The falsifier is material share loss in the scarce node, qualification of an adequate substitute, utilization falling below an economic threshold while capacity expands, order-to-cash conversion breaking, or incremental return falling below the cost of capital. The security-return trap is paying for a decade of scarcity at the moment supply is responding. The portfolio trap is treating TSMC, ASML, NVIDIA, Schneider, Eaton, automation vendors, and data-center landlords as independent when they share hyperscaler capex, East Asian fabrication, power availability, policy, and valuation duration.

Disciplined insurance float, deposit franchises, and other balance-sheet businesses

Balance-sheet franchises correct an important bias in the favored-prior list. Insurance underwriting and deposit banking were omitted from the first model taxonomy because the prior was calibrated on a period that favored asset-light, long-duration quality. In inflation, higher nominal rates, fragmentation, or financial repression, a disciplined underwriter or low-cost deposit franchise can sometimes capture value that a premium-duration franchise loses. The class is favorable only through its own evidence battery; the label "financial" is not a shortcut.

An insurer earns underwriting profit and invests float held before claims are paid. Bargaining power comes from risk selection, distribution, trusted claims handling, data, scarce underwriting expertise, and a willingness to shrink when price is inadequate. Policyholders and brokers gain power when coverage is abundant. Reinsurers, catastrophe-bond investors, and debt providers price tail capital. Skilled underwriters and distribution partners take part of the surplus. Regulators and courts influence capital, coverage terms, and claims. Inflation transfers value toward claimants when reserves were understated. The shareholder benefits only after paid and incurred losses, adverse reserve development, reinsurance cost and recoverability, catastrophe capital, taxes, and share issuance.

A deposit bank earns a spread and fees from a funding franchise, payments, wealth, and credit intermediation. Depositors have bargaining power when they can move instantly or demand market rates. Borrowers gain when banks compete or the state directs credit. Employees can capture franchise value through pay, while regulators claim it through capital and liquidity requirements. The sovereign can suppress rates, direct lending, tax windfalls, or provide a safety net whose value is not freely distributable to shareholders. Wholesale creditors gain power when core deposits are weak. Owner value requires low deposit beta through stress, full-cycle credit losses inside pricing, strong capital without emergency dilution, and cash that can legally be distributed.

Asset management earns fees on assets and advice, but its prior is more uncertain. Clients and platforms possess strong power because products are comparable, fees are visible, and flows can reverse. Distribution platforms, advisers, index licensors, and scarce investment teams can capture much of the economics. Bull markets can make operating profit look recurring even while fee rates compress. Owner cash must be assessed after seed capital, performance compensation, acquisitions, stock compensation, regulatory capital, and cyclic flows. No current named asset-manager example has completed the battery, so the class remains under the no-qualifier rank rather than receiving an assumed favorable status.

The evidence batteries are non-substitutable:

Class	Required evidence	Owner-cash bridge	Primary falsifier
P&C insurance	Ten-year reserve triangles; prior-year development; combined ratio by line and cycle; float cost; catastrophe and casualty concentration; reinsurance recoverables; stressed capital	Underwriting and investment cash less true claims, reserve strengthening, reinsurance, capital retention, tax, and dilution	Cumulative adverse development material to opening equity, structurally inadequate pricing, or capital repair through dilution
Deposit banking	Deposit cohorts and beta; uninsured and operational concentration; loan vintages; non-performing loans and charge-offs; duration and liquidity; stressed CET1	Net interest and fees less normalized credit, deposit cost, operating loss, capital retention, tax, and issuance	Deposit-share loss with rapid repricing, losses above risk pricing, or ROE below cost across a full cycle
Asset management	Net flows by strategy and channel; fee rate; performance persistence; client and platform concentration; seed and regulatory capital	Fees less distribution, compensation, seed losses, acquisition spend, stock compensation, and required capital	Sustained net outflow and fee compression, team loss, or per-share cash lagging market-driven AUM growth

Berkshire Hathaway is a positive research example for low-cost float, retained operating cash, and capital allocation, but succession, scale, energy liabilities, reserve evidence, and sum-of-parts valuation remain open; its

2025 annual report is an evidence source rather than a completed battery. Chubb is the positive P&C example, and its Q2 2026 release begins, but does not complete, the reserve and capital work. DBS is the positive deposit-franchise example; its Q2 2026 release must be supplemented with deposit cohorts, full-cycle losses, property exposure, and stressed capital. All favorable class-battery scores remain provisional and capped at 2 until the named exhibits attach.

Flagstar Financial is the bank counterexample: apparent cheapness relative to book value cannot repair weak funding, commercial-real-estate credit, capital, or supervisory evidence. Nu Holdings is an emerging digital-bank option, not a mature franchise prior, because the record does not yet demonstrate a full Brazilian credit and deposit cycle. These names also remain at zero weight.

The valuation trap is using current combined ratio, net interest margin, or price-to-book without normalizing the cycle. The capital-cycle trap is underwriting abundance when new capital enters insurance after hard pricing, or assuming deposit value remains fixed when market rates and digital transfer change depositor behavior. The portfolio trap is treating balance-sheet franchises as diversifiers without looking through sovereign, property, catastrophe, duration, currency, and liquidity exposure.

Selective brands with cohort desire and controlled distribution

A selective brand earns through real price, repeat purchase, premium mix, scarcity, direct customer relationships, licensing, and sometimes services attached to a product community. The favorable prior belongs to a brand that remains desired by the next cohort, controls distribution, preserves product integrity, and converts scarcity or trust into cash without requiring proportionate advertising and capital. Familiarity, historic prestige, gross margin, or social-media reach alone does not qualify.

Bargaining power shifts quickly in this class. A genuinely scarce brand can ration supply and retain price. Retailers, digital platforms, search engines, and influencers can take distribution and discovery rents. Creative and commercial talent can capture value through pay. Landlords benefit from premium locations. Suppliers may take scarcity rents for inputs. Customers capture value when competition forces innovation or price concessions. Regulators and health policy can constrain marketing or product design. Entrepreneurs may create the initial desire, while public shareholders inherit a mature brand after growth and scarcity are fully priced.

Operating profit must be reconciled to diluted owner cash after marketing required to maintain relevance, wholesale support, inventory provisions, store and digital investment, acquisitions, stock compensation, and buybacks that merely offset issuance. Price/mix must be separated from units. A nominal sales increase can conceal declining unit demand, channel stuffing, promotion, or inflation. Resale value, wait lists, full-price sell-through, and repeat behavior can provide independent evidence of desire, but each can also be manipulated or distorted by temporary scarcity.

Ferrari is the positive quality example and the cleanest warning about the security layer. Controlled supply, personalization, and brand community can support excellent operating economics, yet the current map assigns a failed valuation pillar and zero diagnostic weight. Its corporate reporting page is the source for current operating disclosure, not an entry conclusion. Nike is the counterexample: historic global brand strength does not prove current cohort heat, full-price unit growth, or channel health; current research begins with its results and filings. Both remain zero-weight.

Evidence includes unit volume and price/mix separately, cohort purchase frequency, household penetration, full-price sell-through, returns and promotion, resale or order depth where relevant, geographic and channel mix, marketing efficiency, inventory age, distribution control, and incremental cash return. The falsifier is volume or cohort desire weakening while price, marketing intensity, inventory, or distribution support rises. The valuation trap is paying for permanent cultural relevance. The portfolio trap is assuming consumer diversification when several premium brands depend on the same affluent customer, Chinese demand, platform discovery, and long-duration multiple.

Funded health innovation

Funded health innovation includes therapies, diagnostics, devices, tools, and care-productivity systems with material clinical or economic benefit to a payer able and willing to pay. It earns through net product price, reimbursement, consumables, service, licensing, or shared savings. Aging and disease burden establish need, not collectible revenue. Scientific success and social surplus do not establish that the developer or device owner retains the value.

The value pool has unusually contested bargaining power. A patented therapy with differentiated outcomes can negotiate, but public and private payers use formularies, access, tenders, reference pricing, rebates, and total-cost-of-care analysis. Clinicians and hospitals control adoption. Patients influence adherence. Researchers and specialized labor take value through compensation. Manufacturers and distribution partners can hold scarce capacity. The state grants exclusivity and approval, funds science, regulates safety, and may negotiate price. Competitors can enter with me-too products, biosimilars, or better modalities. Customers and society may capture most of prevention's value through longer or healthier lives even when the provider's margin is modest.

The shareholder bridge starts at net price after rebates, not list price. It deducts manufacturing and distribution, required trials, replacement research and development, post-market commitments, litigation, acquisition spend, stock compensation, tax, and the probability and timing of exclusivity loss. A terminal value cannot bridge a patent cliff by assuming an unidentified replacement pipeline. A product may have high gross margin while the portfolio-level company destroys value through repeated acquisition of external science.

Eli Lilly and Novo Nordisk are valuation-sensitive research examples because funded metabolic innovation may produce material outcomes and revenue, while capacity, access, net price, competition, adherence, pipeline concentration, and duration remain decisive. They are not evidence that disease prevalence automatically creates shareholder returns. Teladoc is a counterexample to equating a useful care channel and growing demand with durable per-share economics; acquisition history, competition, payer power, and impairment matter. Current results should be sourced through the companies' regulatory and investor materials in any full underwrite; this chapter preserves the map's research classification rather than updating it.

Evidence includes indication-level trial and real-world outcomes, adherence, payer access, net price, total cost of care, manufacturing yield and capacity, exclusivity dates, competitive pipeline, replacement-pipeline risk-adjusted value, R&D productivity, acquisition spend, clinician workflow, and diluted cash per share. The falsifier is access or adherence failing to translate efficacy into outcomes, net price falling faster than volume and cost, manufacturing or safety failure, or replacement R&D failing before a material cliff. The valuation trap is capitalizing a current product as a perpetual annuity. The portfolio trap is counting several health names as independent when they share the same modality, payer response, manufacturing constraint, or regulatory policy.

Constructive regulated franchises and funded state procurement

Utilities, grids, water, climate adaptation, defense, and other regulated or state-payer models can have essential demand and long asset lives. They earn through an allowed return on prudent capital, contracted availability, cost-plus or fixed-price programs, equipment sales, services, or maintenance. The prior is favorable only where an enacted rate order, appropriation, or contract demonstrably converts to timely cash and where realized returns compensate all capital and liability. "Regulated" and "backlogged" are not synonyms for safe.

The state or regulator is usually the strongest counterparty. It can grant a monopoly and then cap its economics. Households and voters gain bargaining power through affordability politics. Labor, engineering firms, and material suppliers can capture infrastructure spending through wages and input inflation. Land and power rights can absorb data-center or grid rents. Lenders may claim most of a utility's cash because expansion is debt-financed. Customers can receive resilience and lower outage cost while the owner experiences negative free cash and dilution. Defense primes may possess qualification, clearances, and scarce capacity, yet a sovereign monopsony controls appropriations, specifications, export permission, audits, and renegotiation.

Operating profit is not owner cash if working capital funds a government program, construction overruns are disallowed, fixed-price contracts create charges, or rate-base growth requires repeated equity. For regulated utilities, allowed ROE must be reconciled to realized ROE, financing, tax, maintenance, and diluted free cash per share. For defense, backlog must be decomposed into appropriated funding, contract type, cancellation, customer advances, margin, delivery schedule, and program cash.

Rheinmetall is a Watchlist/Price Sensitive research example for a funded defense bottleneck, not a qualifier: state monopsony, working capital, program conversion, peace reversal, and valuation remain open, and the company's investor-relations materials are an evidence pointer. PG&E is the regulated-utility counterexample because necessary service and rate-base need cannot offset wildfire liability, political affordability, financing, and dilution without a complete jurisdiction and claims underwrite. Boeing also shows how funded defense exposure can coexist with fixed-price overruns and governance failure. No broad state-payer prior follows from fiscal activism or deglobalization.

Evidence includes enacted appropriations or rate orders, payment timing, allowed versus realized return, regulatory lag and disallowances, bill affordability, customer advances, cancellation terms, contract type, construction variance, working capital, capacity funded by customer or owner, realized program ROIC, debt and dilution, export approvals, and state take. The falsifier is funded awards failing to convert, working capital consuming the operating result, realized returns staying below cost, repeated disallowance, or political change reversing procurement. The valuation trap is treating backlog or rate-base plans as cash. The portfolio trap is owning defense, grids, adaptation, and localization as distinct themes when all depend on the same fiscal capacity and political legitimacy.

High-quality serial acquirers - favorable only after the treadmill is charged as reinvestment

A serial acquirer can create value by sourcing small businesses through seller trust, operating them under a disciplined decentralized system, improving stewardship without costly integration, and reinvesting retained cash at attractive prices. It earns through acquired operating cash and subsequent organic improvement. Financing is not the moat. When acquisitions are required to sustain reported growth, acquisition consideration is recurring reinvestment and must be deducted from owner cash.

The acquirer has bargaining power only if its reputation, permanent ownership, speed, and operating model make it a preferred buyer. Sellers gain power when capital is abundant and deal multiples rise. Managers and deal teams can take surplus through compensation. Lenders capture value when leverage replaces operating skill. Minority sellers can retain claims through earn-outs. Customers and employees may gain under better stewardship, but aggressive cost extraction can impair the franchise. Existing shareholders lose when stock issuance, contingent consideration, and rising deal size conceal declining organic economics.

Constellation Software is the positive valuation-sensitive research example for decentralized acquisition and capital allocation. Bausch Health is the counterexample for a debt-heavy healthcare roll-up where enterprise value, legal claims, refinancing, and organic erosion dominate headline franchise EBITDA; its SEC filings are the required starting record. A high historical acquisition return does not create an exception to the current-price or leverage gates.

The evidence reconstruction needs at least five years: price, volume, foreign exchange, acquisitions and disposals; all-in deal cost including shares, assumed debt, earn-outs, retention, restructuring, and integration capex; cohort cash IRR; recurring acquisition spend relative to pre-acquisition free cash; goodwill and intangible changes; payables, factoring, and supplier finance; contingent consideration; leverage, covenants, share count, and same-unit retention. The leading falsifier is declining organic volume combined with acquisition spend above internal cash, rising deal size or leverage, weakening conversion, changing definitions of organic growth, and increasing deferred consideration.

The valuation trap is using an adjusted free-cash-flow number that excludes the reinvestment required to keep the acquisition machine growing. The capital-cycle trap is confusing cheap debt and a rising currency with managerial

skill. The portfolio trap is owning several serial acquirers whose valuations, financing, and seller pipelines all depend on the same low-rate, quality-duration regime.

Emerging and unproven priors

Emerging challengers are not failed mature companies. They may possess a real cost advantage, new distribution, or an accelerating data loop, but they cannot produce multi-cycle evidence they have not yet lived through. The scorecard is intentionally conservative: short history, limited cash generation, customer concentration, incomplete governance evidence, or untested credit and technical cycles constrain ordinary admission. These names belong in research-only or speculative-option work, with total loss assumed and no implication that the policy's speculative cap is an authorized position.

Emerging class	Possible mechanism	Who may capture the early rent	Evidence required before the prior improves	Falsifier or downgrade trigger
Merchant AI capacity and neocloud	Specialized orchestration and scarce accelerators earn a service premium	GPU and networking suppliers, power and land owners, lenders and lessors, hyperscaler customers, scarce engineers; equity receives the residual	Enforceable funded contracts, customer diversification, vintage utilization and rental price, power rights, debt and lease waterfall, economic life, free cash after full capex	Rental price falls while utilization drops; interest and lease cash consume gross profit; refinancing or customer failure
Young vertical software or telemetry	Integrated data becomes an operational system of record and improves outcomes	Cloud and model platforms, employees through stock compensation, customers through productivity, and the application vendor if workflow control is real	Retention after seat/device normalization, lawful data rights, outcome lift, hardware subsidy, compute cost, fully diluted owner earnings	Bundling, low switching cost, NRR deterioration, gross-margin decline, or share count outruns per-share cash
Automation projects and warehouse systems	Proven deployments reduce labor and throughput cost	Anchor customers, systems integrators, component suppliers, project labor and lenders; vendor equity only after milestones convert	Site-vintage throughput and cash ROIC, acceptance tests, cancellation and warranty, customer concentration, project working capital, related parties	Deployment misses contractual output; cumulative project cash stays negative despite revenue and backlog growth
EM digital finance and commerce	Formalization, low servicing cost, logistics and payments density compound engagement	Customers through lower price and access, state through tax and direction, depositors and lenders, local labor, platform, controller, and minority owners in varying shares	Real local FCF/share, unit economics, full credit cycle, deposit beta, regulatory capital, FX funding, related parties, dilution, cash upstreaming and title	Loss-adjusted return below cost, subsidies resume, currency funding mismatch, trapped cash, leakage, or USD per-share cash stagnation
New health platform	Clinical or workflow innovation lowers total cost or improves outcome	Patients and payers may keep most savings; researchers, manufacturers, distributors and state also claim value	Comparative outcomes, reimbursement, adherence, net price, capacity, replacement science, diluted cash	Clinical, safety, access, manufacturing, or pipeline failure

CoreWeave is the map's merchant-capacity example. Its model may earn scarcity economics, but customer concentration, financing, leases, power, utilization, rental price, and economic depreciation can transfer most value to suppliers and creditors. Symbotic is the automation-project example; backlog cannot replace site-level cash evidence. Samsara represents the young vertical-workflow case. Nu and Sea represent EM digital finance, commerce, and payments. BYD represents a scaled but capital-intensive manufacturing challenger where unit growth can coexist with price war, trade barriers, and weak incremental returns. These are research examples only, all at admissible weight zero.

The common mistake is to assume that the entrepreneur or public equity holder captures the value because the product is innovative. During a buildout, scarce labor, chips, transformers, land, power, debt capital, and customer bargaining can take the surplus. Stock compensation can transfer value from existing owners to employees. New equity can fund growth while reducing each existing share's claim. A successful service can lower customer cost and still fail to earn its own cost of capital. The emerging prior improves only when the owner-cash bridge, not the adoption curve, matures.

Currently unsupported or without a qualifying example

This rank prevents absence of evidence from being disguised as either enthusiasm or prohibition. It also corrects the consensus-quality bias by keeping neglected classes in the search even when no current candidate passes.

Natural-resource and infrastructure harvest

Natural resources, shipping, memory, offshore services, and other hard-asset industries can become attractive after an investment bust. The mechanism is not scarcity rhetoric or a price decline. It is financing closure, net capacity exit, survivor balance-sheet repair, restrained replacement capital, stabilized utilization and price without subsidy, and normalized incremental cash return above the cost of capital. A low-cost resource with long reserve life or a royalty can improve the operating prior, but title, depletion, state take, environmental liability, and valuation remain.

Bargaining power begins with geology, permits, logistics, or scarce capacity, but shifts to the state through tax and title, to labor and contractors through cost inflation, to equipment and energy suppliers, to lenders when balance sheets are stretched, and to customers when substitution or recycling becomes credible. In early scarcity, owners may capture rent. High returns then fund entry, expansion, acquisitions, and state claims, moving surplus toward suppliers, labor, customers, and the sovereign. The relevant owner value is after sustaining capex, reserve replacement, reclamation, royalties, hedging, interest, tax, and dilution.

No full harvest candidate currently qualifies. The recorded search examined Micron or memory, Maersk or container shipping, and Transocean or offshore services and rejected them because financing-window closure, confirmed net capacity exit, balance-sheet survivability, and shareholder cash-return gates were not all evidenced. Intel also fails the semiconductor-harvest test because net exit and harvest cash are unconfirmed while state support remains. First Quantum is the commodity-scarcity counterexample: copper demand does not waive jurisdiction, legal title, debt, depletion, or sustaining capex; its financial reporting is a starting source, not proof of capture.

The two-stage test remains unresolved in timing. Financing-side evidence can justify research-candidate status before full harvest: issuance below prior norms, industry capex below depreciation, a material exit, and no state backstop sustaining marginal capacity. But the policy's final map found no candidate that satisfied the full evidence. Waiting may miss some rerating; moving early may fund zombie capacity. That tradeoff remains explicit rather than being silently resolved.

Audit, asset management, broad regulated resilience, and generic EM growth

Public audit has mandated demand, but labor intensity, wage exposure, rotation, conflict, and binary liability weaken the prior. The Big Four are private partnerships, so the current public-company map cannot provide a clean positive or negative pure-play. That structural absence is not an excuse to transfer the ratings prior to audit. Any listed audit-like candidate would require regulator inspection outcomes, independence evidence, partner and staff economics, claims relative to fees, and a severity assumption consistent with de-licensing.

Asset management remains unqualified because no positive or negative name was fully developed despite the class's explicit inclusion in the balance-sheet battery. This is research incompleteness. The next search must examine net flows, fee compression, client and platform concentration, compensation, seed capital, regulatory requirements, and capital allocation through market cycles. Rising AUM in a bull market is not sufficient.

Broad regulated resilience is unsupported as a class because utilities, grids, adaptation, defense, and water do not share one owner-cash mechanism. Each depends on its jurisdiction, contract, rate order, liability, funding, and affordability. Generic EM growth is similarly unsupported. Country GDP, digital users, or nominal local revenue can rise while currency, dilution, related parties, state direction, capital controls, or weak minority rights consume the result. The class prior improves only through audited operational capture, clean governance, and a separately acceptable macro and legal exposure.

Recurring negative priors

Negative priors identify mechanisms that repeatedly convert attractive demand into weak diluted per-share outcomes. They are not permanent prohibitions. A negative prior can be falsified by the required structural evidence, but it cannot be reversed by price decline, market size, management confidence, or a single favorable quarter.

Recurring negative prior	How the model appears to earn	Who usually captures the value instead	Capital-cycle or valuation trap	Evidence needed to reverse the prior	Current research example or counterexample
Financed capacity flood	Volume growth and scarce capacity	Customers through falling prices; equipment suppliers, labor, land/power, and lenders during buildout	Peak margins attract supply; price-to-sales ignores depreciation and financing	Financing closure, net exit, disciplined survivor capex, utilization and price stability, above-cost owner cash	Intel fails current harvest; Plug Power illustrates subsidy and financing dependence
Leveraged merchant compute	Accelerator rentals and orchestration premium	Chip vendors, network suppliers, power and land owners, creditors, and concentrated customers	Short contracts fund long-lived assets; economic life is overstated	Funded take-or-pay coverage, diversified customers, vintage cash returns, conservative asset lives and debt waterfall	CoreWeave remains speculative research only
Acquisition treadmill	Adjusted EPS and acquired growth	Sellers, lenders, management, advisers, and new share recipients	Cheap capital and rising stock currency hide weak organic economics	Organic volume, cohort cash IRR, acquisition spend charged as reinvestment, stable leverage and per-share cash	Bausch Health counterexample; Constellation is favorable only after full reconstruction
Undifferentiated seat software	Recurring subscription and reported retention	Platform, foundation-model provider, customers and employees	AI reduces seats and bundles features while historical NRR supports a premium multiple	Outcome pricing, unique lawful data, workflow control, stable normalized retention, owner cash after compute and dilution	Palantir current Avoid on valuation/capture; Samsara emerging rather than mature
Backlog without funded cash	Announced orders and multiyear demand	State payer, prime contractor, labor and suppliers; equity funds working capital	Appropriations, cancellations, fixed-price overruns and late capacity	Contract funding, cancellation compensation, advances, program cash and normalized ROIC	Rheinmetall remains Watchlist; Boeing is the governance and program counterexample
Capital-intensive turnaround or cheap replacement value	Recovery in utilization and accounting margin	Creditors, pension or legal claimants, labor, suppliers and customers	Low multiple reflects cash burn, obsolescence or required reinvestment	Independent quality repair, balance-sheet survival, positive cash after maintenance, no emergency dilution	Boeing and Intel; low price does not clear C or D
Legacy brand without current cohort desire	Pricing and historic awareness	Retailers, platforms, influencers, labor and marketing channels	Price/mix masks unit decline; terminal growth assumes relevance	Full-price unit and repeat growth, healthy inventory, share recovery without rising marketing intensity	Nike counterexample; Ferrari itself fails valuation despite quality
Commodity scarcity with weak title or leverage	Spot-price NAV and reserve value	Sovereign, labor, contractors, lenders and acquirers at the top	Depletion, sustaining capex, tax, closure and price mean reversion ignored	Conservative mine-level cash, durable title, balance-sheet survival, per-share reserve economics	First Quantum counterexample
High-growth EM platform with weak capture	Users, GMV, payments or local revenue growth	Customers via subsidy, controller and related parties, state, creditors, employees, and new shares	Nominal and currency growth obscure real per-share value	Three-layer capture: real local operations, macro currency cost, governance leakage; title and upstreaming	Adani Enterprises counterexample; Nu, Sea and BYD remain options
Regulated essential service with unfunded recovery	Rate base, essential demand, or public mandate	Customers and state through affordability, lenders through financing, contractors through capex	Allowed ROE is mistaken for realized cash; equity finances the plan	Binding recovery, realized full-cycle return, manageable bills, free cash after capex and dilution	PG&E counterexample
Single-product health duration	High gross margin and large disease burden	Payers, competitors, researchers, manufacturers and state	Patent or exclusivity cliff is hidden in terminal value	Net access, adherence and outcomes, replacement pipeline, R&D productivity and diluted cash	Teladoc as broader health-platform counterexample; Lilly and Novo remain valuation-sensitive

Several named Avoid cases reinforce the principle. Plug Power shows why subsidy, customer aspiration, and a large hydrogen opportunity do not cure financing and cash economics. Adani Enterprises shows why nominal EM

growth does not cure related-party, governance, state, and per-share-capture questions. Flagstar shows why low price-to-book cannot replace a bank battery. Bausch Health shows why enterprise value and recognizable franchises cannot be separated from the debt and legal waterfall. These classifications are dated and falsifiable; none authorizes a short sale or predicts a price decline.

Cross-class value-capture map for 2026-2036

The decade themes matter only through their effect on bargaining and cash distribution. The following synthesis states where capture appears more likely to strengthen, weaken, or remain radically uncertain. It is a research map, not a forecast allocation.

Decade force	Likely scarce input and early bargaining power	Where surplus may migrate as the cycle responds	Conditional implication for model priors	Leading distribution indicators
AI and software diffusion	Accelerators, power, data rights, scarce engineers, foundation platforms, trusted systems of record	Toward customers through lower prices and fewer seats; toward platforms through bundling; toward verification and liability when error cost matters	Strengthen lawful record-to-action and assurance research; weaken generic seat software and leveraged merchant capacity	Price per token or rental, seats and workflow events, outcome price, compute and labor share, capex, diluted FCF/share
Compute and data-center buildout	Advanced equipment, packaging, networking, grid access, land, cooling and power	New capacity can transfer rents to customers; debt and lessors can claim weak projects; state can condition permits and subsidies	Favor scarce qualified bottlenecks but make valuation and capital-cycle tests severe; merchant capacity emerging/unproven	Utilization by vintage, rental price, cancellations, capacity additions, power reservation to actual load, interest and lease burden
Electricity, grids and electrification	Grid connection, transformers, switchgear, engineering labor, permitted generation and transmission	Affordability policy, rate disallowance, local opposition and new supply can move value to customers or state; land and power rents may persist	Favor power-control bottlenecks selectively; regulated owner capture remains mixed	Allowed versus realized ROE, lead times, order conversion, input wages, land/power price, capex and dilution
Automation and labor scarcity	Integration know-how, qualified components, operational data and field service	Skilled labor and integrators can keep the gains; customers retain savings if vendors compete; automation can lower future wage power	Favor reliability distribution and proven workflow; treat project automation and backlog as unproven until site cash appears	Customer payback, throughput, labor hours saved, vendor price realization, warranty, project working capital
Aging and health innovation	Clinical differentiation, evidence, manufacturing capacity and trusted care workflow	Payers and state negotiate savings; clinicians and researchers take rents; competition erodes exclusivity	Funded health innovation favorable but highly valuation-sensitive	Net price, access, adherence, total care cost, capacity, R&D and acquisition spend, per-share cash
Fiscal activism, defense and resilience	Appropriation, clearance, qualification, local capacity and scarce engineers	State monopsony, windfall tax, fixed-price terms and labor/input inflation can absorb supplier economics	State-procurement models require funded contract-to-cash evidence; no generic defense prior	Appropriation to award to cash, advances, cancellation, working capital, realized program ROIC, state take
Deglobalization and economic security	Trusted jurisdiction, duplicate capacity, logistics and inventories	Security benefits accrue to society while utilization falls; suppliers and labor gain during build; taxpayers fund duplication	Favor qualified bottlenecks and reliability selectively; weaken undisciplined capital builders	Utilization, subsidy share, local content cost, inventory, tax burden, supplier concentration, after-capex return
Inflation, real rates, debt and repression	Short repricing lag, low-cost float or deposits, real pricing power	Depositors, claimants, labor and state can reprice; creditors gain from leverage; price controls punish visible rents	Add disciplined balance-sheet franchises to the research canon; weaken levered duration and unfunded regulated models	Deposit beta, reserve development, wage and input share, interest burden, tax take, real price and unit volume
Resource depletion and harvest	Low-cost reserve, permit, logistics and surviving capacity	State, labor, contractors and new entrants capture scarcity; recycling and substitution benefit customers	Keep harvest as a two-stage research option, not a standing favored class	Issuance, capex/depreciation, exits, utilization, sustaining cost, reserve replacement, state take, diluted FCF/share
EM formalization and digital platforms	Local distribution, trust, logistics, deposits and network density	Subsidies benefit customers; state, controller, creditors, employees and currency loss may absorb value	Favor only audited per-share capture; generic country growth unsupported	Real local FCF/share, FX versus inflation differential, dilution, related parties, upstreaming, credit vintages
Climate adaptation and physical resilience	Risk data, insurance capacity, permits, engineering, water and grid assets	Claimants, reinsurers, contractors, landowners and state capture rising need; affordability limits owner return	Favor disciplined insurance and select assurance; regulated resilience remains jurisdiction-specific	Loss trend, reserve development, reinsurance price, rate recovery, project cash, tax and litigation
Boring or plateau decade	Ordinary replacement, compliance, payment, distribution and service needs	Customers retain broad productivity gains; overbuilt thematic capacity loses pricing	Strengthen research on ordinary compounders and harvest survivors; comparator may prevail over all active priors	Organic cash per share, maintenance demand, capex exits, valuation spreads, active results after all costs

The strongest expected capture is therefore conditional, not absolute. It may strengthen where a provider controls a licensed, trusted, high-consequence workflow; where a qualified installed base creates recurring service; where liquidity or acceptance improves user utility; where density lowers the total customer cost; or where disciplined underwriting owns low-cost float. It may weaken where financing, labor, power, land, platforms, or the state own the scarce input; where entry is easy; where customers can multi-home or substitute; where a regulated return requires continuous external equity; or where acquisition and stock compensation consume operating cash. It remains radically uncertain in AI capex, health innovation, state procurement, EM platforms, and resource harvest because the social value and demand case can be strong while the split of value is still changing.

Required evidence and class-level falsification

A prior becomes useful only when it changes what evidence is collected and what would cause the ranking to change. It cannot be defended by adding narrative after the fact. The evidence hierarchy favors filings, contracts, regulators, and independently reproducible calculations, but filings remain management-produced. A material contradiction from credible customer, competitor, regulator, or physical evidence scores zero until resolved. Evidence Coverage below 60 remains research-only; higher coverage scales only the permissible loss budget and never waives a failed pillar or creates a buy signal.

Prior	Minimum evidence packet	36-month class milestone	What would lower or reverse the prior
Installed-base aftermarket	Net base and retirements, attach, service cash, independent alternatives, labor and inventory economics	Service cash tracks qualified base with stable real price and conversion	Base retires, access opens, content falls, or customers capture price
Ratings/select assurance	Mandated share, regulatory quality, conflict and liability, organic fee/volume, staff economics	Quality remains stable while price and cash cover effort and claims	De-licensing, liability, fee cap, entry, or automation benefits retained by customers
Exchanges/rails	Liquidity cohorts, take rate, incentives, losses, regulation, uptime and capital	Network utility persists without deteriorating net yield or tail capital	Share and fee decline together, public-rail substitution, clearing or cyber impairment
Reliability distribution	Same-site organic bridge, density productivity, inventory and receivables, substitution	Density produces share, service and cash gains through a cycle	Working capital and labor absorb gains; direct channel wins
Record-to-action software	Normalized retention, workflow events, rights, outcome lift, compute, diluted cash	Outcome pricing and owner cash persist after AI and seat normalization	Bundling, portability, nonexclusive data and dilution overwhelm workflow
Bottlenecks	Qualification, capacity and utilization, customer ROI, economic lives, FCF after capex	Scarce share and cash persist as supply responds	Substitute qualifies, utilization falls, returns fall below cost
Insurance/banks	Full class battery and stressed capital	Underwriting or deposit advantage survives stress without dilution	Reserve, credit, funding or capital failure
Brands	Unit and price/mix, cohort repeat, inventory, channel, marketing, diluted cash	Full-price cohort demand persists without rising support	Unit and cohort decline masked by price or promotion
Health	Outcome, access, net price, capacity, exclusivity and replacement science	Clinical benefit converts into durable net cash after required R&D	Payer, safety, manufacturing, competition or cliff absorbs value
Regulated/state payer	Enacted order or contract, cancellation, realized return, working capital and funding	Award or rate order converts to owner cash at above-cost return	Disallowance, delay, overrun, peace reversal, affordability or dilution
Serial acquirer	Organic bridge, all-in deal cost, cohort IRR, treadmill and leverage	Organic health and per-share cash persist as financing normalizes	Rising deal size, leverage and deferred cost with organic deterioration
Emerging challenger	Rights, funding, concentration, unit/vintage cash, governance and dilution	At least one full stress or cohort demonstrates self-funded economics	Refinancing, customer loss, price collapse, project miss or per-share leakage
Harvest candidate	Financing closure, capex below depreciation, exit, no zombie backstop	Full exit, utilization and cash-return confirmation	Issuance reopens, state sustains capacity, or cash remains sub-cost after 24 months

Two kinds of falsification must remain distinct. A class can remain operationally excellent while its securities underperform because starting valuations were too high; that falsifies security use or the valuation pillar, not necessarily business quality. Alternatively, high-prior classes can systematically produce weaker real diluted cash per share, more impairment, more restatement, or more distress than lower-prior classes. That would challenge the taxonomy itself. The existing debate's proposed validation must be treated cautiously: three years can test gates and causal milestones, not prove a ten-year moat. A larger, preregistered and class-balanced cohort, pairwise

predictions, retained failures, valuation controls, and independent scoring are necessary before any systematic ranking claims predictive authority.

Positive examples and counterexamples are paired research designs

The company map is most useful when a positive and negative case are studied together. The positive case reveals what evidence an attractive mechanism should produce. The counterexample reveals which apparently similar fact is not enough. None of these pairs implies a long-short trade.

Model question	Positive research example	Counterexample or limiting case	What the pair is meant to teach
Can an installed base create owner cash?	GE Aerospace	Boeing	Fleet and service need matter only after safety, program cash, balance sheet and governance
Can a mandated information toll persist?	S&P Global; Moody's	Audit de-licensing history; Palantir as valuation/capture limit, not an assurance peer	Mandate and workflow can be strong, but legitimacy and price remain separate
Does network density improve economics?	CME; Visa	Coinbase	Volume or user growth is not enough; liquidity quality, net take, regulation, capital and tail risk matter
Does local density beat commodity distribution?	Fastenal	Current named negative example missing	Same-site cash, service levels and working capital are the test; absence must remain disclosed
Can proprietary workflow absorb AI?	RELX	Generic seat software; Palantir current Avoid	Rights, action control, outcome lift and diluted cash matter more than a flywheel diagram
Is a technical bottleneck durable?	TSMC; ASML; Schneider	Intel; NVIDIA's E0 valuation case	Scarcity may be real while capital-cycle, geopolitical and price conclusions differ
Does a balance sheet compound rather than hide leverage?	Berkshire; Chubb; DBS	Flagstar; Nu as short-cycle option	Reserve, float, deposit, credit and capital evidence must constrain favorable as well as weak cases
Does scarcity create a durable brand security?	Ferrari on business quality	Nike on current demand; Ferrari itself on valuation	A-D quality and E valuation can point in different directions
Does health need become per-share cash?	Lilly; Novo as research cases	Teladoc	Clinical or social usefulness is upstream of access, net price, replacement investment and capture
Does state demand become owner return?	Rheinmetall as incomplete Watchlist	PG&E; Boeing	Appropriation, contract type, recovery, working capital and political legitimacy decide capture
Does repeated acquisition create value?	Constellation Software	Bausch Health	Acquisition consideration is reinvestment; leverage and organic health decide the owner result
Has a hard-asset industry reached harvest?	No current qualifier	Intel; First Quantum; examined memory, shipping and offshore cases	Price decline and consolidation rhetoric are not financing closure and survivor cash
Does EM growth reach minority shareholders?	MercadoLibre as price-sensitive research; Nu and Sea as options	Adani Enterprises	Local operations, macro currency, governance leakage and legal title must be decomposed

The map's Core, Watchlist, Speculative Option, and Avoid labels expire if not refreshed. They are not permanent attributes of the companies. A positive example can become a counterexample through price, governance, regulation, cycle, or technical change. A counterexample can improve if the preregistered evidence is met. The research process must preserve both the original classification and the evidence that changed it.

Capital-cycle and valuation traps common to the favored list

The favored list is most vulnerable precisely where it appears safest. High returns and admired franchises attract capital, competitive imitation, regulation, and premium valuation. Several recurrent traps require explicit defense.

First, **scarcity is a state, not a terminal value**. Long lead times, queues, and backlogs can encourage customers, states, and competitors to finance substitutes. The owner must model post-scarcity pricing, not merely current

demand. Second, **recurrence is not eternity**. An installed base retires, a regulatory mandate changes, a seat disappears, a patent expires, and a brand cohort ages. Third, **asset-light can conceal liabilities**. Ratings and audit face claims; payments and clearing face operational tails; software may issue stock and acquire products; insurers and banks carry leverage by design. Fourth, **reported free cash flow can omit required reinvestment**. Acquisitions, stock compensation, model training, customer incentives, reserve strengthening, regulatory capital, and right-to-operate investment may all be economically recurring.

Fifth, **quality classes can share one valuation factor**. A payments network, exchange, data vendor, vertical software company, branded scarcity franchise, and equipment monopolist can all be long-duration securities priced for stable high returns. A rise in real rates or a no-multiple-recovery decade can make operational diversification irrelevant. Sixth, **the comparator may already own the factor**. Adding a favored AI or power bottleneck can deepen rather than diversify embedded exposure. The current proxy factor work is unratified, but while original limits are exceeded, active additions with material AI or power sensitivity remain inadmissible regardless of company quality.

Seventh, **the state can socialize demand and privatize little profit**. Industrial policy, defense, grids, health, and adaptation may be socially valuable while procurement, tax, affordability, and local-content rules cap owner economics. Eighth, **debt can own the buildout**. Merchant compute, utilities, resource projects, and roll-ups may show revenue and EBITDA growth while interest, leases, and refinancing absorb the residual. Ninth, **a low valuation does not establish harvest**. Cheap book value, replacement value, or spot-price NAV may simply indicate obsolescence, weak title, capital need, or liability.

Reverse discounted-cash-flow analysis should therefore ask what duration, price, margins, capital intensity, dilution, and terminal advantage the current security already assumes. It should include a boring-decade case in which dramatic themes disappoint, customers capture productivity, ordinary maintenance businesses persist, and the low-cost comparator wins after cost. If a favorable class clears only through multiple expansion, unexplained terminal share, peak scarcity, defensive-capex add-backs, or an unidentified replacement product, the security conclusion fails even if the business prior remains favorable.

What not to infer from the ranks

The ranks do not say that a Strongest Conditional company deserves a larger position than an Emerging company. Position arithmetic occurs later, after evidence, valuation, liquidity, factor, severity, funded state, and approval. The ranks do not say that a Recurring Negative Prior must be shorted. A negative prior can be widely known, priced, difficult to borrow, or vulnerable to a sharp harvest reversal. The ranks do not imply a sector allocation, geographic allocation, hedge, timing signal, or view that active research is economically justified.

They also do not remove the policy's bias. Eight of the eleven current Core Research Candidates resemble the consensus-quality canon, and the map is North America-heavy. The counter-canon search produced balance-sheet financials but no full harvest or asset-manager qualifier. This composition is evidence about the screen's lens, not evidence of forward returns. Annual refresh must explicitly search non-US, financial, asset-management, deep-value, and verified-harvest candidates and report named rejections rather than leaving blank cells that look like conclusions.

The ranks are not certified. Company Evidence Coverage remains provisional, no named company has completed the independent adversarial-corroboration pass, and investor-specific facts remain missing. Research can be suspended entirely if the active-viability equation fails, if the operator cannot maintain the workload, if the comparator cannot be scored reproducibly, or if liability closure remains unmeasured. In that state, the appropriate use of these priors is a paper journal or annual research map, not capital deployment.

How these priors should change research, not authorize ownership

Use the priors to sequence questions. Begin where a durable value-capture mechanism is plausible, but write down before company work who owns the scarce input, which counterparty can walk away, who finances the growth, who retains the productivity saving, and how operating profit becomes parent cash and then diluted per-share value. Pair every positive example with a counterexample. Record the capital-cycle state, a normalized reverse DCF, the simplest comparator or abstention alternative, the required class battery, and at least one observable that could reverse the prior within 36 months.

A strong prior should raise the standard of skepticism, not lower it, because the most admired classes are precisely those most exposed to crowded valuation and narrative persistence. A negative prior should specify the structural evidence that could overturn it, especially financing closure and verified harvest, so that skepticism does not become dogma. Missing evidence keeps the work research-only; Evidence Coverage below 60 keeps it research-only by rule; a failed pillar, legal-title problem, factor breach, or absent professional review keeps admissible weight at zero regardless of narrative quality.

Nothing in this chapter authorizes ownership. Every named company remains at current admissible weight zero until a separate current security underwrite, verified investor facts, portfolio look-through, tax and legal review where applicable, Evidence Coverage audit, factor headroom, and owner approval all pass. The practical output of a prior is a better research queue and a clearer falsifier - never a buy list, target weight, order, or portfolio mutation.

Chapter 8 - From Worldview to Investment Themes and Business-Model Winners and Losers

This chapter is the bridge between the dated 2026-2036 worldview and security research. It does not convert a worldview into a portfolio. Its purpose is narrower: to explain why a force that may persist for a decade changes the prior probability that a business model will create and retain diluted per-share owner cash. The bridge ends before expected return. A plausible model must still pass the five non-compensating pillars, current evidence coverage, legal and payer gates, conservative valuation, common-factor limits, portfolio fit and owner approval. Every named company remains at current admissible weight zero.

The distinction between social success and shareholder success is controlling. A trend is socially successful when it raises productivity, security, health, resilience or consumer welfare. A business can help produce that success and still disappoint its shareholders. Customers may receive the benefit through lower prices. Labor may capture it through scarce-skill wages. Landlords, power suppliers and other rent claimants may raise their charges. Lenders may absorb cash through interest and covenants. The state may use taxes, procurement power, price controls or mandated reinvestment. New entry can bid away an early margin. A parent can fail to receive cash from an operating subsidiary, and dilution can prevent total corporate growth from becoming growth per share. The useful chain is therefore:

> *observed condition -> causal persistence -> economic need or value pool -> payer -> bargaining-power split -> capital-cycle response -> operating-company cash -> parent-company cash -> diluted per-share owner cash -> price paid -> portfolio role.*

This chapter labels its propositions as follows. **Durable policy** means the translation discipline and non-authorization boundary. **Dated worldview hypothesis** means a proposition about 2026-2036 whose evidence cutoff is 22 August 2026. **Business-model prior** means a favored or disfavored hypothesis to research, not a conclusion about a security. **Illustrative company** means a model example from the dated company map, not a current factual update, ranking or recommendation. **Unresolved** means the policy preserves competing mechanisms because the evidence does not settle them.

The master crosswalk

The table compresses the bridge. “Favored” means that the model has a plausible route to bargaining power and per-share capture if evidence confirms the chain. “Disfavored” means that the model is unusually exposed to weak payers, competitive supply, economic depreciation, state leakage, dilution or a price-taking position. Neither label establishes valuation or expected return.

Worldview force	Persistent economic need or value pool	Payer and likely early surplus claimant	Favored business-model hypothesis	Business models likely to disappoint	Capital-cycle and valuation warning	Leading indicators and falsifiers
AI and software diffusion	Lower task cost, faster decisions, governed automation and error reduction	Enterprises and consumers pay; customers, scarce technical labor, foundation platforms and workflow incumbents contest the gain	Record-to-action workflows, lawful unique data, verification, security, liability transfer, outcome-priced vertical software	Undifferentiated applications, seat-only pricing, generic content tolls, services that cannot retain productivity savings	Low distribution cost invites entry; compute and model costs can transfer value upstream; “AI” multiples can discount the full diffusion	Production use, customer savings, retention after seat normalization, real price, error/liability, inference cost; falsified by commoditization or no customer return
Compute, data centers and semiconductors	Training and inference capacity, advanced fabrication, memory, networking and cooling	Hyperscalers and enterprises pay; equipment bottlenecks, power owners, lenders and chip designers may capture before capacity owners	Qualified equipment oligopolies, trusted foundry, scarce interconnect and service installed bases where full-cycle returns persist	Levered merchant capacity, undifferentiated hosting, builders using long accounting lives, late entrants	High returns attract enormous capex; rental prices, utilization and economic lives can collapse before depreciation catches up	Utilization by vintage, rental/token price, cancellations, inventory, impairments, customer AI gross profit, capex/depreciation; falsified by financed flood

Worldview force	Persistent economic need or value pool	Payer and likely early surplus claimant	Favored business-model hypothesis	Business models likely to disappoint	Capital-cycle and valuation warning	Leading indicators and falsifiers
Electricity demand, grids and power constraints	Reliable generation, connection, transmission, distribution, cooling and power quality	Utilities, data centers, industry and households pay; land, interconnection rights, equipment bottlenecks, utilities and the state divide rents	Power management, grid control, qualified components, efficient load, rate-recoverable networks	Unfunded grid plans, speculative generation or connections, fixed-price contractors, rate-base growth without owner cash	Long queues invite duplicate applications and capacity; affordability politics can cap returns while forcing investment	Energized load versus reservations, queue conversion, transformer lead times, allowed-to-earned return, outages, customer bills; falsified by demand cancellation or rapid abundance
Electrification and industrial automation	Substitute electricity and machines for fuel, scarce labor and unreliable manual processes	Industrial customers and households pay; customers may keep savings, while integrators, skilled labor and qualified equipment share value	Mission-critical automation, controls, sensors, service aftermarket, reliability distribution with demonstrated payback	Project integrators without IP, commodity hardware, subsidy-dependent manufacturers, automation vendors whose customer return is unproven	Subsidies and attractive growth invite capacity; interoperability and Chinese or other low-cost supply can compress hardware margins	Customer payback, utilization, service attach, unit economics after subsidy, order-to-cash, inventory; falsified by weak productivity or overcapacity
Aging, healthcare innovation and labor scarcity	More care with fewer workers, prevention, treatment, diagnosis and care productivity	Governments, insurers, employers and patients pay; labor, drug/device owners, providers and payers contest savings	Funded therapies with material outcomes, tools that remove scarce labor hours, diagnostics tied to action, disciplined insurers	Unreimbursed innovation, labor-heavy providers with capped prices, patent-cliff franchises, access platforms without clinical savings	Scientific success attracts competitors; payer monopsony, net-price resets and manufacturing expansion can shift value away from innovators	Net price, access, adherence, outcomes, labor hours saved, total cost of care, exclusivity and replacement pipeline; falsified by poor real-world economics
Fiscal activism, defense and state procurement	Deterrence, strategic capacity, infrastructure and public services	The state pays from taxes or borrowing and acts as monopsony; labor, primes, bottlenecks and taxpayers divide surplus	Scarce qualified components, service installed bases, funded cost-recoverable contracts, disciplined state-capex suppliers	Budget-story companies, fixed-price development, working-capital funders, politically visible surplus, capacity built before appropriation	Backlog attracts capacity and acquisitions; elections, debt stress or peace can strand it; the state can cap margin after suppliers invest	Appropriation-to-cash , contract type, advances, cancellations, cost overruns, cash conversion, capacity utilization; falsified by unfunded plans or state extraction
Deglobalization, resilience and economic security	Redundant supply, trusted technology, inventories, compliance and localized capacity	States and corporations pay; workers, local land, incumbent bottlenecks, logistics providers and protected entrants compete for rents	Testing and certification, trusted foundry, multi-sourced mission-critical distribution, traceability and logistics control	Subsidized duplicate factories, domestic champions without discipline, low-utilization localization, single-route “resilience”	Protection raises returns and simultaneously funds excess local supply; retaliation can strand assets and suppliers	Local-content rules, export licenses, utilization, subsidy share, inventory, price premium, free cash after duplication; falsified by normalization or uneconomic redundancy
Inflation, real rates, debt and financial repression	Preserve real purchasing power, refinance debt, reprice essential goods and match liabilities	Debtors, savers, taxpayers and customers bear the adjustment; banks, the state, labor and pricing-power firms contest nominal gains	Short-cycle repricers, low-cost funding franchises, sound balance sheets, verified real-liability matching, essential distributors	Long-duration fixed nominal claims mismatched to liabilities, levered duration, visible windfall earners, price takers with long working capital	High nominal returns attract leverage; higher real rates expose weak projects; repression can confiscate visible surplus without formal default	Real cash per share, price versus volume, wage/input lag, interest burden, tax/state take, deposit beta, index integrity; falsified by disinflation without repression or inability to pass costs
Natural-resource depletion and harvest cycles	Replace depleted reserves and supply energy, metals, food and materials	Industrial customers pay; resource states, labor, land, equipment suppliers, lenders and low-cost owners divide scarcity rents	Low-cost long-life assets, disciplined survivor producers, royalties only with sound title and tax terms, service firms in confirmed harvest	High-cost or single-asset producers, spot-price NAV stories, serial reserve buyers, levered builders, poor-jurisdiction claims	Scarcity prices finance exploration, substitution and new capacity; depletion and grade decline consume accounting profit	Reserve replacement cost, grade, sustaining capex, capacity exit, inventories, state take, per-share production and cash; falsified by new supply or demand substitution

Worldview force	Persistent economic need or value pool	Payer and likely early surplus claimant	Favored business-model hypothesis	Business models likely to disappoint	Capital-cycle and valuation warning	Leading indicators and falsifiers
Emerging-market formalization and digital platforms	Move commerce, finance, logistics and records into formal digital systems	Consumers, merchants and borrowers pay; platforms, banks, telecoms, the state and local entrepreneurs contest network value	Two-sided platforms with falling unit cost, audited cash, sound credit and clear minority-owner rights	User-growth stories funded by subsidy, related-party groups, foreign-currency leverage, weak title, credit disguised as payments growth	Cheap capital funds customer acquisition and credit; regulation and local competitors can reset take rates; FX can erase local growth	Cohort contribution, take rate, fraud/credit loss, funding currency, dilution, related parties, hard-currency per-share cash; falsified by growth without capture
Urbanization and infrastructure	Agglomeration, mobility, housing, water, sanitation, logistics and reliable urban service	Users, municipalities and taxpayers pay; landowners, builders, financiers, workers, platforms and the state divide the uplift	Small-ticket maintenance, metering, leakage control, enforceable service infrastructure and clear-title systems with funded collection	Fixed-price megaprojects, speculative land banks, foreign-currency-funded local tariffs, concessions built on heroic volume	Infrastructure raises land rent before operator cash; overbuild, congestion and debt can turn population growth into poor returns	Formal jobs and output per worker, commute, access, collection, utilization, housing burden, land-value capture; falsified by urban growth without productivity
Climate adaptation and physical resilience	Reduce damage, downtime, water stress, fire, flood and insurability loss	Property owners, insurers, utilities and governments pay; contractors, skilled labor, land owners, insurers and the state split value	Risk analytics tied to underwriting, efficiency, grid/water controls, funded hardening, insurers able to reprice or withdraw	Unfunded municipal plans, contractors bearing fixed-price risk, uninsurable asset owners, utilities denied cost recovery	Disaster spending attracts emergency entrants and political controls; repeated losses can raise cost of capital faster than revenue	Paid projects, loss avoided, rate/premium adequacy, coverage availability, litigation, real cash after working capital; falsified by weak funding or suppressed pricing
Plateau, muddle and ordinary compounding	Efficient exposure, necessary mundane services and disciplined use of surplus	Broad customers pay; competition keeps gains dispersed and management allocation determines owner capture	Low-cost comparator, boring recurring services, repair/aftermarket, disciplined capital return, select balance-sheet franchises	Theme-pure builders, research-intensive stories needing heroic growth, serial empire builders, high-fee thematic vehicles	The capital cycle may normalize without a dramatic bust; crowded “quality” can still derate	Broad productivity, capex/revenue, theme revenue versus cash, buybacks net of issuance, active-program attribution; falsified if dramatic bottlenecks deliver sustained excess cash

The master table deliberately avoids a single “winner” for each force. A value pool is usually a bargaining game among several claimants. The firm that solves the social problem is not always the scarce claimant. A data-center operator can create capacity while the power interconnection, land owner, lender and chip supplier collect most of the rent. A therapy can transform health while the payer negotiates away price, skilled manufacturing labor captures wages and a patent cliff shortens the owner's cash duration. A resilience project can prevent enormous social loss while an unfunded municipality never becomes a collectible payer.

Interaction map: the forces are not independent themes

The crosswalk is a network, not twelve parallel sector forecasts. The following map identifies the most important transmission channels.

From	Through	Into	Investment consequence
AI diffusion	Compute demand and task automation	Power, chips, software seats and skilled labor	Helps bottlenecks early, but later price declines can shift value to customers and ordinary adopters
Compute buildout	Power reservations, construction and financing	Grids, land, cooling, utilities, banks and materials	Creates one correlated capex complex; different tickers do not ensure diversification
Aging and labor scarcity	Wage pressure and service shortages	Automation, migration policy, healthcare productivity and fiscal spending	Raises value of labor-saving tools, but labor or the state may keep the benefit
Fiscal activism	Procurement, subsidies and debt issuance	Defense, localization, grids, inflation and repression	Funds demand while increasing state bargaining power, tax risk and real-rate sensitivity
Deglobalization	Redundancy, controls and retaliation	Semiconductors, logistics, defense, EM currencies and working capital	Raises nominal orders but may lower system productivity and utilization

From	Through	Into	Investment consequence
Electrification	Load growth and materials intensity	Grids, copper, power equipment and adaptation	Can create bottlenecks, then induce supply, substitution and affordability intervention
Climate stress	Insurance loss and hardening needs	Municipal finance, utilities, migration, land values and state budgets	Benefits repricers and controls only if payers are funded and law permits recovery
Inflation and higher real rates	Cost of capital and refinancing	Every long-duration theme and state budget	Converts revenue growth into poor equity returns when capex and duration are underpriced
EM formalization	Digital records and credit creation	Payments, banks, tax collection, logistics and governance	Expands the formal value pool while giving the state and funding providers more visibility and claims
Plateau/muddle	Slower narrative growth and normalized capex	Comparator, ordinary compounders and capital return	Tests whether thematic research adds value after its costs and whether valuation already captured success

Three common factors deserve explicit aggregation. First, **AI-capex sustainability** links chips, merchant compute, hyperscalers, power equipment, cooling, automation and some industrial distribution. Second, the **state-payer factor** links defense, grids, localization, healthcare, climate adaptation and parts of emerging-market infrastructure. Third, **quality-duration and discount-rate exposure** links software, payments, information tolls, brands, market infrastructure and other high-multiple recurring models. A portfolio can own many businesses and still be exposed to one payer, one financing window or one valuation duration.

This interaction has a current policy consequence, not merely a research implication. The diagnostic comparator proxy already breaches the original AI and joint AI-plus-power factor-loss limits, and the actual comparator study remains incomplete. Accordingly, active additions with material AI or power sensitivity remain inadmissible while the original limits govern. A favorable crosswalk conclusion cannot override that factor bar; raising a cap would acknowledge embedded concentration rather than remove it.

The interaction map also shows why early and late winners can differ. Early scarcity tends to reward bottlenecks, existing capacity and scarce labor. The response - capex, training, migration, entry, substitution and regulation - attempts to remove the scarcity. If it succeeds, customer welfare can keep improving while early supplier margins fall. The later winners may be low-cost users, disciplined survivors or businesses that need the now-cheap input. Crosswalk research must therefore record whether an industry is in scarcity, buildout, competitive flood, consolidation or harvest. It may not extrapolate the economics of one phase across the decade.

How to analyze the distribution of surplus

For every force, the relevant question is “who is scarce after the system responds?” Labor captures more when skills are scarce, organization is difficult, unions or licensing strengthen bargaining and automation is not yet a substitute. Land and rent claimants capture more when location, power access, water, mineral rights or interconnection cannot be replicated. Debt capital captures more when projects are levered, refinancing is frequent and equity funds the residual risk. Equity capital captures more only when operating returns exceed the full cost of debt and equity after replacement capital. The state captures more through taxes, procurement monopsony, regulation, directed investment and ownership. Entrepreneurs and new entrants capture more when incumbents are slow and financing plus distribution are accessible; incumbents capture more when installed base, standards, lawful data, regulatory qualification or customer workflow blocks entry. Platforms capture more when they govern a two-sided market without being bypassed or regulated. Customers capture more when products are substitutable, entry is easy or procurement is concentrated. Bottlenecks capture more while their qualification, capacity or right cannot be reproduced quickly.

Those categories overlap. An entrepreneurial founder can control an incumbent platform; a state can be payer, regulator and shareholder; a utility can be both capital provider and bottleneck; a worker can receive equity. The analysis therefore follows cash and rights rather than labels. Aggregate operating profit is not parent cash when it is trapped behind regulation, minority interests or capital controls. Parent cash is not per-share value when share

count, stock compensation, acquisition consideration or controller leakage grows. Per-share cash is not expected return until the purchase price and duration of that cash are tested.

Indicators of a changing split include wage share and skilled compensation; customer savings and real unit prices; markups and take rates; supplier concentration; land, power and connection cost; interest expense and refinancing spreads; effective tax and state take; capex intensity; working-capital absorption; buybacks net of issuance; and free cash flow per diluted share. A favored model prior should weaken when its supposed bottleneck expands capacity faster than demand, its payer becomes more concentrated, the state increases take, customers can multi-home or substitute, labor consumes incremental revenue, or per-share conversion diverges from operating profit.

AI and software diffusion

Observed condition and persistence. The dated starting condition is rapid experimentation, large infrastructure commitment and uneven production use. Federal Reserve work cited in the source policy distinguishes adoption from aggregate productivity and provides an explicit buildout dashboard (Federal Reserve AI adoption note; Federal Reserve AI buildout dashboard). Persistence is plausible because models, interfaces and compute are general-purpose inputs, while enterprises need time to redesign workflows, permissions, data, controls and responsibility. The structural force is falling machine capability cost and diffusion into work. The cyclical element is the pace of spending and hiring. Any forecast of near-term aggregate productivity is dated and uncertain.

Need, payer and surplus split. The economic need is not “AI” in the abstract. It is lower task cost, faster cycle time, higher throughput, fewer errors and new feasible services. Enterprises, governments and consumers are the payers. The customer will keep much of the gain where models and applications are substitutable. Scarce technical labor and implementation teams can capture wages during the transition. Foundation platforms can capture when compute, distribution or default access is concentrated. Existing workflow owners can capture when they control the system of record, lawful context and the action taken after a model response. Verification, security and liability-transfer providers can capture when error cost is high. The state can capture through privacy, competition and safety regulation or through taxes on visible rents.

This distribution changes the prior. Favored models are record-to-action systems with high-consequence workflows; lawful, unique and fresh data that improves outcomes; security, identity, audit and verification; and vertical software able to price by outcome or retained value rather than seats alone. Incumbents are favored only where AI deepens switching costs rather than bypassing them. Disfavored models are undifferentiated wrappers, generic content tolls, labor-arbitrage services that cannot move up the value chain, seat-only pricing exposed to headcount compression and products whose “proprietary data” is exportable or recreatable. S&P Global and RELX illustrate the information-workflow hypothesis; Microsoft and Alphabet illustrate incumbent-platform and defensive-capex questions; Palantir illustrates why admiration of a capability does not cure valuation or per-share-capture gates. These are dated research-map examples, not recommendations.

Capital cycle, disappointment and counterexamples. Software's low marginal distribution cost attracts new applications rapidly. Open models, bundling and customer-built tools can erode features before reported retention responds. At the infrastructure level, cheaper inference can increase volume while reducing unit price; the resulting value may accrue mainly to users. A socially successful diffusion - more productive workers and cheaper services - can therefore coincide with disappointing application margins. Conversely, a broad AI disappointment can leave an ordinary workflow company healthy if the product was necessary without AI and management did not overinvest.

Valuation, indicators and falsifiers. The valuation risk is duration and narrative crowding. A company can retain a strong model but produce a poor security return if price assumes long high growth, stable margins and no substitution. Research should track production rather than pilot use; customer savings net of implementation; gross and net retention normalized for seats and price; outcome or usage pricing; model and compute cost; data rights and portability; error, insurance and liability; sales efficiency; stock compensation and diluted free cash per share. The favored prior is falsified when customer return remains unproven, renewal depends on bundling, real price falls

faster than cost, switching becomes easy, liability grows faster than revenue or the vendor cannot retain any share of customer savings.

Compute, data centers and semiconductor capital cycles

Observed condition and persistence. The starting condition is an unusually large, interconnected buildout across accelerators, memory, networking, fabrication, power, cooling and data-center structures. Persistence is plausible because useful AI workloads, cloud growth, sovereign capacity and replacement cycles require physical assets. Yet the capital cycle is not structurally favorable merely because demand persists. The structural need for computation can coexist with cyclical overbuilding, falling unit prices and short economic lives. The policy therefore separates merchant capacity, internal hyperscaler capacity and semiconductor or equipment oligopolies.

Need, payer and surplus split. Hyperscalers, AI developers, enterprises and governments fund the chain, sometimes with debt and long commitments. Scarce equipment suppliers, qualified foundries, advanced chip designers, memory or networking bottlenecks, power interconnections, land owners and specialized construction labor may collect cash before the data-center owner earns a return. Lenders capture through interest and covenants when merchant capacity is levered. Customers gain when competing capacity lowers rental or token prices. The state can alter the split with subsidies, export controls, local-content rules and tax incentives. Entrepreneurship can earn early profits in orchestration or new architectures, but platform buyers can internalize the capability or exercise purchasing power.

The favored prior is therefore not “compute.” It is a qualified equipment or process bottleneck with full-cycle pricing discipline; a trusted foundry with high qualification barriers and per-share returns above its capital cost; critical interconnect, cooling or service installed bases where customer downtime exceeds the fee; and internal capacity whose owner can show that AI preserves or expands an existing franchise after economic depreciation. ASML and TSMC illustrate bottleneck and trusted-foundry questions. NVIDIA illustrates strong operating economics that can still fail the valuation pillar. CoreWeave illustrates the merchant-capacity severity and financing question. Intel illustrates why state support, replacement value or a hoped-for harvest cannot substitute for evidence of capacity exit, utilization and owner cash.

Capital-cycle response and models likely to disappoint. Scarcity rents invite capacity, financing and vertical integration. Orders can be double counted, power reservations can be speculative, and customers can design around a bottleneck. Merchant hosting with concentrated customers, short contracts, high leverage and unhedged residual value is disfavored. So are builders that capitalize recurring replacement, extend accounting lives while rental economics deteriorate or depend on supplier financing. Even an equipment oligopoly can disappoint when customer capex peaks, export controls remove a market, inventories rise or buyers qualify alternatives. The counterexample is the socially useful capacity flood: abundant compute accelerates diffusion and customer welfare while destroying the equity of late capacity builders.

Valuation, indicators and falsifiers. Track utilization and contribution margin by vintage, rental and token prices, contract funding and cancellation, customer concentration, equipment retirement and resale, impairments, inventories, lead times, book-to-bill, export licenses, cash capex versus depreciation and free cash after leases and interest. At the customer layer, track incremental AI gross profit and persistence of the protected franchise. The bottleneck prior is falsified by rapid qualification of alternatives, falling price without sufficient volume response, financed additions exceeding demand, repeated asset-life extensions, weak customer monetization or per-share cash that remains negative after replacement capital. The unresolved minority view - that hyperscaler capex defends existing rents even when stand-alone AI returns are low - remains live and must be tested rather than assumed.

Electricity demand, grids and power constraints

Observed condition and persistence. The observed starting condition is rising connection demand alongside slow transmission, transformer, permitting and interconnection processes. The policy cites the IEA's finding that grid investment and queues have become material constraints while also warning that queue totals contain duplicate and speculative projects (IEA, Electricity 2026 - Grids). Persistence is plausible because electrification, data centers, industrial localization and replacement of aging infrastructure all require reliable power, while large grid projects take years. The structural elements are the physical network, permitting and reliability need. The cyclical elements are data-center reservations, equipment lead times and capital spending. Forecast load is not the same as energized load.

Need, payer and surplus split. Utilities, data centers, industrial firms and ultimately households or taxpayers pay. Early scarcity rent can accrue to existing connection rights, land near power, generators, transformer and switchgear suppliers, power-management providers, construction labor and regulated networks. A data-center landlord may discover that the utility, land owner and equipment bottleneck have more bargaining power than the building. Customers capture value when efficiency avoids a costly connection or when abundant generation lowers all-in bills. The state determines permitting, allowed returns, affordability, market design and sometimes ownership; it can require investment while suppressing recovery.

Favored priors include power management, grid control, protection, sensing and efficiency whose fee is small relative to avoided downtime or capacity; qualified components with service content; and regulated networks only where the allowed return reliably becomes earned cash after financing and dilution. Schneider Electric and Eaton illustrate power-control questions; PG&E illustrates why essential demand and rate-base need do not waive liability, financing and regulatory-capture risks. The disfavored group includes speculative connection developers, fixed-price engineering contractors exposed to material and labor inflation, merchant projects dependent on one forecast load, and utilities whose capex grows faster than collectible cash per share.

Capital-cycle response, valuation and falsifiers. Scarcity attracts factories, alternative technologies, demand response, storage and political reform. Equipment margins can normalize well before physical need disappears. New generation can create energy abundance but still require costly networks; conversely, network reform can release capacity without the forecast construction. Track queue-to-interconnection conversion, energized megawatts rather than reservations, equipment lead times and cancellations, capacity additions, power prices including network and storage, outage intensity, utility allowed versus earned return, customer bill growth, working capital, debt and diluted cash per share. The scarcity prior fails when demand is cancelled, connections accelerate, equipment price and backlog fall without service offsets, affordability blocks cost recovery or capital needs repeatedly force dilution. A socially successful grid build can still be a poor shareholder outcome when customers receive reliability, labor receives scarce wages and the regulator limits owner return.

Electrification and industrial automation

Observed condition and persistence. Electrification substitutes controlled electrical systems for combustion and manual processes; automation substitutes machines and software for scarce, unsafe or inconsistent labor. Persistence is plausible because many industrial assets turn over slowly, reliability requirements rise and aging workforces make some tasks harder to staff. The structural force is the search for lower lifetime cost, throughput and reliability. Subsidy schedules, order surges and inventory cycles are cyclical. Forecast adoption rates remain dated.

Need, payer and surplus split. Factories, warehouses, buildings, utilities, transport buyers and households pay. Customers retain most value when equipment is standardized and suppliers compete on price. Automation providers capture more when integration, safety qualification, embedded workflows and service installed bases make failure expensive. Skilled technicians and integrators can capture wages when deployment is labor constrained. Power and grid claimants capture if connection cost rises. The state shifts economics through

emissions rules, subsidies, tariffs and local-content mandates. Low-cost entrants can transfer early incumbent rents to customers.

Favored models are controls, sensing, safety and mission-critical automation with measurable customer payback; qualified equipment with a recurring aftermarket; reliability distribution whose density lowers downtime; and integrators that own reusable intellectual property rather than selling hours. Fastenal illustrates reliability distribution, while GE Aerospace illustrates the related installed-base service logic in a different domain. Speculative automation providers can be informative, but short history, project concentration and customer finance raise severity. Disfavored models are commodity hardware, bespoke project labor without repeatability, subsidy-dependent manufacturing, equipment vendors whose “backlog” consumes working capital and any company that counts customer capex aspiration as collectible demand.

Capital cycle and indicators. High growth invites factories, inventory and vendor finance. Standardization and interoperability can commoditize hardware; a dominant customer can dual-source after the supplier funds qualification. Automation can also weaken its own payer: if customers struggle financially, the projected savings may not become a funded order. Track payback periods using customer cash, not vendor estimates; utilization, throughput and defect reduction; service attachment and installed-base retirements; organic units and real price; inventories, receivables and customer advances; subsidy share; returns by project or cohort; and cash after maintenance and integration labor. The prior fails when productivity gains do not cover total implementation cost, customers delay acceptance, service alternatives proliferate, capacity outruns units or supplier per-share cash lags reported earnings. Social success can mean cheaper industrial output, not high automation-vendor margins.

Aging, healthcare innovation and labor scarcity

Observed condition and persistence. Population aging, chronic disease, care demand and a constrained care workforce create a decade-long need for treatment, prevention and labor productivity. Persistence rests on demographic stocks, training lags and the time required to change care pathways. Scientific and manufacturing progress can expand feasible treatment. The structural forces are demography, disease burden and worker scarcity. Product cycles, reimbursement decisions and individual pipelines are cyclical or company-specific. A forecast of any therapy's adoption or price is not structural.

The demographic force is geographically divided rather than one global “aging trade.” Near-term cohort arithmetic is relatively durable in older East Asian and European systems, but migration, later retirement and participation can materially alter effective labor supply in North America and elsewhere. Young African and South Asian populations create potential workers and consumers only where education, health, power and formal jobs turn population into productive income. The UN's country projections and the OECD's employment scenarios support this regime distinction, while remaining projections rather than security evidence (UN Population Division, World Population Prospects 2024; OECD Employment Outlook 2025, demographic change). The favored prior must therefore be local: the same health or automation model can face a funded social insurer, an out-of-pocket household, a rationing state or an unfunded need.

Need, payer and surplus split. Governments, insurers, employers and patients fund care. Patients receive health and time; employers can gain attendance and productivity; payers can retain avoided total cost through price negotiation; innovators can earn intellectual-property rent; providers and licensed labor can capture scarcity wages; manufacturers can capture when capacity or quality is scarce; distributors and benefit managers may collect tolls; the state can use monopsony, compulsory discounts and tax. Longer life can also create additional downstream care rather than pure savings.

Unpaid family care is an economic claimant even though it never appears as provider revenue. When formal capacity is unavailable, households absorb the shortage through waiting, travel and reduced paid work; public or employer-funded care can create value by releasing that labor. Local land and housing can also absorb nominal wage gains when scarce workers must live near hospitals, laboratories or productive cities. Housing supply, credential reform, migration, later retirement and task redesign are therefore capital-cycle responses to labor

scarcity just as factory construction is a response to equipment scarcity. A “health demand” thesis that omits unpaid care, housing cost and the public tariff can mistake a social burden for a corporate value pool.

The favored prior is funded health innovation: material clinical outcomes plus demonstrated access, adherence, net price, manufacturing economics, exclusivity and replacement pipeline. Tools that remove scarce labor hours without increasing errors can be favored, as can diagnostics linked to a reimbursed action and insurers that price risk with sound reserves and capital. Eli Lilly and Novo Nordisk illustrate the innovation, access, manufacturing and cliff questions; Chubb and Berkshire illustrate disciplined underwriting in other risk pools; Teladoc illustrates why access and adjusted metrics do not prove clinical savings or diluted cash. Disfavored models include labor-heavy providers with capped reimbursement, undifferentiated telehealth, single-product franchises priced as if exclusivity never ends and services whose customer savings are asserted rather than measured.

Metabolic treatment is the clearest current stress test of the distinction. Effective chronic treatment can improve health, labor participation and employer productivity while reducing some downstream procedures. Yet the surplus can be divided among patients, employers, public and private payers, intellectual-property owners, scarce manufacturers, distributors and clinicians. Indication-specific access, persistence, net price and total budget impact matter more than the gross eligible population. A US Congressional Budget Office illustration found that near-term Medicare medical savings from expanded anti-obesity-drug coverage were small relative to direct drug cost under its assumptions, showing why avoided events do not automatically fund broad premium-price access (CBO, budget effect of Medicare anti-obesity-drug coverage). Competitive mechanisms and capacity can later shift a socially successful treatment from early IP and manufacturing rents toward payers and patients.

Capital cycle, valuation and falsifiers. Scientific success attracts competing mechanisms, manufacturing capacity and payer scrutiny. Labor scarcity raises provider wages, but automation, migration, scope-of-practice reform and care redesign can reverse that rent. The socially best outcome - effective prevention at broad access - may shift value to patients and payers while reducing downstream procedure volumes and drug prices. Track indication-level outcomes, real-world adherence and persistence, net rather than list price, denial and access, total cost of care, labor hours per episode, manufacturing yield and capacity, research productivity, exclusivity dates, replacement pipeline, acquisition spending and diluted cash. The favored prior fails if payer savings do not materialize, access narrows, net price falls faster than volume and cost, adherence is weak, the replacement pipeline cannot bridge a cliff or labor cost absorbs incremental revenue.

Fiscal activism, defense and state procurement

Observed condition and persistence. Defense requirements, industrial policy, infrastructure renewal and economic-security objectives have increased the state's role as payer, rule-maker and capital allocator. Persistence is plausible because procurement cycles, security commitments and industrial capacity take years, while geopolitical rivalry and aging public assets do not reverse quickly. Enacted contracts and installed obligations can be structural for their term. Political targets, annual budgets and election promises are dated. Spending surges and emergencies are cyclical.

Need, payer and surplus split. The state pays through current taxes or debt, but it is often a concentrated buyer with power to specify cost, local content, audit, cancellation and profit. Taxpayers and future savers ultimately finance it. Scarce qualified suppliers, cleared labor and installed service bases can bargain when the state urgently lacks alternatives. Primes can shift risk to subcontractors. Lenders capture interest when suppliers fund working capital. The state can recapture apparent profit through renegotiation, excess-profit taxes, required reinvestment or ownership. Customers receive security and infrastructure benefits that have no direct commercial price.

Favored models are scarce, qualified bottlenecks with funded programs; recurring maintenance, training and aftermarket tied to an installed base; cost-recoverable or appropriately escalated contracts; and suppliers whose balance sheets can survive political delay. Rheinmetall illustrates the positive defense-research hypothesis but remains subject to appropriation, contract and valuation work. Boeing illustrates how safety, fixed-price development and working-capital claims can overwhelm backlog. Regulated resilience models can qualify only

when rate or appropriation becomes cash. Disfavored models are budget-story companies, fixed-price development at technical frontiers, contractors that finance the payer, politically visible windfall earners and localization builders that add capacity before contracts are enforceable.

Capital cycle, valuation and falsifiers. Backlog and high margins attract factories, acquisitions and labor. The buildout can arrive after a ceasefire, election or fiscal constraint changes demand. A peace dividend can strand late capacity. Conversely, a bottleneck can persist while the state caps margin. Track enacted appropriations, contract type and escalation, advances, cancellation compensation, backlog conversion, milestone acceptance, working capital, cost overruns, supplier concentration, production capacity, cash taxes and free cash per diluted share. Falsifiers are repeated unfunded renewals, cancellation, rising backlog with negative cash, margin compression after capacity investment, state take that absorbs returns or production additions that outrun funded demand. Social success - deterrence or resilient domestic capacity - is not proof that minority shareholders earn above the cost of capital.

Deglobalization, supply-chain resilience and economic security

Observed condition and persistence. Export controls, sanctions, tariffs, local-content rules and corporate resilience programs encourage duplicated production, inventories and routes. Persistence is plausible because trust and security restrictions can remain after the triggering crisis and because new plants and supplier qualification are long lived. The structural element is a higher value placed on security of supply. The cyclical elements are emergency inventory, freight prices and subsidy races. A complete reversal to frictionless globalization is a counterthesis, not the base assumption.

Need, payer and surplus split. Corporations and states pay an insurance premium for redundancy. Local labor, land and utilities can capture the investment. Testing, certification, logistics and trusted suppliers can capture when multiple regimes increase complexity. Incumbent bottlenecks gain when they are already qualified within a bloc. New entrants can receive subsidies yet destroy industry utilization. Customers and taxpayers bear higher costs. Retaliating states can remove markets or trap capital. A domestic champion can have national value while minority owners receive poor returns because investment, price limits and political obligations dominate.

Favored priors include verification and standards across jurisdictions; traceability; trusted foundry and secure components with disciplined capacity; reliability distributors that can multi-source; and logistics or service networks that improve resilience without betting on one route. TSMC and ASML illustrate essential-node questions but also concentrated geography, customer and policy risk. Adani Enterprises illustrates why national buildout, project scope and privileged access do not replace audited entity-level debt, guarantees, related parties and per-share cash. Disfavored models include low-utilization duplicate factories, subsidy-dependent national champions, single-route infrastructure marketed as resilience, inventory builders funded with short debt and suppliers dependent on a license that can disappear.

Capital cycle, indicators and falsifiers. Protection initially raises local prices and margins, then attracts protected supply that may never reach efficient scale. Redundancy can be socially rational insurance even when commercial return is below the cost of capital. Track control lists, license approvals, local-content requirements, subsidy share, utilization, inventory turns, freight and insurance, price premiums, customer willingness to pay, retaliation, cash taxes and free cash after duplication. The favored prior fails when policy normalizes, customers refuse the resilience premium, duplicate capacity remains underused, state support is withdrawn, retaliation removes the assumed export pool or suppliers cannot generate per-share cash without continuing subsidy.

Inflation, real rates, debt and financial repression

Observed condition and persistence. High public and private debt, investment needs and an uncertain inflation process create incentives for nominal growth, taxation and policies that lower the real burden of debt. Persistence is plausible because debt stocks and fiscal commitments adjust slowly. Inflation and real rates are regime variables, not one-way forecasts. Financial repression can take the form of administered rates, captive demand, directed credit, capital controls, reserve requirements, price controls or taxes on visible surplus. Disinflation, fiscal consolidation and higher productivity form the reversal path.

The crosswalk preserves a wide regime distribution. Productive normalization could raise real income and ease debt arithmetic. Fiscal repair with expensive capital could reward creditors while impairing levered equity. An inflation-repression grind could preserve nominal solvency by transferring value from savers and visible corporate rents. Sovereign repricing could damage both banks and public contractors. A debt-deflation or theme bust could lower prices while raising real debt burdens and eventually create a harvest only after capacity and balance sheets clear. The IMF's April 2026 Fiscal Monitor documents the rising debt and interest claim, while BIS work emphasizes the sovereign-financial linkage; neither source resolves which regime will dominate (IMF Fiscal Monitor, April 2026; BIS Annual Economic Report 2026, Chapter II).

Need, payer and surplus split. The system must allocate the real loss from debt adjustment. Borrowers gain from unexpected inflation; savers lose when nominal returns lag. Banks can gain from low-cost deposits or lose through deposit beta, duration mismatch and directed credit. Labor can recapture purchasing power through wages, while firms with genuine short-cycle pricing may preserve margin. Customers bear price increases until substitution or regulation binds. The state collects bracket creep, seigniorage, taxes and captive financing. Land and commodity owners can collect rent, but visible windfalls invite state take. Lenders capture more when real rates and spreads rise; equity becomes the residual after refinancing.

Creditor capture is contractual, not unconditional. Floating coupons transfer cash stress to borrowers and can turn spread into default; fixed coupons lose real value under surprise inflation; captive sovereign holdings can impose a hidden tax on banks, insurers and pension savers. Repression may stabilize the state while crowding out entrepreneurship because politically preferred borrowers receive credit before higher-return but less protected entrants. Thus “higher rates favor finance” and “inflation favors real assets” are both inadequate priors. Funding mix, maturity, collateral, regulation, tax, title and the borrower's replacement-capital burden decide who retains the gain.

Favored priors are short-cycle necessary distributors with working-capital discipline; low-cost and sticky funding franchises proven through a full credit cycle; insurers with adequate repricing and reserves; sound balance sheets; and businesses whose price increase is supported by unit retention and low reinvestment. DBS illustrates the deposit-franchise question, Chubb the underwriting and float question, and Fastenal short-cycle repricing and distribution. These examples retain class-specific evidence burdens. Disfavored models are levered long-duration projects, businesses with fixed revenue and floating inputs, long working-capital cycles, nominal growth mistaken for real growth and politically visible scarcity profits. Verified inflation-linked liability matching belongs to strategic design, not this list of corporate “winners.”

Capital cycle, valuation and falsifiers. Easy nominal financing can sustain poor projects; higher real rates expose them through refinancing and lower terminal values. Financial repression can support nominal earnings while reducing real per-share value. Track real price versus units, wage and input lags, inventory funding, interest coverage, refinancing maturity, deposit beta, credit losses, effective tax and state take, capital controls, index integrity and real free cash per diluted share. The favored pricing-power prior fails when volumes contract, labor or inputs absorb price, working capital consumes cash or regulation caps return. The repression hypothesis weakens if debt ratios fall through fiscal consolidation, real rates remain durably positive without financial instability and capital allocation becomes less directed. A correct inflation forecast can still produce poor shareholder results through leverage, tax or valuation.

Natural-resource depletion and harvest cycles

Observed condition and persistence. Mines, wells, reservoirs and biological stocks deplete, while permitting, discovery and construction take time. Electrification, power and resilience can raise demand for selected materials. That structural need does not settle the cycle. High prices promote exploration, substitution, recycling, efficiency and new capacity, while grade decline and sustaining capital can consume the apparent rent. Resource scarcity is therefore an input to a capital-cycle analysis, not an ownership thesis.

Need, payer and surplus split. Industrial customers and consumers pay for materials. Resource states collect royalties, taxes, local-content claims and sometimes ownership. Labor, land and equipment suppliers capture during scarcity. Lenders can claim much of the cash from levered projects. Low-cost reserve owners capture only after sustaining capital, closure and state take. Customers gain from substitution and new supply. Entrepreneurs can earn by discovering or improving deposits, but serial acquisition can transfer value to sellers rather than continuing shareholders.

Favored priors are long-life, low-cost assets with credible title; disciplined survivor producers in a confirmed harvest; royalties or infrastructure claims only where contracts, counterparty and state terms are sound; and service franchises tied to a surviving installed base. Disfavored priors are single-asset or high-cost producers, spot-price net-asset-value stories, serial reserve buyers, projects with large closure liabilities, levered builders and assets in jurisdictions where title or cash upstreaming is uncertain. First Quantum illustrates why copper demand and a high spot price cannot waive jurisdiction, sustaining-capital, debt and title gates. The policy records no fully qualified natural-resource or harvest candidate, which is a supported absence from a limited search, not a negative sector allocation.

Capital-cycle response, indicators and falsifiers. Scarcity prices finance the response that erodes scarcity. Harvest becomes plausible only after financing closes, high-cost capacity exits, survivors can fund maintenance, utilization improves and management restrains replacement capex. A price collapse, impairment or merger is not capacity exit; state support can keep an uneconomic unit alive. Track reserve life and replacement cost, grade, recovery, all-in sustaining capital, closure, project delays, inventories, cost-curve position, balance-sheet survival, net capacity exit, state take, per-share production and cash. The scarcity prior fails with large economic discoveries, substitution, recycling, demand destruction or state extraction. The harvest prior fails if consolidation is debt funded, nominal discipline masks new capacity, maintenance is deferred or survivors still cannot return cash per share. An unconfirmed harvest candidate expires after 24 months rather than rolling forward by narrative.

Emerging-market formalization and digital platforms

Observed condition and persistence. Digital payments, logistics, identity, commerce and credit can move activity from informal cash systems into recorded, scalable networks. Persistence is plausible because urbanization, smartphone use, tax administration and network density reinforce formalization. Yet “emerging markets” are not one cycle. Funding conditions, institutional quality, demographics, currency and competitive supply differ widely. User and gross domestic product growth are observed activity, not foreign-shareholder capture.

Need, payer and surplus split. Consumers and merchants pay fees, spreads or prices; borrowers pay interest. Platforms capture when two-sided density reduces cost and improves utility. Banks and telecoms capture funding and access. The state gains tax visibility, regulatory control and sometimes data. Local entrepreneurs can gain distribution; global capital can fund growth but demand return and create foreign-currency exposure. Customers can receive lower price and inclusion, while competition and instant-payment infrastructure compress private take rates. Credit losses can reveal that apparent payment growth was balance-sheet risk.

Favored priors are platforms with cohort contribution, rising utility, falling service cost, clear property rights, audited cash and disciplined credit; payment or logistics rails with take rates justified by measurable service; and deposit franchises with sound full-cycle loss and capital evidence. MercadoLibre illustrates the mature platform and

per-share-capture question; Nu, Sea and BYD illustrate higher-severity optionality; Adani illustrates governance and entity-map failure. Disfavored models are subsidy-funded user acquisition, related-party conglomerates, foreign-currency borrowing against local revenue, payments businesses that hide credit, and structures in which cash cannot move to the listed parent or diluted shares grow faster than owner cash.

Capital cycle, valuation and falsifiers. Cheap capital funds discounts, merchant subsidies and unsecured credit, producing network metrics before economics. Success attracts local competitors and state rails. Regulation can increase trust while capping fees. Currency depreciation can erase attractive local growth in the owner's liability currency. Track retained cohorts, contribution after incentive, take rate and customer savings, fraud, charge-offs and vintages, funding currency, regulatory capital, related parties and guarantees, tax and state take, capital controls, dilution and hard-currency free cash per share. The favored platform prior fails when network growth requires persistent subsidy, take rate collapses without unit-cost relief, credit losses rise, the state disintermediates the toll, related-party leakage appears or local cash cannot become parent per-share value.

Urbanization and infrastructure as a cross-cutting force

Observed condition and persistence. Urban concentration can raise productivity because workers, customers, suppliers and knowledge interact at density, while much future urban infrastructure in younger economies has not yet been built. The mechanism can persist through 2036 because buildings, transport, water, grids and land use adjust slowly. Urban population growth is not the same as productive agglomeration: congestion, informal employment, unreliable service and unaffordable housing can absorb the gain. Construction booms and property credit are cyclical; a city population projection is a planning input, not operator revenue.

Need, payer and surplus split. Households, enterprises, municipalities and taxpayers pay through tariffs, fares, property charges, taxes and rent. Landowners often receive the first uplift when transport, drainage, power or safety improves access. Construction labor, material suppliers and lenders are paid before an operator's residual equity. The state can recapture some land appreciation through tax, leases or development charges, or leave the uplift to private and political owners while carrying maintenance cost. Platforms can benefit from delivery and matching density, yet warehouses, drivers, congestion and regulation rise with it. Customers capture the social gain when reliable services, shorter journeys and competitive cities lower their total cost.

Favored priors include small-ticket maintenance, metering, leakage detection, inspection and installed-base service; enforceable service infrastructure with conservative currency and debt; and digital identity, payment, land-record and logistics systems that raise collection and reduce transaction cost. Disfavored models are fixed-price megaproject contractors, speculative land banks dependent on rezoning, foreign-currency debt against local tariffs, concessions with heroic traffic and utilities whose losses or receivables make nominal asset growth uncollectible.

Capital cycle, valuation and falsifiers. Visible need attracts public plans, property credit, concessions and industrial parks. Upstream inputs can collect early while occupancy, tariffs and complementary links arrive late. Track formal jobs and output per worker, commute and service reliability, housing burden, property and user-fee collection, project utilization, operation and maintenance coverage, currency match, claims, debt and dilution. The favored prior fails when population grows without formal productivity, infrastructure stands alone or underused, arrears rise, land rent absorbs wage gains, congestion worsens or asset-base growth does not become after-maintenance diluted cash. Socially successful infrastructure can properly allocate most value to users or the public and still be a disappointing listed-owner model.

Climate adaptation and physical resilience

Observed condition and persistence. Heat, flood, fire, storm, drought, water stress and insurability constraints increasingly affect assets, supply chains and public budgets. Persistence is plausible because physical assets and settlement patterns adjust slowly, while hazard and replacement costs can change faster. The structural need is to avoid loss and maintain essential service. Individual disaster years, emergency appropriations and insurance pricing cycles are cyclical. Large estimates of adaptation need are not evidence of funded demand.

Need, payer and surplus split. Property owners, insurers, utilities, municipalities, states and taxpayers may pay. Contractors and skilled labor capture during reconstruction. Land values shift between resilient and exposed locations. Insurers capture only when law allows adequate price and capital survives correlated loss; otherwise policyholders receive subsidized cover and shareholders bear the shortfall. Utilities can earn on hardening only when rate recovery is timely. The state can fund, tax, litigate, mandate or socialize loss. Customers and communities receive the avoided-damage benefit, which may greatly exceed the supplier's fee without guaranteeing supplier bargaining power.

Favored priors include risk analytics embedded in underwriting and capital allocation; power, water and efficiency controls with measurable avoided loss; funded hardening contracts with cost protection; reliability distribution; and insurers able to select, reprice or withdraw while maintaining sound reserves. Regulated resilience is conditional, not automatically favored. PG&E illustrates essential need overwhelmed by liability, financing and recovery risk; Chubb illustrates why underwriting quality requires reserve and capital evidence rather than one favorable period. Disfavored models are unfunded municipal-plan suppliers, fixed-price contractors in volatile inputs, owners of increasingly uninsurable assets, utilities forced to invest without earned return and insurers growing exposure where regulation prevents adequate pricing.

Capital cycle, indicators and falsifiers. High premiums and disaster spending attract entrants, capital and political intervention. Repeated losses can raise reinsurance, debt and equity cost faster than nominal revenue. Adaptation can also reduce the future value pool by successfully lowering loss - a social success that may normalize contractor demand and insurance pricing. Track paid projects rather than announced needs; hazard and loss per unit of exposure; premium or tariff adequacy; non-renewal and coverage availability; reinsurance; litigation; working capital; allowed versus earned return; avoided outage or water loss; and diluted cash after capital. The favored prior fails when payer funding remains aspirational, rate suppression persists, litigation transfers value, loss trends outrun pricing or the company must repeatedly raise capital.

Verification, standards and market infrastructure as cross-cutting models

Verification, standards, identity, testing, clearing and trusted records recur across AI, trade fragmentation, health, finance and state procurement. Their persistent need is to make claims interoperable and enforceable when systems become more complex. The payer may be a regulated issuer, market participant, enterprise or state. Incumbent assurance providers can capture through accreditation, legal embedding, historical data and workflow integration. Platforms can capture through two-sided liquidity. Customers capture when multiple providers are interchangeable or public rails cap fees. The state can license, mandate and withdraw permission, making the same regulation both a moat and a binary tail.

This supports a favored prior for assurance and market-infrastructure models only when mandate, accuracy, liability, real price and switching are evidenced. S&P Global, CME, Visa and LSEG illustrate different versions; Coinbase illustrates why a “rail” label does not resolve custody, token-cycle, capital and regulatory questions. The capital-cycle response may be new standards, open data, public payment rails, AI-assisted verification and client multi-homing. Track real fee, volume quality, share on both sides, error and liability, cyber resilience, regulatory action, client-retained savings and free cash per share. The prior is falsified when the standard is bypassed, take rates fall without offsetting density, a liability event breaks trust or competition turns the mandate into a low-return utility. Socially valuable verification can be delivered at regulated or competitive prices that leave modest shareholder rent.

The plateau/muddle counter-scenario and ordinary compounders

Observed condition and persistence. The strongest counterthesis is not collapse. It is that dramatic themes become useful but ordinary: AI diffuses slowly, power queues clear unevenly, inflation and real rates stabilize, defense and localization persist without accelerating, and nominal growth continues with mediocre productivity. Several theme builders can suffer a sector bust without causing an economy-wide depression. Persistence comes from offsetting forces and institutional adaptation rather than one dominant shock.

Need, payer and surplus split. Customers continue paying for mundane reliability, repair, payments, insurance, distribution and workflow. Competition disperses much of the social gain. Management's capital allocation becomes decisive because reinvestment opportunities are finite. Businesses able to return surplus without dilution can convert modest growth into per-share value. The low-cost comparator captures broad corporate progress with little research burden. The user keeps the benefit of low fees. Active research must earn more than its hours, taxes, data and errors.

Favored priors are ordinary necessary services, installed-base aftermarket, reliability distribution, disciplined insurers and banks, recurring workflows that did not require a heroic AI outcome, selective brands with current cohort evidence and companies that repurchase shares only below value or pay sustainable dividends when reinvestment is poor. CME, Visa, Fastenal, Berkshire and Chubb illustrate versions of mundane recurrence or disciplined capital use, while Ferrari illustrates that a scarce brand can still fail valuation. Disfavored models are theme-pure builders that need accelerating growth, high-fee thematic vehicles, serial empire builders and research programs whose modeled contribution does not clear twice their fully loaded cost.

Capital cycle, valuation and falsifiers. In muddle-through, capacity can normalize gradually rather than through dramatic bankruptcy. This denies both the extrapolator and the premature harvest investor. “Quality” can remain operationally strong while derating for a decade, so no-multiple-recovery cash yield matters. Track capex relative to revenue, broad productivity, theme-specific revenue versus free cash, ordinary-service retention, acquisition spending, interest, buybacks net of issuance, diluted cash per share and active-program attribution against the named comparator. The counter-scenario is falsified if multiple dramatic bottlenecks sustain scarcity, customer value and per-share excess cash through entry and normalization. Until then, it disciplines the whole map: abstention and the comparator are investable alternatives in concept, although this chapter authorizes neither implementation nor allocation.

Where value capture may strengthen, weaken or remain radically uncertain

The policy's dated prior is that value capture is most likely to **strengthen** at qualified bottlenecks whose scarcity cannot be financed away quickly, in record-to-action workflows with lawful context and real switching cost, in short-cycle necessary distribution, in installed-base service where retirements and independent alternatives are controlled, and in balance-sheet franchises with demonstrably disciplined underwriting or funding. Even here, “strengthen” is a business-quality prior. Customers, regulators and entrants react, and valuation may already capitalize the advantage.

Capture is most likely to **weaken** in undifferentiated compute or software, seat-only applications, fixed-price complex projects, subsidy-dependent manufacturing, labor-heavy services with capped reimbursement, long-duration nominal or levered models under higher real rates, spot-price resource stories, user-growth platforms dependent on continuing subsidy, and regulated or state-funded businesses unable to turn an allowed budget into cash. These models often create real social output while their scarce inputs - labor, power, debt or state permission - collect the surplus.

Capture remains **radically uncertain** in internal hyperscaler AI spending, defense buildout, regulated grid investment, healthcare prevention, emerging-market platforms, climate insurance and any early capital-cycle scarcity. In each, the end demand can be strong while the claimant split changes abruptly. The relevant research question is not simply whether the theme occurs. It is whether incremental value survives wages, rent, power,

interest, tax, procurement, entry, maintenance capital, trapped cash and dilution, and whether it then exceeds what the current security price already assumes.

Research protocol for using the crosswalk

For any worldview-linked company, write one page that completes the full chain. State the observed condition and evidence date. Separate the structural mechanism from the cycle and the dated forecast. Name the economic need, the actual payer and the path from announced demand to collected cash. List every material claimant on the surplus: customers, labor, land and power, debt and equity capital, the state, entrepreneurs, incumbents, entrants, platforms and bottlenecks. Explain which claimant is scarce, how bargaining power is enforced and what could reverse it. Record the capital-cycle state and the capacity response. Identify favored and disfavored business models, a positive illustration and a counterexample. Then state the valuation assumptions, common-factor exposure, leading indicators, two primary milestones and hard falsifiers.

Only after that causal page exists does security underwriting begin. The five pillars do not average: demand and lock-in; industry structure, cycle and payer; capital economics; governance and per-share capture; and valuation plus falsifiability must each clear the required minimum. Evidence coverage measures completeness, not favorability. Legal title, state and payer collectibility, economic depreciation, financing, dilution and portfolio factor headroom remain separate gates. A business-quality prior is not a security-return conclusion, and a security-return conclusion is not a portfolio-role conclusion.

No model in this chapter is a buy list, target allocation or trade instruction. Company references are illustrations drawn from a dated research map whose classifications expire and whose current evidence and valuations are incomplete. The plateau/muddle case, the low-cost comparator and abstention remain live competitors to every thematic idea. If the bridge cannot identify the payer, the scarce claimant, the capital-cycle response and the route to diluted per-share cash, the theme has not produced an investable research hypothesis.

The bridge is a research hypothesis, not expected-return proof

The crosswalk therefore allocates research attention only. A security still needs a complete evidence file, conservative expected-return case, loss budget, portfolio-fit analysis, and owner approval before it can acquire any live role.

Part IV - Security Research and Operating System

Chapter 9 - From Business-Quality Prior to Security-Return Conclusion

The purpose of the canonical business-model scorecard is not to identify companies that sound admirable. It is to decide, in a disciplined order, whether an active investment thesis is understandable, evidenced, economically attractive, owned by the shareholder in substance as well as form, and available at a price that can beat a named comparator. Only after those questions have been answered may the policy calculate how much permanent loss the portfolio is allowed to tolerate. Even then, the result is a ceiling, not a target weight, a recommendation, or an instruction to trade.

That order matters. Many bad investment decisions reverse it. They begin with a desired position size, then assemble a narrative about market growth, competitive advantage, or management quality to support it. This policy begins with five non-compensating pillars, tests the coverage and quality of the evidence, applies hard gates, establishes a conservative permanent-impairment severity, and only then converts the permitted loss into a maximum diagnostic weight. A strong answer in one area cannot cancel a fatal weakness in another. Excellent products do not compensate for an uncollectible payer. High returns on capital do not compensate for value leaking through dilution or related parties. A wonderful business does not compensate for a price at which even the upside case fails.

The five-pillar profile, not an additive quality score

The canonical scorecard has five pillars:

Pillar	Governing question	Score 0 anchor	Score 2 anchor	Score 4 anchor
A - Demand and lock-in	Is demand recurring, necessary, switching-costly, and measured in retained cohorts or installed base?	Promotional or undifferentiated	Mixed retention or need	Mission-critical with a measurable ratchet
B - Structure, cycle, and payer	Is industry structure rational, is the payer funded and collectible, and is the cycle state favorable?	Price-taking, weak payer, or financed flood	Mixed	Oligopoly or mandated demand, funded payer, and harvest discipline
C - Capital economics	Does incremental after-tax cash return exceed its cost after economic depreciation, dilution, leases, and maintenance capital expenditure?	Value destructive	Near the hurdle	Durable high incremental return and cash conversion
D - Governance and capture	Does value reach each share after pay, issuance, acquisitions, related parties, and state leakage?	Leakage or opacity	Mixed	Aligned and transparent per-share capture
E - Valuation and falsifiability	Does a reverse discounted-cash-flow conservative return clear the comparator, and are hard disconfirmers preregistered?	The upside case fails	The base case clears	The conservative case clears by 1 percentage point

The table prints anchors at 0, 2, and 4 because those are the decisive reference points. The company map also uses 1 and 3 as intermediate ordinal judgments. They are not separate policy gates or precision estimates. A 1 lies between adverse and mixed; a 3 lies between mixed and exceptional. The score is a disciplined description of evidence, not a probability with two significant figures.

There is deliberately no additive Business-Model Score. Earlier debate work used a long list of measures and a weighted total, but that total over-rewarded correlated forms of apparent quality. Pricing power, switching costs, recurrence, mission-criticality, and installed-base economics can all be different views of the same customer lock-in. Adding each as though it were independent lets one attractive characteristic collect several votes. It also lets a company in a financed capacity flood, or one with a poor price, hide that weakness beneath a large quality total.

The final policy therefore treats the five-pillar profile as the decision object. There is no rule that adds A through E, averages them, or allows an A4 to offset an E0. Ordinary admission requires every pillar to score at least 2. A single pillar below 2 blocks ordinary active admission. A pillar at 0 is a hard failure and the diagnostic maximum weight is zero. This minimum logic is the core anti-narrative control in the scorecard.

The profile also separates three questions that investors frequently blur. Business durability belongs mainly in A and B. Economic value creation and shareholder receipt belong in C and D. Whether the security can earn enough from today's price belongs in E. A company can be a superb business and a poor security. Conversely, a mediocre historical industry entering a verified harvest phase can deserve research, although it still must clear all five pillars before ordinary admission.

Pillar A - Demand and lock-in

Pillar A asks whether customers repeatedly buy because the product or service is necessary and difficult to replace, and whether that claim can be observed in retained cohorts or a qualified installed base. The scoring direction runs from promotional and undifferentiated demand at 0, through mixed retention or mixed need at 2, to a mission-critical relationship with a measurable ratchet at 4.

"Recurring" is not synonymous with "subscription." A subscription can disappear when seats fall, usage declines, a contract is rebid, or a customer discovers that a bundled substitute is good enough. Conversely, a transaction or service business can be highly recurring when the same operational need returns, the provider remains embedded, and customers continue to pay through stress. The analyst should therefore trace logo retention, gross and net revenue retention where relevant, renewal behavior after price increases or outages, usage intensity, cancellation rights, implementation and migration costs, and the economics of defectors and win-backs. Management's claim that revenue is recurring is not enough.

Necessity is also a cash concept rather than an adjective. A product is mission-critical when failure creates downtime, liability, non-compliance, lost revenue, or some other cost that is large relative to the fee, and when the customer continues to prioritize the spend during budget pressure. A product that is frequently praised but easily deferred remains discretionary. The cost of failure should be supported by incident records, contracts, regulation, customer continuity plans, or observed purchasing behavior during stress.

Switching costs can be operational, data-related, regulatory, contractual, or based on retraining and workflow redesign. They are strongest when several operate together and when customers remain because the product creates value, not because a punitive auto-renewal traps them temporarily. Network effects and scale advantages belong here only when another user or another unit of activity improves utility or economics for existing users. Mere size, brand recognition, or a large dataset does not qualify. The useful test is causal: after controlling for price and marketing, does density, accumulated use, or unique lawful data improve customer outcomes, fill rates, cost, accuracy, or retention?

The installed-base element must be measured net of retirements. Shipments are not an installed base, and an installed base is not automatically an annuity. The analyst should connect beginning units, additions, retirements, attach rates, service revenue per unit, third-party alternatives, and service cash after support capital. Aging can strengthen an aftermarket for a time, but it can also signal a fleet that will disappear. Open architecture, right-to-repair rules, independent parts, or a technology transition can break the link between installed units and owner cash.

The word "ratchet" at the 4 anchor means the relationship becomes observably more valuable or harder to displace as use accumulates. Examples include a growing qualified base that produces recurring, high-consequence service events; a system of record whose lawful feedback data measurably improves later decisions; or a two-sided rail whose liquidity improves execution without requiring a falling take rate. The ratchet must be measurable. A presentation slide showing a flywheel is not evidence.

AI absorption modifies this pillar. For software and data businesses, the test is whether artificial intelligence strengthens a record-to-action workflow or commoditizes the feature being sold. Evidence should cover data rights, uniqueness and freshness, customer exportability, outcome lift against a credible foundation-model substitute, workflow control, compute cost, seat contraction, retention, and who keeps the productivity savings. A product with apparently high historic retention can still receive a weak A score if AI lowers migration costs or bundles away the feature. While defensive value from heavy AI spending remains unidentified, the policy does not permit a convenient defensive-capex add-back to make the economics pass.

Hypothetical example: a compliance platform reports 96% annual customer retention. If contracts renew automatically, usage is falling, customers can export their data in a day, and a general platform now includes similar functionality, the headline retention does not support A4. A2 may be appropriate if the workflow remains useful but substitution is credible. A0 or A1 becomes possible if the product is promotional, rapidly displaced, or supported only by uncorroborated claims. Evidence that would improve the conclusion includes stable real pricing and usage after the substitute launches, defector evidence showing costly migration, and measured outcome improvement from lawful proprietary feedback data.

Pillar B - Industry structure, cycle, and payer

Pillar B asks whether the company operates in a structure that permits rational returns, whether the ultimate payer has both the obligation and the means to pay, and whether current industry capacity supports rather than destroys economics. The scoring direction runs from price-taking, a weak payer, or a financed capacity flood at 0; through a mixed position at 2; to oligopoly or mandated demand combined with a funded payer and harvest discipline at 4.

Industry structure is not market share alone. A large participant can remain a price taker when customers are concentrated, suppliers own the scarce input, products are interchangeable, or competitors are willing and able to add capacity. A smaller firm can have bargaining power when it controls a trusted standard, scarce qualification, dense route, critical component, or local service relationship. The analyst should map who keeps an incremental dollar of end-customer price and who funds an incremental dollar of capacity. Sustainable bargaining power requires counterparties to retain enough surplus to keep the chain healthy; extracting all surplus invites substitution, regulation, or vertical integration.

Cycle state is symmetric. High margins and full order books can describe the most dangerous point in a capital cycle if financing is abundant and supply is accelerating. Conversely, an industry with an ugly history can become researchable after capacity exits and survivor economics improve. A score of 0 is appropriate when financed supply is flooding the market, utilization and pricing are falling, or subsidies and easy credit keep uneconomic capacity alive. A score of 4 requires observed harvest discipline: credible net capacity exit or mothballing, closed entrant financing, survivor balance sheets able to fund maintenance without emergency equity, disciplined replacement capital expenditure, stabilized utilization and price without subsidy dependence, and normalized incremental cash returns above the cost of capital. A price collapse by itself is not evidence of harvest.

Payer quality prevents the policy from capitalizing demand that never becomes collectible cash. The ultimate payer should be tested for legal obligation, funding source, cancellation or renegotiation power, payment timeliness and credit quality, and the supplier's demonstrated full-cycle real return. A government announcement, political target, rate-base plan, backlog figure, or customer purchase aspiration is not the same as an enacted appropriation, binding contract, collectible rate order, or paid invoice.

Private payers are not automatically better than sovereigns. A diversified prepaid customer base can be excellent; a venture-financed customer with negative free cash flow can be weak. A funded cost-plus defense contract may protect downside but limit upside. A fixed-price sovereign development contract can be dangerous despite the payer's credit because the supplier absorbs overruns. A utility rate order can be strong if recovery is binding, timely, and inflation-aware, yet weak if affordability politics, disallowances, and financing needs prevent the allowed return from becoming per-share cash. An insurer or reimbursement system can pay reliably while using monopsony power or retrospective denial to prevent attractive supplier capture.

Customer and supplier concentration, regulation, and common-factor dependence sit naturally in this pillar. Several nominally different customers may rely on the same state budget, credit channel, cloud provider, commodity price, or AI capital-expenditure cycle. The score must look through to that shared source. Regulation may protect entry while capping price or creating binary liability. A mandate is valuable only while it has political legitimacy, manageable liability, and demonstrated economics.

Hypothetical example: an equipment supplier has a five-year backlog linked to a national infrastructure program. If only the first year is appropriated, cancellation rights are broad, one ministry is the sole payer, working capital is funded by the supplier, and new competitors are building subsidized plants, the backlog does not support B4. The same announced demand can coexist with B0 or B1 because payer and capacity conditions are adverse. The conclusion could improve if binding funded contracts, customer advances, documented cancellation compensation, competitor exits, and realized cash returns become visible.

Pillar C - Capital economics

Pillar C asks whether the next unit of investment produces an after-tax cash return above its economic cost after charging all the resources required to sustain the business. The score runs from value destruction at 0, through returns near the hurdle at 2, to durable high incremental returns with strong cash conversion at 4.

The word "incremental" is essential. A company can report a high average return on old assets while new projects earn poorly. The analyst should compare changes in after-tax operating profit and owner cash with the capital invested to produce them, using suitable lags and cohorts rather than a single-year ratio. Reinvestment runway is valuable only when there is a repeatable pool of projects at attractive incremental economics. A large total addressable market does not establish such a runway; customer budgets, competitive supply, local execution, and mature cohort returns do.

Economic depreciation takes precedence over accounting convenience. Physical life, revenue life, and economic life are different. An asset may still work physically after customers migrate to a more efficient generation, rental prices collapse, support ends, or power and maintenance costs make the old unit uneconomic. The analyst should bridge productive units by vintage from beginning units through additions, retirements, and impairments, then associate each vintage with utilization, revenue, gross profit, operating cost, resale value, and support. Conservative owner earnings use the shortest life relevant to the modeled cash flow.

This rule is especially important in cloud, merchant computing capacity, semiconductors, automation, and other rapidly changing capital goods. An accounting life extension can lift reported profit without improving economics. Leases, special-purpose vehicles, supplier financing, take-or-pay commitments, capitalized development, and customer-funded assets must be traced to the party that ultimately bears replacement and obsolescence risk. Idle reserve capacity earns option value only when it is explicitly designated and valued no higher than the observable external cost of equivalent short-notice capacity. It may not be turned into a longer depreciation life.

Working capital and cash conversion are part of the economic return. Growth that absorbs receivables and inventory, depends on stretched payables, or repeatedly sells receivables may be less valuable than reported profit suggests. A negative cash-conversion cycle can be excellent when customers durably prepay and refunds are unlikely, but fragile when it rests on supplier abuse or reverses as soon as growth slows. The long-run reconciliation should connect adjusted income to cash after leases, maintenance capital expenditure, restructuring, capitalized costs, and recurring acquisition consideration.

Dilution appears in both C and D for different reasons. In C, issuing shares, stock-based compensation, or acquisition consideration is an economic input required to generate the reported growth. In D, the same issuance determines how much of the total company value reaches each existing share. Counting the cash benefit of stock compensation while ignoring the new shares is not allowed.

Recurring acquisitions require especially careful treatment. If deals are the normal method of sustaining growth, acquisition consideration is reinvestment, not a discretionary use of free cash after the fact. The analysis should

reconstruct organic price and volume, foreign exchange, acquisitions and disposals; calculate all-in deal cost including shares, assumed debt, earn-outs, retention awards, restructuring, and integration capital; compare acquisition-cohort cash returns with organic alternatives; and track leverage and per-share cash. Warning signs include deteriorating organic volume while deal size and leverage rise, acquisition spending exceeding pre-acquisition free cash, payables stretch or factoring, growing contingent consideration, and changes to the definition of organic growth.

Hypothetical example: a data-center operator reports a 20% accounting return on invested capital. If the calculation assumes six-year equipment lives, excludes long-term power commitments, and treats recurring replacement spending as growth capital, a conservative three-year economic life could reduce the true incremental return below the cost of capital. Pillar C would then be 0 or 1 even if current reported margins are high. It could improve if vintage rental prices, utilization, resale values, funded customer contracts, and after-replacement cash support the longer economic life.

Pillar D - Governance and per-share capture

Pillar D asks whether value created by the operation reaches the owner of each share after management pay, issuance, acquisitions, related-party transfers, controller behavior, state intervention, taxes, trapped cash, and the legal path from operating subsidiary to listed security. Its scoring direction runs from leakage or opacity at 0, through mixed capture at 2, to aligned and transparent per-share capture at 4.

The governing unit is the share, not the company. Revenue, operating profit, total free cash flow, national gross domestic product, or industry growth can rise while free cash flow per diluted share stagnates. The analysis should reconcile total cash creation to fully diluted shares over a long period; review stock-based compensation, buybacks net of issuance, dividends, acquisitions, related-party transactions, voting and economic rights, controller incentives, audit quality, restatements, and cash upstreaming. For controlled or emerging-market structures, legal title, pledged shares, variable-interest entities, capital controls, funding currency, and state or promoter extraction can dominate otherwise attractive local economics.

Capital allocation belongs here because management decides which owners receive the fruits of the business. The relevant question is not whether an acquisition is accretive to earnings per share under management's assumptions. It is whether the all-in cash and share cost earns more than the cost of capital and the available organic, debt-repayment, dividend, or repurchase alternatives. Buybacks that merely offset stock compensation, especially at an unattractive valuation, are not automatic capture. A candid post-audit of prior projects and deals is stronger evidence than a new synergy target.

Opacity is itself adverse when it prevents the shareholder from tracing the cash. The policy's financial-complexity control applies to material financial intermediation, receivables finance, or off-balance-sheet funding that is ancillary to a non-financial company, and to segment recasting that obscures such activity. It is designed for the industrial company whose financing arm can manufacture or smooth the apparent economics of the core operation. Principal balance-sheet businesses are not rejected merely because their balance sheet is central; they instead face the class-specific batteries described later in this chapter.

For emerging-market exposure, operating capture, macro currency cost, and governance leakage should not be confused. Real local-currency cash per share can reveal whether operations create value. Currency depreciation and purchasing-power-parity effects belong in the country and portfolio decision. Share issuance, related-party flows, trapped cash, and weak minority rights belong in Pillar D. The security still must survive all three, but separating them identifies the mechanism rather than blaming management for every macro move or excusing genuine leakage as country risk.

Hypothetical example: a controlled platform grows local-currency revenue by 25% a year and reports positive free cash flow. If diluted shares grow 12% annually, related parties supply key services at untested prices, cash cannot be upstreamed, and the listed owners lack a clear claim on the operating subsidiaries, total-company growth does

not support D4. D0 may be appropriate when owner title or leakage is materially impaired. A clean ownership map, independent related-party review, conservative funding, stable diluted share count, and a record of cash reaching minority owners could change the conclusion.

Pillar E - Valuation and falsifiability

Pillar E asks two linked questions: what operating path is already embedded in the price, and what evidence would disconfirm the thesis on a useful decision clock? The score runs from 0 when even the upside case fails, through 2 when the base case clears, to 4 when the conservative case clears the comparator by 1 percentage point and the thesis has hard preregistered disconfirmers.

A reverse discounted cash flow, or reverse DCF, begins with the market price and solves for the revenue growth, margins, reinvestment, dilution, economic depreciation, fade, and terminal assumptions required to justify it. This reverses the usual temptation to choose optimistic forecasts and then calculate an attractive value. The policy requires conservative, base, and upside owner-cash cases, compared with the named low-cost implementable comparator on a standardized after-tax and after-cost basis. The comparator is not risk-free or neutral, but it is the burden-of-proof alternative that active complexity must beat.

The distinction among the anchors is strict. At E0, plausible success is not enough: even the upside case fails to clear the hurdle, the implied path is internally inconsistent, or the downside trips a hard gate. At E2, the base case clears, but the conservative case does not earn the extra margin of safety required for a 4. At E4, the conservative return clears the comparator by at least 1 percentage point and the thesis can be disproved by specified evidence. Scores 1 and 3 express intermediate states; they do not alter the ordinary admission requirement of E>=2.

Falsifiers should be causal and observable. "The share price falls" is not a falsifier. Examples include retained cohorts weakening after a new substitute, service revenue diverging from the qualified installed base, capacity growing while utilization and rental prices fall, cumulative project cash remaining negative after milestones, reserve development becoming adverse, or diluted cash per share failing to grow despite reported expansion. A sound falsifier states the metric, the relevant period or event, and the consequence for the thesis. It is preregistered so that management and the analyst cannot replace it after unfavorable evidence arrives.

Valuation is an admission and hurdle test, not another multiplier in the permanent-loss formula. The final policy deliberately contains no Vmult. Earlier multiplication charged the same uncertainty through severity, evidence coverage, and valuation, allowing several moderate discounts to reduce a position mechanically after the security had already passed its price gate. The canonical rule keeps the concepts separate: price decides whether the candidate may enter; evidence, liquidity, factor headroom, and severity decide the maximum tolerated permanent loss if it does.

Hypothetical example: an admired payments rail scores A4, B4, C4, and D4. At the current price, the base case clears the comparator, but the conservative case does not. It may score E2 and pass the ordinary pillar minimum, subject to every other gate and a full current review. If even the upside case fails, E0 makes maximum weight zero regardless of the other four 4s. This is how the scorecard enforces "quality is not price."

Evidence coverage and the standard of proof

Pillar scores describe the direction and strength of the evidence. Evidence Coverage, abbreviated EC, measures how much of the required file has actually been verified:

EC = verified required fields / applicable required fields

The ratio is weighted toward Tier 1 evidence. In this policy, Tier 1 means the most direct and binding public-company evidence available, such as audited statements and footnotes, filings, contracts where available, and regulator or legal materials. Tier 2 and Tier 3 evidence provide official or independent primary data and triangulation from customers, suppliers, competitors, or other sources capable of contradiction. A management

assertion, an investor-relations link, or the mere existence of a filing is not the same as extracting and verifying the relevant disclosure.

The coverage schedule is exact. EC below 60 is research-only. EC from 60 through 74 maps to an epistemic multiplier of 0.50. EC from 75 through 84 maps to 0.75. EC of at least 85 maps to 1.00. These ranges are percentage points of evidence coverage, not pillar scores. An EC of 84 receives 0.75; an EC of 85 receives 1.00.

EC is not a favorability score. A complete file full of adverse evidence can have EC 100 and pillar scores of 0. A sparse file containing only good news can have low EC. Keeping the two separate prevents the analyst from calling a story well evidenced merely because the available fragments are positive.

The company-map EC values and A through E scores are explicitly provisional analyst assumptions, tagged [A]. No named company has completed the independent adversarial-corroboration pass. If Tier 1 evidence conflicts materially with Tier 2 or Tier 3 evidence, the affected item scores zero until the conflict is resolved; the analyst may not average the conflict into a comfortable 2. The policy treats contradiction as information, not noise.

Evidence can only reduce admissible loss. It cannot waive a severity floor, manufacture a conviction tier, or turn a failed pillar into a pass. The 0.50 floor in the loss formula applies only to the epistemic multiplier after admission. It does not allow EC below 60 into ordinary active admission. Nor does it soften liquidity or factor penalties.

Class overlays: asking the question appropriate to the model

The five pillars are universal, but their decisive evidence differs by business model. The class overlays prevent a generic scorecard from giving the same benefit of the doubt to economically different structures.

For assurance businesses, including ratings, audit, testing, inspection, and certification, the overlay is binary license and liability stress. Mandated demand can create durable recurrence, yet a conflict, failed inspection, de-licensing event, or catastrophic liability can destroy a single firm quickly. A mandate supports Pillar A or B only when accreditation is legitimate, quality is sustained, and the fee remains defensible. The binary tail also drives severity toward the policy's 80%-100% range even if the balance sheet looks unlevered.

For installed-base aftermarket businesses, the overlay is the age and survival of the qualified base and the availability of independent alternatives. Required evidence connects net installed units after retirements with service attach, parts or shop-visit frequency, service revenue and cash per unit, warranty obligations, third-party substitution, and customer economics. A large but retiring fleet or a right-to-repair change can make historical recurrence a wasting asset.

For exchanges, clearing systems, payment rails, and other market infrastructure, the overlay examines both sides of liquidity, take rate, and regulation. More volume should improve utility, reliability, or execution without requiring fee surrender or creating unmanageable clearing, cyber, or conduct risk. Multi-homing and public or regulated alternatives can weaken the apparent network effect.

For vertical software and workflow data businesses, retention must be normalized for seats and AI. Gross and net retention, module attachment, workflow events, customer outcomes, data rights, compute costs, and replacement drills matter more than the subscription label. If agents reduce seats while the vendor preserves outcome value and captures part of the savings, the model may strengthen. If a foundation platform absorbs the feature and customers export context cheaply, it may weaken.

For reliability distribution, the overlay tests density, working capital, and substitution. Route, branch, vending, or on-site density should improve availability and customer productivity while inventory, receivables, labor, and supplier terms remain sound. Revenue scale alone is not a moat, and thin margins can make small errors in working capital or price transparency decisive.

For AI and cloud, three layers must not be blended: merchant capacity, internal hyperscaler capacity, and semiconductor or equipment oligopolies. Merchant capacity is tested through funded contracts, utilization by

vintage, rental prices, debt, leases, customer concentration, and power. Internal hyperscaler investment is tested against franchise persistence, AI gross profit, economic depreciation, cannibalization, and the possibility that capital expenditure is empire-building rather than defense. Equipment oligopolies are tested through node or tool share, qualification, customer concentration, inventory, export control, capacity, research and development, and full-cycle returns. Funding loops and useful-life error are explicit concerns in every layer.

For regulated or state-funded capital expenditure, the overlay follows appropriation or rate case all the way to cash. Announced budgets, awarded backlog, or allowed returns receive no automatic owner-value credit. The analysis must test enacted funding, contract form, cancellation and disallowance, regulatory lag, affordability, cost overruns, working capital, financing, dilution, and realized return.

For brands, the overlay separates cohort desire from reported price and mix. Repeat purchase, full-price units, promotion, household or client depth, resale or residual evidence where relevant, distribution power, customer-acquisition economics, and private-label or cultural substitution matter. Price/mix can rise while unit demand and brand relevance erode.

For healthcare, the overlay tests net price, payer access, exclusivity cliffs, and the replacement pipeline. Clinical value does not automatically become shareholder value. The model must include indication-level outcomes, reimbursement and denial, persistence, net rather than list price, manufacturing economics, patent or regulatory exclusivity dates, and replacement research and development. Terminal growth cannot bridge a patent cliff by assumption.

For natural resources, the overlay covers depletion, position on the cost curve, jurisdiction, and cycle. Reserve life and replacement, grade decline, all-in sustaining cost, reclamation, taxes and state take, title, conservative commodity prices, balance-sheet survival, and per-share production and cash are central. A scarcity narrative or spot-price net asset value cannot waive depletion or jurisdiction risk.

For serial acquirers, the overlay requires the organic bridge, contingent consideration, and leverage. Recurring deal spending is charged as reinvestment. Cohort returns, organic price and volume, goodwill and intangibles, earn-outs, restructuring, payables, factoring, dilution, and acquisition-definition changes determine whether the operating system compounds or merely hides organic decay.

Balance-sheet financials use class-specific batteries instead of the ancillary-finance complexity flag. Insurers require reserve-development history, float cost, and capital evidence. Banks require deposit beta, full-cycle credit, non-performing-loan behavior, and capital. Asset managers require net flows, fee compression, client concentration, and capital. Missing or failed evidence blocks a Core classification. Favorable battery scores remain provisional and are capped at 2 until the named exhibits, such as long reserve triangles, deposit cohorts, full-cycle loss evidence, and stressed-capital analysis, are attached. The overlay is meant to be equally capable of rejecting and supporting a financial company; it is not a blessing for the class.

From an admitted thesis to an allowed loss

Only after every admission gate passes does the policy calculate an allowed permanent loss:

AllowedLoss = BaseIssuerLoss(FS) * max(0.50, ECmult) * Lmult * Fmult

Each symbol has a distinct job. AllowedLoss is the maximum portfolio percentage the policy permits to be permanently impaired by that issuer under the diagnostic assumptions. BaseIssuerLoss(FS) is the starting loss budget determined by funded status, abbreviated FS. Funded status is investable assets after known taxes and transaction costs divided by the probability-weighted present value of liabilities plus an explicit safety margin. The four base issuer loss budgets are 0.25% when FS is below 1.00, 0.50% when FS is 1.00-1.24, 0.75% when FS is 1.25-1.74, and 1.00% when FS is at least 1.75. These are conditional policy bands because the investor facts remain unverified.

ECmult is the evidence multiplier: 0.50 for EC 60-74, 0.75 for EC 75-84, and 1.00 for EC at least 85. The expression max(0.50, ECmult) makes explicit that only epistemic coverage has a 0.50 floor after admission. It was scoped this way to avoid multiplying the same uncertainty into irrelevance while still requiring the evidence gate. It does not rescue a research-only EC below 60.

Lmult is the stress-liquidity multiplier. It equals 1.00 when the position can be sold within 5 stress days, 0.75 at 6-30 days, 0.50 at 31-365 days, and 0.25 beyond one year. These are stressed exit periods, not ordinary-market impressions. The multiplier recognizes that an illiquid position can remain trapped while the thesis breaks, the portfolio needs cash, or legal and market conditions deteriorate.

Fmult is the common-factor headroom multiplier. It equals 1.00 with full headroom, 0.75 when factor use is above 75% of the applicable limit, 0.50 when factor use is above 90%, and zero if the addition would breach a hard cap. Factor use is measured on a look-through basis, including shared revenue, capital-expenditure, payer, sovereign, currency, technology, financing, custody, and other material clusters, including exposure already embedded in the comparator. The factor penalty is portfolio context: the same company can have a lower maximum weight in a portfolio already crowded into its underlying risk.

The final formula contains neither a valuation multiplier nor a probability-of-loss term. Valuation has already determined admission through Pillar E. A comforting probability estimate cannot rescue a severe permanent impairment. Liquidity and factor penalties are not epistemic, so their full force remains. No floor rescues illiquidity, and Fmult = 0 makes AllowedLoss = 0 arithmetically.

The conversion from loss to weight is:

MaxWeight = AllowedLoss / S

S is bear-case permanent-impairment severity expressed as a fraction or percentage of the position. It is not mark-to-market volatility. A 50% drawdown that later recovers is not necessarily a 50% permanent impairment; a security that remains quoted but loses half its sustainable owner cash may be. Because severity is judgmental and easy to game, the policy imposes floors.

For an ordinary single equity, S>=40%. For a levered financial, cyclical company, or roll-up, S>=60%. For binary-license, pre-cash-flow, or flagged emerging-market exposure, S>=80%-100%. When more than one class applies, the more severe applicable floor governs. A lower severity requires written evidence and an independent challenge. Reputation, low historical share-price volatility, or confidence in management is not sufficient.

Hypothetical example 1: assume FS is 1.25-1.74, so the base issuer loss is 0.75%. The candidate has EC 82, giving ECmult = 0.75; stress liquidity is 6-30 days, giving Lmult = 0.75; and existing factor use is above 75% but not above 90%, giving Fmult = 0.75. Allowed loss is 0.75% * 0.75 * 0.75 * 0.75 = 0.3164% after rounding. If the company is a cyclical and therefore uses at least S = 60%, maximum weight is about 0.3164% / 60% = 0.527% of the total portfolio. The arithmetic can produce a ceiling well below a generic issuer limit because several independent portfolio constraints matter.

Hypothetical example 2: assume the best funded-state band, EC at least 85, sale within 5 stress days, and full factor headroom. Allowed loss is 1.00%. For an ordinary equity at the minimum S = 40%, maximum weight is 2.50%. At S = 50%, it is 2.00%. This is the deliberate best-case ceiling created by the policy, not a normal target.

Hypothetical example 3: keep every favorable input but suppose adding the company would breach a hard common-factor cap. Fmult becomes zero, allowed loss becomes zero, and maximum weight is zero. The 0.50 epistemic floor does not revive the position. This example captures the final correction to the formula.

Hypothetical example 4: a fully evidenced but highly illiquid levered roll-up in the best FS band has ECmult = 1.00, Lmult = 0.25 because exit takes beyond a year, Fmult = 1.00, and at least S = 60%. Allowed loss is 0.25%, and maximum weight is about 0.42%. Full evidence does not neutralize illiquidity or severity.

Reading the implied maximum-weight table

The canonical table shows the best-case ceilings when liquidity is within 5 stress days and factor headroom is full. The unparenthesized number assumes the full evidence multiplier; the number in parentheses shows only the 0.50 epistemic floor. It does not include an additional liquidity or factor penalty.

S	FS below 1.00, 0.25% loss	FS 1.00-1.24, 0.50% loss	FS 1.25-1.74, 0.75% loss	FS at least 1.75, 1.00% loss
100%	0.25 (0.13)	0.50 (0.25)	0.75 (0.38)	1.00 (0.50)
80%	0.31 (0.16)	0.63 (0.31)	0.94 (0.47)	1.25 (0.63)
60%	0.42 (0.21)	0.83 (0.42)	1.25 (0.63)	1.67 (0.83)
50%	0.50 (0.25)	1.00 (0.50)	1.50 (0.75)	**2.00 (1.00)**
40%	0.63 (0.31)	1.25 (0.63)	1.88 (0.94)	**2.50 (1.25)**

The entries are percentages of the total portfolio. Rounding explains, for example, why half of 0.25 is displayed as 0.13. Every entry remains a policy limit, never a recommendation. A real candidate can be lower because of liquidity, factor use, a severity above the floor, another portfolio cap, tax or legal constraints, or a failed implementation gate.

The table also makes the policy's concentration decision visible. The legacy generic 5% and 8% single-name rows are void. There is no generic conviction tier. With an ordinary-equity severity floor of 40%, even the best funded state, full evidence, full liquidity, and full factor headroom produce 2.50%. At a 50% severity, the ceiling is 2.00%. This closes the D-1 concentration dispute for the current policy in favor of a deliberately bounded single-issuer approach.

The contrary concentrated-value view is preserved as a rejected-with-reasons minority rather than erased. It may be reopened only after a preregistered five-year record covering at least 20 fully underwritten decisions, independently audited after-tax attribution, verified liability surplus, and independent challenger sign-off. Until that demanding but achievable test is met through a formal amendment, an analyst may not recreate concentration by choosing an implausibly low S or invoking conviction.

Speculative and pre-cash-flow positions use $S = 100\%$. The per-name speculative market-value cap of 0.25%-0.75% binds before any larger diagnostic result from the formula. These positions are assumed capable of total loss. Their separate sleeve status does not waive evidence, factor, legal, or implementation controls, and it does not make a short-history challenger comparable with a mature admitted company.

How the gates interact in practice

The most useful way to apply the scorecard is as a sequence, because each stage answers a different question.

First, define the applicable required evidence fields and class overlay. A bank, a regulated utility, a vertical software company, and an equipment maker do not need identical exhibits even though each ultimately receives A through E scores.

Second, score the five-pillar profile from current evidence. No additive total is calculated. Any pillar below 2 blocks ordinary admission, and a pillar at 0 produces a zero diagnostic maximum. A beautiful A-D profile cannot rescue E0.

Third, calculate EC independently of whether the evidence is favorable. EC below 60 keeps the candidate research-only. Resolve any material Tier 1 versus Tier 2 or Tier 3 conflict at zero rather than averaging it.

Fourth, check the hard legal, title, fraud, de-licensing, financing, factor, and implementation gates elsewhere in the policy. The loss formula operates only after all admission gates pass. If a hard factor cap would be breached, $Fmult = 0$ ends the calculation.

Fifth, set severity from the bear-case permanent impairment and apply the relevant anti-gaming floor. Establish stress liquidity and current factor headroom. Then calculate allowed loss and divide by severity.

Sixth, compare the result with every other applicable portfolio ceiling and use the lower limit. The result remains a maximum diagnostic weight. It does not supply approval, timing, account suitability, tax treatment, custody, or an order.

This sequence prevents several common forms of rule-shopping. An analyst cannot use high EC to compensate for adverse evidence, use quality to compensate for valuation, use a low probability to compensate for a severe tail, use the 0.50 evidence floor to override illiquidity, or set severity low enough to resurrect the void 5% and 8% tiers.

The scorecard also governs deterioration after initial research. A pillar falling to 0, a hard-cap breach, two missed milestones, or EC below 60 requires a full re-underwrite within 30 days and forbids an add. A hard title, fraud, or de-licensing event freezes action immediately and is resolved on the faster policy clock. Valuation deterioration alone does not force a mechanical sale or trim: E below admission forbids adding and requires re-underwriting, while an intact thesis can remain only if the standardized after-tax comparison still clears. A share-price decline alone is never a business-model falsifier.

Boundaries and failure modes of the scorecard itself

The scorecard is a governance instrument, not a discovery oracle. Its numbers are analyst judgments constrained by evidence, and its class overlays reflect mechanisms rather than statistically proven return forecasts. The policy itself acknowledges that a mature-company evidence standard can exclude young challengers, that quality businesses can share a crowded duration factor, and that the evidence may arrive after value is lost.

Most importantly, a scorecard that cannot be maintained should not be approximated casually. The complete policy limits the number of live full underwrites and can suspend active work when the research program does not justify its cost. An unmaintained EC audit, stale reverse DCF, or expired company classification cannot be treated as though the work still exists.

For the current policy document, every named company's admissible weight is zero. Investor facts, independent EC audits, current valuation, factor headroom, and implementation approval are incomplete. A company-map score or Wmax is an optimistic diagnostic ceiling based on provisional assumptions, not a completed underwrite. Nothing in this chapter authorizes an order, target weight, entry price, message, or portfolio change.

Chapter 10 - Research Priority, Classification, and the Eight-Underwrite Capacity Rule

This chapter governs the allocation of research capacity. It does not rank securities by expected return, identify holdings, recommend purchases, or authorize portfolio action. Its central problem is narrower: a one-person process cannot keep 42 company files current at full-underwrite depth, so it must decide which uncertain questions deserve scarce analytical time. The answer is a queue with no more than eight live full underwrites, and fewer - including zero - whenever the hours, evidence, comparator work, or active-program viability tests do not support eight.

The 42 company notes remain the dated research inventory as of 22 August 2026. Their preliminary classifications, pillar profiles, Evidence Coverage estimates, severity assumptions, missing exhibits, and falsifiers are useful for triage. They are not completed evidence-and-valuation underwrites. Every named company retains current admissible weight zero. A place in the research queue means only that a question has been recorded. A place in the live-underwrite set means only that the question has won a temporary claim on research capacity. Neither status is a portfolio role, a recommendation, an approval, or a prediction.

This distinction is especially important because the old label "Core Research Candidate" can be mistaken for "core holding." The revised labels remove that semantic hazard without changing any underlying classification. The old classes are mapped mechanically into clearer research labels. No company is upgraded, downgraded, promoted, recommended, or re-underwritten by the change in vocabulary. The mapped labels describe why a file may deserve attention, not what the investor should own.

The policy's research-capacity rule has four parts. First, preserve the 42-note inventory and the original classification record. Second, estimate the expected information value of additional work, including whether the work is likely to change a recorded decision. Third, compare that value with the hours and complexity required, while correcting for duplication, staleness, geographic and model bias, and comparator-embedded exposure. Fourth, fill no more live slots than can be maintained to the full standard. An empty slot is a valid result. There is no obligation to use all eight merely because the ceiling exists.

Classification vocabulary and explicit legacy mapping

The revised classification vocabulary is deliberately about research state. It does not use words normally associated with portfolio construction, such as core, overweight, starter, or position. The mapping is one-for-one:

Original classification in the 22 August 2026 map	Revised research classification	Count	Meaning of the mapping	What it does not mean
Core Research Candidate	Priority Underwrite Candidate	11	The original file presented a plausible mature-business prior deserving comparative research attention, while valuation, independent evidence, portfolio fit, and approval remained open.	Not a core holding, buy, preferred security, expected-return rank, target weight, or automatic live-underwrite slot.
Watchlist/Price Sensitive	Conditional/Price-Sensitive Research Candidate	13	The original file identified a plausible model with an unresolved gate, expectation, price, evidence, cycle, or payer issue.	Not "buy on weakness," not a price alert, and not evidence that a lower quotation alone would clear valuation.
Speculative Option	Speculative Research Option	6	The original file required short-history, emerging-market, pre-cash-flow, governance-sensitive, or other high-severity treatment under an assumed-total-loss discipline.	Not permission to invest 0.25%-0.75%, not a starter position, and not a path around ordinary admission gates.

Original classification in the 22 August 2026 map	Revised research classification	Count	Meaning of the mapping	What it does not mean
Avoid/Counterexample	Rejected/Counterexample	12	The original file illustrated one or more non-compensating failures and recorded demanding evidence that could falsify the rejection.	Not a short recommendation, sell instruction, permanent condemnation, or exclusion from future evidence review.

"Priority" in Priority Underwrite Candidate refers only to eligibility for research prioritization. It does not guarantee that the candidate enters the live set. There are 11 such files and no more than eight live slots across all classes; therefore the label cannot logically mean live status. A Conditional/Price-Sensitive Research Candidate can outrank a Priority Underwrite Candidate for a temporary slot if new evidence makes the conditional question more decision-relevant. A Rejected/Counterexample can receive bounded work when testing it would challenge an important model prior or protect against groupthink. A Speculative Research Option can receive work when the information option is unusually valuable, but it cannot borrow the mature-company evidentiary standing of the first class.

The following table preserves every original company-level classification and shows only the mechanical label translation. It is not a reordered queue. Sequence is inherited from the original map so that the table cannot be mistaken for a fresh ranking.

No.	Company note	Original classification	Revised classification	Classification effect
1	S&P Global	Core Research Candidate	Priority Underwrite Candidate	Label translation only; no promotion or recommendation.
2	CME Group	Core Research Candidate	Priority Underwrite Candidate	Label translation only; no promotion or recommendation.
3	Visa	Core Research Candidate	Priority Underwrite Candidate	Label translation only; no promotion or recommendation.
4	Fastenal	Core Research Candidate	Priority Underwrite Candidate	Label translation only; no promotion or recommendation.
5	Berkshire Hathaway	Core Research Candidate	Priority Underwrite Candidate	Label translation only; no promotion or recommendation.
6	Chubb	Core Research Candidate	Priority Underwrite Candidate	Label translation only; no promotion or recommendation.
7	DBS Group	Core Research Candidate	Priority Underwrite Candidate	Label translation only; no promotion or recommendation.
8	Taiwan Semiconductor Manufacturing	Core Research Candidate	Priority Underwrite Candidate	Label translation only; no promotion or recommendation.
9	ASML	Core Research Candidate	Priority Underwrite Candidate	Label translation only; no promotion or recommendation.
10	Schneider Electric	Core Research Candidate	Priority Underwrite Candidate	Label translation only; no promotion or recommendation.
11	RELX	Core Research Candidate	Priority Underwrite Candidate	Label translation only; no promotion or recommendation.
12	GE Aerospace	Watchlist/Price Sensitive	Conditional/Price-Sensitive Research Candidate	Label translation only; no promotion or recommendation.
13	Microsoft	Watchlist/Price Sensitive	Conditional/Price-Sensitive Research Candidate	Label translation only; no promotion or recommendation.
14	Alphabet	Watchlist/Price Sensitive	Conditional/Price-Sensitive Research Candidate	Label translation only; no promotion or recommendation.
15	NVIDIA	Watchlist/Price Sensitive	Conditional/Price-Sensitive Research Candidate	Label translation only; no promotion or recommendation.
16	Eaton	Watchlist/Price Sensitive	Conditional/Price-Sensitive Research Candidate	Label translation only; no promotion or recommendation.
17	London Stock Exchange Group	Watchlist/Price Sensitive	Conditional/Price-Sensitive Research Candidate	Label translation only; no promotion or recommendation.
18	Moody's	Watchlist/Price Sensitive	Conditional/Price-Sensitive Research Candidate	Label translation only; no promotion or recommendation.

No.	Company note	Original classification	Revised classification	Classification effect
19	Constellation Software	Watchlist/Price Sensitive	Conditional/Price-Sensitive Research Candidate	Label translation only; no promotion or recommendation.
20	Eli Lilly	Watchlist/Price Sensitive	Conditional/Price-Sensitive Research Candidate	Label translation only; no promotion or recommendation.
21	Novo Nordisk	Watchlist/Price Sensitive	Conditional/Price-Sensitive Research Candidate	Label translation only; no promotion or recommendation.
22	MercadoLibre	Watchlist/Price Sensitive	Conditional/Price-Sensitive Research Candidate	Label translation only; no promotion or recommendation.
23	Ferrari	Watchlist/Price Sensitive	Conditional/Price-Sensitive Research Candidate	Label translation only; no promotion or recommendation.
24	Rheinmetall	Watchlist/Price Sensitive	Conditional/Price-Sensitive Research Candidate	Label translation only; no promotion or recommendation.
25	CoreWeave	Speculative Option	Speculative Research Option	Label translation only; total-loss research discipline retained.
26	Samsara	Speculative Option	Speculative Research Option	Label translation only; total-loss research discipline retained.
27	Symbotic	Speculative Option	Speculative Research Option	Label translation only; total-loss research discipline retained.
28	Nu Holdings	Speculative Option	Speculative Research Option	Label translation only; total-loss research discipline retained.
29	Sea Limited	Speculative Option	Speculative Research Option	Label translation only; total-loss research discipline retained.
30	BYD	Speculative Option	Speculative Research Option	Label translation only; total-loss research discipline retained.
31	Palantir Technologies	Avoid/Counterexample	Rejected/Counterexample	Label translation only; no short, sell, or permanent exclusion implication.
32	Intel	Avoid/Counterexample	Rejected/Counterexample	Label translation only; no short, sell, or permanent exclusion implication.
33	Boeing	Avoid/Counterexample	Rejected/Counterexample	Label translation only; no short, sell, or permanent exclusion implication.
34	Coinbase	Avoid/Counterexample	Rejected/Counterexample	Label translation only; no short, sell, or permanent exclusion implication.
35	PG&E	Avoid/Counterexample	Rejected/Counterexample	Label translation only; no short, sell, or permanent exclusion implication.
36	Plug Power	Avoid/Counterexample	Rejected/Counterexample	Label translation only; no short, sell, or permanent exclusion implication.
37	Nike	Avoid/Counterexample	Rejected/Counterexample	Label translation only; no short, sell, or permanent exclusion implication.
38	Teladoc Health	Avoid/Counterexample	Rejected/Counterexample	Label translation only; no short, sell, or permanent exclusion implication.
39	Flagstar Financial	Avoid/Counterexample	Rejected/Counterexample	Label translation only; no short, sell, or permanent exclusion implication.
40	Bausch Health	Avoid/Counterexample	Rejected/Counterexample	Label translation only; no short, sell, or permanent exclusion implication.
41	First Quantum Minerals	Avoid/Counterexample	Rejected/Counterexample	Label translation only; no short, sell, or permanent exclusion implication.
42	Adani Enterprises	Avoid/Counterexample	Rejected/Counterexample	Label translation only; no short, sell, or permanent exclusion implication.

The preserved map contains 11 Priority Underwrite Candidates, 13 Conditional/Price-Sensitive Research Candidates, 6 Speculative Research Options, and 12 Rejected/Counterexamples. Those counts reconcile to 42.

The translation does not change provisional A-E profiles, Evidence Coverage, severity, falsifiers, valuation status, missing evidence, factor restrictions, or current admissible weight. If a later current review changes a classification, it must record the new evidence, the date, the old and new labels, the reason for the change, and the fact that reclassification still does not authorize a transaction.

Why 42 notes cannot become 42 live underwrites

A research note is an indexed question. A live full underwrite is a maintained analytical obligation. The latter requires a current evidence file, application of the relevant business-model battery, a five-pillar profile, an independent or adversarial corroboration pass, Evidence Coverage, a current reverse discounted-cash-flow analysis, a conservative expected-return comparison with the named implementable comparator, a bear permanent-impairment severity, stress liquidity, factor headroom, legal and tax work where relevant, preregistered disconfirmers, decision records, and monitoring. Treating a short note as equivalent to that package is evidence laundering.

The capacity rule recognizes this difference. Recurring governance is estimated at about 40 hours a year before company-specific deep work. Each live underwrite is estimated at roughly 15-25 hours a year. Eight live names therefore imply about 160-240 hours annually, before one-off comparator scoring, the first independent factor-scoring pass, custody work, unusual legal review, or a new complex event. Those hours are not a forecast of efficiency. They are an honesty requirement: the live roster must fit the operator's precommitted time using realistic complexity estimates.

The maximum live count is determined before selecting names:

LiveCapacity = min(8, HoursCapacity, ViabilityCapacity, MaintenanceCapacity)

HoursCapacity is the number of full underwrites that fit after recurring governance and one-off required work are reserved. ViabilityCapacity is zero when the prospective active-program contribution does not reach 2 times fully loaded annual cost and the program therefore remains in paper-journal mode; paper mode may still support a bounded research exercise, but it authorizes zero active capital. MaintenanceCapacity is the number of files whose evidence, valuation, and monitoring can be refreshed to standard rather than carried as stale labels. The expression is a scope rule, not a portfolio formula.

The eight-name ceiling is absolute for one operator under this policy. It is not an average, a target, or a safe harbor. A ninth apparently compelling question waits, displaces a lower-value live file through an explicit review, or remains an annual-check note. It does not enter by being called a partial underwrite, a quick update, a shadow model, or a watch item that in practice receives full cadence. Shadow underwrites count toward capacity if they consume the exhibits, monitoring, or decision responsibility of a live underwrite.

The annual hours table makes the capacity burden visible. The range includes the 40-hour recurring governance estimate and 15-25 hours for each live name, but excludes one-off comparator, independent-scoring, custody, tax, and legal projects. Those one-off commitments must be added before capacity is approved.

Live full underwrites	Company hours at 15-25 each	Recurring governance hours	Baseline annual total before one-off work	Capacity interpretation
0	0	about 40	about 40	Valid default when research is uneconomic, unmaintainable, or blocked.
1	15-25	about 40	about 55-65	One live question only; all other notes remain queue or annual-check items.
2	30-50	about 40	about 70-90	Small program; no entitlement to add more because hours happen to remain.
3	45-75	about 40	about 85-115	Capacity still depends on name-specific complexity and one-off work.

Live full underwrites	Company hours at 15-25 each	Recurring governance hours	Baseline annual total before one-off work	Capacity interpretation
4	60-100	about 40	about 100-140	A reasonable ceiling only if the files can be refreshed and challenged.
5	75-125	about 40	about 115-165	Requires explicit precommitment and a displacement rule.
6	90-150	about 40	about 130-190	Complexity concentration becomes a material operational risk.
7	105-175	about 40	about 145-215	Leaves little resilience for litigation, regulatory, or factor events.
8	120-200	about 40	about 160-240	Absolute maximum, not quota; one-off work can make this infeasible.

If only 120 total hours are reliably available, the operator does not start eight names and hope that each takes 10 hours. The operator reserves the governance and one-off burden, estimates each candidate's complexity, and selects only the count supported by the remaining hours. If a complex bank, insurer, cross-border controlled company, regulated payer, merchant-capacity company, or litigation-sensitive assurance business requires more than the ordinary estimate, its actual budget governs. Complexity is not solved by lowering the evidentiary standard.

Expected information value, not expected share-price return

Research priority is based on expected information value. Expected information value asks what a specified piece of work could teach, how likely it is to resolve the uncertainty, how likely that resolution is to change a recorded decision, how important that decision is, and what the work costs. It is not expected investment return. A deeply unattractive security can have high information value if resolving it tests an important model prior or prevents a recurring analytical error. An excellent business can have low information value if the missing evidence is unavailable, the valuation is far from any plausible hurdle, the exposure merely duplicates the comparator, or no realistic finding would change the current abstention.

The queue uses the following conceptual expression:

Expected Net Information Value = P(resolve) x P(decision changes | resolved) x DecisionConsequence + CrossCaseLearning + PortfolioOrComparatorLearning - FullyLoadedResearchCost

The terms are estimates, not facts with spurious decimal precision. P(resolve) asks whether the named work can actually answer the question within the proposed research window. P(decision changes | resolved) asks whether an answer would change a recorded research or policy decision rather than merely enrich the narrative. DecisionConsequence is the importance of that change for liability coverage, comparator design, factor headroom, a live abstention, or a future admission decision. CrossCaseLearning credits evidence that adjudicates a model class across several notes. PortfolioOrComparatorLearning credits work that exposes duplication or a hidden risk in the strategic comparator. FullyLoadedResearchCost includes operator time, data, external advice, tax and legal work, and the attention displaced from other files.

The expression need not be converted into dollars when the inputs cannot support that precision. Its practical function is to force a written comparison. Each proposed live file must state the discrete question, the evidence expected to resolve it, the decision that could change, the probability of change, the consequence of being wrong, the hours and external cost, and the next-best use of those resources. A proposal that cannot identify a decision-changing output belongs in the queue or annual-check file, not the live set.

This is deliberately different from the retired additive Business-Model Score. The business-model pillars judge a company after evidence is assembled, and a failure in one cannot be compensated by another. The

research-priority assessment decides where to gather evidence. High research priority therefore cannot waive a pillar, Evidence Coverage minimum, valuation hurdle, legal gate, factor cap, severity floor, or approval boundary. The queue may use comparative estimates because scarce attention must be allocated; the underwrite remains governed by non-compensating admission tests.

The eight selection dimensions

Every candidate for a live slot is described across eight dimensions. These dimensions are recorded separately before an overall queue judgment is made. Keeping them separate prevents a famous company, a strong prior, or a dramatic theme from silently dominating the choice.

Selection dimension	High research-priority condition	Low research-priority condition	Guardrail against misuse
Expected information value	A specific, obtainable exhibit can resolve a material uncertainty with consequences beyond narrative improvement.	The work would restate known facts, collect inaccessible data, or add detail without changing any recorded threshold.	Information value is about learning and decisions, not expected stock return.
Model-prior strength	A causal business-model prior is plausible, externally grounded, and important enough to test, with a clear counterthesis.	The prior is merely thematic, promotional, circular, or already contradicted by a hard failure that the proposed work cannot cure.	Strong prior earns investigation, not admission; correlated quality traits do not count as independent support.
Valuation distance	Current expectations appear close enough to the conservative comparator hurdle that bounded work could adjudicate E, or a new event materially changed the required assumptions.	Even an optimistic but defensible case remains far from the hurdle, or no current valuation can be constructed.	A lower price alone is not evidence; valuation distance never converts a candidate into a buy.
Portfolio/comparator uniqueness	The answer could expose a hidden comparator factor, provide genuinely different economics, improve a capped/blended comparator study, or test a liability-relevant risk not covered elsewhere.	The note duplicates an already crowded economic exposure or another live file answers the same question more efficiently.	Different ticker, country of listing, or sector label does not establish uniqueness.
Evidence freshness	A new filing, contract, regulatory action, cohort, cycle turn, or other dated primary exhibit makes the uncertainty resolvable now.	The file is stale with no decision-relevant new evidence, or management commentary is the only update.	Staleness creates expiry, not automatic priority; a source pointer earns no Evidence Coverage by existing.
Counter-canon and non-US balance	The work tests a neglected model, region, payer, financial structure, cycle, or legal context and can correct the map's quality-duration or North America bias.	Inclusion is sought only to make the roster look diverse, with no adequate question or evidence path.	Balance is a search obligation, not a quota and never forces a weak name into live status.
Complexity and capacity cost	The question is bounded, the required exhibits are identifiable, and the likely learning justifies the hours and specialist cost.	The work requires open-ended legal, technical, tax, data, or multi-entity reconstruction with a low chance of resolution.	Complexity is normally a cost penalty; it cannot be ignored because the company is interesting.
Probability the work changes a decision	Plausible outcomes would cross an explicit research threshold, produce a journaled abstention, alter comparator/factor treatment, or change which file deserves the next slot.	All plausible outcomes leave the same decision in place and produce no important cross-case learning.	A changed opinion or richer narrative is not enough; the changed decision must be named in advance.

Expected information value

The first dimension begins with a question, not a ticker. "Update Company X" is not a question. "Determine whether independently measured customer retention remains above the threshold after a bundled substitute launched" is a question. So is "determine whether reserve development and float cost remain favorable through the selected stress period," "trace an announced state budget through appropriation, cancellation rights, working capital, and collected cash," or "estimate whether the actual comparator already consumes the factor headroom that this candidate would appear to provide."

A question has high expected information value when three conditions coincide. The uncertainty is material, the missing evidence is obtainable, and different plausible answers lead to different decisions. A question has low expected information value when the evidence is not likely to arrive, the outcome cannot cross a stated threshold, or the answer matters only to the elegance of the story. The operator should prefer the smallest decisive exhibit over the broadest possible company study. If one contract cohort, reserve triangle, deposit cohort,

utilization-by-vintage schedule, or reverse-DCF sensitivity can resolve the key issue, the queue should not authorize an open-ended industry survey first.

Model-prior strength

Model-prior strength asks whether there is a plausible causal reason to spend time before the current full evidence exists. It is not a preliminary quality score. A useful prior connects demand, competitive structure, capital economics, per-share capture, and failure paths in a way that can be tested. It also names the counterthesis. A weak prior is a theme label - artificial intelligence, aging, defense, scarcity, electrification, digital payments, or emerging-market growth - without a demonstrated route to diluted per-share owner cash.

The relationship between prior strength and research priority is not monotonic. A very weak prior usually does not deserve scarce full-underwrite time. A strong but already well-understood prior may also have low incremental information value. The most productive zone is often a plausible prior with a decisive unresolved link: a payer question, economic-depreciation question, cohort-retention question, valuation assumption, shareholder-capture question, or cycle-state question that current evidence can address. The queue must record both the reason the prior may be right and the observation that would show it is wrong.

Valuation distance

Valuation distance is the gap between what the current price appears to require and what the policy's conservative comparator-relative hurdle could support. It belongs in research prioritization because a full valuation underwrite is expensive and perishable. A company can be a plausible business yet too far from a conservative return case for more work to change a near-term decision. Conversely, a new price, earnings base, capital structure, or operating disclosure can bring the required assumptions close enough that a current reverse DCF has high information value.

Valuation distance is not a mechanical price trigger. A large price decline may coincide with deteriorating economics, dilution, a broken payer, or a new hard gate. A high multiple may be supported by cash durability, while a low multiple may conceal permanent impairment. The triage question is whether a bounded valuation exercise could change the recorded research decision after updated A-D evidence, not whether a quote looks cheaper than its own history. E0 and E1 remain admission failures in the underlying company map. Research priority cannot relabel them as investable.

Portfolio and comparator uniqueness

The marginal value of research depends on what is already known and what the comparator already owns. A candidate is not unique because its ticker differs. Several technology, semiconductor-equipment, electrical-equipment, cloud, and platform companies can depend on the same artificial-intelligence capital-spending and power-availability chain. Several defense suppliers can depend on the same state procurement and working-capital mechanism. Several financial companies can depend on the same rates, credit, property, or sovereign outcome. Duplicative research consumes hours without diversifying understanding.

Uniqueness receives credit when the answer can reveal a genuinely different cash-flow mechanism, legal exposure, payer, cycle, currency, capital structure, or failure mode. It also receives credit when the work improves the actual comparator study. Because comparator exposure consumes factor headroom before active additions, a company file that helps score embedded exposure may have more decision value than a seemingly attractive duplicate security. Comparator-design learning and single-name selection value must be attributed separately at the annual review.

Evidence freshness

Freshness is a decision-timing variable. A classification expires after 12 months without refresh and reverts to Unclassified Research Note. That rule stops stale favorable and adverse labels from masquerading as current conclusions. Expiry does not, however, grant an automatic live slot. A stale file with no important new evidence may remain unclassified or annual-check because refreshing it would not change a decision.

Fresh evidence raises priority when it makes a previously unresolved question answerable. Examples include a long-awaited regulator decision, a contract with enforceable terms, a new reserve or credit cohort, a cycle turn that tests pricing discipline, a mature acquisition cohort, an economic-useful-life disclosure, a legal-title development, or enough operating history to test a preregistered milestone. Ordinary quarterly commentary and management guidance do not automatically qualify. Conflicting evidence is preserved and investigated; it is not averaged into comfort.

Counter-canon and non-US balance

The original map openly reflects the 2010-2025 consensus-quality canon and is North America-heavy. The queue must challenge that lens. Each annual refresh therefore includes a counter-canon screen and a non-US screen. For any claimed model-class absence, at least three named candidates must be examined and rejected with the failing gates recorded. The purpose is to prove that neglected areas were searched, not to force representation.

Balance operates as a challenge step, not a portfolio quota. Before finalizing the live roster, the operator compares the marginal proposed live file with the best unresolved counter-canon question and the best unresolved non-US question. If one of those alternatives has comparable or higher expected net information value, it receives the slot. If neither clears the minimum question, resolvability, and decision-change tests, the operator records the examined rejections and may leave the slot empty. A weak non-US or counter-canon file is not made stronger by the map's bias, but the bias may reveal that an otherwise comparable file has greater learning value.

Complexity

Complexity is a real cost and an error source. It includes the number of legal entities, jurisdictions, payers, customer cohorts, capital cycles, reserve or credit vintages, financing instruments, acquisition cohorts, related parties, regulatory regimes, tax questions, and factor links that must be reconstructed. It also includes the probability that specialist work is necessary. The queue estimates complexity before granting a slot and revises the estimate when scope expands.

Complexity can create information value when a misunderstood structure is consequential, but it does not automatically do so. A very complex file with a small chance of resolution may be dominated by two simpler files that can change decisions. The correct response to excess complexity is to narrow the decisive question, seek the required specialist evidence, defer the file, or leave the slot empty. The response is not to call the work complete after reading management materials.

Probability of changing a decision

The proposed decision change must be stated before work begins. Acceptable decision changes include moving a file between live and annual-check status; replacing vague interest with a journaled abstention threshold; determining that a model prior is contradicted; determining that a current valuation question deserves or does not deserve a complete underwrite; changing the factor treatment of a comparator or research line; or identifying that a legal, payer, evidence, or complexity gate makes further work uneconomic. A future change in a research classification is also possible, but only through a dated current review and never by inference from queue rank.

"I may understand the company better" is not a decision change. Neither is "the share price may rise" or "the market may notice." The probability estimate should compare plausible result branches. If every branch ends with the same abstention and no cross-case learning, the work has low decision-change probability even if the company

is important. If one obtainable exhibit could move the file from open-ended monitoring to a clear abstention, that can be valuable even though it does not create a positive classification.

The queue construction procedure

The queue is rebuilt at least annually and whenever a live slot opens after a material completion, expiry, hard-gate event, or scope change. It is not reranked continuously from price movements or headlines. Continuous reranking would consume capacity, amplify recency bias, and turn research allocation into trading by proxy.

The procedure has nine steps.

1. **Freeze the inventory and dates.** Begin with the 42 notes, their original and mapped classifications, last substantive refresh, evidence cutoff, missing exhibits, preregistered falsifiers, and current status. Expired classifications revert to Unclassified Research Note before scoring. The historic label remains in the audit trail.
2. **Reserve governance and one-off capacity.** Deduct the roughly 40 recurring governance hours and budget the actual comparator, independent scoring, custody, tax, legal, and other one-off obligations. Estimate each candidate at 15-25 hours or more when complexity requires it. Calculate the live capacity before choosing names.
3. **Apply program-level stop rules.** If the prospective 2 times fully loaded cost threshold fails, active capital remains zero in paper-journal mode. If the process has two consecutive negative deployed viability years, new active work is suspended until re-justified. If comparator scoring or Evidence Coverage maintenance cannot be reproduced in the available hours, reduce live scope or default to the named comparator plus the investor liquidity floor and verified nominal and real liability matching.
4. **Write a decisive research question for each contender.** Identify the smallest evidence package that could resolve it, the counterthesis, the branches of possible results, the recorded decision each branch would change, and the expected completion date. A ticker without a decision-changing question is not eligible for a live slot.
5. **Assess the eight dimensions separately.** Record expected information value, model-prior strength, valuation distance, portfolio/comparator uniqueness, freshness, counter-canon/non-US contribution, complexity, and probability of changing a decision. Do not hide a low value behind an overall adjective. A famous company gets no automatic advantage.
6. **Compare expected net information value per constrained hour.** Prioritize questions that are material, resolvable, decision-changing, and efficient. Also consider cross-case learning: one insurer reserve study, bank deposit study, economic-life study, or payer-quality study may improve several archived notes. Do not double count the same learning across multiple live files.
7. **Run duplication and balance challenges.** Check the actual comparator and proposed live roster for shared economic drivers. Compare the marginal candidate with the best counter-canon and non-US alternatives. Document why the selected question adds more learning or why no alternative clears the threshold. This is a bias audit, not a regional allocation exercise.
8. **Fill up to the capacity, not up to eight.** Select the highest expected net information-value questions that fit the hours and maintenance standard. Leave every unsupported slot empty. Record the marginal selected question and the best omitted question so a later displacement decision is explicit.
9. **Set a sunset and next review.** Each live file has a question-specific completion target, a 12-month classification-expiry date, and the applicable 36-month milestone or theme-sunset boundary. Work that produces neither a journaled decision nor an explicit abstention threshold within 36 months moves to annual-check status, with hours and external cost recorded.

The operating queue table is as follows:

Queue state	Maximum population	Required content	Review cadence	Exit or transition
Dated research inventory	42 in the current map	Original and mapped class, date, evidence pointers, provisional profile, missing work, falsifier, current admissible weight zero	Annual inventory reconciliation	Remains historical record; a stale class reverts to Unclassified Research Note.
Ranked research queue	Up to all unresolved notes	Decisive question, eight-dimension assessment, rough hours, next evidence event, decision that could change	Annual and when a live slot opens	Moves to live only if capacity and expected net information value justify it.
Live full-underwrite set	No more than 8; may be 0	Full evidence file, class battery, A-E, EC, reverse DCF, comparator hurdle, S, liquidity, factor headroom, disconfirmers, decision record	Event-driven, quarterly breakage review, comprehensive annual refresh	Completion, hard failure, displacement, expiry, or sunset returns the file to an appropriate non-live state.
Paper-journal work	Bounded by available research capacity; active capital 0	Preregistered hypothetical decision, evidence, expected outcome, cost and attribution	Annual viability review	Can support learning; cannot be treated as deployed evidence or authorization.
Annual-check status	No live-underwrite claim	Preserved history, explicit abstention or revival threshold, annual evidence check	Annual	New decision-relevant evidence can return the question to the ranked queue, not directly to ownership.
Rejected or archived counterexample record	No live-underwrite claim unless a new high-value question is approved	Original failure, falsifier for the rejection, evidence provenance, hours already spent	Annual map and bias review	Strong new evidence may trigger re-triage; rejection does not become permanent dogma.

A practical selection ledger

The selection decision should fit on a concise ledger even though the underwrite itself is deep. Each row records: company or cross-case question; original classification; mapped classification; last refresh; live or non-live status; decisive question; evidence expected; model prior and counterthesis; valuation distance; comparator and portfolio overlap; evidence freshness; counter-canon or non-US contribution; estimated hours and specialist cost; probability of resolution; probability that resolution changes a named decision; expected net information-value judgment; next review; expiry; and the best alternative displaced.

The ledger uses verbal bands such as high, medium, and low with a short reason rather than a mechanically precise total. If the operator uses numerical aids, the underlying notes and branch estimates must remain visible. A single summed score cannot erase a low probability of resolution, a hard capacity problem, or a missing decision link. The purpose of the ledger is comparative accountability, not statistical theater.

The marginal-selection record is especially important. Suppose the live roster has room for one more file. One candidate has a prestigious model prior but duplicates the comparator, remains far from the conservative valuation hurdle, and lacks new primary evidence. Another is a less familiar non-US or counter-canon model with a specific new contract, reserve, credit, capacity, or governance exhibit that could decide an abstention. The second may have higher information value even if its company classification appears less favorable. Selecting it would not promote the security. It would recognize that learning is most valuable where a material uncertainty can actually be reduced.

The reverse can also occur. A neglected class should not receive a live slot merely to make the roster look balanced. If the required evidence is inaccessible, the structure cannot be understood within capacity, or every plausible outcome leaves the same decision, the file remains in the queue or annual-check status. The operator then documents the examined rejection. Search balance is mandatory; favorable classification balance is not.

Wmax must be removed from queue logic

Wmax is especially likely to confuse research priority with portfolio priority. The queue should therefore not display it as a selection input, sort key, recommended size, or expected position. It does not measure attractiveness, confidence, or the value of doing research. A large theoretical ceiling can accompany a low-information question, while a zero ceiling can accompany a high-value counterexample or factor study.

Where the notation is retained for traceability, it should be printed only as:

Theoretical maximum weight conditional on every admission, evidence, valuation, investor-fact, liquidity, factor, legal, tax, portfolio, approval, and implementation gate passing; not a current weight, target, or recommendation.

This long label is intentional. The shorthand Wmax should not appear alone in a queue table, company heading, dashboard card, or sorting field. If space requires a short label, use Theoretical all-gates-pass ceiling and place the complete boundary immediately beside the number. The number belongs downstream of admission work, not upstream in research-capacity allocation.

The calculation itself remains governed by the controlling policy. After every gate passes, allowed permanent loss is based on funded state, the Evidence Coverage multiplier, stress liquidity, and factor headroom; the resulting allowed loss is divided by bear-case permanent-impairment severity. The 0.50 floor applies only to the epistemic multiplier. It does not rescue illiquidity, and a hard factor breach makes the factor multiplier zero. A pillar at zero, including E0 under the final correction, produces a diagnostic ceiling of zero. Ordinary single-equity severity cannot generally be set below 40%; levered financial, cyclical, or roll-up cases use at least 60%; binary-license, pre-cash-flow, or flagged emerging-market exposures use 80%-100%. Speculative cases use total-loss discipline and the separate per-name cap before any larger raw output.

Even when the theoretical calculation is nonzero, the present company map has not supplied verified investor facts, current complete evidence audits, current reverse DCFs, actual-comparator scoring, factor headroom, relevant legal and tax review, or implementation approval. Every current admissible company weight therefore remains zero. Presenting Wmax only as an all-gates-pass theoretical ceiling is not a drafting preference. It is necessary to prevent a provisional arithmetic output from becoming a shadow target allocation.

Valuation distance remains a legitimate research-priority input, but it is not derived from Wmax. It asks whether a current reverse-DCF exercise could change the research decision. Research capacity should be allocated before any final allowed-loss arithmetic and without assuming a portfolio weight. This ordering preserves the difference among business quality, security return, research value, and portfolio role.

Current-state consequence: this chapter does not appoint eight names

The preserved source does not support naming a current live roster of eight. The existing 42 rows are provisional, the independent Evidence Coverage audits are incomplete, current valuations remain unverified, the actual comparator and factor look-through are certification blockers, investor facts are missing, and the active-program capacity and prospective 2 times cost tests have not been completed for a specific operator budget. Selecting eight named companies in this chapter would therefore invent a fresh investment conclusion from a dated research map.

The correct current result is procedural. The 11 mapped Priority Underwrite Candidates form one input to the queue, not a preselected roster. The 13 Conditional/Price-Sensitive Research Candidates, 6 Speculative Research Options, and 12 Rejected/Counterexamples remain eligible for bounded triage when new evidence or cross-case learning raises expected information value. Before any name receives live status, the operator must calculate capacity, state the decisive question, compare alternatives, test duplication and balance, and record why the work is likely to change a decision.

This non-selection is not indecision. It respects the difference between an inherited classification and a current capacity allocation. Once the required current inputs exist, the selection ledger can identify up to eight live

underwrites without implying that any one of them is admitted to a portfolio. Until then, the map remains a 42-note research queue with zero current admissible weight for every named company.

Refresh, displacement, expiry, and sunset

A live slot is temporary. It is not a tenure right earned by a favorable old classification. Every live file must defend its continuing claim on attention against the best omitted question. The review should avoid constant churn, but it must permit displacement when expected net information value changes materially.

Displacement is appropriate when the decisive question is answered; the evidence becomes inaccessible; complexity expands beyond capacity; a hard failure makes further work unable to change a decision; the candidate duplicates another live file; an important new primary exhibit makes another question substantially more resolvable; or a comparator, liability, legal, or factor event raises the consequence of another file. A price move alone does not force displacement. Nor does publicity. The record states why the old file leaves, what status it enters, which file replaces it, and which hours are released.

Any classification not refreshed within 12 months reverts to Unclassified Research Note. The historical original and mapped classifications remain in the audit trail, but neither may be presented as current. The annual review updates every live-underwrite scorecard and reverse DCF. A stale non-live file does not receive an automatic refresh; it competes for attention through expected information value. This prevents the 12-month rule itself from recreating a 42-name maintenance burden.

The 36-month research sunset applies to a theme or line that produces neither a journaled decision nor an explicit abstention threshold. Such a line moves to annual-check status, with hours and external cost recorded. Sunset changes attention, not history. Evidence is archived, falsifiers remain visible, and new decision-relevant evidence can return the question to the ranked queue. Renewed cultural prominence is not enough.

The annual map review also reports the counter-canon and non-US searches. If no candidate from a model class or region enters the live set, that is not automatically a defect. The review shows what was examined, records at least three named rejections for a claimed class absence, and states whether the cause was failed gates, low resolvability, valuation distance, duplication, excessive complexity, or low probability of changing a decision. This makes absence inspectable without manufacturing representation.

Capacity governance and failure modes

The first capacity failure is treating eight as a quota. It produces filler work, hides the zero-active alternative, and encourages the operator to underprice time. The remedy is to calculate capacity before selecting names and leave unsupported slots empty.

The second failure is keeping shadow underwrites outside the count. A watch file with current models, repeated evidence extraction, quarterly monitoring, legal work, and a decision obligation is functionally live regardless of its label. It counts. The operator should classify work by actual burden, not by the folder name.

The third failure is equating favorable classification with high information value. A mature admired company may be well understood, far from the valuation hurdle, or economically duplicated in the comparator. Its incremental research value can be low. Conversely, a rejected counterexample may reveal why a popular model fails. Classification and queue priority answer different questions.

The fourth failure is treating evidence staleness as urgency. Expiry means the label cannot remain current. It does not mean the operator must spend scarce hours refreshing every expired file. The queue asks whether new work is resolvable and decision-changing. Otherwise the correct label is Unclassified Research Note or annual-check status.

The fifth failure is valuation anchoring. A falling price attracts attention even though the business, payer, capital structure, legal title, or per-share capture may have worsened. Valuation distance must be considered alongside current A-D evidence and the comparator hurdle. Price alone never supplies a thesis.

The sixth failure is diversity theater. A live roster can look geographically or stylistically balanced while every file has low information value. The policy requires a genuine counter-canon and non-US search, not forced positive representation. It also requires looking through legal domicile to economic geography and common factors rather than counting exchange listings.

The seventh failure is complexity admiration. Difficult companies can feel more worthy because the work is sophisticated. Complexity is normally a cost, delay, and error penalty. It earns a slot only when the expected consequence and probability of resolution justify it.

The eighth failure is double counting cross-case learning. Several live files may all be justified as ways to understand the same AI-capex, power, defense-payer, obesity-treatment, financial, or emerging-market mechanism. The operator should identify the most efficient representative question and credit the shared learning once. Other files wait unless they add a distinct mechanism.

The ninth failure is using Wmax to rank the queue. That converts a theoretical all-gates-pass ceiling into a portfolio aspiration and reverses the policy's sequence. Wmax remains downstream, conditional, and absent from queue sorting.

The tenth failure is allowing research completion to imply approval. A completed full underwrite can conclude with abstention, rejection, unresolved evidence, or a request for professional review. Even a favorable research conclusion does not verify investor facts, liability closure, comparator fit, factor headroom, tax and legal suitability, owner approval, or execution authority. Completion earns a decision record, not a transaction.

If these controls cannot be maintained, the policy reduces scope. It does not relax evidence standards, carry stale classifications, assume missing factor exposure is zero, or let a research pointer stand in for extracted evidence. The valid outcomes include fewer live files, paper-journal mode, suspension, annual-check status, and the strategic default of a named comparator plus the investor liquidity floor and verified nominal and real liability matching. None is a failure merely because it produces no active security decision.

Research priority is not portfolio priority

The research queue answers which uncertainty deserves the next scarce hour. A portfolio decision answers whether a security, in a specific investor context, passes every evidence, valuation, liability, liquidity, factor, legal, tax, governance, approval, and implementation gate. Those are different decisions. The first may rationally prioritize an uncertain, rejected, foreign, complex, or currently expensive case because the learning is valuable. The second may still and often should remain at zero.

Accordingly, the revised labels are research labels only; the 42-note order is not an expected-return rank; a live-underwrite slot is not a holding; a completed underwrite is not approval; a theoretical all-gates-pass ceiling is not a target; and an empty slot is not a missed opportunity. No more than eight live full underwrites may be maintained, fewer must be used whenever capacity or viability requires it, and every named company remains at current admissible weight zero until a separate current process completes all gates and supplies the required authority. Nothing in this chapter authorizes a trade, order, message, account action, target allocation, or portfolio change.

Chapter 11 - The Canonical Decision Workflow and a Complete Worked Example

This chapter turns the policy into one continuous decision path. It begins with the investor rather than a security, moves through liabilities, funded status, liquidity, strategic design and comparator look-through, and reaches company research only after the portfolio has established that it can bear active judgment. It then carries one fictional company through the five non-compensating pillars, Evidence Coverage, valuation, allowed-loss arithmetic, independent challenge, separate owner approval, monitoring, and an eventual evidence-based exit.

Everything in the worked case is fictional. The investor, family, companies, funds, prices, accounts, returns, contracts, tax facts, evidence exhibits, decisions, and events do not describe a real person or security. The numbers are assumptions constructed to demonstrate the rules. They are not advice, a forecast, a recommendation, a target allocation, an entry price, or authorization. Nothing here authorizes a trade, order, hedge, transfer, message, account opening, comparator change, or portfolio mutation. Even the fictional owner's approval is part of the illustration, not authority for any reader. In the live policy, a completed analysis produces eligibility for the next governed handoff, never self-executing permission.

The case also illustrates the policy's fail-closed character. A missing input is not filled by optimism. A failed non-compensating gate is not offset by a strong score elsewhere. A hard factor breach makes the relevant active addition zero. A diagnostic maximum weight is a ceiling, not a target. Independent challenge tests the work but does not authorize it. Owner approval applies only to the fictional case and cannot manufacture evidence or waive a hard gate. If any required real-world fact remains missing, the live result remains non-authorization and current admissible weight zero.

The canonical sequence

The order is part of the policy. Later work cannot cure an earlier stop.

```text START

                                                
```

v Select a conditional research prior, not a buy list

```
-- active program fails 2x fully loaded cost test? --> paper mode; active capital zero
```

v Underwrite company: class overlay, A-E, EC, reverse DCF, falsifiers

```
-- any pillar <2, EC <60, or hard gate fails? --> research-only; weight = 0
```

v Calculate AllowedLoss and MaxWeight; aggregate EB and derating budget

```
-- any portfolio limit tighter? --> lower limit governs
```

v Independent challenge

```
-- unresolved material objection? --> remediate, reject, or remain at zero
```

v Separate owner decision and implementation authority

```
-- approval absent? --> no transaction
```

v Monthly survival / quarterly breakage / annual reset

```
-- integrity event? --> immediate freeze; minimum risk-monotone response in 0-5 days
-- A-D pillar 0, persistent cap, or failed full gate? --> close/reduce review in 5-30 days
-- valuation only below comparator? --> no add; tax/cost review; no price-only sale
-- price decline alone? --> monitor evidence; do not infer thesis break
```

v EXIT, HOLD, OR CONTINUE RESEARCH ONLY UNDER THE APPLICABLE GOVERNED DECISION ```

The flow has three distinct conclusions that should never be collapsed. A business-quality prior asks whether a business model deserves research. A security-return conclusion asks whether the rights available at the current price clear the comparator after tax, cost, risk, and uncertainty. A portfolio-role conclusion asks whether the security fits the actual investor's liabilities, liquidity, factor headroom, loss budgets, and operating capacity. A company can pass the first question and fail the second; it can pass both and fail the third. None of the three supplies execution authority.

Case file: the entirely fictional Northbridge household

The fictional investor is the Northbridge household. Its home-liability currency is the fictional Crown dollar, abbreviated CD. The household has two adult decision makers, one successor agent under a current incapacity mandate, no recourse investment leverage, and a same-jurisdiction second custodian. One adult owns a cyclical logistics business, so household income is materially correlated with equity and industrial downturns. Essential payments due within three years are inflexible. These facts will shorten the error and market-drawdown budgets by one funded state even if the balance-sheet ratio initially appears stronger.

The case file uses the policy's epistemic labels:

- • [F] means a stipulated fictional fact in the case record.
- • [A] means an assumption needed to run a diagnostic.
- • [I] means an inference from the stipulated facts.
- • [FC] means a forecast or scenario input, not a prediction.
- • [C] means a research classification, not an approval.

• [U] means unresolved until the specified evidence appears.

The household's gross financial assets are CD3,180,000 [F]. Known taxes payable if all assets were made available are CD120,000 [F], and known transaction, settlement, and exit friction is CD60,000 [F]. No private-company value, residence value, unvested compensation, or inaccessible trust value is counted. Adjusted investable assets are therefore:

A* = CD3,180,000 - CD120,000 - CD60,000 = CD3,000,000

That subtraction is not prudence theater. Counting a gross balance would raise every downstream capacity number with assets that the mandate cannot spend. The case assumes that account ownership, settlement access, and the deductions have been reconciled to current records. If they had not been, A* would remain unresolved and active discretion would stop here.

The liability register

The household builds liabilities before discussing an asset allocation. Every obligation has a date range, currency, nominal or real character, probability treatment where relevant, and an explicit safety margin. The present values below already reflect the case's stipulated discounting and probability weights. They are not market forecasts.

Liability component	Character and horizon	Present value in CD	Epistemic status
Essential living and care spending	Real, years 1-15	1,250,000	[F] schedule; inflation path [A]
Housing and fixed contractual payments	Nominal, years 1-8	300,000	[F]
Health contingency	Probability-weighted real claim, years 4-12	350,000	probability and inflation [A]
Known taxes and fixed commitments	Nominal, mostly years 1-3	150,000	[F]
Incapacity, transfer, and estate administration reserve	Mixed, years 1-10	100,000	[A]
Explicit model and omission safety margin	Added after probability weighting	150,000	[A]
Base PVL*		**2,300,000**	**case input, not a certificate**

The liability schedule is also mapped by currency and matching character. CD1,600,000 is treated as real Crown-dollar liability, CD550,000 as fixed nominal Crown-dollar liability, and CD150,000 as the explicit cross-category safety margin [A]. The household has no verified foreign-currency essential liability in the case. If it did, a home-currency cash reserve alone would not close the mismatch.

The withdrawal-path file separately lists annual spending and dated payments. That file matters because an endpoint stress cannot prove how long a portfolio survives. For compactness, this chapter prints only the present-value summary and near-term cash schedule, but the fictional case stipulates that the year-by-year schedule reconciles to PVL*. In a real review, a missing withdrawal path would remain a certification blocker even if the starting ratio looked comfortable.

Funded status and the mandatory liability sensitivities

Base funded status is:

FSbase = A* / PVL* = CD3,000,000 / CD2,300,000 = 1.3043

Viewed alone, 1.3043 falls in the 1.25-1.74 funded-state band. The policy forbids stopping at that favorable point estimate. PVL* must be shown at base and at -10%, +10%, -20%, and +20%. If the +/-10% range changes the band, the lower state governs.

PVL* sensitivity	PVL* in CD	FS = 3,000,000 / PVL*	Apparent state
-20%	1,840,000	1.6304	1.25-1.74
-10%	2,070,000	1.4493	1.25-1.74
Base	2,300,000	1.3043	1.25-1.74
+10%	2,530,000	1.1858	1.00-1.24
+20%	2,760,000	1.0870	1.00-1.24

The +10% case crosses the 1.25 boundary. The binding state is therefore 1.00-1.24 [C], not the more permissive base state. This one rule changes the strategic and opportunity ranges, base issuer loss, EB, illiquidity, and active tracking-error ceiling. The household may not call 1.3043 the "real" answer and the +10% result a pessimistic footnote.

For the binding state, the policy range is 90-95% strategic, 5-10% opportunity, EB 3-5%, illiquidity 0-8%, and active tracking error 0-3%. These are ranges and ceilings, not required holdings. Unused opportunity capacity remains in the strategic sleeve. The household does not have to fill a 5% opportunity sleeve merely because the table prints 5-10%.

The second axis then tightens the case. Essential liabilities within three years are inflexible [F], and the logistics-business income is highly correlated with industrial and equity stress [I]. MDB and EB are shortened by one funded state. The case committee therefore records an EB of 2.5%, within the more conservative 0-3% band, equal to CD75,000. The separately assessed MDB is CD300,000, or 10% of A*, after the one-state shortening. The policy does not provide a universal MDB number; CD300,000 is a fictional capacity conclusion supported by the stipulated cash-flow, behavior, and liability file. Neither EB nor MDB is an amount to spend.

Promotion is also fail-closed. Even if the next annual review places the household securely in 1.25-1.74, the more permissive state would require a second consecutive annual confirmation. Death, incapacity, income loss, leverage, litigation, custody impairment, or newly inflexible spending would demote the architecture immediately. Favorable change waits; verified deterioration does not.

The Investor Liquidity Floor

The ILF is calculated before market exposure. The fictional near-term schedule provides:

- • D36 = CD300,000 of dated liabilities within 36 months [F].
- • Tax12 = CD45,000 of known tax due within 12 months [F].
- • Commit = CD60,000 of fixed undertaken commitments not already counted as spendable slack [F].
- • Essential spending beyond D36 is CD90,000 per applicable horizon unit [A].
- • Dependable income covering that same later interval is CD55,000 [A].
- • The documented income-interruption horizon is H = 2 [A], reflecting concentrated cyclical income.

The uncovered later amount is max(0, CD90,000 - CD55,000) = CD35,000. Only spending beyond D36 enters this term, so the first 36 months are not counted twice. The target is:

ILF target = CD300,000 + CD45,000 + CD60,000 + 2 x CD35,000 = CD475,000

The separately calculated hard floor consists of CD45,000 of six months' uncovered essential spending, CD30,000 of twelve months' fixed commitments, and CD45,000 of known taxes:

ILF hard floor = CD45,000 + CD30,000 + CD45,000 = CD120,000

The CD355,000 difference between target and hard floor is not available risk capital. CD475,000 remains the planned reserve. CD120,000 is the emergency level below which no risk purchase may be funded. A fall below

target requires restoration; a fall below the hard floor is a survival breach. The household does not reach for yield inside either amount to disguise the real insurance cost.

The ILF is held in short-duration Crown-dollar instruments with verified title and same-day or next-day access [F]. Beneficiary designations and the incapacity mandate have been tested [F]. A same-jurisdiction Tier-A second custodian exists, and the annual transfer test moved a nominal CD1,000 through the intended path successfully [F]. A Tier-B foreign legal-system structure is not needed. A proposed Tier-C trust wrapper remains at zero because the full tax, reporting, title, transfer, sanctions, estate-situs, and legal gauntlet has not been performed [U].

The D24 election is unavailable. The binding funded state is below 1.75, regardless of how flexible some optional spending may be. The household must use D36. Wanting less cash drag is not a qualifying election.

Strategic architecture and the grind test

With survival protected, the household names its strategic baseline. The fictional comparator is the Northbridge Policy Comparator, or NPC. It consists of the fictional low-cost Global All-Market Fund - Capped Blend share class, a Crown-dollar nominal liability ladder, a Crown-dollar inflation-linked liability ladder, and the ILF. The equity vehicle, fee, rebalance rule, tax treatment, custody, and holdings file are stipulated as implementable in the fictional case. A real policy could not use this invented fund.

The starting design is 100% strategic while active work is being evaluated. It reserves an opportunity capacity ceiling of 6%, but unused capacity stays in the comparator and matching sleeves.

Strategic component	CD amount	Portfolio weight	Role
Global All-Market Fund - initial broad version	1,025,000	34.167%	Productive comparator equity, subject to look-through
Crown-dollar nominal ladder	600,000	20.000%	Matches verified nominal liability dates
Crown-dollar inflation-linked ladder	900,000	30.000%	Matches verified real liabilities, subject to basis integrity
ILF	475,000	15.833%	Near-liability and forced-sale protection
Total	**3,000,000**	**100.000%**	**Initial strategic state**

Liability-matched nominal duration is not counted as directional rates FEW. Its residual basis risk is reported separately: reinvestment gaps, a two-month weighted-average duration mismatch, tax timing, and the legal behavior of the Crown inflation index [A]. The inflation-linked ladder receives matching credit only because currency, index, law, market integrity, duration, title, and tax treatment are stipulated as verified in the case. The label "inflation-linked" would not be enough.

The case applies the policy's default grind stresses: equity -45% real, matched nominal duration -25% real, verified duration-matched inflation-linked bonds 0% real, and ILF cash -10% real. These are assumptions, not forecasts. The weighted loss is:

- • Equity: 34.167% x 45% = 15.3752% of the portfolio.
- • Nominal matching: 20.000% x 25% = 5.0000%.
- • Verified real matching: 30.000% x 0% = 0.0000%.
- • ILF: 15.833% x 10% = 1.5833%.

Total weighted real loss is 21.9585%. Stressed end assets are:

CD3,000,000 x (1 - 0.2195845) = CD2,341,246.50

Against base stressed real liabilities of CD2,300,000, endpoint Coverage is:

Coverage = CD2,341,246.50 / CD2,300,000 = 1.0179

The immediate endpoint therefore passes by CD41,246.50 [I], conditional on the real match being genuine. It is a narrow pass. If only 25% rather than 30% qualified for the 0% real matching treatment and the other 5% had to receive the -25% nominal stress, an additional 1.25% of the portfolio, or CD37,500, would be lost. The cushion would fall to CD3,746.50. A modest omitted withdrawal or basis error would reverse it. The policy does not round a narrow pass into broad safety.

The case's year-by-year withdrawal file also has to pass. It is stipulated to preserve Coverage above 1 under the same assumptions, with the lowest annual point at 1.006 [A]. If that file were missing, the immediate endpoint would not certify the grind. No positive return is assigned to commodities, trend, or tail protection to make the result pass. Lower gross risk remains the crash null; verified inflation-linked matching is the structural grind instrument for verified real liabilities.

The architecture records a 6% opportunity capacity, a 2.5% EB, a 10% MDB, a 0-8% illiquid ceiling that the household sets to 0% because no illiquid implementation qualifies, and a 3% active tracking-error ceiling. The opportunity number is capacity, not a quota. Before any company work, active exposure is zero and all capital remains strategic.

Comparator look-through and the forced choice

The initial broad equity vehicle is not treated as neutral. Its 20 largest fictional constituents are archived individually; the remaining 15.067% of total portfolio value is assigned to documented sector buckets. Weights below are percentages of the total portfolio, not merely percentages of the fund. Scores use 0, 0.25, 0.50, 0.75, and 1.00. The table abbreviates AI-capex sustainability as AI, state-payer exposure as State, China/export-control exposure as China, directional rates/refinancing/quality duration as Rates, and power availability/cost as Power.

Holding or bucket	Weight	AI	State	China	Rates	Power
Atlas Compute	3.200%	1.00	0.25	0.25	0.75	0.75
Northstar Cloud	2.600%	0.75	0.25	0.25	0.75	0.50
Polaris Semi	2.000%	1.00	0.00	0.75	0.50	0.75
Meridian Power	1.500%	0.50	0.50	0.25	0.50	1.00
Harbor Workflow	1.300%	0.50	0.25	0.00	0.75	0.25
Greenline Grid	1.100%	0.25	0.75	0.00	0.50	1.00
Civic Defence	0.900%	0.25	1.00	0.25	0.25	0.25
Orient Foundry	0.800%	0.75	0.25	1.00	0.50	0.50
Aster Health	0.750%	0.00	0.50	0.00	0.50	0.25
Granite Bank	0.700%	0.00	0.25	0.00	0.75	0.00
Keystone Exchange	0.650%	0.25	0.25	0.25	0.50	0.25
Summit Consumer	0.600%	0.00	0.00	0.25	0.50	0.00
Bluewater Insurer	0.550%	0.00	0.25	0.00	0.50	0.25
Prairie Resources	0.500%	0.00	0.25	0.50	0.25	0.25
Verity Assurance	0.450%	0.25	0.25	0.25	0.50	0.25
Redwood Automation	0.400%	0.50	0.25	0.25	0.50	0.50
Lighthouse Pharma	0.350%	0.25	0.50	0.00	0.50	0.25
Metro Rail	0.300%	0.25	0.75	0.00	0.50	0.50
Juniper Telecom	0.250%	0.25	0.25	0.25	0.75	0.50
Orchard Brands	0.200%	0.00	0.00	0.25	0.50	0.00
Developed ex-top-20 technology bucket	3.000%	0.50	0.25	0.25	0.75	0.50
Healthcare bucket	2.500%	0.10	0.50	0.10	0.50	0.20
Industrials bucket	2.400%	0.25	0.50	0.20	0.50	0.50
Financials bucket	2.200%	0.05	0.25	0.10	0.75	0.10

Holding or bucket	Weight	AI	State	China	Rates	Power
Consumer bucket	2.000%	0.10	0.10	0.20	0.50	0.10
Emerging-markets bucket	1.200%	0.20	0.50	0.75	0.50	0.25
Other bucket	1.767%	0.10	0.25	0.20	0.50	0.20
Total equity weight	**34.167%**					

The first scoring pass produces FEW of 13.4267% for AI, 11.0918% for State, 8.5909% for China, 20.0460% for Rates, and 14.3109% for Power. The case applies common exposed-portion stress rates of 45%, 35%, 50%, 25%, and 35%, respectively [A]. The corresponding FSL values are 6.0420%, 3.8821%, 4.2955%, 5.0115%, and 5.0088% of the total portfolio.

Factor	FEW	Preferred FEW reference	FSL	Governing hard FSL cap	Result before redesign
AI capex	13.4267%	15% total portfolio or 25% of risk assets	6.0420%	5.0%	**Breach by 1.0420%, or CD31,260**
State payer	11.0918%	20% total; 15% for one sovereign	3.8821%	5.0%	Pass
China/export control	8.5909%	15% total; 10% for one irreplaceable node	4.2955%	5.0%	Pass
Directional rates	20.0460%	25% net	5.0115%	6.0%	Pass; matched duration excluded
Power availability/cost	14.3109%	20% total	5.0088%	4.0%	**Breach by 1.0088%, or CD30,264**
Joint AI plus power	n/a	n/a	8.4100%	8.0%	**Breach by 0.4100%, or CD12,300**
Joint state plus rates	n/a	n/a	6.7200%	8.0%	Pass

The independently repeated six-month pass returns AI FSL of 5.79%-6.31%, power FSL of 4.73%-5.22%, and joint AI-plus-power FSL of 8.05%-8.76% under the required +/-0.25 score perturbations [A]. The exact point estimates differ, but the breach conclusion is robust. The worksheet is refreshed at material events and annually.

The breach has appeared for four consecutive quarters. A renewable exception is forbidden. The fictional owner must choose in writing between an implementable capped or blended comparator, with costs, tracking error and abandonment point specified, or a formal amendment raising the caps after re-underwriting the risk. Raising the cap would accept more risk; it would not remove it.

The fictional challenge panel rejects a cap increase because the evidence supports concentration, not a superior liability outcome [I]. The owner separately approves the case's capped-blend branch. The fictional blend reduces the largest AI and power exposures while preserving the 34.167% equity allocation. Its modeled FSL becomes 4.60% for AI, 3.70% for Power, and 7.40% jointly. Expected tracking error is 1.2% annually [FC], incremental all-in cost is 0.08% of the equity sleeve [A], and the abandonment point is a governed review - not automatic reversal - if rolling three-year tracking error exceeds 4% or implementation cost exceeds 0.20% for two consecutive years. Tax cost is stipulated as immaterial in the case. The decision and separately assumed implementation cure the embedded breach before company admission proceeds.

Had the owner not made and implemented one of the two permitted choices, the policy would remain uncertified at comparator look-through. Material AI/power additions would have Fmult = 0, and no analyst confidence could revive them. A fixed temporary diagnostic ceiling could not exceed twice the original cap, could not float with index concentration, and would auto-revert after 12 months; the case does not use one.

Selecting a research prior without creating a buy list

With liabilities, liquidity, architecture, and factor headroom established, the operator chooses where scarce research time might have value. A decade theme is not expected-return evidence. It is a conditional prior that identifies a causal mechanism, a capital-cycle risk, a value-capture question, and falsifiers.

The case compares four fictional research lines:

Conditional research prior	Why it might deserve research	Strongest counterthesis	Portfolio interaction	Initial decision
AI equipment bottlenecks	Qualification and tool scarcity may support cash returns	Overbuild, shorter asset lives, export controls, customer bargaining	Comparator already near AI/power limits	Annual-check only; no active queue
State-funded grid contractors	Appropriation and rate recovery may convert scarcity into cash	Budgets may not be collectible; working capital and state take may absorb value	Power factor is constrained	Research note only
Reliability distribution	Dense local inventory may save customer downtime and support recurring replenishment	Working-capital drag, price transparency, supplier disintermediation, cyclicality	Low modeled AI/power sensitivity	One full underwrite authorized for research
Ordinary low-cost comparator	Broad ownership may win if dramatic themes disappoint	Embedded concentration and valuation-duration risk	Already strategic baseline	Remains burden-of-proof comparator

Reliability distribution is selected as the research prior [C]. This means only that its mechanism is suitable for a bounded full underwrite. It does not mean the industry will outperform, that the selected company is superior, or that a security is cheap. The capital-cycle response is explicitly included: apparent high returns can attract local branches, supplier direct sales, digital price comparison, and inventory financing. Value must travel from customer need to collected revenue, then through working capital, leases, maintenance investment, tax, pay, dilution, and acquisition spending before it reaches each share.

The operator is limited to no more than eight live full underwrites. In this case, only one is opened. The active program must also pass the prospective viability test before active capital can exist. The fictional five-year paper record supports an evidenced active premium of 4.0% and an evidenced hit rate of 70% [A]. With A* = CD3,000,000, an intended active sleeve ceiling of 6%, and fully loaded annual cost of CD2,200, expected program value is:

CD3,000,000 x 6% x 4.0% x 70% = CD5,040

Twice fully loaded cost is 2 x CD2,200 = CD4,400. The fictional calculation clears by CD640. Because the margin is thin, the challenger requires quarterly hour and cost recording and no expansion beyond one live company. If the threshold failed, the work would remain in paper-journal mode with zero active capital. If deployed capital later produced two consecutive negative realized-viability years, new active work would be suspended to comparator-plus-ILF-plus-matching status until re-justified.

The annual research-universe refresh also performs the required counter-canon and non-US screens. If the operator claims that no company qualifies in a model class, at least three named fictional candidates must be examined and rejected with their failing gates recorded. "None found" is not evidence of absence. The rule challenges research bias; it does not force representation or ownership.

Full company underwrite: fictional Alder Route Systems plc

Alder Route Systems plc, abbreviated ARS, is entirely fictional. It distributes critical maintenance parts through dense local branches and customer-site inventory units. The case classifies it as reliability distribution with cyclical exposure [C]. It is not a real issuer, and CD42 is an invented price.

The class overlay requires evidence on installed-site density, customer downtime economics, replenishment frequency, retention, supplier and independent alternatives, inventory turns, receivables, labor productivity, and substitution. Scale alone is not treated as a moat. Thin margins make working-capital errors material.

Evidence Coverage

The research plan defines 50 applicable weighted evidence fields before reading the outcome. Forty-one are verified, producing:

EC = 41 / 50 = 82%

EC of 82 maps to ECmult = 0.75. EC is completeness, not favorability. An adverse field can be verified. A favorable assertion with no extractable evidence remains unverified. Material conflict between a Tier 1 filing or contract and a Tier 2 or Tier 3 source scores zero until resolved; it is not averaged.

Evidence block	Verified / applicable fields	What the file contains	Main unresolved item
Demand and installed-site cohorts	9 / 10	Five-year customer and site cohorts, cancellation logs, service frequency	One regional cohort lacks full history
Structure, cycle, payer, and substitution	7 / 9	Competitor exits, supplier contracts, customer payment history	Digital marketplace response incomplete
Capital economics	10 / 12	Inventory, receivables, leases, maintenance capex, acquisition bridge	Two small acquisition cohorts immature
Governance and per-share capture	7 / 10	Pay, dilution, related parties, capital allocation	Succession and one incentive metric unresolved
Valuation and falsification	8 / 9	Reverse DCF, tax/cost comparison, milestone definitions	Terminal reinvestment range remains wide
Total	**41 / 50 = 82%**	**ECmult = 0.75**	**No unresolved item is a hard conflict**

The five non-compensating pillars

The score uses the canonical 0, 2, and 4 anchors. There is no additive total. The weakest pillar controls the description, and every pillar must be at least 2.

Pillar	Case evidence and reasoning	Score	Gate result
A - Demand and lock-in	96% customer retention, 93% site retention, replenishment linked to downtime-sensitive maintenance, and a five-year cohort ratchet are verified [F]. Customers can switch, but doing so requires inventory remapping and service interruption.	4	Pass
B - Structure, cycle, payer	Local density is helpful, but the industry remains competitive and end demand is cyclical. Customer payment history is sound; no state budget is credited as cash. Digital price transparency is a live counterthesis.	2	Pass, but severity reflects cyclicality
C - Capital economics	Incremental after-tax cash return is 18% against a 9% case cost of capital after leases, maintenance capex, economic depreciation, working capital, and dilution. Five-year cash conversion is 92%.	4	Pass
D - Governance and per-share capture	Dilution averages 0.6% annually, pay is disclosed, and acquisitions are charged as reinvestment. Two immature cohorts and succession uncertainty prevent the top score.	2	Pass, continuing monitor
E - Valuation and falsifiability	At invented price CD42, the conservative reverse DCF yields 8.4% after tax and cost versus the 6.9% comparator replacement return, a 1.5 percentage-point spread. Two primary disconfirmers are preregistered.	4	Pass at the fictional date only

The weakest pillars are B and D at 2. The company is not described as a 16-out-of-20 business because additive scoring would hide those weaknesses. If any pillar were below 2, ordinary admission would fail. If E were 0, printed diagnostic maximum weight would be zero regardless of A-D. If EC were 59, the case would remain research-only even if every available item were favorable.

Reverse DCF, conservative return, and falsifiers

The invented price is CD42 per share. The market case is represented by a current multiple of 25 times normalized owner cash [A]. The conservative case uses 18 times, low-single-digit real sales growth, a two-point operating-margin decline, working capital reverting to the weak end of the cycle, full lease treatment, continuing 0.7% annual dilution, and acquisition spending charged as reinvestment [FC]. It produces an 8.4% standardized after-tax expected return [FC] against the NPC replacement return of 6.9% [FC]. The 1.5 percentage-point margin clears the case's comparator hurdle and the Pillar E top-anchor requirement that the conservative case clear by at least 1 percentage point.

This is not a forecast that 8.4% will occur. It is the output of stipulated assumptions. A large addressable market, a past price decline, management reputation, or the research prior receives no valuation credit. If the current price or owner-cash assumptions change, E must be refreshed.

The two primary disconfirmers are registered before the decision:

1. Site retention below 92% for two consecutive quarters, accompanied by verified losses to independent or digital alternatives, would falsify the demand-lock-in mechanism and set Pillar A to 0.
2. Trailing-twelve-month cash conversion below 75% for two consecutive measurement dates after documented cycle normalization would falsify the claimed capital economics and set Pillar C to 0.

Secondary milestones include inventory turns, receivables aging, per-share dilution, organic revenue per installed site, and acquisition-cohort return. The 36-month causal horizon does not let the company postpone all evidence for ten years. Two missed primary milestones trigger no add and a full re-underwrite within 30 days.

Allowed-loss arithmetic and portfolio budgets

All admission gates have passed in the fictional file at the decision date. Only now may the case calculate allowed permanent loss.

The binding funded state is 1.00-1.24, so BaseIssuerLoss(FS) = 0.50%. EC is 82, so ECmult = 0.75. The case's archived stress-liquidity study says the position can exit within five stress days, so Lmult = 1.00 [A]. The complete liquidity schedule is 1.00 for no more than five stress days, 0.75 for 6-30 days, 0.50 for 31-365 days, and 0.25 beyond one year. After the capped comparator has been separately implemented, ARS has full headroom in every factor it materially uses, so Fmult = 1.00 [A]. The complete factor schedule is 1.00 with full headroom, 0.75 when applicable factor use is above 75% of its limit, 0.50 above 90%, and zero if the addition would breach a hard cap. Its AI and power scores are zero under the case rubric because its customer demand, capital spending, and cash flow do not materially depend on those shared scenarios. Missing factor exposure would not have been scored zero; the conclusion is based on the stipulated customer and supplier evidence.

AllowedLoss = 0.50% x max(0.50, 0.75) x 1.00 x 1.00 = 0.375%

In currency:

AllowedLoss = CD3,000,000 x 0.375% = CD11,250

ARS is cyclical, so the applicable severity floor is at least 60%, not the 40% ordinary-equity floor. For completeness, the anti-gaming floors are at least 40% for an ordinary single equity, at least 60% for a levered financial, cyclical, or roll-up, and 80%-100% for binary-license, pre-cash-flow, or flagged emerging-market exposure; the more severe applicable floor governs. A lower severity requires written evidence and independent challenge. The analyst initially uses S = 60%. Diagnostic maximum weight is:

MaxWeight = 0.375% / 60% = 0.625%

CD3,000,000 x 0.625% = CD18,750

The 0.625% is an optimistic diagnostic ceiling before independent challenge. It is not a target and does not require that capital be used. There is no generic 5% or 8% conviction tier. The legacy rows are void. High confidence cannot lift this result.

The analyst's fictional proposal is 0.50%, or CD15,000, below the diagnostic ceiling. At 60% severity, modeled permanent impairment is CD9,000, or 0.30% of the total portfolio. That fits inside CD11,250 of issuer allowed loss and inside the aggregate CD75,000 EB. It also leaves most of the 6% opportunity capacity unused in the strategic sleeve.

Liquidity and factor penalties retain full force. If stress liquidity were 31-365 days, Lmult would be 0.50 and the maximum at 60% severity would fall to 0.3125%. If an addition breached a hard factor cap, Fmult = 0, making allowed loss and maximum weight zero. The 0.50 epistemic floor would rescue neither result. Valuation is not multiplied again because it already controls admission through E. Speculative and pre-cash-flow positions instead use S = 100%; each is normally 0.25%-0.75% of portfolio market value and assumed to suffer total loss. In this binding funded state, speculative loss is capped at 2%-3% inside the opportunity sleeve, not added as a third sleeve. The household admits none.

The derating-loss budget

The position's current multiple is 25 and the conservative multiple is 18. At the proposed 0.50% weight, prospective derating loss is:

0.50% x (1 - 18 / 25) = 0.50% x 28% = 0.14% of portfolio

CD3,000,000 x 0.14% = CD4,200

Active derating loss must fit within 50% of MDB after the strategic sleeve's own stress is recognized. The case MDB is CD300,000, so the default active derating allowance is CD150,000. Existing active derating loss is zero. The ARS proposal would use CD4,200, or 2.8% of that allowance. It passes. The household does not use the optional path up to 75% of MDB; that would require a written amendment and independent challenge. An ex-post overshoot with an intact thesis would cause a sizing-process post-mortem, not an automatic trim.

Constraint	Case amount before challenge	Proposed use	Result
Aggregate EB	CD75,000	CD9,000 at S=60%	Pass
Issuer allowed loss	CD11,250	CD9,000	Pass
Diagnostic issuer weight	CD18,750 / 0.625%	CD15,000 / 0.50%	Below ceiling
Active derating allowance	CD150,000	CD4,200	Pass
Opportunity capacity	CD180,000 / 6.0%	CD15,000 / 0.50%	Pass; unused capacity stays strategic
Active tracking-error ceiling	3.0%	0.50% direct contribution before covariance	Pass in case model
Illiquid ceiling	Household chooses 0%	Security is <=5 stress days	No illiquid use

Independent challenge and separate fictional owner approval

The analyst cannot approve the analyst's own assumptions. The independent challenger receives the investor-facts pack, liability schedule, FS sensitivities, ILF and custody evidence, strategic design, comparator worksheet, company evidence extracts, reverse DCF, falsifiers, and arithmetic. The challenger is asked to disprove each link, not to provide a second narrative.

The challenge produces four findings:

1. The funded-state selection is correct. The +10% liability case crosses the boundary, and the lower 1.00-1.24 state governs. No promotion can occur without two annual confirmations.

2. The comparator cure is acceptable only after the capped blend is actually implemented through separate authority. A decision memo alone does not create factor headroom. Until implementation, any active addition materially sensitive to the breached factors remains zero. ARS is not materially AI/power-sensitive under the case evidence, but the strategic certification blocker still has to be removed before the owner relies on the full architecture.
3. Pillar D should remain 2, not 4, because succession and two acquisition cohorts are unresolved. This does not fail admission, but it prevents quality language from outrunning evidence.
4. S = 60% is too low for the combined industrial cycle, customer concentration, and supplier-rebate risk. The challenger requires S = 70%. This is stricter than the 60% floor. No special rule is needed to raise severity.

With S = 70%, allowed loss remains CD11,250, but maximum diagnostic weight falls to:

MaxWeight challenged = 0.375% / 70% = 0.535714%

In currency, CD3,000,000 x 0.535714% = CD16,071 after rounding. The proposed 0.50%, or CD15,000, remains below the challenged ceiling. Modeled permanent impairment becomes 0.50% x 70% = 0.35% of the portfolio, or CD10,500, still within CD11,250 of allowed loss. The challenged result, not the analyst's earlier 0.625%, governs.

The challenge does not authorize anything. It only resolves specified objections. The owner then receives a decision packet that distinguishes the three conclusions:

Question	Fictional conclusion	What it does not mean
Business-quality prior	Reliability distribution merits one bounded live underwrite	The class will outperform
Security-return conclusion	ARS conservative return clears the NPC by 1.5 percentage points at invented price CD42	8.4% will be realized
Portfolio-role conclusion	A 0.50% case position is below the challenged loss, factor, liquidity, EB, MDB, and tracking-error limits	The ceiling is a target or the security is suitable for a real investor

The fictional owner signs two distinct decisions. First, the owner approves and separately implements the capped comparator design, accepting its stated cost and tracking-error risk. Second, after verifying that the case's legal, tax, custody, title, and authority conditions are complete, the owner approves a maximum initial ARS case weight of 0.50%, with no automatic averaging down, the two preregistered disconfirmers, and the normal monitoring clocks. The owner declines the analyst's unused 0.035714 percentage points of challenged ceiling. Capacity need not be filled.

For pedagogical continuity only, the fictional state ledger assumes that a separately authorized implementation process establishes a 0.50% ARS holding at CD42 with no pending information event. Because all gates clear and no scheduled filing or preregistered milestone is pending, full deployment is the statistical default in the case. If a material information event had been pending, any pacing would have needed to be event-anchored and complete by day 180. Calendar-only tranches could not roll indefinitely. None of this paragraph is a live order or permission for a reader.

The post-decision portfolio has 0.50% in the fictional active company and 99.50% in strategic assets, with the active amount funded from the equity comparator. Although the architecture allowed up to 6% opportunity capacity, only 0.50% is used. The ILF, nominal match, and real match are untouched. No risk is bought from the emergency floor.

Decision journal at the point of the fictional approval

The decision journal records the conclusion that existed before the outcome was known.

Journal field	Entry
Decision date and evidence cutoff	Fictional day 0; all evidence dated and archived

Journal field	Entry
Governing FS	1.00-1.24 because +10% PVL* crosses the band
Second-axis treatment	EB and MDB shortened one state for inflexible liabilities and correlated business income
ILF	Target CD475,000; hard floor CD120,000; both verified and protected
Comparator	Capped blend separately approved and assumed implemented; top-20 file and second pass archived
Research classification	Fully underwritten active candidate for the fictional date only
Pillars and EC	A4, B2, C4, D2, E4; EC82; no hard gate
Valuation	Conservative after-tax return 8.4% versus comparator 6.9% [FC]
Allowed loss	0.375%, or CD11,250
Challenged severity and ceiling	S70%; 0.535714%, or CD16,071 maximum diagnostic weight
Fictional approved case weight	0.50%, or CD15,000; not a target for any real portfolio
Derating loss	0.14% of portfolio, or CD4,200
Primary falsifiers	Site retention below 92% for two quarters; normalized cash conversion below 75% twice
Prohibited responses	No automatic add after a price decline; no tax rescue of a hard gate; no exception rollover
Authority boundary	Analysis and challenge complete specified work; separate fictional owner approval assumed; chapter authorizes nothing

Monitoring the case without converting noise into action

Monitoring is hierarchical. Monthly review covers ILF, custody, title, and hard-gate alerts only. Quarterly review reconciles holdings, issuer and factor headroom, pillar-breaking evidence, tax lots, milestones, and the exception log. The annual review recalculates FS and PVL* sensitivities, refreshes the comparator top 20 and buckets, audits ILF drag, updates every live scorecard and reverse DCF, performs the transfer test, tests active-program viability, and considers amendments.

The fictional post-approval record unfolds as follows:

Review point	New information	Policy interpretation	Governed result
Month 1	ILF remains CD475,000; custody and title unchanged	Survival controls intact	No action
Quarter 1	Retention 95%; cash conversion 90%; factor headroom intact	Thesis and gates remain intact	Continue monitoring; no routine rescore
Quarter 2	Market price falls 35% while retention, site economics, dilution, and cash conversion remain on plan	Mark-to-market loss is not evidence of permanent impairment; E may improve	No price-only sale and no automatic add
Quarter 3	Site retention is 91.5%, first miss; independent alternatives gain in one region	One primary milestone miss, not yet the two-miss trigger	Log contrary evidence; refresh A; no story-based dismissal
Quarter 4	EC refresh is 76%; second-custodian transfer passes; price partly recovers	ECmult remains 0.75; classification refreshed before 12-month expiry	Continue, with A at risk; annual reset completed
Quarter 5	Site retention is 90.8% for a second consecutive quarter; customer losses to alternatives verified	Second primary milestone missed; no add and full re-underwrite due within 30 days	Freeze additions immediately; open governed re-underwrite
Re-underwrite day 18	Cohort evidence shows the local-density switching-cost mechanism no longer holds in two major regions	Pillar A becomes 0; this is a causal break, not a price signal	Owner receives close-or-reduce recommendation under 5-30-day hierarchy
Re-underwrite day 23	Separate fictional owner approves closing the case position; legal, tax, and liquidity checks complete	Hard business-pillar failure cannot be rescued by tax deferral or a recovered price	Fictional case records full exit within the governed window

At Quarter 2, the 35% price fall creates a mark-to-market loss of about CD5,250 on the initial CD15,000 case position before other changes. The price movement is real, but it does not by itself set A-D to zero. The monitor checks evidence, factor headroom, valuation, and the registered milestones. It does not average down because a cheaper price is not evidence. Any addition would require a current thesis, updated EC and valuation, intact portfolio budgets, factor headroom, ILF protection, and new separate approval.

At Quarter 5, the second missed primary milestone triggers the exact 30-day re-underwrite and no-add rule. The subsequent A0 finding is a non-compensating failure. The governed exit is based on verified cohort deterioration in the mechanism that supported demand and lock-in. A later price recovery would not repair it. Tax deferral could not turn A0 into a pass. The owner-approved case exit on day 23 falls within the five-to-30-day close-or-reduce review for an A-D pillar at zero.

Had the event instead been credible fraud, de-licensing, a title failure, or a financing-covenant threat, the affected discretion would have frozen immediately and the minimum separately authorized risk-monotone response would have occurred within zero to five trading days. An emergency exception would have been recorded within 24 hours, independently reviewed within seven days, and expired within 30 days. It could not have been used to add risk or establish a permanent precedent.

Had the thesis remained intact while the prospective return fell below the comparator after standardized tax and cost, the immediate result would have been no add, not automatic sale. The deferral comparison would use the applicable statutory tax rate, a fixed ten-year horizon, and the comparator replacement return. A hold relying chiefly on tax deferral at two annual reviews would have required an independent challenger. After a governed decision to act, the next liquid window would mean five trading days absent a security-specific integrity event. A proposed 90-day tax wait would have been costed on day 0. That valuation-only branch did not decide this case because the actual trigger was Pillar A at zero.

The case has no material Policy Exception Event. If the household had used the ILF floor to fund risk, taken an unjournaled panic action, or missed a mandatory review without written deferral, the event would have entered the log. Two material Policy Exception Events in a rolling 12 months would freeze new discretion pending review regardless of whether the deviations happened to make money.

Closure, expiry, and research sunset

The fictional exit does not erase the evidence file. ARS moves from live underwrite to archived counterexample. Its classification would expire after 12 months without refresh and revert to Unclassified Research Note. Expiry is not an investment view; it means the old classification is no longer current.

The reliability-distribution research line remains active only if it produces decision value. If, over 36 months, it produces neither a journaled decision nor an explicit abstention threshold, it moves to annual-check status. Hours and external cost are recorded. The historical ARS file remains archived so a future analyst does not relearn the same failure. Sunset reduces attention; it does not rewrite history or claim that the entire class is permanently unattractive.

The active-program viability test also continues. A single successful or unsuccessful outcome is not enough to infer edge. Annual value is recalculated as active sleeve times evidenced premium times evidenced hit rate and compared with fully loaded hours, data, tax, and external cost. Below the prospective 2x threshold, capital returns to paper mode. Two consecutive negative realized-viability years suspend new active work. The default then contracts to the named comparator, ILF, and verified nominal and real liability matching until the active program is re-justified.

The annual portfolio reset does not mechanically refill the exited company or its theme. Broad strategic sleeves may rebalance when more than 5 percentage points or more than 20% relative from an approved band, subject to ILF and tax. A broken active thesis is never refilled merely because its portfolio weight fell. Any future ARS research would start again with current evidence and could not borrow approval from the expired case.

Control ledger: where the case would fail closed

The following table makes explicit that every apparent permission is conditional.

Stage	Required evidence or decision	Fail-closed result
Investor facts	Reconciled A*, taxes, liabilities, currencies, income correlation, authority, succession	Stop affected discretion; do not infer a funded state
FS	Base and PVL* +/-10% and +/-20%; hysteresis	Use lower state on ambiguity; no one-year promotion; immediate verified demotion
ILF	Target, hard floor, home-liability currency, title, access, Tier-A transfer	No new discretion if hard floor, title, access, or transfer fails
Grind	Default-admissible assets and withdrawal path	Coverage below 1 requires liability/design remedy; no unearned diversifier gain
Architecture	Named comparator, matching, MDB, EB, illiquidity and tracking error	Placeholder cannot persist beyond 12 months; missing budget blocks active admission
Look-through	Top 20, remainder buckets, five factors, joint shocks, second pass	Hard-cap-creating addition has Fmult 0; four-quarter embedded breach requires choice
Research prior	Causal mechanism, counterthesis, value capture, capital cycle, viability	Theme remains research only; failed 2x cost threshold means paper mode
Company gates	Class overlay, A-E all at least 2, EC at least 60, no hard gate	Any failure means research-only and current admissible weight zero
Valuation	Current reverse DCF and conservative comparator hurdle	E below admission means no add and re-underwrite; quality cannot compensate
Loss arithmetic	Correct FS base loss, ECmult, Lmult, Fmult, severity floor	Lower applicable ceiling governs; hard factor breach arithmetically gives zero
Derating budget	Current and conservative multiples; aggregate active loss <=50% MDB	No admission above default; up to 75% needs amendment and independent challenge
Challenge	Independent review of facts, gates, severity, valuation and arithmetic	Unresolved material objection means remediate, reject, or remain zero
Owner and implementation	Separate authority, tax/legal/custody completion, decision journal	Analysis does not self-authorize; absence of approval means no transaction
Monitoring	Monthly survival, quarterly breakage, annual reset, current classification	Missed evidence or clocks tighten or freeze discretion; no headline substitution
Exit	Causal trigger and correct decision clock	Price alone never compels sale; hard evidence cannot be postponed by narrative or tax

What the arithmetic does and does not say

The case's arithmetic is internally connected:

• A* is CD3,000,000 after CD180,000 of known tax and exit friction.

• Base PVL* is CD2,300,000, producing FS 1.3043, but +10% liabilities reduce FS to 1.1858, so the 1.00-1.24 state binds.

• ILF target is CD475,000 and hard floor CD120,000. The target is included in the strategic architecture and never funds the active case.

• The initial grind design loses 21.95845% under stipulated real stresses, leaving CD2,341,246.50 against CD2,300,000 of liabilities, a CD41,246.50 conditional endpoint cushion.

• Initial comparator look-through breaches AI FSL by 1.0420%, Power FSL by 1.0088%, and the joint cap by 0.4100%. The four-quarter breach is cured by an explicit fictional capped-blend choice, not a renewable exception or a floating cap.

• ARS has A4, B2, C4, D2, E4 and EC82. The business prior, security return, and portfolio role are separately stated.

• AllowedLoss = 0.50% x 0.75 x 1 x 1 = 0.375%, or CD11,250.

• At the analyst's S60%, the diagnostic ceiling is 0.625%. The challenger raises severity to S70%, lowering the ceiling to 0.535714%, or CD16,071.

• The fictional proposal is 0.50%, or CD15,000. At S70%, its modeled permanent impairment is CD10,500, below CD11,250.

• Derating from 25 times to 18 times at a 0.50% weight creates a 0.14% portfolio loss, or CD4,200, within the default 50%-of-MDB active derating allowance.

These calculations establish only consistency with the fictional inputs at the dated decision point. They do not establish that the inputs are true for any real investor, that the stresses have correct probabilities, that the invented company would earn 8.4%, or that a 0.50% position is appropriate elsewhere. Loss arithmetic bounds a modeled error; it does not predict the outcome. The grind pass is narrow and conditional; it is not a guarantee. The comparator redesign exchanges concentration for cost and tracking error; it does not create a free improvement.

The governing boundary after a fully worked case

A complete example can create a dangerous illusion of certification. The Northbridge file looks complete because all missing elements were stipulated into existence. Real files rarely arrive that way. The controlling policy remains dated 22 August 2026, and its numeric bands are policy hypotheses awaiting actual investor facts. A source pointer proves only that a disclosure existed; it does not verify a management claim. A research classification expires after 12 months. A favorable challenge satisfies only the challenge condition. An owner or operator cannot waive evidence by signing a page.

For any real company, current admissible weight remains zero until the actual investor facts, liability closure, FS sensitivities, trait tolerance, MDB, ILF and access, actual named comparator, top-20 look-through, factor headroom, current evidence audit, class overlay, five pillars, valuation, tax and legal review where relevant, loss budgets, independent challenge where required, and separate implementation approval have all been completed. Even then, the result is specific to that separate process. This chapter itself authorizes nothing.

Open disputes also remain open. The worked case chooses a capped comparator only inside its fictional facts; it does not settle the policy-wide capped-versus-uncapped comparator dispute. It uses no private assets, so it does not resolve the private/illiquid question. It assumes Crown-dollar liabilities, so it does not settle the long-horizon FX hedge range. It does not prove that AI capex will disappoint, that inflation-linked instruments always match, or that active research has durable edge. The concentration minority remains closed under the current policy unless a formal future amendment is supported by a preregistered five-year record of at least 20 fully underwritten decisions, independently audited after-tax attribution, verified liability surplus, and challenger sign-off.

What the example demonstrates - and what it cannot establish

The example demonstrates the operating logic of the policy: liabilities select capacity; sensitivity can bind more tightly than the base case; the ILF protects near claims; real matching, not an optimistic diversifier, closes the grind; the comparator consumes factor headroom; research priors organize inquiry without proving return; five pillars do not compensate for one another; EC controls the epistemic multiplier but cannot waive admission; liquidity and factor penalties retain full force; severity converts allowed loss into a diagnostic ceiling; independent challenge can tighten the result; owner approval is separate from research; and exit follows causal evidence and the correct clock rather than price alone.

It cannot establish that any real investor has the Northbridge facts, that any real asset matches a liability, that the invented comparator is implementable, that any actual company resembles ARS, that the scenario inputs will occur, that active management will add value, or that a trade should be made. It is not advice, a forecast, or authorization. Where real evidence is missing, stale, contradictory, unmaintainable, or unapproved, the correct outcome remains fail-closed: research-only, paper mode, frozen discretion, or current admissible weight zero.

Chapter 12 - Monitoring, Falsification, Sell Discipline, and Research Sunset

Monitoring exists to keep a decision aligned with current evidence, not to create a reason to trade. That distinction is central to this policy. A price can move sharply while the underlying claim on future cash remains sound, and a business can deteriorate while its quoted price is rising. The monitoring system therefore watches survival risks, evidence quality, thesis milestones, portfolio limits, valuation, taxes, and the cost of research on different clocks. It does not treat every market move as new information, and it does not allow a long-term story to postpone recognition of contrary evidence.

This chapter governs research and decision discipline. It does not authorize an order, target allocation, transfer, hedge, message, or portfolio change. Even when an exit rule says that action is required on a stated clock, an actual transaction may occur only through the separate implementation process and with whatever owner approval, verified investor facts, tax or legal review, and execution authority that process requires. The present company map remains research-only, and every named company has a current admissible weight of zero until all policy gates and a complete current underwrite pass.

Monitoring is a hierarchy, not a stream of alerts

The policy deliberately assigns different questions to monthly, quarterly, and annual review. This prevents two opposite errors. The first is neglect: allowing a custody problem, broken covenant, stale evidence file, or failed business premise to persist because the investor calls everything "long term." The second is over-monitoring: repeatedly rescoring a ten-year thesis from headlines, price changes, and quarterly noise until the process becomes expensive, emotional, and procyclical.

The monthly review is narrow. It covers the investable liquidity floor, or ILF, custody and legal-title alerts, and hard-gate events only. The ILF is the liquid, high-integrity reserve designed to cover near liabilities and protect against forced selling. A monthly ILF check asks whether that protection still exists; it is not a monthly return contest. Custody and title monitoring asks whether the investor still has an enforceable, accessible claim on the assets. Hard-gate monitoring looks for events such as evidenced fraud, de-licensing, a financing-covenant threat, or another condition that could turn delay itself into permanent loss. These are survival and integrity questions. They deserve the fastest cadence because the damage from waiting may be discontinuous.

The monthly cadence is not an instruction to refresh every company model. Routine earnings commentary, market sentiment, and a share-price decline do not belong in this survival check unless they contain evidence of a hard-gate event. The one-page operating rule elsewhere in the policy states the idea plainly: monthly means survival alerts, not a complete rescore. This boundary preserves attention for the events where speed matters.

The quarterly review covers holdings, issuer and common-factor headroom, evidence capable of breaking a scorecard pillar, tax lots, and the policy exception log. Holdings must be reconciled because the monitored portfolio must be the portfolio actually owned, including changes caused by distributions, corporate actions, or market movement. Issuer headroom measures distance from a single-issuer limit. Factor headroom measures aggregated exposure to risks that can appear in many holdings and in the comparator at once, such as AI capital spending, power infrastructure, rates, a sovereign, a currency, or a financing chain. The purpose is look-through control: ten individually plausible names do not create diversification if they depend on the same economic outcome.

Quarterly pillar review is also selective. The five non-compensating pillars are A, demand and lock-in; B, industry structure, cycle, and payer quality; C, capital economics after economic depreciation and dilution; D, governance and per-share capture; and E, valuation and falsifiability. The quarterly task is to identify pillar-breaking evidence, not to manufacture tiny score changes from every new data point. Examples include a binding regulatory action that changes the payer economics, reserve development that impairs capital, customer retention that contradicts

the demand premise, or new dilution that prevents operating value from reaching each share. Evidence that does not alter a causal link can be logged without forcing a new decision.

Tax lots appear quarterly because taxes can affect implementation cost and the relative merit of holding or selling, but tax is not allowed to become the business thesis. The exception log appears quarterly because repeated deviations can reveal that the policy is not executable as written. Under the wider policy, two material Policy Exception Events in 12 months freeze new discretion pending review. That safeguard is behavioral and operational: an exception is not a renewable substitute for an amendment.

The annual review is the comprehensive reset. It recalculates funded status, or FS, using the probability-weighted present value of liabilities plus safety margin, or PVL, at the base estimate and at plus or minus 10 percent and plus or minus 20 percent. It refreshes the actual comparator's top 20 constituents and buckets the remainder, audits ILF drag, updates every live-underwrite scorecard and reverse discounted cash-flow analysis, performs the jurisdiction transfer test, evaluates the economic viability of the active research program, and considers policy amendments.

Each annual item answers a different form of drift. FS and PVL sensitivity identify whether the investor's capacity has changed. The comparator refresh prevents active holdings from being assessed against a stale or falsely neutral benchmark. ILF drag makes the cost of survival liquidity visible rather than tempting the investor to reach for yield inside the reserve. A reverse discounted cash-flow analysis, or reverse DCF, starts from price and asks what operating and cash-flow assumptions the market price requires; refreshing it prevents an old valuation conclusion from surviving a new price or capital structure. The transfer test checks that cross-jurisdiction access works in practice, not merely in account documentation. The active-program review asks whether the expected contribution of active work justifies its fully loaded time, data, tax, and external costs.

That annual viability test is prospective as well as retrospective. Active capital is enabled only when portfolio value multiplied by the active sleeve, evidenced premium, and evidenced hit rate is at least two times fully loaded annual cost. Below that threshold, the active program stays in paper-journal mode with zero active capital. If capital has been deployed, two consecutive negative realized-viability years suspend new active work. Monitoring is therefore allowed to falsify the value of the monitoring apparatus itself.

Evidence has an expiry date

A research label is not a permanent property of a company. Any company classification that has not been refreshed within 12 months expires and reverts to "Unclassified Research Note." The reversion is substantive, not clerical. It means that a former Core Research Candidate, Watchlist/Price Sensitive name, Speculative Option, or Avoid/Counterexample cannot continue to claim current standing from old work. The evidence may remain useful history, but the classification no longer represents a current judgment.

This rule controls map decay. A research map with 42 named companies can easily become an archive that looks like a live portfolio of judgments. The policy accepts that one operator cannot maintain 42 full underwrites and permits no more than eight live full underwrites. The 12-month expiry keeps the visible label aligned with the work that can actually be maintained. It also prevents favorable classifications from enjoying implicit permanence while adverse classifications are refreshed only when the price becomes interesting.

Evidence Coverage, abbreviated EC, is the percentage of applicable required fields supported by verified evidence, weighted toward primary filings, contracts, and regulators. EC below 60 percent is research-only under the scorecard. In monitoring, EC below 60 triggers a full re-underwrite within 30 days and forbids any add. The same 30-day re-underwrite and no-add rule applies after two missed milestones. These are separate tests: a thesis can have broad documentary coverage yet miss the outcomes it predicted, or it can appear operationally on track while the evidence file has decayed below the minimum needed to rely on that appearance.

Milestones must be causal and preregistered. A decade thesis is not expected to reach its terminal outcome in three years, but it must produce observable intermediate evidence within 36 months. A thesis about installed-base

economics might require rising service revenue per installed unit and stable retention. A thesis about a capital-cycle harvest might require verified capacity exit, repaired financing, and actual shareholder cash returns. A thesis about AI monetization might require incremental gross profit to cover incremental depreciation, power, and financing costs. Missing two primary milestones cannot be excused by repeating the ten-year narrative. The policy's governing idea is that a long horizon changes what evidence should appear, not whether evidence must appear.

The 30-day review clock is a governed window for rebuilding the evidence record, reassessing the causal thesis, updating the scorecard and valuation, and deciding the permitted status. It is not an automatic 30-day grace period for continued buying. No add applies immediately once the trigger is recognized. Nor does the review guarantee restoration. If the evidence cannot bring EC back to at least 60 or restore the missed causal links, the name cannot retain a fully underwritten standing.

The annual company-map refresh must also challenge the composition of the research universe. It must report a counter-canon screen, a non-US screen, and at least three named rejected candidates for any claimed class absence. A counter-canon screen deliberately examines areas outside the fashionable quality consensus. A non-US screen tests whether geography is being mistaken for quality or data availability. If the policy says that no company qualifies in a model class, "none found" is insufficient: at least three named candidates must have been examined and rejected, with the failing gates recorded. The requirement does not force representation, and it does not turn a rejected candidate into a recommendation. It makes an absence an evidenced result rather than an unexamined assumption.

Falsification separates an investment thesis from a story

A falsifier is an observable that would materially disconfirm the causal chain supporting a classification or holding. It is not merely a disappointing quarter, and it is not a price threshold. Good falsifiers connect business reality to per-share owner economics. Examples in the research map include recurring growth below a stated rate, adverse reserve development above a stated share of opening equity, a combined ratio structurally above 100, non-performing loans above a stated level, sustained utilization below a stated threshold absent recession, or cash conversion remaining below a stated level through a full cycle. The exact company-specific tests belong in each current underwrite; the monitoring rule is that they must be stated before the evidence arrives.

Preregistration matters because the same fact can be rationalized after the event. If an investor writes in advance that falling utilization would show overcapacity, the investor cannot later call falling utilization an intentional land grab without new evidence and a documented thesis change. The policy permits re-underwriting; it does not permit an ex-post rewrite that makes the thesis impossible to falsify.

Pillar scores make this discipline non-compensating. A strong brand does not offset destructive incremental capital economics. High growth does not offset governance leakage. A low price does not restore legal title. Evidence can improve EC and therefore the amount of loss that may be tolerated elsewhere in the policy, but it cannot waive a hard gate. Likewise, a favorable narrative cannot compensate for a pillar at zero.

Consider a clearly hypothetical example, not a forecast or live position. A distributor is classified on the premise that embedded replenishment sites create density, customer retention, and better cash conversion. Its two primary milestones are that new sites produce measurable share gains and that cash conversion improves above the preregistered floor. The share price falls 25 percent during an industrial slowdown, but retention, site economics, working-capital normalization, and per-share cash remain consistent with the thesis. That price move alone is not falsification. If, however, two milestone dates pass without share gains and cash conversion remains below the stated floor even after cycle normalization, the two-missed-milestone trigger applies. Adds stop and the full underwrite must be reopened within 30 days, even if the share price has recovered.

Now alter the example. Suppose the distributor's EC falls from 78 percent to 58 percent because key cohort data can no longer be reconciled and a conflict between primary and secondary sources remains unresolved. The stock

price and reported earnings are irrelevant to the trigger. EC below 60 independently requires the same 30-day re-underwrite and no add. The correct question is not whether the investor still feels confident; it is whether the evidence required to bear risk remains verifiable.

Price volatility is not thesis breakage

The policy maintains three separate budgets: permanent impairment, mark-to-market drawdown, and decade-long real grind. Chapter 9 is especially concerned with keeping the first two separate. Mark-to-market drawdown is the decline in quoted value. Permanent impairment is damage to the underlying owner claim, such as loss of title, fraud, a broken license, insolvency, destructive financing, failed economics, or governance leakage. A large drawdown may occur without permanent impairment, while permanent impairment can begin before the market price acknowledges it.

For that reason, mark-to-market drawdown alone never forces a sale. This is not a promise to hold through every decline. It is a prohibition on using price alone as evidence that Pillars A through D have failed. If those pillars and the registered milestones remain intact, a 30 to 40 percent multiple derating can be painful without being a thesis break. At the lower price, valuation Pillar E may improve because the prospective return is higher. Conversely, a rising price can leave E below admission because the expected return has fallen beneath the comparator.

The distinction also blocks thesis-free averaging down. A lower price is not, by itself, permission to add. Any addition would still require a fresh evidence and valuation record, adequate EC, intact pillars, issuer and factor headroom, sufficient ILF, compliance with the loss and valuation-drawdown budgets, and separate implementation authorization. Where the current policy bars an exposure because a factor cap is breached, a cheaper price does not remove that bar. The investor may not convert a falling quote into evidence that the business is safer.

A second hypothetical illustrates the point. Assume a company was underwritten on stable customer retention, incremental cash returns above its cost of capital, and disciplined per-share capital allocation. Its market multiple falls from 30 times to 20 times while retention, cash returns, dilution, and milestones remain unchanged. The quoted loss is real, but the thesis has not necessarily broken. Monitoring should update E and the reverse DCF, check common-factor exposure, and compare expected return with the comparator. It should not invent a fundamental failure from the chart. If instead retention collapses, the company issues equity to fund structurally negative returns, or a regulator removes the license on which demand depends, the problem is no longer volatility. One or more causal pillars or hard gates have failed, and the faster exit hierarchy applies.

This approach carries a trade-off. Waiting for causal evidence can allow further price loss before impairment is fully verified, especially when obsolescence or political action moves faster than financial reporting. Acting on price alone can avoid some losses by chance, but it also creates a systematic tendency to sell intact claims after declines and to retain impaired claims while prices remain strong. The policy resolves that trade-off by using fast clocks for title, fraud, licensing, and financing threats; governed review clocks for pillar failures; and no price-only sell rule.

The exit hierarchy and its decision clocks

Exit decisions follow four levels. The ordering reflects the reversibility of the problem and the cost of delay.

First, a legal-title, fraud, de-licensing, or financing-covenant threat requires the minimum risk-monotone action within five trading days, if that action has been authorized under the separate implementation process. "Risk-monotone" means the permitted response must reduce or contain the affected risk; it cannot use an emergency as an excuse to increase gross risk or buy a dislocation. "Minimum" means no broader action than necessary to protect title, solvency, or compliance. The five-day ceiling is not a required wait. A security-specific integrity event can justify a faster response, subject to market integrity and implementation authority.

The wider emergency protocol remains narrow. Such an exception is recorded within 24 hours, independently reviewed within seven days, and expires within 30 days. Unaffected policy rules continue to apply. This boundary is important: the research operator may identify the threat and start the governed process, but this chapter does not itself grant transaction authority.

Hypothetically, suppose credible primary evidence shows that a custodian's asset segregation may have failed. The monthly alert becomes an immediate hard-gate review. If the separately authorized implementation process permits a transfer, the response should be the minimum transfer needed to reduce title risk, not a wholesale change in market exposure. The action clock is within five trading days, potentially sooner if delay threatens recovery. The event is documented and independently reviewed on the emergency timetable. A favorable market outlook cannot override the title problem.

Second, if any of Pillars A through D becomes zero, a hard cap persists, or a full-underwrite gate fails, the position is closed or reduced under a governed review lasting from five to 30 days. A zero in A through D signals failure in demand, structure or payer, capital economics, or governance and capture. These are owner-claim failures, not ordinary valuation disagreements. A failed full-underwrite gate also includes evidence or control conditions that make fully underwritten status inadmissible. The range allows enough time to verify complex evidence, taxes, liquidity, and legal consequences while refusing indefinite delay.

Persistent cap breaches cannot become renewable exceptions. For an embedded issuer or factor hard-cap breach that continues for four consecutive quarters, the broader policy forces a written choice between a compliant capped or blended comparator and a cap raised by formal amendment; there is no renewable exception. For an active holding whose hard cap persists, the five-to-30-day close-or-reduce hierarchy applies. A formal amendment changes the governing policy and therefore requires its written process and approval; an analyst cannot amend a limit by continuing to hold and calling the breach temporary.

Third, a different rule applies when the business thesis is intact but expected return falls below the named implementable comparator after standardized tax and cost. This is a valuation and opportunity-cost failure, not necessarily a business failure. The immediate consequence is no add. If tax deferral is the chief reason for holding in two annual reviews, an independent challenger is required. Once the governed decision is made, action occurs within five trading days. The separation between review and execution matters: evidence may support "hold, not add" for a period, but once the decision is to act, "next liquid window" means five trading days absent a security-specific integrity event.

The tax comparison uses the applicable statutory rate, a fixed ten-year horizon, and the comparator replacement return. A proposed 90-day tax wait must be costed on day 0. More generally, every tax delay is priced at day 0, meaning the journal must recognize immediately the expected opportunity cost and risk of waiting rather than pretending the wait is free. Tax deferral can support retention only when its value plausibly exceeds the valuation or thesis shortfall. It cannot preserve a hard-gate failure, and it can never support an add to a position that no longer clears admission.

Consider another hypothetical. A company continues to generate the cash flows expected at purchase, but its price rises enough that the standardized conservative return is now below the comparator after tax and research cost. The thesis is intact, so the position is not treated as fraudulent or fundamentally broken. Adds stop. If selling would trigger a material tax, the annual tax-blind review evaluates the holding as if unowned and applies the standardized deferral comparison. A first documented hold may be justified by deferral. If the next annual review again relies chiefly on tax deferral, an independent challenger must test the case. After the authorized decision to sell or reduce, action is due within five trading days. If someone proposes waiting 90 days for a tax window, the cost of that wait must be recorded immediately. None of these rules authorizes the trade in this hypothetical.

Fourth, a mark-to-market drawdown alone never forces a sale. This final level is as binding as the faster rules above. It prevents a price decline from being smuggled into the exit hierarchy as though it were evidence of fraud, failed economics, or a broken milestone. It also prevents the inverse error: the absence of a drawdown does not excuse a genuine pillar break.

The valuation-drawdown budget is set before purchase

Because price volatility does not create an automatic exit, plausible derating loss must be budgeted before any purchase that may be authorized elsewhere. The file must archive both the current valuation multiple and a conservative multiple, together with the relevant reverse-DCF outputs. The position's derating loss is:

w x (1 - Mconservative / Mcurrent)

Here, w is the position weight as a fraction of the total portfolio, Mcurrent is the current valuation multiple, and Mconservative is the conservative-case multiple. The expression in parentheses is the proportionate valuation decline if the multiple moves from the current level to the conservative level while the selected operating measure is unchanged. Multiplying by w translates that security-level decline into a portfolio-level mark-to-market loss.

For a hypothetical calculation, assume an otherwise admissible position would be 2 percent of the portfolio, the current multiple is 30, and the conservative multiple is 20. The proportional derating is 1 - 20/30, or one-third. The portfolio derating loss is 0.02 x 1/3, approximately 0.0067, or 0.67 percent of total portfolio value. This is a planning estimate, not a forecast, target weight, or authorization.

The calculation is aggregated across all active holdings because several expensive quality companies can derate together. After the strategic sleeve's own stress is recognized, the sum of active derating losses must be no more than 50 percent of the portfolio's mark-to-market drawdown budget, abbreviated MDB. Suppose, only for illustration, that MDB were 8 percent of portfolio value after applying the relevant strategic stress. The default active derating allowance would then be 4 percent. The hypothetical 0.67 percent consumes part of that allowance alongside every other active holding's derating loss.

The allowance may rise as high as 75 percent of MDB only through a written amendment and independent challenge. In the illustrative 8 percent MDB case, that outer level would be 6 percent. It is not an analyst discretion, a normal range, or an after-the-fact excuse. The amendment must make the additional concentration of valuation risk explicit, and the independent challenger must test the assumptions. Above 75 percent is not admitted by this rule.

The budget is calculated at purchase because the purpose is to prevent oversizing before a drawdown occurs. If the sum later overshoots because correlations rise, multiples fall differently than expected, or the initial assumptions prove wrong, the consequence is a sizing-process post-mortem. It is not a mechanical trim of an intact thesis. The journal should examine why the conservative multiple, aggregation, or strategic-sleeve interaction failed and improve the future process. Turning the ex-post breach into an automatic sale would recreate the very volatility error the budget was designed to prevent.

This is not leniency. A post-mortem applies only to an ex-post valuation-budget overshoot when the underlying thesis remains intact. If A through D becomes zero, a hard cap persists, title is threatened, or another full-underwrite gate fails, the corresponding exit rule still governs. The valuation budget cannot shield a broken owner claim.

Research sunset is an attention rule

The policy applies falsification to themes as well as companies. A theme that produces neither a journaled decision nor an explicit abstention threshold within 36 months moves to annual-check status. The annual review must record the hours and external cost spent on that research line. This is the research sunset.

A journaled abstention counts as a decision only when it states the evidence or valuation threshold that would justify reconsideration. "Interesting, keep watching" is not an abstention threshold. A valid record might say that the line remains outside active work until independently verified customer cash returns exceed a stated hurdle, financing dependence falls below a stated condition, or a specified payer contract becomes binding. The threshold keeps the research line falsifiable and tells a future reviewer what new information would matter.

Moving to annual-check status does not delete the evidence or declare the theme false. The debate that produced this rule was explicit: sunset attention, not history. Archiving protects against relearning old lessons, preserves base rates and rejected hypotheses, and allows later evidence to be compared with the original record. What changes is the claim on scarce recurring attention. The line no longer receives live-underwrite cadence merely because its story remains culturally prominent.

This rule is designed for the plateau or muddle-through scenario in which dramatic themes generate little decision value and the low-cost comparator quietly wins after time and error. It also controls recency bias. Themes tied to AI, power, health, defense, and emerging-market growth can dominate the research agenda because they produce a constant flow of news. A 36-month sunset asks whether that flow produced an actionable, governed decision or a clear reason to abstain. If not, the research program must stop treating activity as value.

There is a real trade-off. Reducing attention can leave the investor less prepared for a later dislocation. Maintaining every line can preserve option value, but it can also consume hundreds of hours, encourage trading, and turn the policy into decision theater. Annual-check status is the compromise: retain the evidence and a low-frequency review, while reserving live work for lines that create demonstrable decision value. Revival should be based on new decision-relevant evidence, not renewed publicity. The authoritative final policy requires the annual-check move and hours/cost record; it does not force a purchase, a permanent rejection, or deletion of the archive.

Failure modes the discipline is meant to prevent

The first failure mode is headline monitoring. If everything is reviewed continuously, temporary commentary is mistaken for structural evidence. The monthly-survival, quarterly-breakage, and annual-reset clocks prevent that collapse of timescales.

The second is evidence staleness. A once-careful classification can look current long after filings, valuation, management, capital structure, regulation, or factor exposure have changed. The 12-month expiry makes staleness visible by reverting the label to Unclassified Research Note.

The third is narrative immunity. Without preregistered milestones and hard falsifiers, every miss can be described as a delay in a ten-year opportunity. Two missed milestones and the 36-month causal horizon force the investor to confront whether the promised mechanism is appearing.

The fourth is price-driven behavior. Selling solely because a quote fell converts volatility into permanent loss; averaging down solely because it fell converts price into unsupported evidence. The four-level exit hierarchy and the separate valuation-drawdown budget block both reactions.

The fifth is tax entrenchment. Low turnover and deferral are useful only when they improve after-tax outcomes. They become dangerous when they preserve a holding that would not otherwise survive comparison with the comparator. Standardized tax assumptions, the two-annual-hold challenger, the five-trading-day post-decision clock, and day-0 pricing of any delay keep tax in its proper role as friction rather than thesis.

The sixth is exception drift. A persistent limit breach cannot be renewed indefinitely as temporary. It must be cured by a compliant design or made explicit through a governed amendment. The distinction preserves decision rights: the analyst supplies evidence and records the breach; the owner-approved process decides whether policy itself changes.

The seventh is research theater. A rich map can create the appearance of edge even when no decisions improve and the hours exceed the possible value added. The 36-month sunset, annual cost record, 12-month map expiry, counter-canon and non-US screens, and active-program viability review turn attention into a resource that must earn its use.

No control eliminates judgment. Evidence can arrive late, conservative multiples can be wrong, a five-day implementation window can coincide with impaired liquidity, and an independent challenger can still share the same bias. The policy's answer is not omniscience. It is traceability: state in advance what would change the

conclusion, use the clock appropriate to the type of risk, record contrary evidence and cost, and preserve the boundary between research judgment and transaction authority.

Part V - Governance, Risks, and Certification

Chapter 13 - Risks, Unresolved Disputes, and Portfolio-Level Falsifiers

This policy is designed to remain useful when forecasts fail, but it can also fail as a governing system. That distinction is the purpose of this chapter. A company thesis can be wrong while the portfolio process remains sound. A theme can disappoint without invalidating liability matching, the Investor Liquidity Floor (ILF), or the named comparator. Conversely, every individual holding might appear defensible while the portfolio has become ungovernable because its common exposures are hidden, its evidence cannot be maintained, or its active program costs more than it can plausibly add.

The risks below therefore operate at three different levels. Company-level risks test whether one issuer can still meet its demand, industry-structure, capital-economics, governance, and valuation claims. Theme-level risks test whether a causal story such as artificial intelligence, defense, electrification, healthcare innovation, or emerging-market growth creates shareholder cash rather than merely social or customer benefit. Portfolio-level risks test whether the full allocation can fund liabilities, remain within common-factor limits, and be operated with reproducible evidence at tolerable cost. These levels interact, but they are not interchangeable.

This is also why a falsifier is not simply bad news. A falsifier is a preregistered observation that changes the status of a claim or process. It is meant to stop an investor from explaining away every adverse outcome after the fact. The policy preserves uncertainty where evidence cannot yet decide the issue and specifies what would change the view. None of the discussion authorizes a purchase, sale, hedge, account change, or target allocation. The evidence cutoff remains 22 August 2026, all personal calibration remains incomplete, and every named company in the research map remains at an admissible weight of zero until all policy gates pass through a separate authorized process.

Uncertified liability closure

The first risk is the most fundamental: the portfolio has not yet been shown to cover the investor's actual liabilities across the relevant paths. The illustrative grind test begins with investable assets, A^*, of 1.50 million and stressed real liabilities, PVL^*, of 1.10 million. It assumes equities fall 45%, matched nominal duration falls 25%, verified duration-matched inflation-linked bonds, or ILBs, produce 0% real in the modeled terminal calculation, and ILF cash loses 10% in real purchasing power. These are assumptions, not forecasts, and the test hides no withdrawal.

Under those assumptions, a 45% equity, 25% nominal-duration, 20% ILB, and 10% ILF mix ends with 1.0875 million. Coverage is 0.989, so it fails the stressed liability amount by 12,500 immediately. Any spending during the path extends the failure. A 40% equity, 20% nominal-duration, 30% ILB, and 10% ILF mix ends with 1.14 million and passes by 40,000, with coverage of 1.036. But that boundary result is valid only if roughly 30% of the portfolio can genuinely match real liabilities and the inflation index, law, title, tax treatment, currency, and market integrity all hold.

The risk is not merely that the first illustrative mix loses money. The deeper risk is false closure: treating a stylized terminal calculation as proof that the investor can pay a dated stream of obligations. A true liability test needs the timing, currency, indexation, flexibility, tax, and probability of the cash flows. A duration-matched ILB ladder held to maturity is not economically identical to an ILB fund that may have to be sold after real yields rise. A terminal 0% real assumption can coexist with a severe interim mark-to-market loss. If withdrawals force sale during that interval, an apparently matched asset can fail operationally.

The required observable is therefore a verified liability cash-flow schedule, with currencies, inflation index basis, legal title, tax treatment, and the intended implementation vehicle. The policy also needs the withdrawal path, because a grind is defined by time and spending sequence, not only by an endpoint. Evidence that would improve the view is a liability-by-liability matching analysis demonstrating that nominal obligations are matched by nominal duration and verified real obligations by suitable real duration, with enough ILF liquidity to avoid forced interim sale.

Evidence that would worsen the view includes index mismatch, uncertain title, unmodeled tax, a duration gap, a withdrawal schedule that exhausts liquid assets before matched securities mature, or liabilities whose currency differs from the supposed hedge.

Hypothetical example, not a forecast or allocation instruction: suppose an investor owes indexed payments in years four through ten, but owns a long-duration ILB fund and holds only two years of liquid spending. A rise in real yields could depress the fund before year four even if the bonds ultimately deliver the modeled real principal. If the investor must sell units in year three, the matching claim has failed in the only sense that matters: the liability could not be funded without realizing the path loss. The proper response is not to declare ILBs ineffective in general. It is to show that the particular vehicle, duration, currency, and liquidity structure did not match the particular liability.

Until this evidence exists, liability closure remains uncertified and is a certification blocker. It is also one of the explicit portfolio-level falsifiers: if liability closure remains unmeasured at the next annual review, new active work freezes. The fallback is the named comparator plus the ILF plus verified nominal and real liability matching. The fallback does not mean that the missing liability work can be replaced with generic bonds, a new macro hedge, or an improvised defensive asset.

A comparator born in breach

The named low-cost comparator is the burden-of-proof baseline, but it is not neutral. It contains issuers, sectors, currencies, duration, and common economic sensitivities. The diagnostic proxy of 60% MSCI ACWI and 40% liability-matched high-grade fixed income produced an artificial-intelligence factor shock loss, or FSL, of 8.2% and a joint AI-plus-power FSL of 9.0%. Those estimates exceed the original caps of 5% for AI and 8% for the joint exposure.

The proxy estimates are plausible only at an order-of-magnitude level. They are not reproducible enough to ratify a cap because the archived top-20 worksheet and the required independent second scoring pass do not yet exist. A scoring perturbation can determine whether the portfolio is classified as inside or outside a cap. The final policy therefore treats actual-comparator scoring as a certification blocker, retains the original caps until the study is complete, and bars active additions with material AI or power sensitivity while the original limits are exceeded. A temporary diagnostic ceiling may never exceed twice the original cap and must auto-revert after 12 months without ratification.

Calling this a provisional cap amendment does not remove the risk. It discloses acceptance of embedded exposure. Raising a limit changes the rule used to describe the portfolio; it does not diversify the portfolio or make a common shock smaller. The unresolved strategic dispute is whether to keep an uncapped market-weight comparator, accepting concentration and avoiding tracking-error intervention, or adopt a capped or blended comparator, accepting design risk, turnover, tax consequences, and the possibility of underweighting long-lived winners. The historical cases specified for the eventual study point in both directions: Japan in 1989 and Nortel in the Toronto market or Nokia in Helsinki warn against allowing a dominant constituent or cluster to define the benchmark, while the United States from 1992 through 2025, including 2013 through 2025, warns that capping sustained winners can create costly tracking error.

The dispute remains unresolved because there is no universally passive choice. Market weighting lets concentration evolve with price and market capitalization. Capping replaces that rule with a different active rule, whose rebalancing, tax, and abandonment behavior must be specified. The observable that can move the view is the actual investor comparator: archived top-20 holdings, sector buckets for the remainder, two independent scoring passes, sensitivity ranges using plus or minus 0.25 changes in factor scores, and an executable capped or blended alternative with costs, tracking-error budget, and abandonment point. Liability-matched duration must be excluded from directional rates exposure and reported separately with basis risk.

Hypothetical example: assume the actual comparator carries an 8% modeled AI shock loss against a 5% limit. Adding an individually attractive semiconductor name does not diversify that position merely because the security

has a different ticker. If it shares the same customer capex, power, export-control, and valuation-duration drivers, the active addition deepens the common loss channel. The company may pass its own scorecard, but the factor multiplier becomes zero when the addition would breach a hard cap. This is a portfolio-level veto, not a negative judgment about the company.

The active program may be uneconomic

Active research is not free, and the policy refuses to pretend that unpaid owner time has no opportunity cost. Active capital is enabled only when:

portfolio value x active sleeve x evidenced premium x evidenced hit rate >= 2 x fully loaded annual cost

The left side estimates the expected annual value added by the active program. Portfolio value is the capital base. Active sleeve is the proportion assigned to active decisions. Evidenced premium is the net return advantage the process can support rather than merely assert. Evidenced hit rate is the demonstrated probability that the decisions earn that premium. Fully loaded annual cost includes research hours valued at the operator's opportunity cost, data, tax work, professional advice, custody work, and other implementation expenses. The factor of two demands a margin of safety against optimistic assumptions and estimation error.

If the prospective threshold is not met, the program runs in paper-journal mode with zero active capital. Paper mode preserves a record and the ability to study dislocations without paying portfolio-level tuition for an uneconomic process. If the threshold is met and capital is deployed, two consecutive negative realized-viability years still suspend new active work. This prospective gate and retrospective stop address different problems: the first prevents an obviously subscale program from beginning, while the second tests whether a plausible program actually delivered enough value.

Hypothetical example: a 2 million portfolio proposes a 10% active sleeve, a 2% annual premium, and a 50% evidenced hit rate. The estimated contribution is 2,000,000 x 0.10 x 0.02 x 0.50 = 2,000 per year. To clear the 2x rule, fully loaded annual cost must be no more than 1,000. If the real cost is 4,000, active capital remains zero even if the owner enjoys research or believes the next idea could be exceptional. This is not a forecast about active management generally. It is the arithmetic of this specific program at this scale and evidence level.

The workload itself is a risk observable. The illustrative burden is about 40 recurring governance hours plus 15 to 25 hours per live underwrite. Eight live names therefore require about 160 to 240 hours a year before one-off comparator and custody work. Eight is a ceiling, not a target, and even eight may be too many. If the operator cannot precommit the hours, the live set must be reduced, possibly to zero, before research begins. The full 42-company map is a research map, not 42 live dashboards; classifications expire after 12 months without refresh and revert to Unclassified Research Note.

This dispute cannot be settled by a general claim that active management works or fails. One minority view favors zero active work because expected excess return is small and governance cost is real. The counterview values a maintained capability as an option during dislocations, sector busts, or harvest phases. The prospective threshold and paper mode structure the disagreement without pretending to eliminate it. Evidence that would support deployment includes audited, after-tax, decision-level attribution; a reproducible premium and hit rate; sufficient portfolio scale; and an hours budget that the operator actually maintained. Evidence that would favor zero active capital includes repeated negative viability, selective or unstable attribution, high abandonment, exception events, or a paper record that fails the comparator after all costs.

Concentration, quality crowding, and the danger of hidden sameness

The policy closes the main concentration dispute for current purposes by rejecting a generic conviction tier. At the highest funded state and a 50% bear-case permanent-impairment severity, the best-case ordinary position ceiling is about 2%. At the 40% ordinary-equity severity floor, it is 2.5%. A compelling story cannot waive the evidence, liquidity, factor, or severity arithmetic. The minority view remains that rare investors with genuine edge or control can rationally hold much larger positions and that excessive diversification may underuse scarce knowledge.

That minority remains visible because the cost of the adopted rule is real: a genuine edge could be economically underused. But the opposing cost is permanent loss from overestimating edge, plus correlated failure among apparently different high-quality companies. The company map itself is materially correlated with the 2010-2025 consensus-quality factor. Premium software, payments, market infrastructure, semiconductor bottlenecks, and branded compounders may differ operationally yet share duration, crowding, discount-rate, and multiple-compression exposure.

The reopening observable was corrected by the adopted F5-13 amendment. The controlling test is a preregistered five-year journal covering at least 20 fully underwritten decisions, independently audited after-tax attribution, verified liability surplus, challenger sign-off, and demonstrated operator capacity. Earlier wording requiring at least 10 years is retained only in provenance and is not a second live standard. Meeting the corrected evidence threshold permits reconsideration of the rejected concentration minority; it does not create a conviction tier, approve a holding, or authorize execution.

Quality-duration crowding remains unresolved even when every company thesis is intact. The relevant falsifier is not whether prices eventually recover. It is whether the sizing rules preserve liability coverage and expected owner cash through a no-multiple-recovery decade without breaching aggregate factor limits. The aggregate derating budget is therefore set at no more than 50% of the mark-to-market drawdown budget after strategic-sleeve stress; up to 75% requires a written amendment and independent challenge. An ex-post breach causes a sizing-process post-mortem rather than an automatic sale of an intact thesis.

Hypothetical example: eight active holdings are each limited to 2%, and no issuer appears large. Yet all eight depend on low discount rates, AI spending, data-center construction, and premium terminal multiples. The active sleeve can then behave like one 16% position during a factor shock. Issuer caps have not failed arithmetically, but diversification has failed economically. The governance response is look-through factor aggregation, not the addition of a ninth name with the same payer and capex cycle.

Evidence that would change the view includes performance through a full derating without thesis rewriting; stable or growing diluted cash per share while multiples do not recover; reliable factor scoring; and decision-level attribution showing value from security selection rather than exposure to a shared quality-duration tailwind. Evidence against the view includes losses clustering across supposedly unrelated holdings, repeated factor-cap exceptions, or an inability to explain returns without multiple expansion.

Procyclical diversifier evidence and the unresolved insurance question

Trend strategies, explicit tail protection, commodities, and other diversifiers can look effective or ineffective depending on the starting date. A backtest beginning just before a crisis may overstate benefit; one beginning after the crisis may emphasize carry drag. A single favored decade cannot decide the issue. The required observable is net full-cycle performance across rolling start dates spanning inflation, deflation, crash, and calm, including implementation cost, capacity, crisis correlation, tax, and the effect on liability coverage.

The tail-insurance dispute has a split resolution. For crash and withdrawal-sequence risk, lower gross risk plus the ILF is the null hypothesis. Convex protection must show that its premium, roll, basis, and implementation costs buy more useful protection than simply holding less risk. For a long inflationary grind, however, lower nominal risk is not a hedge if cash and nominal bonds lose real value. Verified inflation-linked liability matching is the structural default

for verified real liabilities when index, law, and market integrity pass. Trend, commodities, and tail-premium strategies remain conditional diversifiers rather than substitutes for matching.

The issue remains open because gross return is the wrong comparison. A diversifier can have a low standalone return and still improve goal probability through crisis liquidity, while an attractive backtest can fail when capacity, taxes, or correlations change. The proper comparison holds worst-case-loss reduction constant: does the diversifier improve coverage more than an equal reduction in gross risk? Evidence supporting admission would include rolling-start, full-cycle net results; stable crisis behavior; realistic capacity; and a demonstrable improvement in liability coverage. Evidence favoring rejection would include disappearing crisis correlation, excessive carry, fragile derivatives or counterparties, start-date dependence, or an inability to implement the historical signal.

Narrative and capital-expenditure risk

AI, defense, electrification, health innovation, climate adaptation, and emerging-market demand may all create large real-world value without creating acceptable minority-shareholder returns. Customers can capture lower prices and better service. Labor can capture scarce skills. States can impose price controls, tax windfalls, or demand local investment. Lenders and suppliers can claim cash before equity holders. Competitive entry can turn demand growth into a capacity flood. The existence of an addressable market, backlog, budget, or technical breakthrough therefore does not establish per-share value.

The unresolved AI-overbuild dispute is layer-specific. Equipment oligopolies, internal hyperscalers, and merchant cloud providers do not share the same contracts, financing, utilization, or residual value. The observables are utilization, rental or token prices, order cancellations, customer AI gross profit, capex relative to depreciation, and free cash flow after capex. A decline in rental prices while capacity expands can reveal overbuild earlier in merchant infrastructure; internal hyperscalers may hide vintage economics through shared workloads. Strong demand alone does not resolve whether the shareholder earns the return.

State-directed capex is similarly unresolved by theme. Announced appropriations are not collectible revenue, and allowed regulatory returns are not the same as realized cash returns. The payer must be tested for legal obligation, funding source, cancellation or renegotiation power, payment record, and the supplier's full-cycle real return. Evidence that would improve the view is binding, funded, collectible contracts or rate recovery that produces incremental after-tax cash per diluted share without offsetting dilution, debt, or state leakage. Evidence that would reverse it includes project cancellation, working-capital absorption, political disallowance, cost overruns, or shareholder cash that persistently trails reported earnings.

The entry-deterrence dispute asks whether heavy spending protects an existing franchise or merely expands management's empire. The claim remains unresolved because both can depress near-term return on invested capital and both can be described as strategic. Evidence of defense requires a named threat, a credible no-investment counterfactual, minimum efficient spend, a rival response that does not trigger endless escalation, and owner value from avoided erosion plus new cash exceeding economic depreciation and opportunity cost. Defensive credit remains capped until franchise outcomes and per-share capture corroborate the claim. Persisting rents after spending do not prove that the spending caused them.

Hypothetical example: a data platform doubles infrastructure spending while revenue rises 20%. Management calls the spending essential to defend its moat. If rivals also double spending, customer prices fall, and diluted free cash flow per share declines, industry growth may be real while shareholder capture is absent. The view would improve only if retention or pricing stabilizes relative to a credible low-spend comparison and the present value of protected profit exceeds the true economic cost.

Political, legal-title, custody, and currency risk

Export controls, reimbursement rules, tariffs, rate cases, sanctions, estate situs, and capital controls can dominate otherwise excellent business economics. These risks are discontinuous: a license can disappear, a security can become untransferable, a government can change the payer terms, or an investor can discover that legal title and beneficial access differ. Diversifying operating businesses does not necessarily diversify the legal system, currency, custodian, or sovereign that ultimately controls access.

The relevant evidence is binding law, enforceable contracts, title records, sub-custody chains, and actual transfer tests, not announced policy intentions. Tier A requires a same-jurisdiction second custodian. Tier B uses plain instruments in a second legal system and requires focused review of tax, reporting, title, transfer, sanctions, and estate situs. Tier C entities, trusts, and complex wrappers default to 0% until the full professional gauntlet passes. The structure of this custody response is resolved, but the appropriate Tier B size remains dependent on verified domicile, citizenship, tax, estate, and liability facts.

Long-horizon currency hedging also remains open within a 0-50% conditional range. Dated liabilities are matched first. Beyond them, the relevant questions are the actual liability currencies, carry, basis, tax, crisis covariance, and instrument risk. Currency volatility that does not impair a liability is not automatically a problem; a calm exchange rate that masks convertibility or capital-control risk is not automatically safe. Evidence supporting a hedge is a robust net benefit for the investor's actual currency basket across regimes after all costs. Evidence against it is expensive carry, unstable basis, tax leakage, or a hedge that fails when the liability is due.

Hypothetical example: an investor holds shares in globally diversified firms through one custodian and assumes operational diversification. A sanctions or access event at that custodian can affect every holding at once, independent of business quality. A verified transfer to the second custodian is stronger evidence than account-opening paperwork because it tests the path by which assets would actually move. An annual transfer failure triggers re-underwrite; unresolved title, situs, succession, or tax issues keep the relevant second-jurisdiction allocation at zero.

Obsolescence and dishonest economic lives

Rapid replacement can make reported profit look healthier than economic profit. Physical life asks whether an asset still operates. Revenue life asks whether it can still attract paying use. Economic life asks whether its vintage gross profit remains positive after power, support, maintenance, and opportunity cost relative to replacement. The shortest economically relevant life should govern the cash-flow model. Accounting depreciation is evidence, not a safe harbor.

This matters most when technology changes faster than reporting conventions. A one-year error in useful life can overstate gross margin, free cash flow, and incremental return on invested capital. The error becomes more dangerous when the asset is debt-funded, customer-concentrated, or dependent on falling rental prices. Conversely, assets with diversified workloads, software optimization, service revenue, or resale value may have genuine economic lives longer than a simplistic technical-obsolescence assumption. The dispute therefore cannot be settled by always using management's life or always using the shortest imaginable life.

The required evidence is a vintage bridge: beginning productive units by vintage plus additions, less retirements and impairments, equals ending productive units, with each vintage mapped to utilization, revenue, gross profit, power and maintenance cost, resale value, and software support. Quarterly rental prices and utilization provide early signals; annual retirement, impairment, resale, accounting-life changes, and vintage margins test the longer claim. Customer cash return is essential because a supplier's replacement cycle is durable only if customers can earn enough to keep buying.

This risk interacts directly with the comparator and narrative risks. If AI assets have shorter economic lives than reported, hyperscaler earnings, merchant-cloud collateral, equipment demand, power forecasts, and premium valuations can all weaken together. What looks like five independent theses can be one depreciation error

transmitted through a financing chain. Evidence that would improve the view is stable vintage utilization and margin, credible resale values, customer return sufficient to fund replacement, and cash returns that survive recast depreciation. Evidence that would change it negatively is accelerated retirement, falling rental prices, rising impairments, or capex that continually exceeds the cash generated by prior vintages.

Other unresolved allocation disputes

Some disputes are not direct stop rules but remain important boundaries on future policy changes. Private and illiquid assets retain only conditional funded-state bands. They must show manager- and vehicle-specific cash flows, governance, valuation integrity, and net return that beats a liquid alternative by 4 to 6 percentage points. Actual investor access remains unknown. The dispute persists because private structures may offer control, differentiated access, or an illiquidity premium, but can also conceal stale valuations, fee layers, capital calls, and governance weakness. Evidence must be specific to the vehicle; an asset-class average cannot admit an individual fund.

The bounded speculative sleeve also remains behaviorally unresolved. Its ranges are conditional, speculative positions assume total loss, and the normal per-name market-value cap is 0.25-0.75%. The question is whether a small option-like sleeve improves goal probability and preserves a useful dislocation capability, or instead creates an exception channel that grows when narratives become exciting. Total-loss simulations, exception frequency, time cost, and change in goal probability must be compared with simple abstention. No averaging down is permitted merely because price fell.

The possibility that the modal decade is a plateau or muddle-through regime also remains open. In that scenario, the dramatic themes generate research theater, repeated false starts, and overtrading while a low-cost comparator and disciplined capital return quietly win. The 36-month research sunset contains the cost of this possibility: a theme that produces neither a decision nor an explicit abstention threshold moves to annual-check status and records its hours. Evidence that would change the view is not a new narrative but measurable value from the research queue compared with the comparator and matching work.

The assurance-rent dispute asks whether AI strengthens regulated data, ratings, legal, scientific, and verification businesses or commoditizes their labor and information tolls. The observables are real price, labor hours, errors and liability, entry, retention, and the share of savings kept by clients. Higher volume with falling price and rising liability may destroy rents. Stable pricing, lower delivery cost, controlled error, and retained per-share cash would support the opposite view.

Finally, the policy keeps a research-priority dispute in view: comparator engineering may create more value than single-name selection because the strategic sleeve is normally 75-95% of assets while the active sleeve is smaller and tightly capped. Improving, or avoiding a mistake in, the comparator and liability-matching design by a small amount at large weight can outweigh a larger apparent edge in a small active sleeve. Annual attribution should therefore separate value from comparator and matching decisions from value created by company selection. This does not authorize a redesigned comparator. It identifies the evidence needed to decide where scarce research hours belong.

How risks combine into correlated failure

The principal danger is not any one row above but the way several can activate together. An AI capital-spending boom can increase index concentration, power demand, equipment orders, merchant-cloud leverage, and premium quality multiples at the same time. If customer monetization disappoints, rental prices fall and economic lives shorten. Lenders then tighten, capex is canceled, power orders weaken, and crowded quality multiples derate. A political response such as export restrictions can intensify the same shock. The portfolio may discover that its comparator, active names, power exposure, and currency exposure were all one trade.

A second chain begins with inflation or financial repression. Nominal bonds and cash lose real value, so an investor reaches for commodities, infrastructure, or regulated capex. But if liabilities are not measured, the new diversifiers may not match the spending index or currency. State price controls and rate disallowance can then reduce shareholder capture just as taxes, custody, or capital controls limit transfer. What appeared to be an inflation hedge can fail simultaneously on payer quality, legal title, and liquidity.

A third chain is operational. Too many live underwrites reduce evidence quality. Lower Evidence Coverage, or EC, makes factor scoring and valuation less reproducible. Exceptions then proliferate because the operator cannot complete the work on time. Research costs rise while attribution quality falls, causing the active program to miss its viability threshold. This is why operator capacity is a portfolio risk rather than an administrative inconvenience.

Governance must respond to these chains at the earliest common node. Common-factor look-through comes before adding issuers. Liability measurement comes before inventing hedges. Hours and evidence capacity are set before opening live underwrites. Binding legal and payer evidence comes before crediting budgets or market size. These responses are deliberately risk-monotone: they freeze new discretion or reduce scope. They do not convert an unresolved risk into an urgent trade.

Company-level and portfolio-level falsification

A company-level falsifier tests a specific causal thesis. Examples from the research map include a loss of retention, failure of backlog to convert to cash, falling per-query economics, adverse reserve development, a deterioration in credit vintages, or capex that fails to produce per-share owner earnings. These tests typically operate over a stated 36-month window. Two missed milestones or EC below 60 trigger a 30-day re-underwrite and no add. A legal-title, fraud, de-licensing, or financing-covenant threat receives faster treatment under the separate implementation process. A mark-to-market decline by itself is never a thesis falsifier.

A portfolio-level falsifier tests whether the active architecture can still be trusted and maintained. It does not require every company thesis to fail. One common-factor exposure, an unmeasured liability, or a governance breakdown can invalidate new active work even when individual dashboards look healthy. Likewise, one failed company does not automatically falsify the entire active program if it was correctly sized, documented, and contained within the loss budget.

This difference matters for remedies. A company-level failure normally leads to a specific re-underwrite, a no-add status, or a governed exit decision if authorized elsewhere. A portfolio-level failure freezes new active work and invokes the simple default. It does not demand improvised substitutions and does not, by itself, authorize indiscriminate liquidation of existing positions. Existing holdings still require the exit hierarchy, tax, liquidity, and integrity rules.

The portfolio falsifiers are explicit and exhaustive for this chapter:

1. Freeze new active work if EC audits cannot be maintained.
2. Freeze new active work after two annual active-viability results are negative.
3. Freeze new active work if two material Policy Exception Events occur within 12 months.
4. Freeze new active work if the actual comparator cannot be scored reproducibly within 16 operator hours per year.
5. Freeze new active work if liability closure remains unmeasured at the next annual review.

The 16-hour comparator limit is a test of minimum viable governance, not permission to use a careless proxy. If reproducible scoring requires an institutional apparatus the owner cannot sustain, the design has failed for this operator. Similarly, two exception events are not merely two isolated mistakes. They suggest the rules are either poorly calibrated, poorly understood, or being overridden when inconvenient. The freeze creates space for review before exception becomes habit.

In every one of these states, the stated default is the named comparator plus the ILF plus verified nominal and real liability matching. This is not a claim that the comparator is riskless. Its embedded risks remain subject to look-through and forced amendment. It is the lowest-complexity governed baseline available while active evidence is unavailable. Improvised substitutes are prohibited because a failure of governance is the worst moment to add an untested hedge, private vehicle, thematic fund, or concentrated idea.

What evidence would change the overall view

The policy becomes more confident only when evidence closes the relevant causal and operational loop. For liabilities, that means a dated, currency-specific, tax-aware cash-flow schedule and an executable matching plan. For the comparator, it means archived holdings, reproducible factor scores, independent review, sensitivity ranges, and a real capped or blended alternative. For active work, it means a prospective 2x economic case, maintained evidence coverage, audited attribution, and an hours budget the operator can sustain. For themes, it means incremental after-tax cash per diluted share after true maintenance and economic depreciation. For political and title risks, it means enforceable documents and successful transfer tests. For diversifiers, it means rolling-start, full-cycle, net evidence that improves coverage more than equal lower gross risk.

Evidence can also move the view toward greater simplicity. If the comparator cannot be scored within 16 hours, if the active program fails viability twice, if exceptions cluster, if liability closure remains missing, or if EC audits decay, the proper conclusion is not that the investor needs a more sophisticated model. It is that the active apparatus has failed its own burden of proof. Paper mode, a smaller live set, or zero active capital is an intended policy outcome, not an embarrassment.

Chapter 14 - Governance, Decision Rights, Maintenance, and Certification Readiness

This chapter makes the policy governable without turning governance into implied investment authority. It identifies who must maintain each record, who may challenge it, where approval must come from, what an executor may do, and what happens when any role or evidence item is missing. These are functional roles. The policy does not name the investor, an approver, an adviser, a broker, a custodian, a tax professional, a legal professional, a challenger, or an executor. A role shown here is therefore a required seat in a process, not evidence that the seat has been filled.

The current status is unambiguous. This is a final policy document, but it is not a certified investment plan. Liability closure, actual funded status, trait tolerance, the mark-to-market drawdown budget, tax basis, domicile and custody structure, the actual comparator, factor look-through, the withdrawal-path grind budget, company Evidence Coverage audits, current valuations, implementation approval, and execution authority remain incomplete or unverified. Every named company therefore retains a current admissible weight of zero. No trade, order, hedge, target allocation, entry price, transfer, account opening, external message, execution, or portfolio change is authorized.

The evidence cutoff is 22 August 2026. A source snapshot establishes only that a disclosure existed by that date. It does not verify a management claim, complete an independent underwrite, or prove that the fact remains current. This chapter preserves the policy's adopted rules and clocks; it does not upgrade the evidence, resolve an open dispute, fill an unnamed role, or create an implementation mandate.

Governance rests on an authority hierarchy

When records disagree, the controlling order is the concise final policy, adopted final amendments, integrated appendices and dispute dispositions, explanatory chapters, then working research files and worksheets. Debate rounds and provenance records explain how a rule arose but cannot override the final adopted text. Current evidence files can change an input or show that a gate passes or fails; they cannot silently change a policy rule. An approval can apply a policy to a proposed action; it cannot make missing evidence true. An execution instruction can carry out an approved action; it cannot supply the approval it lacks.

The distinction is important because the policy contains visible corrections. The final F5 amendments are already adopted, not optional commentary. F5-1 limits the 0.50 floor to the evidence-coverage multiplier, leaving liquidity and factor penalties fully effective and making a hard factor breach zero. F5-2 makes an E0 gate print a diagnostic maximum weight of zero. F5-13 replaces the earlier, impossible ten-year prerequisite for reopening the D-1 concentration question with a five-year record and at least 20 audited decisions. The accompanying conditions remain: decisions must be fully underwritten and preregistered, after-tax attribution must be independently audited, liability surplus must be verified, and a challenger must sign off. An older source row once said at least ten years. The current Appendix B and every operative rule use F5-13; the former wording survives only in provenance, never as a second live standard.

The authority hierarchy can be summarized as follows.

Level	Instrument	What it controls	What it cannot do
1	Final concise policy and adopted F5 amendments	Governing rules, thresholds, gates, clocks, prohibitions, and final dispute dispositions	Verify personal facts, approve an account action, or execute a transaction
2	Integrated appendices and current dispute register	Traceability, status of minority views, and adjudicating observables	Override an adopted correction or convert an observable into a finding without evidence
3	Revised and expanded explanatory chapters	Interpretation, workflow, examples, and navigation	Amend a rule, fill a missing decision-rights role, or authorize implementation
4	Current investor, comparator, custody, tax, legal, and research files	Evidence used to apply the policy	Rewrite the policy by implication or treat a source pointer as verified evidence

Level	Instrument	What it controls	What it cannot do
5	Decision packet and independent challenge	A documented proposed conclusion under the rules	Waive a hard gate, cure a missing fact, approve itself, or place an order
6	Separate implementation approval and execution mandate	Approval of a specific action and the authorized means of carrying it out	Backfill research, enlarge the approval, or treat a diagnostic ceiling as a target

Three events must never be collapsed into one. A new fact changes an input. An amendment changes a governing rule. An approval applies the current rule and current facts to a defined proposed action. Execution then implements only the approved action. Research precedes all four but possesses none of their authority.

Functional roles and decision rights

The minimum viable system uses roles rather than assumed personalities. One person may occupy more than one maintenance role if doing so is lawful, disclosed, operationally viable, and not prohibited by a required independence condition. The same person must not claim independence from their own work. Approval and execution must remain separately evidenced even if the external mandate lawfully assigns them to the same natural person. The decision packet must show when the person acted as approver and when the person acted as executor; an analytical file is never the approval record.

The functional roles are:

• **Policy owner.** Holds the purpose hierarchy and accepts or rejects formal policy amendments. The final policy does not name this person. Until an external record names the owner, no amendment or implementation may be assumed approved.

• **Governance operator.** Maintains the ILF check, look-through dashboard, decision journal, exception log, annual funded-status record, calendar, and active-program viability file. Maintenance does not grant investment authority.

• **Research lead.** Builds evidence files, scorecards, reverse DCFs, falsifiers, tax/cost comparisons, and decision packets. A completed underwrite must remain marked "non-buy until approved."

• **Independent challenger.** Tests evidence, assumptions, attribution, exceptions, and specifically required departures. Independence is a condition of the review, not a power to trade.

• **Investor-facts verifier.** Confirms assets, liabilities, taxes, currencies, income correlation, spending flexibility, decision makers, succession, and other personal inputs. The verifier may be the owner or a qualified delegate, but the source and date must be recorded.

• **Tax and legal reviewers.** Establish relevant tax, title, domicile, custody, transfer, sanctions, estate-situs, and account constraints. Their conclusions can be blocking; they are not security recommendations.

• **Implementation approver.** A person or body named outside this policy with lawful authority to approve a specific proposed action after all gates pass. No such holder is named here.

• **Executor.** The person or system separately authorized to transmit and carry out the approved instruction through the approved account and channel. Execution authority is absent until documented externally.

• **Continuity steward.** The person designated outside this policy to preserve access, records, and minimum risk controls during incapacity or absence. This role has no discretionary investment authority unless a separate lawful mandate expressly grants it.

• **Record custodian.** Preserves versioned policy, evidence, decisions, exceptions, approvals, execution receipts, and access instructions. Record custody is not asset custody and confers no authority over assets.

In the RACI tables below, **R** means responsible for preparing or performing the work, **A** means accountable for accepting that work or decision, **C** means consulted, and **I** means informed. "External mandate" means that the named authority must exist in a separate, current, legally effective record. A blank, unavailable, or unverified A seat is a stop condition, not permission for the R seat to self-promote.

Governance activity	Policy owner	Governance operator	Research lead	Independent challenger	Facts / professional reviewer	Implementation approver	Executor	Continuity steward
Maintain master policy and amendment ledger	A	R	C	C	I	I	I	I
Update investor facts and funded-status inputs	A	R	C	C	R	I	I	C
Maintain ILF, title-alert, look-through, journal, and exception records	A	R	C	C	C	I	I	C
Refresh actual comparator and factor map	A	R	R	C; independent second pass when due	C	I	I	I
Prepare a company underwrite or option record	I	C	R	C	C	I	I	I
Independently challenge a specified departure or conclusion	I	C	C	R/A for challenge completion only	C	I	I	I
Approve a formal policy amendment	A	R	C	C or required sign-off	C	I	I	I
Approve a specific implementation	I	C	R for packet	C	C	A	I	I
Translate approval into an executable instruction	I	I	I	I	C	A	R	I
Reconcile execution with approval and portfolio records	A	R	I	C	C	C	R	I
Activate incapacity continuity controls	I	R	I	C	C	A only if external mandate so provides	I or R only if separately authorized	R

The RACI assignment is deliberately asymmetric. A challenger is accountable only for the integrity of the challenge, never for the investment outcome or execution. A tax or legal reviewer is responsible only for the scoped professional issue. The implementation approver is accountable for a specific approval, not for rewriting the policy. The executor is responsible for faithful execution, not for selecting the security, changing the size, choosing a different account, or extending an expired approval.

The decision-rights boundary is even clearer when expressed by action.

Role	May prepare analysis	May reject or stop on its own evidence	May change policy	May approve a portfolio action	May execute
Governance operator	Yes	Yes, by recording a blocker or freeze	No; may propose only	No	No
Research lead	Yes	Yes, by failing a gate or recommending abstention	No; may propose only	No	No
Independent challenger	Yes, as challenge	Yes, by withholding required challenge sign-off	No; may recommend amendment	No	No

Role	May prepare analysis	May reject or stop on its own evidence	May change policy	May approve a portfolio action	May execute
Tax or legal reviewer	Yes, within scope	Yes, where a professional gate is unmet	No	No, unless separately appointed under an external mandate	No, unless separately authorized under an external mandate
Policy owner	Yes	Yes	Yes, through the formal amendment record	Only if separately named as implementation approver	Only if separately named as executor
Implementation approver	May review packet	Yes	No, unless separately acting through formal policy-owner capacity	Yes, only for the specific packet and mandate	No, unless separately named as executor
Executor	No investment analysis required	Must reject an instruction outside the approval	No	No	Yes, only within the verified approval
Continuity steward	May inspect continuity records	Must freeze unsupported discretion	No	Only if expressly granted by a current external mandate	Only if expressly authorized and within that mandate

An executor must be able to answer five questions before acting: What exact action was approved? Who held approval authority at the time? Which current decision packet and policy version supported it? Which account, instrument, size, price or order constraint, time window, and tax/legal conditions were approved? Has any event invalidated the approval? If any answer is missing, execution stops. The executor may reduce scope to comply with a more restrictive approval, but may not enlarge or reinterpret it.

Incapacity continuity is a survival control

Incapacity continuity stands alongside survival, legal title, and near-term essential liabilities, ahead of nominal return. A plan that only the usual operator can access or understand is not operationally viable. Death, incapacity, income loss, leverage, litigation, custody impairment, or inflexible spending can demote the funded-state band immediately. Promotion still requires two consecutive annual confirmations. Continuity therefore must support both access and rapid conservatism.

The policy does not create a power of attorney, trustee appointment, beneficiary designation, brokerage authority, or medical determination. Those must be supplied by valid external documents. The continuity record must identify, without exposing unnecessary private information in the policy itself, the current decision maker, successor or continuity steward, legal authority relied upon, activation method, relevant custodians, emergency contacts, document location, access recovery method, essential liability calendar, ILF location, and the accounts or instruments that must not be changed without specialized review.

On credible notice of death, incapacity, material absence, or loss of access, the continuity sequence is:

1. Preserve title, records, and access. Do not infer authority from possession of passwords or documents.
2. Verify the legal activation condition and the identity and scope of the authorized substitute decision maker.
3. Freeze new discretionary active work and any unsupported execution. Existing automated or contractual activity is reviewed, not casually interrupted.
4. Check the ILF and essential payments, known taxes, and fixed commitments. The emergency floor may never be treated as dislocation buying power.
5. Apply immediate funded-state demotion if the event impairs income, flexibility, access, title, or capacity. Do not wait for the next annual review.
6. Restrict any urgent response to the minimum risk-monotone action within the authority actually held. A title, fraud, de-licensing, or financing-covenant threat follows the 0-5 trading-day action clock if separately authorized.
7. If emergency exception authority is used, record it within 24 hours, obtain independent review within seven days, and let it expire within 30 days.

8. Inform the policy owner, record custodian, relevant professional reviewer, implementation approver, and executor according to the external mandate. Do not send an external instruction merely because this policy says those parties should be informed.

The annual continuity test must verify more than the existence of a successor's name. It must establish that the succession and incapacity documents are current, that the relevant person can locate the policy and liability schedule, that title and beneficiary records have been checked, that the Tier A same-jurisdiction second custodian exists where required, and that the annual transfer test has succeeded. Tier B exposure in a second legal system or plain instruments requires the tax, reporting, title, transfer, sanctions, and estate-situs subset review. Tier C entities, trusts, and complex wrappers remain at 0% until the full legal, tax, title, reporting, sanctions, transfer, and estate gauntlet passes.

The continuity steward's safe default is maintenance, payment continuity, evidence preservation, and risk reduction. The steward may not infer a mandate to rebalance, add active risk, exercise speculative discretion, change the comparator, amend a cap, or open an account. If the external mandate does not clearly allocate approval and execution rights during incapacity, the status is blocked and new discretion remains frozen.

The maintenance calendar

Governance should be event-sensitive without becoming a stream of alerts. The recurring calendar reserves fast attention for survival and integrity, quarterly attention for portfolio and exception drift, and annual attention for full-system renewal. A due date does not postpone an event response: if a hard gate breaks one day after the monthly review, the event-driven clock begins that day.

Monthly cycle

Monthly task	Responsible record	Completion evidence	Escalation
Verify the ILF remains funded, accessible, correctly denominated, short-duration, and above its hard floor and target	Governance operator	Dated ILF reconciliation, source balances, liability coverage note, and any variance	Any shortfall or access problem is escalated immediately; never refill risk from the emergency floor
Review custody, legal-title, fraud, de-licensing, and financing-covenant alerts	Governance operator with relevant reviewer	Dated alert log with source tier and disposition	Credible hard-gate threat freezes immediately and must be resolved within 0-5 trading days
Confirm essential payment and known-tax calendar for the next interval	Governance operator / continuity steward	Liability calendar and access confirmation	Missing access or authority activates continuity review and may demote FS immediately
Check for a material event that makes the comparator, factor map, or underwrite stale	Governance operator / research lead	Event log stating refresh required or no material event	Material event overrides annual cadence

The monthly cycle does not rescore every company, refresh every valuation, or react to price volatility alone. It is a survival and integrity control. Price decline without thesis evidence is not a hard-gate event and never authorizes averaging down.

Quarterly cycle

Quarterly task	Responsible record	Completion evidence	Escalation
Reconcile actual holdings, corporate actions, distributions, and tax lots	Governance operator	Account-to-dashboard reconciliation	Unexplained difference freezes affected decision work until reconciled
Refresh issuer, sovereign, currency, financing-chain, counterparty, and common-factor headroom	Governance operator / research lead	Look-through dashboard and headroom report	Active additions stop at a hard cap; a persistent active-holding breach enters the 5-30-day governed review
Review only evidence capable of breaking a pillar, gate, milestone, or EC threshold	Research lead	Evidence-delta log	Pillar 0, two missed milestones, or EC below 60 causes no add and full re-underwrite within 30 days
Review exception log, expiry, scope, and recurrence	Governance operator / challenger	Exception register with 24-hour, seven-day, and 30-day checks	Two material Policy Exception Events in 12 months freeze new discretion pending review
Count consecutive quarters of an embedded issuer or factor hard-cap breach	Governance operator	Breach-age field in comparator register	Four consecutive quarters force a written compliant comparator design or a cap raised by formal amendment; no renewable exception
Review tax-deferral-led holds and decision clocks	Research lead / tax reviewer	Standardized hold-versus-redeploy worksheet	Two annual holds chiefly justified by deferral require an independent challenger; a decided exit normally acts within five trading days

Quarterly review is selective. It asks whether a causal premise, evidence threshold, limit, tax premise, or exception status changed. It does not manufacture activity from routine news. Strategic sleeves use their approved mechanical rule when outside the approved band by more than 5 percentage points or more than 20% relative, subject to the ILF floor and taxes. A broken theme is not mechanically refilled, and the rule itself is not an authorization to trade.

Annual cycle

Annual task	Responsible record	Completion evidence	Consequence if missing or failed
Refresh assets, liabilities, taxes, currencies, income correlation, spending flexibility, decision makers, succession, domicile, and capacity	Facts verifier / governance operator	Dated verified-facts file	Actual FS and liability closure remain uncertified; new active work freezes if closure is still unmeasured at the annual review
Calculate FS = A*/PVL* at base PVL* and PVL* +/-10% and +/-20%; apply lower state if +/-10% changes the band	Governance operator	Sensitivity worksheet and band record	No promotion without two consecutive annual confirmations; verified deterioration demotes immediately
Reconfirm ILF target and hard floor, D36 or eligible D24 election, fixed commitments, real drag, title, beneficiary and incapacity access	Governance operator / continuity steward	ILF and continuity certificate of process status	Missing access, coverage, or authority is a survival blocker; D24 cannot be presumed
Refresh the actual comparator's top 20 individually and bucket the remainder by sector proxy	Governance operator / research lead	Archived holdings, scores, sensitivity, currency gap, and implementable capped/blended study	If not reproducible within 16 operator hours per year, freeze new active work and default in policy design to comparator + ILF + verified matching
Repeat the comparator scoring independently six months after the first pass in year one	Independent challenger	Second-pass worksheet and reconciliation	Until complete, feasibility caps are not ratified and original caps govern
Audit each live underwrite, scorecard, reverse DCF, falsifiers, EC, valuation, factor headroom, and tax/cost	Research lead / challenger	Current underwrite no older than 12 months	Classification expires to Unclassified Research Note after 12 months; no stale label carries forward
Run counter-canon and non-US screens and name at least three rejected candidates for any claimed model-class absence	Research lead	Screen and rejection log	Claimed absence remains unsupported; this does not force admission
Perform jurisdiction transfer test and custody review	Governance operator / professional reviewer	Successful transfer evidence and Tier A/B/C review	Unverified title or transferability blocks certification and relevant exposure
Run prospective active-program viability and print the hours budget	Governance operator	Viability worksheet and time budget	Below the 2x fully loaded cost threshold, active capital remains zero and the program stays in paper-journal mode
Attribute comparator and matching design separately from company selection	Governance operator / challenger	Decision-level after-tax/cost attribution	Two consecutive negative realized-viability years suspend new active work until re-justified
Audit ILF drag, exceptions, tax deferrals and challenges, diversifier rolling-start evidence, classifications, dispute observables, and amendments	Governance operator / challenger	Annual governance packet	Open or missing work remains open; it cannot be converted into a pass by omission
Rehearse incapacity continuity and verify external mandates	Continuity steward / policy owner	Dated access and authority test	Unclear authority freezes new discretion during incapacity

Recurring governance is estimated at roughly 40 hours per year before company work. Each live full underwrite requires about 15-25 hours per year. Eight live names therefore imply roughly 160-240 hours per year in total before one-off comparator, year-one independent scoring, and custody work. Eight is a maximum, not a quota. The live set must be reduced prospectively, possibly to zero, when the hours cannot be precommitted or the viability test fails.

Event-driven clocks

Trigger	Immediate status	Required clock	Required evidence and governed consequence
Legal-title, evidenced fraud, de-licensing, or financing-covenant threat	Freeze affected discretion immediately	Resolve and take minimum risk-monotone action within 0-5 trading days if separately authorized	Primary evidence, authority check, action rationale, and execution receipt; emergency documentation clocks also apply when an exception is used
Pillar A-D becomes 0; hard cap persists; full-underwrite gate fails	No add; governed close/reduce review	5-30 days	Updated pillar/gate evidence, tax/liquidity/legal effects, challenge where required, and separate implementation decision
Two missed milestones or EC below 60	No add; fully underwritten standing suspended	Full re-underwrite within 30 days	Rebuilt evidence set, updated thesis, scorecard, valuation, and disposition; restoration is not automatic
Valuation E falls below admission while thesis remains intact	No add; no mechanical trim	Re-underwrite on decision clock	Standardized after-tax/cost comparison with comparator; two deferral-led annual holds require challenge
Decision to exit for intact-thesis expected-return or tax/cost reason	Await separate approval	Act within five trading days after the approved decision, absent security-specific integrity event	Approved comparison, applicable tax calculation, instruction, and execution reconciliation
Proposed 90-day tax wait	Decision remains governed	Cost of the wait calculated at day 0	Applicable statutory rate, fixed ten-year horizon, comparator replacement return, and stated opportunity cost
Strategic sleeve outside approved band by >5 percentage points or >20% relative	Rebalance review, subject to ILF and tax	Per approved mechanical process	Current band, measurement, tax and ILF check; do not mechanically refill a broken theme
Proposed entry pacing	No action until separately approved	Tranches tied to filings or preregistered milestones and expire by day 180	Event anchors and expected opportunity cost; otherwise full deployment is only a statistical default after all gates pass, not authorization
Emergency Policy Exception Event	Risk-monotone scope only	Record within 24 hours; independent review within seven days; expire within 30 days	Event, necessity, authority, scope, affected rule, risk reduction, review, and closure; no chain of renewals
Second material Policy Exception Event	Freeze new discretion	On occurrence within any rolling 12 months	Full governance review; exception handling may not become the operating model
Embedded comparator issuer or factor breach	Additions that worsen breach stop	Four consecutive quarters force written comparator redesign or formal cap amendment	Top-20/bucket evidence, feasibility study, currency risk/feasibility gap, and explicit accepted risk
Material comparator or company event	Stale affected record	Refresh when event is recognized, not at next annual date	Dated event analysis and new status
Death, incapacity, income loss, leverage, litigation, custody impairment, or inflexible spending	Possible immediate FS demotion and continuity activation	Immediate; do not await annual review	Verified event, access and authority record, revised capacity and band, essential-payment protection

No event-driven clock is a standing order. It is a deadline for a governed decision and, only where an external mandate supplies authority, a deadline for minimum authorized action.

Shelf lives and the 12-, 24-, 36-month, and five-year horizons

Maintainability requires distinguishing durable rules from dated inputs. Annual review does not mean every durable principle expires annually, and a long policy horizon does not mean every company classification lasts ten years. Each layer has its own shelf life.

Layer or item	Status type	Shelf life or horizon	What happens at expiry or review
Purpose hierarchy, epistemic labels, three risk budgets, liability-first measurement, comparator burden, Shortfall Parity, evidence-to-loss, no-recourse-leverage default, look-through, fast-broad/slow-active, tax-as-friction, amend-not-excuse, minimum viable governance, and research non-authorization	Durable policy	No automatic expiry; reviewed annually and changed only by explicit amendment	Remains controlling unless formally amended; market discomfort or repeated exceptions do not change it
Numerical FS bands, loss budgets, severity floors, ILF formula, multipliers, caps, and viability thresholds	Calibrated or provisional policy	Reviewed annually and on relevant structural event	Await verified investor facts where stated; do not drift with markets; change only through compliant design or formal amendment
Evidence synthesis and source snapshot	Dated	Cut off at 22 August 2026	Later facts require a current review and must not be backfilled as if known at cutoff
Company research classification	Living research	12 months from refresh	Reverts to Unclassified Research Note; expiry is not a sell instruction and does not imply deterioration
Policy exception recurrence	Governance window	Rolling 12 months	Two material events freeze new discretion pending review
Actual comparator	Living strategic record	Annual and on material event	Refresh top 20 and remainder buckets; independent repeat after six months in year one
Tax-deferral-led hold	Living implementation analysis	Annual	Two annual holds chiefly justified by deferral require independent challenge
Entry pacing	Decision-specific	Maximum 180 days	Unused pacing authority expires; milestones cannot be rolled into calendar averaging
Harvest-candidate status	Dated option research	24 months	Expires without verified exit and shareholder cash-return evidence; cannot roll indefinitely
D24 liquidity election	Annual conditional election	One annual cycle	Requires FS at least 1.75 and documented high spending flexibility; otherwise D36 governs
D36 ILF component	Liability-matching horizon	36 months of essential liabilities in the target formula, subject to the formula's non-overlap	Remains part of the annual ILF calculation unless a valid D24 election applies
Company or theme milestones	Dated falsification record	Normally up to 36 months as preregistered	Two missed primary milestones trigger no add and a 30-day full re-underwrite
Theme without a decision or explicit abstention threshold	Attention allocation	36 months	Moves to annual-check status and records hours and cost; it is not declared false
Concentration minority reopening	Long-horizon audit	Five years and at least 20 fully underwritten decisions	May be reconsidered only with preregistration, independent after-tax attribution audit, verified liability surplus, challenger sign-off, and operator capacity; reopening is not approval
Tax-deferral valuation	Standardized analytical horizon	Fixed ten years	Uses applicable statutory rate and comparator replacement return; this ten-year clock does not revive the superseded D-1 rule
Policy worldview	Dated decade frame	2026-2036	Must be revisited as the horizon matures; scenarios remain conditional beliefs, not durable constitutional facts
Open dispute	Unresolved or conditional	Until stated observable adjudicates it or a formal amendment changes status	Remains visible; passage of time and consensus do not resolve it

The 12-month horizon primarily governs freshness and exception recurrence. The 24-month horizon governs the harvest option and, through an annually renewed election, the shorter D24 liquidity base available only at FS at least 1.75 with high spending flexibility. The 36-month horizon governs the default D36 liquidity component and the point at which a theme without a decision or explicit abstention threshold loses recurring attention. The five-year

horizon governs only the possibility of reopening the rejected concentration minority, and it also requires at least 20 audited decisions and all accompanying safeguards. None of these clocks authorizes ownership.

Amendment, compliant design, approval, exception, and execution flow

A maintainable policy needs one route for permanent rule change and another for a temporary emergency. It also needs a clean separation between changing the rule and applying the rule. The following sequence is controlling.

Ordinary evidence and decision flow

1. The governance operator or research lead records the new fact, inference, assumption, forecast, classification, or unresolved item with source, date, and affected record.
2. The responsible role updates the applicable investor-facts file, comparator map, underwrite, liability schedule, tax/legal memorandum, or dispute record. This changes an input, not the policy.
3. The decision packet is tested against all non-compensating gates. Missing evidence remains missing; a favorable result in one gate cannot cure another.
4. An independent challenger reviews wherever the policy requires independence or where the policy owner requests it. Challenge completion satisfies only the challenge condition.
5. The packet states one of the governed outcomes: blocked, research-only, paper mode, no add, re-underwrite, freeze, suspension, abstention, eligible for separate approval, or compliant strategic design. A diagnostic maximum weight remains a ceiling, never a target.
6. If the outcome is eligible for action, the packet moves to the implementation approver named by an external mandate. The approver checks authority, policy version, evidence currency, portfolio look-through, tax/legal conditions, account scope, and whether any event invalidated the packet.
7. Only a specific approval moves to the executor. The instruction must match the approved instrument, account, size or maximum, price/order conditions, timing, and contingencies.
8. The governance operator reconciles the execution receipt against the approval and updates holdings, tax lots, headroom, decision journal, and monitoring clocks.

Formal amendment flow

A formal amendment is required when a governing rule is to change rather than a design being made compliant. The amendment record must identify the exact rule and policy version; the reason for change; whether the driver is an objective, constraint, verified durable fact, structural change, capability failure, or implementation defect; the evidence and source date; the strongest countercase and any surviving minority; effects on liabilities, liquidity, title, tax, currency, concentration, behavior, workload, and implementation; every numerical threshold and clock affected; the language superseded; the effective date; the policy owner's approval; any required independent challenge; and a plan for post-amendment review.

Two policy examples have special conditions. First, active derating loss is limited to 50% of the mark-to-market drawdown budget after strategic-sleeve stress. Allowance above 50% and up to 75% requires a written amendment and independent challenge. The policy provides no inherited informal route beyond 75%. An ex-post overshoot produces a sizing post-mortem, not an automatic sale of an intact thesis. Second, four consecutive quarters of embedded issuer or factor hard-cap breach force a written choice between a compliant capped or blended comparator and a cap raised through amendment. Raising the cap explicitly accepts more risk; it does not make the exposure disappear or ratify the diagnostic proxy.

The D-1 concentration minority has a different kind of amendment gate. It is closed for the current policy. Only after five years and at least 20 fully underwritten, preregistered decisions, independently audited after-tax attribution, verified liability surplus, challenger sign-off, and demonstrated operator capacity may the question be reopened.

Meeting that evidence threshold permits reconsideration; it does not create a conviction tier, amend the policy automatically, approve a holding, or authorize execution.

Compliant design flow

A compliant design change leaves the policy rule intact and changes the separately approved portfolio or process so the rule is met. Examples include using a capped or blended comparator, reducing active scope to fit the hours budget, retiring a stale classification, or reducing a factor exposure through an authorized implementation. The record must show the original breach, the unchanged rule, the proposed compliant design, tax/cost/tracking-error and abandonment risks, independent evidence where required, implementation approval, and the post-execution measurement. A design cannot be called compliant while relying on a placeholder comparator, missing look-through, or an expired exception.

Emergency exception flow

An emergency exception is narrow, temporary, and risk-monotone. It must reduce or not increase the relevant risk, preserve liquidity, protect title, or restore compliance. It may not increase gross exposure, add illiquidity, open a new discretionary position, buy a perceived bargain, or waive unaffected rules. The event and action are recorded within 24 hours; necessity, conflicts, scope, evidence, and authority are independently reviewed within seven days; and the exceptional authority expires within 30 days. Continued need must move to a formal amendment or compliant design. A new 30-day label does not renew the exception.

Two material Policy Exception Events in any rolling 12 months freeze new discretion pending governance review. The freeze applies even when each event appeared individually defensible, because recurrence is evidence that the written policy or operating capability may be misaligned with reality.

Approval and execution separation

Implementation approval must be a dated, specific record outside the research output. At minimum it identifies the authorized approver and source of authority, the current policy and decision-packet versions, the proposed action, account and instrument, maximum scope, tax/legal conditions, validity window, contingencies, and any prohibited substitutions. Execution must be a separate record linking to that approval and showing the actual order or action, timestamp, account, instrument, quantity, price or order terms, channel, outcome, and reconciliation.

If the same person lawfully holds policy-owner, approver, or executor roles, the records must still be sequential and role-specific. The person cannot write an underwrite, call it approved by implication, and transmit an order. An approval must exist before execution, and execution may not exceed or creatively interpret it. Any material change in evidence, valuation, factor headroom, ILF, tax/legal status, title, or investor capacity before execution sends the packet back for review.

Current certification-readiness and blocker register

The readiness statuses in this register are **blocked**, **open**, **conditional**, **expired**, or **process-ready**. "Process-ready" means only that the policy describes a usable process; it does not certify the investor, portfolio, security, or transaction. A blocker is not cured by assigning an optimistic owner. Required evidence must be attached, current, and independently reviewed where specified.

Readiness item	Current status at 22 August 2026	Consequence now	Functional owner	Required evidence to change status
Policy document and adopted amendments	Process-ready as a final document	Governing text may be maintained, but finality is not investment certification	Policy owner / record custodian	Version-controlled policy, adopted F5 ledger, hierarchy note, and preserved supersessions
Implementation approver	Blocked / unnamed	No portfolio action can be approved under this document	External mandate / policy owner to identify	Current lawful mandate naming approver, scope, accounts, limits, succession, and validity

Readiness item	Current status at 22 August 2026	Consequence now	Functional owner	Required evidence to change status
Executor and execution channel	Blocked / unnamed	No order, hedge, transfer, account action, or external instruction may be sent	External mandate	Current execution authority, approved account/channel, controls, and reconciliation process
Approval-execution separation	Blocked until evidenced	Research or maintenance cannot become action	Implementation approver / executor	Separate dated approval and execution templates, role identities, authority source, and test of rejection controls
Verified investor facts	Blocked	All numeric bands remain conditional	Investor-facts verifier	Assets after tax/cost, liability schedule, taxes, currencies, income correlation, spending flexibility, leverage, decision makers, succession, domicile, and source dates
Liability closure	Blocked / uncertified	Strategic design and overall certification cannot close	Governance operator / facts verifier	Verified cash-flow schedule, probability weights, safety margin, currency and index basis, title, tax, timing, and Coverage
Actual funded status and hysteresis	Blocked	No live FS band, promotion, or sleeve range may be treated as personal	Governance operator	A*/PVL* at base and +/-10% and +/-20%, lower-state rule, second-axis assessment; two annual confirmations for promotion
Trait tolerance and capacity	Blocked / unverified	Risk cannot be enlarged on assumed tolerance	Policy owner / facts verifier	Calm-state adherence assessment, capacity record, essential liability flexibility, and income-market correlation
ILF target, floor, and D36/D24 election	Blocked pending personal facts	Emergency liquidity and liability matching are not certified	Governance operator	D36, Tax12, Commit, H and dependable-income inputs; hard-floor proof; FS >=1.75 and high-flexibility finding for D24 election
Withdrawal-path and decade-long grind budget	Blocked / open	Stress snapshots cannot certify slow-failure survival	Governance operator	Scenario paths with withdrawals, liability inflation, taxes, currencies, sequence, duration, and Coverage; verified real matching
Mark-to-market drawdown budget	Blocked / unverified	Derating allowance and active risk capacity cannot be finalized	Policy owner / governance operator	Current MDB based on capacity and trait tolerance, strategic-sleeve stress, and aggregate derating ledger
Tax basis and account constraints	Blocked / unverified	After-tax comparisons and implementation remain conditional	Tax reviewer	Actual basis, lots, statutory rates, domicile/account rules, withholding, estate and transaction consequences
Domicile, title, custody, and transferability	Blocked / unverified	Relevant accounts or structures cannot be certified	Legal/custody reviewer	Ownership and segregation evidence, Tier A second custodian, annual transfer test, Tier B subset review, Tier C full gauntlet where relevant
Actual named comparator	Blocked / missing	Opportunity cost and strategic design are incomplete; no placeholder after 12 months	Policy owner / governance operator	Named low-cost implementable comparator, tax basis, holdings, operational feasibility, and approval
Comparator top-20 and sector-bucket scoring	Blocked / incomplete	Factor multipliers and feasibility caps are not ratified; original caps govern	Governance operator / research lead	Archived top 20, bucket remainder, sensitivity scores, currency gap, capped/blended study, and independent six-month year-one pass
Factor and issuer headroom	Blocked / provisional	Active additions to affected AI/power rows remain zero; a hard breach means Fmult = 0	Governance operator	Actual portfolio plus comparator look-through, FEW/FSL, joint stresses, risk caps, and reproducible dashboard
Company EC audits	Blocked / incomplete	Every named company remains research-only and current admissible weight zero	Research lead / challenger	Verified required fields, evidence tiers, conflict resolution, EC calculation, scorecard, falsifiers, and current date

Readiness item	Current status at 22 August 2026	Consequence now	Functional owner	Required evidence to change status
Current company valuations and reverse DCFs	Blocked / incomplete	No security clears admission; quality cannot substitute for price	Research lead	Current price and rights, conservative reverse DCF, comparator hurdle, tax/cost, severity, liquidity, factor headroom, and derating budget
Active-program prospective viability	Blocked until annual worksheet; paper mode is default	Active capital must remain zero unless expected contribution is at least 2x fully loaded annual cost	Governance operator / policy owner	Portfolio value, active sleeve, evidenced premium, evidenced hit rate, operator opportunity cost, data/tax costs, and hours budget
Operator capacity and live-underwrite count	Conditional / unresolved	Maximum is eight, but viable count may be lower or zero	Governance operator	Precommitted roughly 40 governance hours plus 15-25 hours per live name, one-off work allowance, and maintenance evidence
Active-program realized viability	Not yet established	If deployed later, two consecutive negative years suspend new active work	Governance operator / challenger	Decision-level independently reviewable after-tax/cost attribution separated between comparator design and company selection
Company-map freshness	Dated	Each classification expires 12 months after its own refresh and reverts to Unclassified Research Note	Research lead / record custodian	Dated current underwrite and annual counter-canon/non-US/absence screens
Open disputes	Open or conditional as recorded	No dispute may be presented as settled without its observable or amendment	Policy owner / challenger	Stated adjudicating evidence, dated disposition, surviving minority, and formal amendment if the rule changes
Incapacity continuity	Blocked until external authority and access are verified	New discretion freezes on incapacity; only lawful minimum continuity actions may proceed	Continuity steward / policy owner	Successor authority, activation test, access map, liability and ILF record, custodian contacts, annual rehearsal, and transfer test
Certification as an investment plan	Blocked	No security, portfolio, or transaction is certified; no company is a buy	Policy owner / implementation approver	Joint completion of liability closure, withdrawal/grind path, actual comparator scoring, company EC audits, all personal and professional facts, all gates, and separate implementation approval

The blockers are non-compensating. A completed valuation cannot cure unmeasured liabilities. Verified liabilities cannot cure an actual comparator born in breach. A strong EC audit cannot cure a failed valuation gate. Independent challenge cannot cure a hard factor breach. An approver cannot make a management claim independent, and an executor cannot infer approval from a completed checklist. Certification requires joint completion, not a majority vote among workstreams.

The published grind example demonstrates the boundary but does not close it. Under the example assumptions, the 45% equity / 25% nominal duration / 20% inflation-linked bond / 10% ILF mix fails immediately by 12,500 with Coverage 0.989. The 40/20/30/10 diagnostic passes by 40,000 with Coverage 1.036 only if the 30% real-liability match is genuine. Withdrawals were not supplied, so neither path certifies a real investor or the duration of a slow failure.

The actual comparator is also unresolved. The diagnostic proxy's AI factor stress loss of 8.2% and joint AI-plus-power stress loss of 9.0% exceed the original caps. That diagnosis discloses risk; it does not ratify a cap. The actual top-20 worksheet, remainder buckets, independent pass, taxes, and capped or blended implementation study remain required. Until then, original caps govern and affected active additions remain zero.

Readiness outputs and records

The governance system should produce a small number of durable records, each with an owner, date, status, evidence links, next due date, and retention history. These are not additional gates; they are the evidence that existing gates were maintained.

Record	Minimum contents	Retention and status rule
Master policy and amendment ledger	Version, effective date, adopted amendments, superseded text, unresolved disputes, and owner approval	Permanent policy history; never overwrite a superseded rule without preserving traceability
Investor-facts and funded-status file	Verified assets, liabilities, taxes, currencies, income correlation, flexibility, succession, sensitivities, band, and second axis	Annual snapshots plus event-driven updates; preserve prior bands and demotions
ILF and continuity file	Formula inputs, floor, target, D36/D24 status, custody/access, beneficiaries, successor authority, and rehearsal	Monthly ILF evidence; annual continuity evidence; event records retained
Comparator and look-through workbook	Named comparator, top 20, sector buckets, issuer/factor/currency stress, headroom, feasibility gap, and independent pass	Annual and material-event versions; year-one second pass after six months
Research underwrite	Evidence tiers, A-E, EC, reverse DCF, severity, loss arithmetic, derating budget, falsifiers, tax/cost, and status	Expires as a classification after 12 months without refresh; retain as historical research
Decision journal	Timestamped facts and labels, thesis, countertheme, alternatives, disconfirmers, action clock, proposed status, and "non-buy until approved"	Permanent decision-level audit record
Exception log	Event, affected rule, authority, risk-monotone rationale, 24-hour record, seven-day review, 30-day expiry, and recurrence count	Retain permanently; rolling 12-month recurrence field remains live
Approval and execution register	External authority, specific approval, expiry, instruction, receipt, reconciliation, and deviations	Permanent; approval and execution records remain separate and linked
Annual governance packet	Calendar completion, viability, hours, attribution, ILF drag, tax deferrals, disputes, map decay, screens, custody test, amendments, and blocker register	Annual immutable snapshot; describes process status, not investment certification

The record custodian should preserve human-readable evidence and reproducible calculations. A hash or provenance receipt can show that a file existed in a particular form, but it is not investment evidence. Likewise, a source link can show where a disclosure was found but grants no Evidence Coverage credit until the relevant claim is extracted, assessed, and reconciled.

Failure states and governed defaults

Governance fails when authority is assumed, evidence is stale, exceptions renew themselves, or workload exceeds capacity. Each failure has a narrower default rather than a relaxed standard.

If EC audits cannot be maintained, new active work freezes. If the prospective active-program value does not reach 2 times fully loaded annual cost, capital remains at zero and work stays in paper-journal mode. If a deployed program produces two consecutive negative realized-viability years, new active work is suspended until re-justified. If two material Policy Exception Events occur within 12 months, new discretion freezes. If the actual comparator cannot be scored reproducibly within 16 operator hours per year, or if liability closure remains unmeasured at the next annual review, new active work freezes. The policy design then contracts to the named comparator, ILF, and verified nominal and real liability matching, without improvising substitutes and without treating that default as transaction authority.

If a company file is not refreshed within 12 months, its classification expires. If a harvest candidate lacks confirmed exit and shareholder cash-return evidence by 24 months, candidate status expires. If a theme has produced neither a decision nor an explicit abstention threshold by 36 months, it moves to annual-check status and its hours and cost are recorded. If the one-person process cannot support eight live underwrites, the count is reduced before research begins. Scope yields before evidence standards do.

If approval authority is missing, the result is non-authorization. If execution authority is missing, the result is no execution. If incapacity authority is unclear, new discretion freezes while title, access, ILF, and essential payments are protected through whatever lawful authority actually exists. If a professional tax, legal, title, custody, sanctions, transfer, or estate gate is relevant and incomplete, the affected conclusion remains blocked.

The same discipline applies to disagreement. D-2 comparator design, D-4 private and illiquid assets, D-5 currency hedging, AI overbuild, accelerator lives, state-capex cash capture, quality-duration crowding, diversifier timing, zero-active investing, and comparator-engineering-first remain open or conditional as recorded. Their observables may change a future conclusion. Until then, governance preserves the uncertainty rather than converting the preferred narrative into policy.

Annual certification-readiness statement

The annual governance packet should end with a short statement containing five distinct statuses:

1. **Policy maintenance status:** whether the master policy, amendment ledger, calendar, exceptions, and dispute register are current.
2. **Investor-fact and liability status:** whether facts, FS sensitivities, capacity, ILF, continuity, and the withdrawal/grind path are verified.
3. **Strategic architecture status:** whether the actual comparator, factor look-through, custody, nominal and real matching, and limits are current and implementable.
4. **Active-program status:** paper mode, eligible for separate approval, frozen, suspended, or inactive, together with the 2x viability result, hours, underwrite count, and realized attribution where applicable.
5. **Implementation status:** whether an externally authorized approver and executor exist for any specific proposed action, with approval and execution still separated.

Only the first status can presently be described as capable of being process-ready. The remaining statuses are blocked, incomplete, conditional, or unverified on the evidence date. The annual statement must therefore repeat the practical boundary: final as a document is not certified as an investment plan; no company is a buy; every named company's current admissible weight is zero; and no transaction is authorized.

The policy is maintainable only if responsibility is explicit

The governance design is complete only when every required fact, review, challenge, approval, exception, and execution has a named owner and an auditable record. Until then, the document remains research and governance guidance rather than an implementable portfolio plan.

Chapter 15 - Certification, Ownership, and the Policy Boundary

Final as a document, not certified as an investment plan

Chapter 15 draws the last and most important boundary in the policy. The policy is final in the limited sense that the five rounds of Codex positions, the five adversarial Fable responses, and the adopted final amendments have been assembled into one debate-integrated edition dated 22 August 2026. That editorial finality does not turn the document into a certificate of financial readiness, a completed portfolio design, a recommendation, or proof that any proposed investment is suitable. The exact certification language is:

> *This is the final debate-integrated policy edition, not a certificate. Liability closure, actual FS, trait tolerance, MDB, tax basis, domicile/custody structure, actual comparator, factor look-through and every company valuation remain unverified. Numeric bands are dated policy hypotheses awaiting investor facts. The source snapshots establish only that a disclosure existed as of the evidence date; they do not verify management claims or complete a full independent evidence-and-valuation underwrite.* **No company is a buy; no trade or portfolio change is authorized.**

Every part of that paragraph matters. "Final debate-integrated policy edition" describes the status of the text. "Not a certificate" describes the status of the investor, portfolio, and research. The two statements are not in tension. A policy can be complete as a record of adopted principles and still be unable to certify that the unknown investor facts satisfy those principles. It can also classify companies for research without concluding that any security is investable now.

The policy therefore certifies almost nothing about implementation. What it does establish is documentary. It records which governance principles survived the debate, which final Fable amendments were adopted, which disputes remain open, which numerical rules are conditional, which research classifications were made on a first-pass basis, and where trading authority stops. It also establishes that the edition has an evidence date and that later facts were not part of the synthesis. These are statements about the document's contents and provenance, not assurances about future returns or personal suitability.

The opening language of the full policy reinforces the same distinction. The instrument is a policy and research map, not a portfolio or recommendation. It authorizes no trade, order, hedge, message, account opening, or portfolio change. No theme, class, company, ceiling, or scenario is an instruction to trade. A calculation can be correct, a company can appear in a favorable research category, and a policy ceiling can be above zero, yet the current admissible weight still remains zero. The missing approvals and evidence are not administrative loose ends. They are part of the substantive boundary.

What the edition establishes

The edition establishes a governance architecture for reasoning under uncertainty. It sets the purpose hierarchy, defines the funded-state and loss-budget concepts, distinguishes liability matching from directional risk, imposes evidence and valuation gates, requires look-through analysis, and records how disputes, exceptions, and amendments are treated. The document also establishes a research map. That map organizes questions, classifications, provisional scores, falsifiers, and evidence needs. It is useful because it shows what would have to be investigated, not because it has already delivered an investable conclusion.

The edition further establishes the result of the final adversarial review. Appendix D records F5-1 through F5-14 as adopted. Those amendments correct the epistemic-multiplier floor, align zero gate states with zero diagnostic maximum weight, impose severity floors, harden comparator requirements, separate liability matching from rates exposure, keep favorable financial-battery scores provisional, require counter-canon work, condition active capital on economic viability, expire stale classifications, preserve the current AI/power factor bar, correct evidence and

classification hygiene, identify certification blockers, make the concentration reopening test achievable, and print the true workload of the active program. These adopted dispositions are part of the final policy text. They do not certify that their inputs have been supplied or that the resulting gates have passed.

The edition also establishes the non-authorization boundary itself. That boundary is not a disclaimer placed outside the policy. It is a governing result repeated at the beginning, in the company rows, in the operating checklist, in Section 15, and in Appendix D. The repetition prevents a reader from detaching an attractive company description, a scenario, or a numerical ceiling from the conditions that make it non-actionable.

What remains uncertified

The policy names the uncertified matters directly. Liability closure has not been demonstrated. This means the policy does not yet know that investable assets, liability cash flows, safety margins, currencies, inflation indices, time horizons, taxes, title, and withdrawal demands have been reconciled into a portfolio that can meet the investor's obligations. The illustrative grind analysis is diagnostic only. One mix fails immediately, and a boundary mix passes only if a substantial real-liability match is genuine. The actual withdrawal path remains open. A hypothetical pass cannot substitute for the missing liability schedule.

Actual funded status, or FS, is also unverified. FS is the ratio of investable assets after known taxes and transaction costs to the probability-weighted present value of liabilities plus an explicit safety margin. Because both sides depend on investor-specific facts, the policy's funded-state bands remain conditional. A band printed in the policy is a dated hypothesis about how governance should respond once the facts are known. It is not a declaration that an investor currently occupies that band.

Trait tolerance and the market drawdown budget, or MDB, remain unverified. Trait tolerance concerns the investor's durable ability to live with losses and uncertainty. It is different from financial capacity, and neither can be inferred from a questionnaire answer, a past bull market, or the existence of a numerical ceiling. MDB defines the tolerated mark-to-market drawdown budget. Without verified tolerance, capacity, liability timing, and income correlation, the policy cannot certify that its conditional architecture fits the person who might use it.

Tax basis is unknown. So are domicile and custody structure. Those facts affect after-tax comparisons, legal title, estate situs, transferability, sanctions exposure, beneficiary and incapacity access, and the feasibility of second-custodian or second-legal-system arrangements. A general policy cannot convert a company domicile, exchange listing, or custody concept into a personalized tax, legal, currency, or account conclusion. Relevant tax and legal review remains a gate where the policy calls for it.

The actual comparator is missing. A diagnostic proxy was studied, but its top-20 worksheet, second independent scoring pass, and capped-comparator study were incomplete. Appendix D adopts the conservative result: no feasibility-cap ratification without that work, and the original caps govern in the meantime. The proxy therefore cannot be promoted into the investor's approved strategic portfolio by repetition or convenience. Actual comparator scoring remains a certification blocker.

Factor look-through is likewise unverified. The policy requires exposure to be traced through the comparator and any proposed active holdings rather than inferred from labels. Embedded artificial-intelligence and power sensitivity is especially relevant because the diagnostic proxy breached the original caps. Appendix D's row-level factor bar preserves current admissible additions of zero for sensitive rows. This is a dated risk control, not a settled forecast about artificial intelligence, asset lives, or defensive capital spending.

Every company valuation remains unverified. The research map contains classifications, provisional pillar scores, evidence coverage percentages, loss severities, diagnostic ceilings, and falsifiers. None of these completes a current reverse discounted cash flow analysis, independent evidence audit, factor-headroom check, tax and legal review where relevant, or implementation approval. Favorable financial-battery scores are assumptions capped at 2 until the named exhibits are attached. Evidence Coverage, abbreviated EC, is not a buy signal. A research

pointer is not extracted evidence. A high-quality business description is not proof that the security offers an adequate return at its current price.

Appendix D restates the final boundary exactly:

> **Final boundary:** *liability closure is uncertified; all numeric bands await verified investor facts; the withdrawal-path grind budget is open; actual comparator scoring and EC audits are incomplete; and no company in this map is a buy absent a current, full evidence and valuation review. No transaction is authorized.*

This is not an invitation to certify the document piecemeal by checking one missing item. The blockers interact. A completed company valuation cannot cure an unmeasured liability position. Verified investor facts cannot cure a hard factor breach. A strong evidence audit cannot cure a failed valuation gate. Approval cannot make an unsupported management claim true. The policy is built around non-compensating gates precisely so that one attractive result cannot erase another missing condition.

The controlling document hierarchy

The authoritative source for this chapter is Section 15 of the final policy. The complete policy supplies the definitions and cross-references needed to understand that short certification paragraph. Appendix D supplies the final Fable amendments and the final restatement of the boundary. The debate rounds and traceability material explain provenance and rationale, but they cannot override the integrated final policy. This expanded chapter is subordinate to all of them. Its purpose is explanation, not amendment.

That hierarchy matters whenever two formulations appear to point in different directions. A final adopted amendment controls over the earlier formulation it explicitly corrects. F5-1, for example, says that the 0.50 floor applies only to the evidence-coverage multiplier. Liquidity and factor penalties retain full force, and a hard factor breach means zero. An earlier floor formulation cannot be used to rescue illiquidity or factor concentration. F5-2 similarly requires an E0 gate state to print a diagnostic maximum weight of zero, and it records Nike as blocked by E1. A stale nonzero diagnostic result cannot supersede the final gate-state rule.

F5-13 offers the clearest example of how to read a later correction. It states that five years and at least 20 audited decisions replace the impossible ten-year prerequisite for reopening the D-1 concentration question. The final Section 1 formulation adds the associated conditions: a preregistered record of fully underwritten decisions, independently audited after-tax attribution, verified liability surplus, and challenger sign-off. Any surviving ten-year wording in an earlier register is superseded by this adopted final amendment. This does not reopen concentration today. It only defines the evidence needed before the rejected minority view may be reconsidered.

Appendix A records article and amendment provenance. Appendix B records open disputes, their present status, and adjudicating observables. Appendix C records model and run provenance and expressly says those receipts are not investment evidence. Appendix D records the final Fable amendments. None of those appendices is a brokerage instruction, an account mandate, or an approval sheet. Provenance shows how the policy was produced. It does not prove the truth of an investment claim.

Research files, company disclosures, worksheets, legal memoranda, tax analysis, and current valuations sit below the governing policy in a different sense: they provide the inputs needed to apply it. They can fill missing evidence, reveal that a gate fails, or support a later policy amendment. They cannot silently rewrite the policy. Conversely, a policy rule cannot manufacture missing evidence. Governance and evidence are complementary, not interchangeable.

The evidence cutoff and the meaning of a source snapshot

The evidence date is 22 August 2026. The cutoff means that the edition's factual and research synthesis is bounded by information available as of that date. It does not mean that every disclosure bearing an earlier date was independently verified, complete, or decision-useful. It also does not mean that a fact remained true after the cutoff. A later filing, price, corporate action, legal decision, investor circumstance, or market event is outside this edition until it is brought into a current review.

Section 15 states the narrower meaning of the source snapshots: they establish only that a disclosure existed as of the evidence date. A link to an investor-relations page may show where evidence could be found. It does not itself establish the accuracy of management's statements, the economic interpretation of a filing, or the completeness of the analyst's evidence set. The company map therefore calls such links evidence pointers and grants no EC credit until the cited disclosure is extracted and assessed.

This distinction prevents source laundering. A management presentation does not become independent corroboration because it is linked. A backlog figure does not become collectible cash merely because it was published. A state budget does not become revenue without the payer gate. An index factsheet does not complete factor look-through. A current price cannot be backfilled into a dated reverse DCF and called part of the 22 August 2026 edition. The cutoff is both a temporal boundary and a warning about epistemic status.

The 12-month expiry of company classifications reinforces this limit. A classification that is not refreshed reverts to an Unclassified Research Note. Expiry is not a sell instruction and does not imply that the business deteriorated. It means the dated research label no longer carries even its prior provisional status. Fresh evidence must be reviewed under the current policy before the classification can again support research work.

Ownership, maintenance, challenge, and approval

The policy distinguishes document maintenance from authority to act. It refers to an operator who can maintain the minimum viable governance system, keep the decision journal and exception log, score the comparator, and conduct a bounded number of full underwrites. It concludes that one person cannot maintain the full 42-company apparatus and caps live full underwrites at eight, subject to the hours and viability tests. F5-14 makes the workload explicit at roughly 160 to 240 hours a year for eight live names, before one-off comparator and custody work. F5-8 permits active capital only above the prospective 2 times fully loaded cost threshold; otherwise the program remains in paper mode.

These maintenance responsibilities do not confer trading authority. The person who calculates FS, updates EC, or writes a reverse DCF does not gain permission to place an order by completing the worksheet. The policy refers to verified investor facts and implementation approval, but it does not name an approving person, broker, adviser, committee, account, or execution channel. This chapter must not invent one. Until the separate implementation process supplies the required authority and all gates pass, the policy's answer is non-authorization.

Independent challenge is also not approval to transact. The policy requires challenge in defined situations, including lower-than-floor severity claims, the larger derating-loss allowance, repeated tax-deferral reliance, and any future attempt to reopen the concentration minority. A challenger tests evidence, assumptions, attribution, and process. A favorable challenge can satisfy one governance condition. It cannot verify missing investor facts, waive a hard gate, or authorize execution.

Tax and legal professionals have similarly bounded roles. Where relevant, their review can establish facts or identify constraints concerning tax basis, title, custody, domicile, sanctions, estate situs, or transferability. Their work may be necessary to pass a policy gate. The policy does not present itself as their substitute, and it does not treat a professional review as a recommendation to own a particular security.

The clean ownership interpretation is therefore functional rather than personal. The final policy owns the governance rules. Current research files own the evidence record. The operator owns maintenance only to the

extent assigned outside this document. The investor facts must be verified for the actual person. A separate implementation process must supply approval and execution authority. Because the policy names no holder of that authority, no reader may assume it.

Prohibited uses of the policy

The standing prohibition is broad: no trade, order, hedge, message, account opening, target allocation, entry price, execution, or portfolio change is authorized. "No transaction is authorized" is the shortest reliable summary. That rule also covers indirect attempts to turn research into action, such as treating a diagnostic maximum weight as a target or treating entry-pacing rules as a live order schedule.

The policy's Section-15 prohibition self-audit makes ten misuse cases explicit. Its reference to "Round 4 Section 15" is provenance from the debate, not a competing current section number. Read as policy boundaries, the prohibitions are:

1. There is no buy call, target weight, entry instruction, order, or portfolio mutation. Every named company remains non-buy with current admissible weight zero. A diagnostic maximum weight is a ceiling used to test policy arithmetic, never an instruction.
2. A theme, market size, prior, reputation, price decline, or index membership is not proof of expected return. Artificial intelligence, rearmament, electrification, health innovation, scarcity, and harvest narratives remain causal questions until cash capture and valuation are demonstrated.
3. The retired additive business-model score cannot be used without the five non-compensating pillars, weakest-pillar result, EC, gates, reverse DCF, and dates. A high total assembled from strengths cannot offset a zero gate.
4. A short-history challenger cannot be promoted to Core. The six such names remain Speculative Options, are analyzed under a total-loss assumption, and have only a diagnostic speculative cap. Their current admissible weight is still zero.
5. Comparator-embedded exposure cannot be ignored. Active additions do not begin with a blank factor budget. The comparator must be scored first, and the original AI/power caps continue to govern while the actual study remains incomplete.
6. A management-only claim cannot receive more than the permitted evidentiary status, conflicting evidence cannot be averaged into comfort, and missing look-through cannot be scored as though it were harmless. Backlog, budgets, or management assertions receive no automatic cash credit.
7. A harvest candidate is not a normal-size holding and cannot roll indefinitely after 24 months. The final map claims no fully qualified harvest candidate. Intel fails the relevant gate and has a policy ceiling of zero.
8. State budgets are not collectible revenue without a payer gate. Appropriation, rate recovery, contract terms, cancellation rights, working capital, state take, and per-share cash capture remain relevant.
9. The document cannot be used to state a personalized tax, account, domicile, currency, or custody conclusion. Those conclusions require the missing investor-specific and professional evidence.
10. Disputes and professional gates cannot be silently resolved. A dispute stays open or conditional until the policy's stated observable adjudicates it or an explicit amendment changes its status. The existence of a preferred narrative is not adjudication.

These prohibitions close common routes by which a research document can become a shadow portfolio. They prevent a reader from extracting a ticker, percentage, scenario, or pacing rule while discarding the gates and approvals that surround it.

How final amendments and future amendments should be read

Appendix D's F5 amendments are already adopted into this edition. They are not optional commentary. Where an F5 item corrects an earlier defect, the corrected rule is the final rule, and the earlier error remains visible for traceability. The policy explicitly records that the original F4-1 floor defect came from Fable, was faithfully implemented by Codex, and was then corrected by Fable in Round 5. Keeping that history visible is not a weakness. It prevents later readers from mistaking provenance for infallibility or concealing a correction to preserve appearances.

Future amendments should be understood through the durable principle "amend rather than excuse." A persistent breach must become an explicit amendment or a compliant design change. It cannot be normalized through renewable exceptions or informal reinterpretation. The policy gives concrete examples: four consecutive quarters of an embedded issuer or factor hard-cap breach force a written choice between a compliant comparator design and a cap raised by amendment; allowing active derating loss above the 50 percent default and up to 75 percent of MDB requires a written amendment and independent challenge.

An amendment changes the governing rule explicitly. New evidence changes an input. An approval applies the rule to a proposed implementation. These are distinct events. A new company filing does not amend the severity floor. A policy amendment does not certify the filing. Completion of a worksheet does not authorize an order. Treating these events as interchangeable would erase the boundary Section 15 is designed to preserve.

Emergency exceptions are not a substitute for amendment. The constitution makes them narrow, risk-monotone, recorded within 24 hours, independently reviewed within seven days, and expired within 30 days. An exception does not create a permanent precedent. Two material Policy Exception Events in 12 months freeze new discretion pending review. Persistent conditions must return to the explicit amendment or compliant-design path.

Nor does this expanded prose amend the final policy. If a future integrated edition changes a rule, the new edition should identify the amendment, its effective evidence context, the language it supersedes, and any surviving dispute using the policy's existing traceability approach. That statement describes how the current policy already distinguishes amendments from excuses; it does not add a new approval process. Until such an amendment exists, explanation cannot overrule the governing text.

Boundary failure modes

The first failure mode is certification laundering: calling the edition "final" and then implying that the investor or portfolio has been certified. This skips the unresolved liabilities, personal facts, comparator work, factor look-through, EC audits, and valuations. It is the exact error corrected by the phrase "not a certificate."

The second is evidence laundering. A source pointer, management disclosure, market-size estimate, state budget, backlog figure, or index factsheet is treated as verified investment evidence. The policy allows those items to identify questions and possible sources. It does not allow them to stand in for independent corroboration, payer collectibility, per-share cash capture, or a completed underwrite.

The third is hierarchy inversion. An old debate formulation, provisional proxy, stale appendix row, or explanatory chapter is used against the final adopted rule. F5 corrections, especially the floor rescope, comparator hardening, and five-year concentration reopening test, show why the final disposition must control. Traceability preserves old language to explain history, not to offer a menu of whichever rule permits more risk.

The fourth is converting a ceiling into a target. Diagnostic maximum weights show what loss arithmetic might permit after gates pass. They do not recommend owning the security, and hard gates still reduce the result to zero. Using Wmax as an allocation instruction would ignore the current admissible weight of zero for every named company.

The fifth is proxy substitution. The diagnostic comparator is treated as the actual approved comparator, or provisional factor scores are treated as ratified caps. Appendix D rejects this. The worksheet, second pass, and capped-comparator study remain blockers, and original caps govern meanwhile.

The sixth is stale-policy drift. Facts after 22 August 2026 are blended into the edition without a current review, while old company classifications are left standing past their 12-month expiry. This creates the appearance of contemporaneous research without the required update. The correct status after expiry is Unclassified Research Note, not automatic renewal.

The seventh is assumed authority. A researcher, operator, model, challenger, adviser, or reader assumes the ability to trade because the analytical work looks complete. The document grants no such authority and does not identify an approver. The separate implementation process remains outside the policy and must not be fabricated from implication.

The eighth is amendment by attrition. A breach persists, exceptions are rolled, and the original rule quietly stops being enforced. The policy requires the opposite: amend explicitly or make the design compliant. A cap increase accepts more risk; it does not make the exposure disappear or certify the reason for taking it.

The ninth is disagreement erasure. Open questions about comparator design, private assets, currency hedging, AI overbuild, asset lives, quality-duration crowding, zero-active investing, and comparator engineering are presented as settled. Appendix D preserves these minority views and the observables that can adjudicate them. Structuring a disagreement is not forcing consensus.

The tenth is operational overload. One person tries to maintain all 42 company files as live underwrites, fails to refresh evidence or score the comparator reproducibly, and continues using the outputs anyway. The policy's bounded version limits full live underwrites to eight and may reduce that number to zero if the hours or prospective viability test fails. If EC audits cannot be maintained, the actual comparator cannot be scored within the stated hours, or liability closure remains unmeasured at the next annual review, new active work freezes and the default contracts to the named comparator, the investor liquidity floor, and verified nominal and real liability matching.

Annex A - Full Company Research Map

A map of questions, not a portfolio

This chapter turns the policy's company list into a readable research map. It does not turn the list into a portfolio, a ranking of expected returns, or a set of recommendations. The distinction is fundamental. A company appears here because its business model helps test one or more of the decade themes, counterthemes, and model classes developed elsewhere in the policy. Inclusion says that a question is worth investigating. It does not say that the security is attractive at its current price, that the evidence has been independently verified, that the security fits the investor's liabilities or domicile, or that any amount of it may be owned. The current admissible weight of every named company is zero.

The map contains exactly 42 named companies. Eleven are classified as **Core Research Candidates**, thirteen as **Watchlist/Price Sensitive**, six as **Speculative Options**, and twelve as **Avoid/Counterexamples**. Those four classifications describe the state and purpose of research. They are not four rungs on an automatic promotion ladder. A Core Research Candidate is a business for which quality appears plausible enough to justify priority work, while valuation and the complete evidence file remain open. A Watchlist/Price Sensitive company may have a plausible model, but at least one gate, expectation, or evidence problem is unresolved. A Speculative Option is a short-history, pre-cash-flow, emerging-market, governance-sensitive, or otherwise high-severity case that can be studied only under an assumed-total-loss discipline; even the policy's diagnostic 0.75% speculative cap is not an authorization. An Avoid/Counterexample is included to show how an attractive narrative can fail one or more non-compensating gates. "Avoid" is therefore a policy classification at the evidence date, not a claim that the share price must fall, and each row supplies evidence that could falsify the classification.

Legal domicile and economic exposure must be kept separate. On the map's stated legal-domicile count, North America accounts for 27 companies, Europe for 9, and Asia or emerging markets for 6. That count does not imply that all revenue, assets, liabilities, regulation, or political exposure sits in the legal domicile. MercadoLibre, for example, is legally domiciled in Delaware but is economically exposed to Latin America. Chubb is listed in the United States but identified here with Switzerland. TSMC combines a Taiwan operating and geopolitical exposure with listings in Taiwan and New York. First Quantum is Canadian-listed while its critical jurisdiction exposures include Zambia and Panama. The map records such distinctions because title, taxation, sanctions, state power, transferability, currency, and operating economics can point to different places.

The composition of the list is itself evidence about the screen. It is materially correlated with the 2010-2025 consensus-quality factor: recurring revenue, high margins, capital-light tolls, installed bases, regulated or mission-critical workflows, and established franchises are conspicuous. This does not establish that those models will outperform from 2026 to 2036. It reveals what the research method tends to surface. The unresolved quality-duration crowding dispute, U3-2, therefore remains open. Eight of the eleven Core names resemble the quality canon that performed strongly in the preceding era, and the map remains North America-heavy. The policy preserves this bias openly rather than claiming that a five-pillar score can remove it. Annual refreshes must run a counter-canon screen and a non-US screen and report what they find, including rejections.

How to read the common fields

Each company entry uses the same conceptual grammar. The five letters A through E are provisional scores on non-compensating pillars. **A** asks whether demand is recurring, necessary, and protected by measurable lock-in. **B** asks whether industry structure, cycle position, payer funding, and collectibility are sound. **C** asks whether incremental after-tax cash returns exceed the true cost of capital after economic depreciation, dilution, leases, and maintenance capital. **D** asks whether value actually reaches each share after compensation, issuance, acquisitions, related parties, state leakage, and other forms of capture. **E** asks whether a conservative reverse discounted-cash-flow case clears the comparator and whether hard disconfirming evidence has been preregistered. A score of 0 is a hard failure, 2 is the ordinary minimum for admission, and 4 represents unusually strong evidence. No strength in one pillar compensates for a zero in another.

The scores in this chapter are analyst assumptions, denoted [A], not completed findings. Evidence Coverage, or **EC**, is the proportion of applicable required evidence fields that has been verified, weighted toward filings, contracts, and regulators. An investor-relations or filings link is only an **evidence pointer**. Its presence shows where a disclosure may be obtained; it earns no EC credit until the disclosure is extracted and tested. The map has not completed the independent adversarial-corroboration pass. A conflict between management disclosure and credible external evidence must score zero until resolved rather than being averaged away. EC below 60% is research-only; EC of 60-74% carries a 0.50 epistemic multiplier; 75-84% carries 0.75; and at least 85% carries 1.0. Even high EC cannot cure a failed pillar or make a recommendation.

S is the assumed bear-case permanent-impairment severity. It is not expected volatility. It estimates the fraction of a position that could be permanently lost under the relevant bear case. The policy floors S at 40% for ordinary single equities, at 60% for levered financials, cyclicals, and roll-ups, and at 80-100% for binary-license, pre-cash-flow, single-project, or governance-flagged emerging-market exposures. **Wmax** is produced by dividing the loss budget by S after evidence, liquidity, and factor penalties. In this chapter it is an optimistic diagnostic policy ceiling. It is never a target weight, recommended size, or permission to act. If any pillar is zero, or if a hard factor cap would be breached, Wmax is zero. The six Speculative Options are additionally subject to the 0.75% per-name speculative cap even where a raw diagnostic formula might print a larger number.

Valuation is deliberately kept separate from business quality. A high A-D profile can still fail E because the market price already assumes too much. NVIDIA and Ferrari make that point explicitly: they are admired businesses on the Watchlist, but their E scores are zero and their Wmax figures are therefore zero. Palantir is the clearest admired-company Avoid case for the same broad reason, compounded by per-share capture concerns. Conversely, a low headline valuation does not repair weak capital economics, governance, title, payer quality, or financial structure. Intel, Flagstar, Bausch Health, and First Quantum illustrate why book value, replacement value, enterprise value, or spot-price net asset value cannot stand in for owner cash.

Every company has a 36-month falsifier. For a Core or Watch name, the falsifier identifies observable deterioration that would break the working thesis or force a re-underwrite. For an Avoid name, it identifies the demanding evidence that could overturn the negative classification. These are research triggers, not automatic trades. A company classification also expires after 12 months without refresh. On expiry it reverts to **Unclassified Research Note**, rather than silently preserving a stale label. The evidence date for this edition is August 22, 2026. Nothing in this chapter updates facts beyond that date.

Core Research Candidates - quality plausible, valuation and full evidence still open

The Core group receives research priority because its business-model fit appears plausible, not because it has passed a complete underwrite. All eleven entries retain valuation sensitivity, missing evidence, a falsifier, and an explicit non-buy boundary. The label cannot be cited as a forecast of superior returns.

1. S&P Global - SPGI, NYSE; United States

Business-quality prior: Favorable and plausible, but provisional - RP-1 (legacy: Core Research Candidate).
Security-return conclusion: Not established: E2, but the current valuation and full underwrite are incomplete.
Portfolio-role conclusion: No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Core Research Candidate. S&P Global represents the ratings, indices, data, and workflow branch of assurance and information tolls. The causal thesis is that regulation, debt issuance, benchmark adoption, and embedded workflow can support recurring revenue with limited incremental capital, while the Mobility separation may make the remaining portfolio more focused. That is a forecast to test, not a fact. The dated evidence pointer is the company's Q2 2026 release, which reported revenue growth of 11%, adjusted-margin expansion of 200 basis points, adjusted EPS growth of 23%, and the July 1 separation of Mobility.

The provisional screen is A4/B3/C4/D3/E2, EC 78%, S 60%, and Wmax 1.67%. Those figures say demand and capital economics look strong enough for priority research, while structure, governance, and the conservative-return case still need work. The 60% severity reflects ratings liability and issuance cyclicality rather than treating a capital-light model as harmless. Valuation is sensitive to a long-duration multiple and debt-cycle volume; a conservative case must normalize issuance rather than extrapolate a favorable quarter. The failure routes are artificial-intelligence substitution in workflow, loss of the ratings franchise through regulation or liability, and weaker information-toll economics. Required work includes segment retention, a price-versus-unit bridge, litigation exposure, and a reverse DCF. The 36-month falsifier is recurring growth below 5%, margin erosion greater than 300 basis points, or a material de-licensing or liability event. This is a candidate only. No full current valuation or legal review exists, so admissible weight remains zero.

2. CME Group - CME, Nasdaq; United States

Business-quality prior: Favorable and plausible, but provisional - RP-1 (legacy: Core Research Candidate).
Security-return conclusion: Not established: E2, but the current valuation and full underwrite are incomplete.
Portfolio-role conclusion: No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Core Research Candidate. CME Group is the map's exchange and clearing rail, where two-sided liquidity and the recurring need to manage risk may form a durable network. The thesis is that macro uncertainty and continued electronification can deepen liquidity, and that additional contract volume can be monetized at high cash margins. The thesis is constrained by fee regulation, concentration in major products, operational resilience, and clearing risk. The evidence pointer is management's Q1 2026 commentary, which reported record quarterly revenue and adjusted earnings alongside continued volume growth.

The provisional screen is A4/B4/C4/D3/E2, EC 78%, S 50%, and Wmax 2.00%. The scores express a strong working view of demand, structure, and capital economics, while leaving governance and valuation short of the highest grade. The current multiple must be tested against normalized, not exceptional, volatility and against the resilience of both fees and volume. Rival liquidity, mandated fee reductions, a clearing default, or a serious cyber outage could break the economics quickly. Research therefore needs product-level open interest and cohort economics, concentration analysis, stressed-liquidity evidence, and a reverse DCF. The 36-month falsifier is a loss of share in core contracts combined with fee compression greater than 10%, or a material impairment of the clearing franchise. CME is a research candidate, not a recommendation, and its admissible weight is zero.

3. Visa - V, NYSE; United States

Business-quality prior: Favorable and plausible, but provisional - RP-1 (legacy: Core Research Candidate).
Security-return conclusion: Not established: E2, but the current valuation and full underwrite are incomplete.
Portfolio-role conclusion: No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Core Research Candidate. Visa represents a global acceptance, data, and payment rail whose relevance rests on scale, density, reliability, and four-party network economics. The causal thesis is that the continued shift toward electronic payments and the growth of value-added services can outpace underlying consumption while remaining capital-light. The model is not immune to law, incentives, or technological substitution. The dated evidence pointer is Visa's fiscal Q3 2026 release, which supplied the quarter's results and operating data.

The provisional screen is A4/B4/C4/D4/E2, EC 82%, S 50%, and Wmax 2.00%. Visa is one of the strongest apparent A-D profiles in the map, but the E score remains only 2 because a current conservative reverse DCF is absent. The valuation behaves like a long-duration asset and depends on cross-border mix, the normalized cost of client incentives, and whether service growth represents genuine incremental owner value. Regulatory reductions in interchange, account-to-account displacement, cyber failure, or settlement impairment are central failure paths. Work still needed includes net yield, incentive cohorts, token and service economics, and antitrust exposure. The 36-month falsifier is payment-volume share loss together with net-revenue yield falling by more than 100 basis points, or a structural regulatory reset of the take rate. The classification is not an entry judgment; admissible weight remains zero.

4. Fastenal - FAST, Nasdaq; United States

Business-quality prior: Favorable and plausible, but provisional - RP-1 (legacy: Core Research Candidate).
Security-return conclusion: Not established: E2, but the current valuation and full underwrite are incomplete.
Portfolio-role conclusion: No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Core Research Candidate. Fastenal is the positive case for reliability distribution. Its dense routes, on-site inventory, vending, and low-ticket but mission-critical products may become embedded in customer operations. The causal thesis is that replenishment density and on-site service improve customer productivity, deepen share, and support attractive economics. Working capital is the central cost of the model, and digital price transparency or customer insourcing may weaken its advantage. The company's Q2 2026 release reported daily sales growth of 14.7%, a 21.0% operating margin, and operating cash flow equal to 69% of net income.

The provisional screen is A3/B3/C3/D4/E2, EC 80%, S 50%, and Wmax 2.00%. The screen is solid rather than flawless: customer need, industry structure, and capital economics all require cohort-level proof, while per-share capture appears relatively strong. A premium valuation assumes resilient organic growth and margins, so the underwrite must normalize the industrial cycle and inventory cash use. Failure could come from customers internalizing procurement, online pricing eroding service value, or route-density investments diluting returns. Required evidence includes site and cohort returns, price and unit bridges, inventory obsolescence, and a reverse DCF. The 36-month falsifier is failure of on-site growth to produce share gains while cash conversion remains below 75% through a full cycle. No valuation or portfolio-fit review has been completed; admissible weight is zero.

5. Berkshire Hathaway - BRK.B, NYSE; United States

Business-quality prior: Favorable and plausible, but provisional - RP-1 (legacy: Core Research Candidate).
Security-return conclusion: Not established: E2, but the current valuation and full underwrite are incomplete.
Portfolio-role conclusion: No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Core Research Candidate. Berkshire combines insurance float, decentralized operating companies, a large balance sheet, and capital-allocation optionality. The thesis is that low-cost float, retained operating cash, and disciplined redeployment can compound per-share value across regimes. Scale narrows the opportunity set, succession matters, and a conglomerate can hide both value and risk. The evidence pointer is the 2025 annual report filed in February 2026, which contains insurance, railroad, energy, manufacturing, and liquidity disclosures.

Because insurance is a principal balance-sheet business, Berkshire uses a class battery rather than a generic complexity flag. Reserve development is provisionally 2, float cost 2, and capital 2. All are analyst assumptions and remain capped at 2 until ten-year reserve triangles, float-cost analysis, and stressed-capital exhibits are attached. The broader screen is A3/B3/C4/D3/E2, EC 82%, S 40%, and Wmax 2.50%. That 2.50% is the map's optimistic diagnostic ceiling for a fully liquid ordinary-equity severity floor, not a size suggestion. Valuation requires a sum-of-parts view, normalized underwriting, and explicit treatment of cash drag, tax, and succession. Failure modes include poor post-succession allocation, catastrophe reserve error, liabilities at Berkshire Hathaway Energy, and a persistent conglomerate discount. The 36-month falsifier is persistent adverse reserve development combined with per-share intrinsic-value growth below the comparator after tax and cost. The insurer and conglomerate review is incomplete; admissible weight remains zero.

6. Chubb - CB, NYSE; Switzerland

Business-quality prior: Favorable and plausible, but provisional - RP-1 (legacy: Core Research Candidate).
Security-return conclusion: Not established: E2, but the current valuation and full underwrite are incomplete.
Portfolio-role conclusion: No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Core Research Candidate. Chubb is the positive property-and-casualty underwriting and float franchise in the map. The causal thesis is that disciplined risk selection, broad global distribution, and positive float economics can compound book value through pricing cycles. Insurance quality cannot be inferred from one favorable combined ratio; reserve development, reinsurance recoverability, capital resilience, and the possibility of long-tail casualty inflation require multi-cycle evidence. The company's Q2 2026 release reported an 83.8% combined ratio, 12.3% growth in net investment income, core operating return on equity of 15.3%, and 12.3% year-on-year growth in book value per share.

The class battery is reserves 2, float cost 2, and capital 2, all provisional [A] and capped until ten-year reserve triangles, float-cost exhibits, and stressed-capital evidence are attached. The broader screen is A3/B3/C4/D4/E2, EC 84%, S 40%, and Wmax 2.50%. Valuation must normalize catastrophe losses and relate sustainable return on equity to price-to-book. Latent casualty inflation and catastrophe-model error are the principal failure routes. Needed work includes the reserve triangles, reinsurance recoverables, and climate concentration. The 36-month falsifier is cumulative adverse development above 5% of opening equity or a combined ratio structurally above 100. No current price-to-book and return-on-equity underwrite has been completed; admissible weight is zero.

7. DBS Group - D05, SGX; Singapore

Business-quality prior: Favorable and plausible, but provisional - RP-1 (legacy: Core Research Candidate).
Security-return conclusion: Not established: E2, but the current valuation and full underwrite are incomplete.
Portfolio-role conclusion: No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Core Research Candidate. DBS represents the deposit-franchise bank model and provides an important counterweight to a list otherwise tilted toward capital-light quality franchises. The causal thesis is that low-cost regional deposits, digital service, payments, and wealth fees can sustain returns above the cost of capital as Asian wealth deepens. A bank's apparent quality can disappear through funding repricing, credit losses, property exposure, operational failure, or state direction, so a class-specific balance-sheet battery is mandatory. The Q2 2026 release reported net profit of S$3.08 billion, first-half return on equity of 17.5%, an NPL ratio of 1.0%, and deposit growth of 11%.

The provisional battery is deposit beta 2, full-cycle credit 2, and capital 2. Those scores cannot rise until deposit-cohort data, full-cycle loss evidence, and stressed common-equity-tier-one capital exhibits are attached. The broader screen is A4/B3/C4/D3/E2, EC 82%, S 50%, and Wmax 2.00%. Valuation must normalize net interest margin, credit costs, and regional property exposure. Failure could arise from deposit repricing, China or Southeast Asia credit losses, an operational outage, or policy direction that diverts value from minority shareholders.

Research needs stressed CET1, credit vintages, deposit cohorts, and related-party and jurisdiction work. The 36-month falsifier is NPL above 2.5%, material deposit-share loss, and return on equity below the cost of capital across six quarters. It is not a completed bank valuation or domicile-fit review; admissible weight is zero.

8. Taiwan Semiconductor Manufacturing - TSM/2330, NYSE/TWSE; Taiwan

Business-quality prior: Favorable and plausible, but provisional - RP-1 (legacy: Core Research Candidate). **Security-return conclusion:** Not established: E2, but the current valuation and full underwrite are incomplete; admission is also blocked by the current hard factor gate. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Core Research Candidate. TSMC is the trusted-foundry, leading-process, and advanced-packaging bottleneck in the artificial-intelligence equipment ecosystem. The causal thesis is that process complexity, yield learning, and customer trust can sustain share and returns. Geopolitical severity is irreducible: operating excellence does not diversify blockade, conflict, export controls, or cross-strait legal and title risk. The evidence pointer is TSMC's Q2 2026 results page, which published audited or management results and guidance.

The provisional screen is A4/B4/C3/D3/E2, EC 80%, S 80%, and Wmax 1.25%. The 80% severity is essential to the reading: business quality does not justify treating a Taiwan-centered foundry as an ordinary 40% impairment case. Valuation must normalize artificial-intelligence demand, utilization, economic depreciation, and overseas-fab returns, and no terminal multiple may ignore Taiwan risk. Failure could come from conflict or blockade, export controls, node execution, or overbuilding. Required evidence includes wafer price and unit bridges, packaging economics, customer concentration, overseas-fab returns, and a reverse DCF. The 36-month falsifier is leading-node share loss greater than 10 percentage points, sustained sub-cost overseas returns, or advanced-node utilization below 70% absent recession. In addition, while the unratified comparator proxy remains above the original AI and power risk caps, any active addition is zero regardless of a later underwrite. Geopolitical, legal, factor, and valuation work is incomplete; admissible weight is zero.

9. ASML - ASML, Euronext Amsterdam/Nasdaq; Netherlands

Business-quality prior: Favorable and plausible, but provisional - RP-1 (legacy: Core Research Candidate). **Security-return conclusion:** Not established: E2, but the current valuation and full underwrite are incomplete; admission is also blocked by the current hard factor gate. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Core Research Candidate. ASML represents the lithography-equipment oligopoly and the service economics of a critical installed base. The picks-and-shovels thesis is that rising patterning complexity, scarce technical capability, and an expanding installed base can sustain rents and recurring service revenue. Customer concentration, export policy, technical substitution, equipment useful lives, and semiconductor-capex digestion constrain the case. The Q2 2026 filing exhibit reported sales of EUR9.3 billion, a 54% gross margin, net income of EUR2.9 billion, and a 2026 sales outlook of EUR43-45 billion.

The provisional screen is A4/B4/C4/D3/E2, EC 82%, S 70%, and Wmax 1.43%. The elevated severity reflects the binary influence of export rules and the cyclicality and concentration of advanced-equipment demand. Valuation is exceptionally sensitive to terminal extreme-ultraviolet and High-NA economics; bookings and China mix must be normalized rather than projected mechanically. Failure paths include tighter China restrictions, customer digestion, technical substitution, and shorter tool lives. Research needs installed-base margins, backlog cancellation terms, customer return on investment, and a reverse DCF. The 36-month falsifier is material loss of leading-edge lithography share or six quarters in which service revenue diverges negatively from the installed base. While the comparator proxy breaches the original AI and power risk caps, active addition is zero whatever a future underwrite concludes. Quality does not clear price, policy, or factor exposure automatically; admissible weight is zero.

10. Schneider Electric - SU, Euronext Paris; France

Business-quality prior: Favorable and plausible, but provisional - RP-1 (legacy: Core Research Candidate). **Security-return conclusion:** Not established: E2, but the current valuation and full underwrite are incomplete; admission is also blocked by the current hard factor gate. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Core Research Candidate. Schneider Electric is the power-management, automation, and data-center electrical bottleneck in the map. The causal thesis is that electrification and constrained power availability increase demand for efficient control hardware, software, and service supported by a broad channel. The same growth can attract competing capacity, create working-capital needs, and produce expectations that run ahead of cash returns. Schneider's financial-results hub reported first-half revenue of EUR21.2 billion, organic growth of 14%, adjusted EBITA of EUR4.1 billion, a 19.3% margin, and free cash flow of EUR1.6 billion.

The provisional screen is A4/B3/C3/D3/E2, EC 78%, S 50%, and Wmax 2.00%. Valuation may capitalize data-center demand too far into the future, so working capital, price versus volume, mix, and cycle normalization matter. Capacity glut, channel destocking, integration problems, or Chinese competition could reverse the thesis. Needed work includes organic price and volume, precise data-center exposure, installed-service economics, and economic depreciation. The 36-month falsifier is a break in orders-to-revenue conversion while capital expenditure rises and incremental return on invested capital falls below its cost. The breached AI and power factor state makes any active addition zero regardless of future company work until the comparator issue is resolved. The expectations audit is incomplete; admissible weight remains zero.

11. RELX - REL, LSE/NYSE; United Kingdom

Business-quality prior: Favorable and plausible, but provisional - RP-1 (legacy: Core Research Candidate). **Security-return conclusion:** Not established: E2, but the current valuation and full underwrite are incomplete. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Core Research Candidate. RELX represents embedded legal, scientific, and risk workflow supported by data and authoritative content. It is an information toll for which artificial intelligence can be both an enhancer and a substitute. The causal thesis is that authoritative datasets, workflow integration, and high renewal can sustain per-share compounding if the company captures more value from AI than it loses through content commoditization. The 2025 results release published in February 2026 reported continued underlying growth and gave a 2026 outlook.

The provisional screen is A4/B3/C4/D4/E2, EC 77%, S 50%, and Wmax 2.00%. A premium multiple assumes durable AI monetization, continued renewal, and successful reinvestment; acquisitions and content investment must be normalized. Generative substitutes could weaken pricing or retention, while data regulation could limit product economics. Research needs product cohorts, AI attachment economics, an organic-versus-acquired growth bridge, and a reverse DCF. The 36-month falsifier is weakening renewal or usage accompanied by organic growth below 4%, or AI-related cost causing margin erosion greater than 300 basis points. RELX remains a research candidate only, not a recommendation, and admissible weight is zero.

Watchlist/Price Sensitive - model plausible, gate or expectation unresolved

The Watchlist is not a waiting room in which the passage of time produces approval. Each of these thirteen companies has a plausible business-model connection to the policy, but a present gate, valuation, legal, payer, cycle, or evidence issue blocks Core treatment or admission. Several are deliberate demonstrations that excellent reported growth and powerful competitive positions can still fail E. The correct use of this group is to define the missing work and wait for evidence, not to anticipate promotion.

12. GE Aerospace - GE, NYSE; United States

Business-quality prior: Plausible with an unresolved gate or price - RP-2 (legacy: Watchlist/Price Sensitive). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E1 is below the ordinary admission minimum; the current valuation and full underwrite are incomplete. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Watchlist/Price Sensitive. GE Aerospace is the positive installed-engine aftermarket and aviation-bottleneck case, but it also carries safety, supply-chain, warranty, and ancillary-finance complexity. The causal thesis is that a growing engine installed base can create long-duration shop-visit and service cash flows. That operating leverage can reverse if fleet availability, supplier execution, safety, or economic-life assumptions disappoint. The Q2 2026 release reported orders of $16.5 billion, up 17%; revenue of $13.3 billion, up 21%; and free cash flow of $3.0 billion, up 43%.

The provisional screen is A4/B3/C3/D2/E1, EC 72%, S 60%, and Wmax 1.67%. The diagnostic number does not overcome E below 2, and the incomplete complexity work is independently blocking. Current expectations require durable shop-visit growth, while engine lives, warranties, customer advances, and supplier or receivables financing must be recast into owner economics. Failure modes include a safety event, fleet grounding, supplier bottleneck, or hidden finance exposure. Required work includes engine-vintage service cohorts, warranty evidence, receivables finance, and owner earnings. The 36-month falsifier is shop-visit conversion missing plan by more than 15% while cash conversion remains below 70% without a recession. GE Aerospace is not a buy; admissible weight is zero.

13. Microsoft - MSFT, Nasdaq; United States

Business-quality prior: Plausible with an unresolved gate or price - RP-2 (legacy: Watchlist/Price Sensitive). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E1 is below the ordinary admission minimum, the current valuation and full underwrite are incomplete, and the current hard factor gate independently blocks admission. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Watchlist/Price Sensitive. Microsoft combines an installed enterprise workflow, cloud platform, developer ecosystem, and internal hyperscaler. The causal thesis is that distribution and ecosystem strength can monetize artificial intelligence across software and cloud. The unresolved question is per-share capture after massive capital expenditure, depreciation, power, model costs, and possible cannibalization of existing products. The company's fiscal 2026 release reported revenue of $331.8 billion, up 18%; operating income of $155.2 billion, up 21%; and net income of $133.7 billion.

The provisional screen is A4/B4/C3/D4/E1, EC 74%, S 50%, and Wmax 2.00%. The map treats Microsoft as admired consensus quality whose price may embed a long duration of AI success. Because a conservative reverse DCF has not been completed, the E gate fails. The critical failure is AI revenue that does not cover incremental depreciation and power while also weakening the pricing or value of existing software. Research needs disaggregated AI revenue, depreciation lives, capital-expenditure cohort returns, and cloud-backlog quality. The 36-month falsifier is AI-linked gross profit remaining below incremental depreciation and power for six quarters. The current comparator proxy already breaches the original AI and power risk caps, so admissible active addition is zero regardless of a later company underwrite. No entry conclusion follows from franchise quality; admissible weight is zero.

14. Alphabet - GOOGL, Nasdaq; United States

Business-quality prior: Plausible with an unresolved gate or price - RP-2 (legacy: Watchlist/Price Sensitive). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E1 is below the ordinary admission minimum, the current valuation and full underwrite are incomplete, and the current hard factor gate independently blocks admission. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Watchlist/Price Sensitive. Alphabet links an intent auction, dominant distribution, cloud, and an internal hyperscaler. The causal thesis is that proprietary demand intent and distribution can defend advertising economics while cloud and AI create new revenue. The countertheme is direct: answer engines can reduce clicks and change the unit economics of search before AI or cloud profits fill the gap. Alphabet's investor site is the evidence pointer for its Q2 2026 Form 10-Q and results package.

The provisional screen is A4/B3/C4/D3/E1, EC 72%, S 60%, and Wmax 1.67%. A proper valuation must separate Search, Cloud, and Other Bets in a reverse DCF and must test capital expenditure and depreciation lives. The failure routes are query monetization declining faster than AI and cloud contribution grows, and antitrust remedies changing distribution or economics. Research requires AI-query revenue and cost cohorts, traffic-acquisition costs, depreciation, and remedy scenarios. The 36-month falsifier is four quarters in which both Search revenue per query and operating profit decline while capital expenditure as a percentage of sales remains elevated. The current AI and power factor-cap state independently sets admissible active addition to zero. Valuation and legal arithmetic remain incomplete; admissible weight is zero.

15. NVIDIA - NVDA, Nasdaq; United States

Business-quality prior: Plausible with an unresolved gate or price - RP-2 (legacy: Watchlist/Price Sensitive). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E0 is a hard failure, the current valuation and full underwrite are incomplete, and the current hard factor gate independently blocks admission. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Watchlist/Price Sensitive. NVIDIA is the accelerated-compute equipment and ecosystem oligopoly. Its market thesis is powerful: integrated hardware and software, rapid product cadence, developer adoption, and scarce compute can preserve rents if customers earn sufficient cash returns to keep replacing systems. The policy refuses to infer shareholder return from demand or technological leadership alone. NVIDIA's fiscal 2026 release reported revenue of $215.9 billion, up 65%; Data Center revenue of $193.7 billion, up 68%; and a 71.1% gross margin.

The provisional screen is A4/B3/C4/D4/E0, EC 74%, and S 70%. **Wmax is zero because E0 fails admission.** This consistency matters: the raw severity arithmetic would otherwise print a diagnostic ceiling, but a failed pillar cannot be rescued by admired quality. The conservative case has not cleared, cyclically exceptional margins may not persist, and customer concentration makes the terminal case fragile. Failure could arrive through poor end-customer monetization, custom silicon, export controls, financing loops, or faster obsolescence. Needed evidence includes end-customer cash return on investment, equipment-financing relationships, useful lives, rental prices, and normalized margins and replacement cycles. The 36-month falsifier is falling rental prices and customer AI gross profit while supply and capital expenditure rise, or Data Center growth below 10% with an inventory correction. The factor-cap state is also independently blocking. This is the required admired consensus-quality rejection by E and severity arithmetic; admissible weight is zero.

16. Eaton - ETN, NYSE; Ireland

Business-quality prior: Plausible with an unresolved gate or price - RP-2 (legacy: Watchlist/Price Sensitive). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E1 is below the ordinary admission minimum, the current valuation and full underwrite are incomplete, and the current hard factor gate independently blocks admission. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Watchlist/Price Sensitive. Eaton is an electrical distribution and control franchise exposed to the grid, electrification, and data-center bottleneck. The causal thesis is that scarce power and electrification can support backlog, pricing, and attractive incremental economics. The countertheme is that visible scarcity attracts capacity and acquisitions, while duplicated orders can make backlog appear stronger than realized demand.

Eaton's Q2 2026 release reported sales growth of 21%, including 14% organic growth; Electrical Americas order growth of 41%; backlog growth of 43%; and segment margins of 23.1%.

The provisional screen is A4/B3/C3/D3/E1, EC 72%, S 50%, and Wmax 2.00%. The multiple may be extrapolating backlog and peak margins, so research must translate orders into cash and test whether price, cost, and capital intensity normalize unfavorably. Failure could come from order duplication or cancellation, a flood of new capacity, or integration problems. Required evidence includes backlog conversion, a price-cost bridge, organic-versus-acquired growth, and incremental return on invested capital. The 36-month falsifier is reversing backlog growth accompanied by negative incremental margins through four quarters. Eaton also falls under the current AI and power factor bar, which makes active addition zero whatever a later underwrite finds. It is a price-sensitive research watch, not a recommendation; admissible weight is zero.

17. London Stock Exchange Group - LSEG, LSE; United Kingdom

Business-quality prior: Plausible with an unresolved gate or price - RP-2 (legacy: Watchlist/Price Sensitive). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E1 is below the ordinary admission minimum; the current valuation and full underwrite are incomplete. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Watchlist/Price Sensitive. LSEG combines market infrastructure, clearing, indices, data, and workflow. The causal thesis is that embedded data and workflow plus clearing can compound recurring income if Microsoft-linked product investment produces genuine organic per-share capture. Integration, capitalized development, data disintermediation, and clearing regulation make the model less simple than a pure information toll. The H1 2026 release called for 7-7.5% organic total-income growth, about 100 basis points of margin expansion, and at least GBP2.7 billion of free cash flow.

The provisional screen is A4/B3/C3/D3/E1, EC 73%, S 50%, and Wmax 2.00%. E remains below admission because integration and AI expectations have not been tested in a segment reverse DCF. Failure routes include data disintermediation, a regulatory change to clearing, continued Refinitiv integration drag, or a capital loss in the clearing operation. Research needs annual-subscription-value retention, product cohorts, capitalized-development analysis, and clearing stress. The 36-month falsifier is annual-subscription-value growth below 4%, a margin miss greater than 300 basis points, or a clearing capital loss. LSEG is not admitted; admissible weight is zero.

18. Moody's - MCO, NYSE; United States

Business-quality prior: Plausible with an unresolved gate or price - RP-2 (legacy: Watchlist/Price Sensitive). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E1 is below the ordinary admission minimum; the current valuation and full underwrite are incomplete. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Watchlist/Price Sensitive. Moody's combines ratings assurance and analytics and is structurally similar to S&P Global, with sensitivity to debt issuance and binary regulatory or liability risk. The causal thesis is that regulatory embedding and proprietary credit-risk datasets can support compounding. The policy withholds Core treatment because price, evidence, and assurance severity remain unresolved. The company's financials page is the evidence pointer for its 2026 quarterly releases and SEC reports.

The provisional screen is A4/B3/C4/D3/E1, EC 70%, S 70%, and Wmax 1.43%. The severity is higher than for an ordinary information business because a ratings-liability or licensing event could permanently damage the franchise. Valuation must normalize debt issuance and test analytics retention; no conservative return has yet cleared. Ratings reform or liability, private-credit opacity, and AI-driven data substitution are the main failure paths. Research requires segment price and volume, retention, litigation, and issuance sensitivity. The 36-month falsifier is materially falling analytics retention combined with declining ratings share or fees through a full issuance recovery. Assurance severity and E block Core status and admission; admissible weight is zero.

19. Constellation Software - CSU, TSX; Canada

Business-quality prior: Plausible with an unresolved gate or price - RP-2 (legacy: Watchlist/Price Sensitive). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E1 is below the ordinary admission minimum; the current valuation and full underwrite are incomplete. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Watchlist/Price Sensitive. Constellation Software is the decentralized vertical-market-software serial acquirer. The causal thesis is that many small, disciplined acquisitions of durable niche workflows can compound per share if the deployment runway, culture, and organic quality persist. The map treats acquisition as reinvestment, not as an automatic source of growth, and therefore demands vintage economics and a clean organic bridge. The company's press-release and filings page is the evidence pointer for its 2026 shareholder materials.

The provisional screen is A4/B3/C4/D3/E1, EC 69%, S 60%, and Wmax 1.67%. The valuation needs acquisition-cohort internal rates of return, dilution-adjusted owner cash, leverage, contingent consideration, and organic-versus-acquired growth. Failure may emerge through deployment at falling returns, concealed organic decay, increased leverage, or dilution of key-person and cultural advantages. The 36-month falsifier is three acquisition vintages earning below the cost of capital or organic revenue remaining negative for six quarters. Constellation is an admired compounder, but the evidence and price gate remain open; admissible weight is zero.

20. Eli Lilly - LLY, NYSE; United States

Business-quality prior: Plausible with an unresolved gate or price - RP-2 (legacy: Watchlist/Price Sensitive). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E1 is below the ordinary admission minimum; the current valuation and full underwrite are incomplete. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Watchlist/Price Sensitive. Eli Lilly represents patented health innovation and a manufacturing bottleneck. The causal thesis is that incretin and oncology innovation, scaled manufacturing, and broader treated populations can create large clinical and economic value. The payer, not addressable patient count, determines collectibility; net price, access, safety, franchise concentration, patent life, and replacement pipeline dominate the shareholder case. Lilly's Q2 2026 release reported revenue of $23 billion, up 48%; a reported price effect of negative 13%; four acquisitions; and $4.5 billion of manufacturing investment.

The provisional screen is A4/B3/C3/D3/E1, EC 71%, S 60%, and Wmax 1.67%. The valuation embeds substantial flagship growth and terminal share, while capacity expansion and acquisitions create additional capital at risk. Failure could come through a safety event, access or net-price reset, capacity misbuild, or a pipeline cliff. Research requires indication-level net price, patient persistence, payer mix, patent and exclusivity mapping, and a replacement pipeline. The 36-month falsifier is volume failing to offset net-price erosion while lead-franchise contribution peaks without an adequate replacement. E and concentration remain unresolved; admissible weight is zero.

21. Novo Nordisk - NOVO-B, Nasdaq Copenhagen; Denmark

Business-quality prior: Plausible with an unresolved gate or price - RP-2 (legacy: Watchlist/Price Sensitive). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E1 is below the ordinary admission minimum; the current valuation and full underwrite are incomplete. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Watchlist/Price Sensitive. Novo Nordisk is the diabetes and obesity franchise with specialized peptide manufacturing. The causal thesis is that chronic-disease penetration can compound if efficacy, supply, net pricing, and pipeline renewal withstand competition. The market size alone proves nothing about reimbursement, competitive share, manufacturing return, or value reaching each share. Novo's financial-results page published Q2 and first-half 2026 results on August 4.

The provisional screen is A4/B3/C3/D3/E1, EC 71%, S 60%, and Wmax 1.67%. Valuation must normalize GLP-1 market share, net price, capital expenditure, and patent timing rather than assume terminal franchise persistence. Failure routes are share loss, payer restrictions, safety, manufacturing diseconomy, or failure of the next generation of products. Research needs country- and indication-level net price, persistence cohorts, capacity returns, and exclusivity dates. The 36-month falsifier is a material decline in both obesity share and net price while next-generation assets fail noninferiority. A full clinical and valuation review is absent; admissible weight is zero.

22. MercadoLibre - MELI, Nasdaq; United States legal domicile, Latin America economic exposure

Business-quality prior: Plausible with an unresolved gate or price - RP-2 (legacy: Watchlist/Price Sensitive). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E1 is below the ordinary admission minimum; the current valuation and full underwrite are incomplete. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Watchlist/Price Sensitive. MercadoLibre combines commerce, logistics, payments, and credit across Latin America. Its causal thesis is that density and trust lower transaction and fulfillment friction, allowing commerce and payments to reinforce one another. Credit makes the model a balance-sheet and funding-cycle problem as well as a platform problem, while foreign exchange and jurisdiction determine how local success becomes US-dollar value per share. The company overview reported 131 million last-twelve-month unique buyers, 2.9 billion last-twelve-month items delivered, and 88 million monthly active users in Q2; the company filings page reported first-quarter local-currency revenue growth of 45.5%.

The provisional screen is A4/B3/C3/D3/E1, EC 70%, S 70%, and Wmax 1.43%. Valuation requires a segment sum of the parts, foreign-exchange normalization, and credit-vintage economics. Failure could come from credit losses, regulation, logistics intensity, dilution, or other leakage. Research needs customer-cohort returns, credit-vintage losses and funding, and US-dollar per-share cash capture. The 36-month falsifier is credit loss exceeding risk pricing while US-dollar free cash flow per share lags regional nominal activity for two years. The emerging-market and financial battery work is incomplete. MercadoLibre's Delaware domicile must not be confused with its Latin American economics; admissible weight is zero.

23. Ferrari - RACE, NYSE/Euronext Milan; Netherlands/Italy

Business-quality prior: Plausible with an unresolved gate or price - RP-2 (legacy: Watchlist/Price Sensitive). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E0 is a hard failure; the current valuation and full underwrite are incomplete. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Watchlist/Price Sensitive. Ferrari is the scarcity-luxury-brand case, built around controlled supply, personalization, resale values, and a durable owner community. The causal thesis is that enforced scarcity and brand desire can support price and mix while preserving high returns. The policy deliberately separates this exceptional apparent quality from the share's expected return. Ferrari's corporate page reported 2025 net revenue of EUR6.7 billion, adjusted EBITDA of EUR2.6 billion, and industrial free cash flow of EUR1.0 billion.

The provisional screen is A4/B4/C4/D4/E0, EC 70%, and S 50%. **Wmax is zero because E0 fails admission.** The conservative reverse DCF has not cleared, and electric-vehicle transition, research and development, personalization, client depth, and terminal duration are sensitive. Failure can arise through diluted scarcity, brand damage, weaker Chinese demand, or capital intensity. Needed work includes order-book cancellation, client cohorts, resale values, model-level economics, and a current price-based reverse DCF. The 36-month falsifier is unit volume growing faster than client and order depth while residual values and personalization decline. Ferrari is an admired quality business rejected at the current unknown arithmetic; admissible weight is zero.

24. Rheinmetall - RHM, Xetra; Germany

Business-quality prior: Plausible with an unresolved gate or price - RP-2 (legacy: Watchlist/Price Sensitive). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E1 is below the ordinary admission minimum; the current valuation and full underwrite are incomplete. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Watchlist/Price Sensitive. Rheinmetall is the state-funded defense bottleneck and ammunition-capacity case. The causal thesis is that replenishment and European rearmament can support a multi-year funded backlog. State budgets are not cash, however. Appropriation, contract funding, cancellation rights, advances, working capital, state monopsony, capacity returns, and the possibility of a peace dividend determine shareholder capture. Rheinmetall's investor-relations page is the evidence pointer for its 2026 reports and order and backlog disclosures.

The provisional screen is A4/B2/C3/D3/E1, EC 68%, S 60%, and Wmax 1.67%. The B score of 2 reflects that the state payer and procurement structure is only minimally adequate for further research, not proven. Valuation must translate backlog into funded cash and normalize margins and working capital. Failure may come from procurement delay or cancellation, cost overruns, political control, or peace-driven reversal. Research needs contract funding, cancellation terms, program-level cash conversion, and returns on new capacity. The 36-month falsifier is funded awards or backlog conversion below 70% of plan while working capital absorbs more than 50% of cumulative EBIT. Payer and valuation gates are incomplete; admissible weight is zero.

Speculative Options - research optionality only, with total loss assumed

The six Speculative Options are not smaller versions of Core candidates. Their short histories, capital structures, governance exposures, cyclicality, or emerging-market risks prevent ordinary concentration treatment. Each uses S of 100%, meaning total loss is the planning assumption. The diagnostic Wmax shown is capped at 0.75% per name before any larger formula result, and the actual admissible weight remains zero. The purpose is to preserve disciplined learning about challengers without letting an exciting narrative evade the mature-company evidence standard.

25. CoreWeave - CRWV, Nasdaq; United States

Business-quality prior: Emerging and unproven, under a total-loss assumption - RP-3 (legacy: Speculative Option). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E0 is a hard failure. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Speculative Option. CoreWeave is the merchant or neocloud AI-capacity case, distinct from an internal hyperscaler and from an equipment oligopoly. The causal thesis is that specialized infrastructure, software, and customer relationships might monetize scarce accelerated compute. The countercase is a financing and collateral loop: falling rental prices or utilization can impair both operating cash and the residual equity beneath debt and lease commitments. The company's investor-relations page is the evidence pointer for 2026 results, debt, leases, and customer disclosures.

The provisional screen is A3/B1/C1/D1/E0, EC 55%, S 100%, and Wmax 0.75% only because the speculative cap binds before any larger diagnostic result. EC below 60% and E0 independently keep the name research-only. Conventional sales multiples are inadmissible. Valuation must begin with the debt and lease waterfall, contract enforceability, equipment useful lives, normalized rental prices, power commitments, and the residual claim of equity. Failure routes include falling utilization or rental prices, refinancing pressure, customer concentration, and GPU-collateral concentration. Research needs related funding, power leases, covenants, depreciation, and customer contracts. The 36-month falsifier is rental prices falling more than 30% while utilization stays below 70%, or interest and lease cash exceeding gross profit. CoreWeave is research-only; admissible weight is zero.

26. Samsara - IOT, NYSE; United States

Business-quality prior: Emerging and unproven, under a total-loss assumption - RP-3 (legacy: Speculative Option). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E0 is a hard failure. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Speculative Option. Samsara is a connected-operations vertical workflow and telemetry challenger. The causal thesis is that integrated fleet, safety, and asset data can become a system of record and support additional modules. The hard question is whether hardware subsidies, competition, and stock-based compensation allow that product value to become fully diluted cash flow per share. The company's quarterly-results page holds its fiscal 2027 filings and shareholder materials.

The provisional screen is A3/B2/C2/D1/E0, EC 58%, S 100%, and Wmax 0.75% under the speculative cap. The company remains below the 60% EC threshold, D shows a per-share capture problem, and E fails. Valuation requires durable net retention after normalizing seats and devices, as well as fully diluted owner cash rather than adjusted headline measures. Failure may come through commoditized hardware, platform bundling, or stock-compensation leakage. Needed evidence includes cohort retention, device subsidies, competitive wins and losses, and fully diluted owner earnings. The 36-month falsifier is dollar retention below 105%, gross-margin decline greater than 500 basis points, or share count growing faster than free cash flow per share. It is speculative research, not an entry; admissible weight is zero.

27. Symbotic - SYM, Nasdaq; United States

Business-quality prior: Emerging and unproven, under a total-loss assumption - RP-3 (legacy: Speculative Option). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E0 is a hard failure. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Speculative Option. Symbotic is a warehouse-automation bottleneck with concentrated anchor customers. The causal thesis is that automation can lower fulfillment cost and create a repeatable installed system. The policy treats backlog and revenue with skepticism because project accounting, installation milestones, cancellation rights, working capital, customer concentration, and related parties determine actual cash capture. The company's quarterly-results page is the evidence pointer for 2026 filings and operating updates.

The provisional screen is A3/B1/C1/D1/E0, EC 52%, S 100%, and Wmax 0.75% under the speculative cap. EC is well below 60%, and the B-E pillars expose several hard weaknesses. Backlog is neither revenue nor cash, so valuation must use milestone margins, site-vintage return, and cumulative project cash. Failure routes include deployment delays, cost overruns, concentrated contracts, throughput shortfalls, and related-party complexity. Research needs cancellation rights, warranty obligations, site-vintage returns, a project cash bridge, and a complete related-party map. The 36-month falsifier is installed sites missing contractual throughput or cumulative project cash remaining negative despite reported revenue growth. Evidence gaps are hard and current admissible weight is zero.

28. Nu Holdings - NU, NYSE; Cayman Islands/Brazil

Business-quality prior: Emerging and unproven, under a total-loss assumption - RP-3 (legacy: Speculative Option). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E0 is a hard failure. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Speculative Option. Nu Holdings is the digital deposit, credit, and payments franchise in an emerging market. The causal thesis is that low servicing cost and strong engagement can broaden access and create attractive customer economics. The company has not yet demonstrated that underwriting, funding,

governance, foreign exchange, and per-share capture can survive a full Brazilian credit cycle. The official results center is the evidence pointer for 2026 regulatory and SEC filings.

As a principal balance-sheet financial, Nu uses a battery: deposit beta 2, full-cycle credit 1, and capital 2. Its broader provisional screen is A4/B2/C2/D2/E0, EC 57%, S 100%, and Wmax 0.75% under the speculative cap. The low full-cycle credit score and sub-60 EC prevent the strong demand score from becoming admission. Valuation must translate customer growth into loss-adjusted, fully diluted US-dollar cash per share. Failure may arrive through credit losses, a funding shock, regulation, currency depreciation, or governance leakage. Needed evidence includes credit-vintage curves, deposit beta, regulatory capital, and cohort contribution. The 36-month falsifier is loss-adjusted return on equity below the cost of capital while deposits reprice and US-dollar free cash flow per share stagnates. The full-cycle record is too short; admissible weight is zero.

29. Sea Limited - SE, NYSE; Singapore/Southeast Asia

Business-quality prior: Emerging and unproven, under a total-loss assumption - RP-3 (legacy: Speculative Option). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E1 is below the ordinary admission minimum. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Speculative Option. Sea combines emerging-market commerce, payments, and digital entertainment. The causal thesis is that logistics density and payments can reinforce commerce, while entertainment may help fund the ecosystem. The countertheme is that subsidy cycles, competition, credit losses, and cross-segment capital allocation can consume value even when gross merchandise volume or users grow. Sea's investor portal is the evidence pointer for 2026 filings and results.

The provisional screen is A3/B2/C2/D2/E1, EC 58%, S 100%, and Wmax 0.75% under the speculative cap. EC below 60% makes it research-only even before valuation. A consolidated sales multiple is inappropriate; valuation needs a sum of the parts with normalized commerce subsidies, fintech credit, gaming cash, and fully diluted US-dollar per-share capture. Failure routes include renewed customer-acquisition spending, credit loss, gaming decline, and governance weakness. Research needs country cohorts, fulfillment-unit economics, credit vintages, and dilution. The 36-month falsifier is commerce contribution margin becoming negative again without share gain, or US-dollar free cash flow per share failing to grow. Sea is a speculative option only; admissible weight is zero.

30. BYD - 1211, HKEX / 002594, Shenzhen; China

Business-quality prior: Emerging and unproven, under a total-loss assumption - RP-3 (legacy: Speculative Option). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E1 is below the ordinary admission minimum. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Speculative Option. BYD is a vertically integrated electric-vehicle, battery, and energy-hardware challenger. The causal thesis is that manufacturing integration and scale can lower cost and support global expansion. The opposing mechanism is equally important: brutal price competition, capital intensity, export barriers, tariffs, dealer inventory, and state or governance leakage can destroy returns despite rapid unit growth. BYD's investor page is the evidence pointer for its 2026 interim and exchange filings.

The provisional screen is A3/B2/C2/D2/E1, EC 59%, S 100%, and Wmax 0.75% under the speculative cap. The name remains below the 60% research threshold and carries the maximum severity assumption. Valuation must normalize vehicle pricing, subsidies, dealer inventory, capital expenditure and depreciation, utilization, tariffs, and the economics of overseas plants. Unit growth alone is not owner value. Failure may come through a prolonged price war, trade barriers, a technology or safety recall, or governance and state leakage. Research needs model-level margins, dealer inventory, overseas-capital returns, related-party analysis, and US-dollar per-share cash. The 36-month falsifier is auto gross margin remaining below cost-of-capital economics while capacity

utilization worsens. Emerging-market, state, governance, and cycle gates remain incomplete; admissible weight is zero.

Avoid/Counterexamples - appealing narratives that fail the gates

The twelve Avoid/Counterexamples prevent the map from becoming a collection of admired winners. Each company expresses a theme that might sound attractive in isolation, then demonstrates why demand, industry structure, payer quality, capital economics, governance, valuation, title, leverage, or per-share capture must remain non-compensating. The entries also give demanding falsifiers so that "Avoid" does not become a permanent opinion immune to evidence. Policy ceiling and current admissible weight are zero in every case.

31. Palantir Technologies - PLTR, Nasdaq; United States

Business-quality prior: Negative/counterexample - RP-4 (legacy: Avoid/Counterexample). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E0 is a hard failure. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Avoid/Counterexample. Palantir is the admired mission-critical analytics and AI-software counterexample to the vertical-workflow thesis. The thesis offered by the market is that deployments compound through an ontology, deep integration, and switching costs. The policy asks a different question: after procurement concentration, competition, stock-based compensation, and the market's expectations, does value reach each diluted share at a conservative return? The company's quarterly-results page is the evidence pointer for 2026 results and fully diluted share data.

The provisional screen is A4/B3/C3/D1/E0, EC 68%, S 80%, and Wmax zero because E is a hard gate. Strong apparent demand does not compensate for weak per-share capture or a failed conservative return. Even high growth can produce a poor outcome when the starting price and dilution are demanding. Failure paths are procurement concentration, commoditization, dilution, and price compression. Research would need customer-cohort cash returns, stock-compensation-normalized free cash flow per share, contract concentration, and a reverse DCF. The 36-month falsifier for the Avoid classification is fully diluted free cash flow per share compounding above 20% while a conservative expected return clears the comparator by at least 1 percentage point. Unless that demanding combination appears, the classification remains. Palantir is the explicit admired-consensus arithmetic failure; policy ceiling and admissible weight are zero.

32. Intel - INTC, Nasdaq; United States

Business-quality prior: Negative/counterexample - RP-4 (legacy: Avoid/Counterexample). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E1 is below the ordinary admission minimum. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Avoid/Counterexample. Intel is the semiconductor-foundry buildout and proposed harvest counterexample. The offered thesis is that state support, manufacturing investment, and capacity rationalization can restore foundry economics. The policy's two-stage harvest discipline requires evidence of capacity exit, closed financing windows, survivor balance-sheet strength, disciplined replacement capital, and shareholder cash returns. State support and low reported book value do not demonstrate that harvest. Intel's filings page is the evidence pointer for continuing manufacturing investment and government-linked support.

The provisional screen is A2/B1/C0/D2/E1, EC 72%, S 100%, and Wmax zero. Fab or book value is not owner value without competitive yields, external customers, and cash returns after sustaining capital expenditure. Failure routes are further node delay, inability to win external customers, subsidy dependence, and persistent negative free cash flow. Required research includes external-foundry customer cohorts, capital expenditure versus depreciation, capacity exits, yields, and subsidy terms. The 36-month falsifier for Avoid requires two independent external customers scaling on leading nodes, capital expenditure divided by depreciation below 1, and segment

cash return above the cost of capital. Intel therefore fails the harvest rule: confirmed net capacity exit and shareholder cash returns are absent while the state backstop continues. Its ceiling and admissible weight are zero.

33. Boeing - BA, NYSE; United States

Business-quality prior: Negative/counterexample - RP-4 (legacy: Avoid/Counterexample). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E1 is below the ordinary admission minimum. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Avoid/Counterexample. Boeing is the counterexample to a seemingly attractive installed-base aftermarket and defense oligopoly. The offered thesis is that a duopoly backlog, production recovery, and services will restore cash generation. The policy refuses to treat backlog or customer advances as owner cash when safety, certification, manufacturing quality, supplier finance, leverage, fixed-price defense programs, and rework remain unresolved. Boeing's SEC filings page is the evidence pointer for safety, certification, cash, and contract risks.

The provisional screen is A4/B2/C0/D0/E1, EC 80%, S 100%, and Wmax zero. Demand can be strong while capital economics and governance fail completely; this is exactly why the five pillars do not add into an average. Valuation must normalize rework, customer advances, debt, supplier support, and program charges. Failure routes include another safety or certification event, a deficient quality culture, defense overruns, and leverage. Research would need independent safety evidence, unit cash by program, supplier-finance analysis, and fixed-price exposure. The 36-month falsifier for Avoid is independently verified quality stabilization plus positive cumulative free cash flow excluding advances and no new material program charge. Until that evidence exists, governance and capital gates fail; ceiling and admissible weight are zero.

34. Coinbase - COIN, Nasdaq; United States

Business-quality prior: Negative/counterexample - RP-4 (legacy: Avoid/Counterexample). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E0 is a hard failure. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Avoid/Counterexample. Coinbase is the crypto exchange and custody-rail counterexample. The offered thesis is that regulated scale, custody, and institutional adoption can make the company durable financial infrastructure. The policy asks whether the business can earn owner cash through a full crypto cycle while customer assets retain clear legal title and custody integrity. Coinbase's SEC filings page is the evidence pointer for transaction volatility, custody, and regulatory exposure.

The provisional screen is A2/B1/C2/D2/E0, EC 70%, S 100%, and Wmax zero. Valuation must normalize token prices, trading volumes, take rates, and stablecoin economics rather than capitalize cycle-peak revenue. Failure routes include fee compression, adverse regulation, custody loss, title uncertainty, and a token bear market. Required evidence includes through-cycle subscription margin, asset segregation and title, regulatory capital, and customer concentration. The 36-month falsifier for Avoid is non-transaction owner earnings covering fixed costs across a full crypto drawdown while legal title is independently validated. Cycle, title, and E gates currently fail. No infrastructure narrative changes the zero policy ceiling or admissible weight.

35. PG&E - PCG, NYSE; United States

Business-quality prior: Negative/counterexample - RP-4 (legacy: Avoid/Counterexample). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E1 is below the ordinary admission minimum. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Avoid/Counterexample. PG&E is the regulated-grid and resilience counterexample. The offered thesis is that necessary hardening and grid capital expenditure enter rate base and earn an allowed return. The policy distinguishes an allowed return from realized owner cash after wildfire mitigation, financing, insurance,

political disallowance, bill-affordability constraints, and catastrophic liability. PG&E's quarterly-results page is the evidence pointer for regulatory filings, wildfire exposure, capital plans, and financing.

The provisional screen is A4/B1/C1/D1/E1, EC 76%, S 100%, and Wmax zero. Necessary demand cannot compensate for a weak payer and liability structure. Valuation needs a rate-base-to-cash bridge, funding and dilution analysis, and probability-weighted catastrophic and regulatory outcomes. Failure routes are another catastrophic liability, an adverse rate case, costly financing, or material dilution. Required evidence includes insurance and recovery, hardening outcomes, bill affordability, and realized cash return on invested capital. The 36-month falsifier for Avoid is no catastrophic loss, a funding plan without material dilution, and realized cash return reaching the allowed return. Payer, liability, and capture gates fail; ceiling and admissible weight are zero.

36. Plug Power - PLUG, Nasdaq; United States

Business-quality prior: Negative/counterexample - RP-4 (legacy: Avoid/Counterexample). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E0 is a hard failure. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Avoid/Counterexample. Plug Power is the hydrogen, energy-abundance, and capital-bottleneck counterexample. The offered thesis is that scale and subsidies lower the cost of green hydrogen and unlock a large market. The policy treats addressable market and subsidies as hypotheses, not owner cash. Plant utilization, unsubsidized contribution, subsidy collectibility, financing, covenants, and dilution determine whether the business has economic value. Plug Power's SEC filings page is the evidence pointer for liquidity, losses, and financing requirements.

The provisional screen is A1/B1/C0/D1/E0, EC 68%, S 100%, and Wmax zero. Valuation must begin with plant-level cash cost and fully financed capital expenditure, not a sales multiple or theoretical market size. Failure routes include inability to fund operations, plant-execution problems, customer concentration, and negative unit economics. Required evidence includes plant utilization, subsidy collection, debt covenants, liquidity, and dilution. The 36-month falsifier for Avoid is two plants sustaining positive unsubsidized contribution while consolidated operating cash becomes positive without equity issuance. Capital and evidence gates fail. The policy ceiling and admissible weight remain zero.

37. Nike - NKE, NYSE; United States

Business-quality prior: Negative/counterexample - RP-4 (legacy: Avoid/Counterexample). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E1 is below the ordinary admission minimum. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Avoid/Counterexample. Nike is the global-consumer-brand counterexample. The offered thesis is that brand heat, product innovation, and a repaired direct and wholesale channel can restore premium growth. The policy requires cohort desire and full-price unit evidence rather than relying on brand history, price increases, or a product-reset narrative. Nike's results and reports page is the evidence pointer for fiscal filings documenting the product and channel reset.

The provisional screen is A2/B2/C3/D3/E1, EC 75%, S 60%, and Wmax zero because E1 fails admission. This corrects any implication that A is the hard failure: A2 meets the minimum, while E does not. Valuation must separate price and mix from units and account for inventory aging, promotions, customer-acquisition cost, wholesale support, and marketing intensity. Failure routes include cultural-relevance loss, channel conflict, promotions, and weakness in China. Needed evidence includes purchase-frequency cohorts, full-price sell-through, inventory aging, and customer-acquisition economics. The 36-month falsifier for Avoid is two years of full-price unit growth, stable gross margin, and share recovery without rising marketing intensity. Demand evidence and the conservative return remain insufficient; ceiling and admissible weight are zero.

38. Teladoc Health - TDOC, NYSE; United States

Business-quality prior: Negative/counterexample - RP-4 (legacy: Avoid/Counterexample). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E1 is below the ordinary admission minimum. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Avoid/Counterexample. Teladoc is the healthcare-access and platform counterexample. The offered thesis is that virtual care and chronic management gain scale and density. The policy requires evidence that payer savings, clinical outcomes, retention, and utilization create fully diluted owner cash. Acquired goodwill, adjusted EBITDA, or access to a large patient pool cannot substitute for those economics. Teladoc's quarterly-results page is the evidence pointer for segment, impairment, and customer-economics disclosures.

The provisional screen is A2/B1/C1/D1/E1, EC 72%, S 100%, and Wmax zero. Weak payer evidence, uncertain capital economics, acquisition impairment, and stock-based compensation all matter. Failure routes are employer churn, commoditization, continued acquisition impairment, and failure to prove payer value. Required research includes retention cohorts, utilization-to-outcome links, independent payer-savings evidence, and stock-compensation-normalized free cash flow. The 36-month falsifier for Avoid is independently verified clinical savings and retention above 95% producing positive diluted free cash flow per share. Until that combination appears, payer and capture gates fail; ceiling and admissible weight are zero.

39. Flagstar Financial - FLG, NYSE; United States

Business-quality prior: Negative/counterexample - RP-4 (legacy: Avoid/Counterexample). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E1 is below the ordinary admission minimum. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Avoid/Counterexample. Flagstar Financial is the deposit and commercial-real-estate balance-sheet counterexample. The offered thesis is that restructuring and deposit stabilization can normalize returns. The policy rejects a simple price-to-book argument because loan quality, funding stability, uninsured deposits, capital, supervision, and dilution determine whether the reported book belongs to shareholders. The quarterly-results page is the evidence pointer for capital, deposits, credit, and commercial-real-estate exposure.

The class battery is deposit beta 1, full-cycle credit 0, and capital 2. The broader provisional screen is A2/B0/C0/D2/E1, EC 78%, S 100%, and Wmax zero. This is the blocking direction of the financial battery in action: adequate disclosure coverage and some capital evidence cannot average away failed credit and industry-structure pillars. Valuation is meaningless without normalized losses, funding, and stressed capital. Failure routes include commercial-real-estate losses, uninsured-deposit flight, a capital raise, and supervisory constraints. Research needs loan vintages, criticized assets, deposit cohorts, and stressed CET1. The 36-month falsifier for Avoid is normalization of NPLs and charge-offs, persistence of core deposits, and return on tangible common equity above its cost without dilution for eight quarters. The class battery blocks admission; ceiling and admissible weight are zero.

40. Bausch Health - BHC, NYSE/TSX; Canada

Business-quality prior: Negative/counterexample - RP-4 (legacy: Avoid/Counterexample). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E1 is below the ordinary admission minimum. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Avoid/Counterexample. Bausch Health is the legacy roll-up and healthcare-franchise counterexample. The offered thesis is that asset separation and debt reduction can reveal valuable franchises. The policy starts with the enterprise-value and legal waterfall because debt maturity, refinancing, litigation, organic

erosion, and entity-level claims can absorb value before it reaches common shareholders. Bausch's SEC filings page is the evidence pointer for leverage, separation, legal, and operating risks.

The provisional screen is A2/B1/C0/D0/E1, EC 75%, S 100%, and Wmax zero. Reported brand EBITDA does not compensate for failed capital economics or governance and capture. Valuation requires debt by entity, covenants, contingent liabilities, separation mechanics, and an organic cash-flow bridge. Failure routes are refinancing, litigation, continuing organic decline, and governance weakness. The 36-month falsifier for Avoid is net leverage below 3 times funded from retained cash, a clean separation, and positive organic free cash flow per share. Leverage and roll-up gates fail; ceiling and admissible weight are zero.

41. First Quantum Minerals - FM, TSX; Canada/Zambia-Panama exposure

Business-quality prior: Negative/counterexample - RP-4 (legacy: Avoid/Counterexample). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E1 is below the ordinary admission minimum. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Avoid/Counterexample. First Quantum Minerals is the copper-scarcity and natural-resource counterexample. The offered thesis is that electrification increases copper demand and scarcity rents. The policy requires mine-level economics after depletion, sustaining capital, cost inflation, tax, closure, legal title, and jurisdiction risk. A high spot copper price or net asset value based on it does not answer those questions. The company's financial-reporting page is the evidence pointer for operations, debt, and jurisdiction matters including Cobre Panama.

The provisional screen is A3/B1/C1/D1/E1, EC 73%, S 100%, and Wmax zero. Valuation must stress conservative copper prices and incorporate reserves, grade, sustaining capital, closure obligations, debt, tax, and expropriation or title outcomes. Failure routes are jurisdiction and title impairment, leverage, cost inflation, and a commodity downcycle. Research needs mine-level reserve economics, legal outcomes, the debt ladder, and a full copper-price stress. The 36-month falsifier for Avoid is a durable Panama resolution combined with balance-sheet survival at a conservative copper price and positive free cash flow after sustaining capital. Title and cycle gates fail; ceiling and admissible weight are zero.

42. Adani Enterprises - ADANIENT, NSE; India

Business-quality prior: Negative/counterexample - RP-4 (legacy: Avoid/Counterexample). **Security-return conclusion:** Current security-return gate fails and admission is blocked: E0 is a hard failure. **Portfolio-role conclusion:** No live portfolio role; current admissible weight is zero; research-only; no trade authorization.

Classification: Avoid/Counterexample. Adani Enterprises is the emerging-market infrastructure and serial-incubation counterexample. The offered thesis is that privileged execution across airports, energy, logistics, and other projects can capture India's buildout. The policy does not apply a conglomerate growth multiple to a complex structure. It requires entity-by-entity debt, guarantees, related-party relationships, payer contracts, ownership, auditor evidence, cash waterfalls, and US-dollar per-share capture. Adani's investor downloads page is the evidence pointer for exchange filings, annual reports, and its governance framework.

The provisional screen is A3/B1/C1/D0/E0, EC 55%, S 100%, and Wmax zero. EC below 60% makes the name research-only, and D0 and E0 are hard failures. Failure routes include governance leakage, refinancing, state dependence, project overruns, guarantees, and cash trapped away from minority shareholders. Research needs an independently verified ownership and guarantee map, cash waterfalls, project payer contracts, debt by entity, and auditor evidence. The 36-month falsifier for Avoid is a simplified audited structure, declining consolidated leverage, no related-party leakage, and positive US-dollar free cash flow per share across two years. Governance and evidence coverage fail; ceiling and admissible weight are zero.

Reading across business-model classes

The map is meant to be read laterally as well as company by company. A positive representation shows how a model might work; an Avoid or blocked Watch representation shows how a similar story can fail. The pairing is part of the research design. It makes the investor compare mechanisms rather than accumulating unrelated company narratives.

In ratings, audit, testing, inspection, and certification assurance, S&P Global and Moody's provide the listed-company ratings cases. Both retain binary ratings-liability and licensing stress. The absence of a listed global audit partnership is structural rather than an investment conclusion: the Big Four audit networks are private partnerships. Their absence cannot be used to claim that audit is an unattractive model, nor can the ratings companies be treated as perfect substitutes for audit economics.

In installed-base aftermarket, GE Aerospace is the positive research representation and Boeing is the counterexample. Both have large installed fleets and powerful service narratives. The difference the research must test is whether safety, warranty, program accounting, working capital, supplier finance, and cumulative unit cash permit the installed base to create shareholder value. An installed base is not itself a moat if liabilities or capital claims absorb the annuity.

In exchanges, clearing, and payment rails, CME, Visa, and LSEG are the positive research representations, while Coinbase is the counterexample. Two-sided liquidity, network density, clearing integrity, take rates, regulation, title, and stress capital are the common questions. The class contains different economic mechanisms: Visa is a four-party acceptance and data rail; CME monetizes contract liquidity and clearing; LSEG combines infrastructure with data and workflow; Coinbase depends much more heavily on token cycles, custody title, and transaction economics. The label "rail" does not make those risks equivalent.

In vertical workflow and data, RELX, Microsoft, and Constellation Software are positive research representations, while Palantir is the Avoid counterexample on valuation and per-share capture. The central question is whether the product remains a record-to-action system with measurable retention and customer value after AI changes seats, interfaces, content scarcity, and compute cost. AI can strengthen workflow, bypass it, or transfer savings to customers. Brand names and user enthusiasm cannot decide among those outcomes.

Fastenal is the positive reliability-distribution representation. The search for a robust listed counterexample or second positive was incomplete, and that incompleteness is not a negative sector call. Grainger, WESCO, and MSC Industrial were named and rejected from this edition because the necessary current valuation, leverage and integration, organic-growth, and cycle evidence was not completed. The correct conclusion is "search incomplete," not "Fastenal is uniquely attractive" and not "the other distributors are unattractive."

The semiconductor, foundry, cloud, and equipment layer is deliberately split. TSMC and ASML are positive Core research representations; CoreWeave is a merchant-cloud Speculative Option; Intel is the Avoid counterexample; and NVIDIA is an admired but E-failing Watch name. Microsoft and Alphabet belong to the internal-hyperscaler layer rather than being treated as equivalent to an equipment vendor or merchant capacity provider. This separation forces the research to identify who funds capacity, who bears utilization and useful-life risk, who owns the customer relationship, and where AI cash profit ultimately appears. The current comparator factor bar sets active additions to zero for the named AI and power-sensitive rows, so the richness of the research map cannot be used to enlarge an already breached factor exposure.

In power management, grid, and regulated capital expenditure, Schneider and Eaton are the positive research representations; PG&E and Plug Power are counterexamples. The common narrative is physical scarcity and necessary investment. The differentiating evidence is contract or rate recovery, payer collectibility, capital intensity, working capital, liability, subsidies, and cash per diluted share. A mandatory project or public budget is not collectible revenue, and an allowed return is not the same as realized cash return.

Ferrari represents consumer scarcity and brands on the positive-research side, but it remains Watchlist/Price Sensitive and fails E. Nike is the Avoid counterexample because historical brand strength has not yet been replaced by current cohort desire, full-price unit growth, and a conservative valuation case. The pairing shows why the policy does not confuse an admired brand with an investable security. Ferrari's controlled scarcity can still be overcapitalized; Nike's brand legacy cannot substitute for current demand evidence.

Eli Lilly and Novo Nordisk represent funded healthcare innovation and manufacturing. Teladoc is the services and access counterexample. The common discipline is the payer gate: patient need, clinical promise, or addressable population does not establish reimbursement, net price, collectibility, or owner value. Lilly and Novo also cannot assume terminal growth through a patent or exclusivity cliff without a replacement pipeline. Teladoc cannot substitute adjusted EBITDA or goodwill for verified clinical savings, retention, and diluted free cash flow.

Rheinmetall is the positive research representation for defense and state procurement, with Boeing also acting as a counterexample through safety and fixed-price program risk. In both cases, government demand must be traced from announced budget to appropriation, funded contract, cancellation terms, working capital, program cost, and shareholder cash. The defense theme remains exposed to monopsony, political control, cost overruns, and a peace dividend.

No full natural-resource or harvest candidate qualifies. That is a supported absence, but only in the limited sense that named searches failed their gates; it is not a permanent judgment against the class. Micron or the memory cycle, Maersk or container shipping, and Transocean or offshore services were examined and rejected from this edition because financing-window closure, net capacity exit, balance-sheet survivability, and shareholder cash-return gates were not all evidenced. First Quantum is the natural-resource counterexample, and Intel also demonstrates the absence of a confirmed semiconductor harvest. A price collapse, consolidation financed by debt, announced capacity discipline, or state support is not enough. The search remains incomplete and must be rerun.

Constellation Software is the positive serial-acquirer representation. Bausch Health is the direct roll-up counterexample, while Adani's complex incubation model also tests acquisition, project, and conglomerate governance. The comparison turns on organic bridges, acquisition-vintage returns, contingent consideration, leverage, guarantees, related parties, and value reaching each share. Revenue assembled through transactions is not evidence of successful reinvestment.

MercadoLibre is the more mature emerging-market platform and per-share-capture research representation, while Nu, Sea, and BYD are Speculative Options and Adani is the Avoid counterexample. Legal domicile must not erase economic geography. The common research burden includes country cohorts, funding, credit, foreign exchange, regulation, related parties, state influence, capital controls, and US-dollar value per diluted share. A growing local platform can be socially useful and operationally strong while providing poor foreign-shareholder capture.

Balance-sheet financials are represented rather than excluded by a quality-duration screen. Berkshire, Chubb, and DBS are Core research candidates; Nu is a Speculative Option; and Flagstar is the failed bank-battery case. Favorable battery scores for the Core names remain provisional and capped at 2 until named reserve, deposit, credit, and capital exhibits are attached. The asset-manager search is incomplete, not a negative call. BlackRock, T. Rowe Price, and Invesco were named and rejected from this edition because current net-flow, fee-compression, client-concentration, capital, and valuation work was not completed. Future refreshes must keep the class visible rather than allowing the existing quality canon to crowd it out.

How the map should be used

The map should allocate research attention and make uncertainty explicit. A reader can begin with a model class, compare its positive and negative cases, identify the causal mechanism, and assemble the missing evidence file. The five-pillar screen provides a common language, but the business-model overlays determine what evidence is actually required. Insurance needs reserve triangles, float cost, and stressed capital. Banks need deposit cohorts, credit vintages, nonperforming-loan and loss evidence, and stressed CET1. AI capacity needs utilization, rental prices, customer cash returns, funding loops, and equipment retirement curves. State procurement needs appropriations, enforceable contracts, cancellation rights, and cash conversion. Healthcare needs indication-level net price, payer access, persistence, exclusivity, and a replacement pipeline.

The map should also support falsification. Every working thesis is paired with a 36-month observable. A falsifier is not a forecast that must occur and not an automatic sell order. It is a preregistered point at which the analyst must stop telling the original story and conduct a fresh evidence review. For Avoid names, the falsifier prevents a negative label from becoming dogma: sufficiently strong evidence can force reclassification. For Core and Watch names, the falsifier prevents a prestigious business from retaining research priority solely because its past reputation remains attractive.

The map should be refreshed rather than passively carried forward. Twelve months after the last complete refresh, a classification expires to Unclassified Research Note. The annual refresh must report a counter-canon screen and a non-US screen, and any claimed class absence must name at least three examined and rejected candidates together with the failed gates. This is a bias control. It does not require admitting a weak company merely to diversify the list; it requires showing that neglected classes and regions were actually searched.

The map should remain subordinate to the full policy architecture. Before any security could move beyond research, it would need a current complete evidence and valuation review, an adversarial corroboration pass, investor facts, liability and liquidity fit, tax and legal review where relevant, portfolio look-through, factor headroom, and owner approval under a separate implementation process. None of that work is completed by this chapter. At most eight names may be maintained as full live underwrites by one operator, subject to the policy's hours and active-program viability tests. The other entries are a queue or annual-check file, not forty-two continuously updated investment theses.

How the map must not be used

The classifications must not be translated into buy, hold, sell, rank, target weight, or expected-return conclusions. Core does not mean highest expected return. Watch does not mean buy after a price decline. Speculative Option does not mean that 0.75% may be invested. Avoid does not mean short, sell, or permanent condemnation. Wmax must not be copied into an allocation sheet. It is a best-case policy ceiling produced from provisional assumptions, and every current admissible weight is zero.

The map must not be used as proof of a theme. Orders, backlogs, market size, installed base, regulation, patent protection, state budgets, user growth, and technological leadership are inputs to research. None proves that value reaches each diluted share. Artificial-intelligence spending may accrue to customers, labor, equipment suppliers, energy providers, or lenders. Defense budgets may be delayed, cancelled, or captured by working capital and cost overruns. Electrification spending may invite capacity and regulation. Healthcare innovation may face net-price and access resets. Emerging-market growth may be offset by currency, governance, or state leakage.

The dated company releases must not be treated as independently verified facts beyond their narrow role. They are evidence pointers to management or regulatory disclosure available by August 22, 2026. The map does not update them beyond the evidence cutoff, and it does not claim that a link proves the thesis. A later reader must extract the relevant disclosure, reconcile it against external and historical evidence, and lower a score to zero if credible sources conflict.

The scores must not be averaged into a quality ranking. A company with A4 and C4 can still be blocked by D0 or E0. A high EC percentage cannot rescue a failed gate. A low valuation cannot rescue weak capital economics, governance, title, or payer quality. A strong narrative cannot lower S below the policy severity floor. Liquidity and factor penalties retain full force, and a hard factor-cap breach arithmetically sets allowed loss to zero.

Finally, the map must not conceal its own historical and geographic bias. It surfaces many businesses admired during 2010-2025. That may reflect genuinely durable models, but it may also load the research queue with long-duration expectations at the wrong point in a cycle. The proper response is not to reverse every quality judgment mechanically. It is to keep U3-2 open, run no-multiple-recovery and aggregate-derating tests elsewhere in the policy, search the counter-canon and non-US universe, and require per-share cash and conservative valuation evidence before any classification can have implementation significance.

Annex B - Detailed Operating Checklist

The one-page operating checklist in the original policy remains the compact operating checklist. This chapter is its explanation, not its replacement. The checklist is deliberately short because it is meant to be used at the point of decision. The fuller policy supplies the definitions, worksheets, legal and tax work, evidence files, and dispute record that sit behind it. This chapter explains why the ten items occur in their exact sequence, what evidence completes each item, where work must stop, and what must be handed to the next stage.

The sequence matters. It begins with the investor's facts, not with securities. It then protects survival, designs the strategic portfolio, exposes hidden concentrations, and only after those steps considers active investments or speculative options. It next creates the decision record, separates light recurring monitoring from deep annual governance, and finishes with blockers and stop rules. Reversing that sequence invites a familiar error: researching an attractive company first and inventing a portfolio rationale later.

The checklist also inherits the policy's standing boundary. It is a governance device, not a trade authorization. Completing a worksheet, passing a gate, or calculating a diagnostic maximum does not authorize an order, target allocation, entry price, message, account opening, hedge, or portfolio change. The policy itself states that every named company has a current admissible weight of zero until verified investor facts, a current complete evidence and valuation review, portfolio look-through, relevant tax and legal review, every applicable gate, and a separate implementation approval are complete. "Pass" therefore means eligible to move to the next governed handoff. It never means "buy."

The bounded checklist exists because the debate reached an adverse conclusion about one-person executability. Forty-two company dashboards, five pillars for every company, factor scoring, custody, taxes, and full evidence maintenance would amount to institutional work. The accepted operating version permits at most eight live full underwrites, with the actual number reduced in advance if time is inadequate. Deep exhibits remain in research files. Items outside the live set remain research notes or annual-check items and cannot borrow credibility from the compact checklist.

1. Annual facts

> **Annual facts:** *update assets, liabilities, taxes, currencies, income correlation, decision makers, succession, domicile; calculate FS at PVL base/+/-10/+/-20 and apply hysteresis.*

This is the first item because every later limit is conditional on the investor rather than universal. Investable assets, written as A^*, are assets available for the mandate after known taxes and transaction costs. PVL^* is the probability-weighted present value of liabilities plus an explicit safety margin. Funded status is $FS = A^*/PVL^*$. The four funded-state bands are below 1.00, 1.00-1.24, 1.25-1.74, and at least 1.75. They affect strategic and opportunity sleeves, the error budget, illiquidity, tracking error, and base issuer loss. Using stale inputs can therefore contaminate every downstream calculation.

"Update assets" means reconciling ownership, availability, tax basis where relevant, and deductions needed to reach A^*; a headline brokerage balance is not enough. Liabilities should be dated, probability weighted where uncertain, expressed in the currencies in which they will be paid, and assigned nominal or real characteristics. Tax obligations and fixed commitments must not be hidden inside a broad spending estimate. Income must be tested for durability and for correlation with portfolio losses. An apparently high income is weak protection if it is likely to decline during the same market shock. Decision makers, succession, domicile, and incapacity arrangements are investment facts because they determine who can act, under which law, and whether assets can actually fund liabilities.

The required evidence is a dated balance-sheet reconciliation; a liability cash-flow schedule; tax and commitment schedules; currency assignments; documented income sources and a reasoned correlation classification; legal domicile and citizenship facts relevant to ownership; and a current list of authorized decision makers, beneficiaries,

and successors. Every material assertion should be labeled fact, inference, assumption, forecast, classification, or unresolved. A repeated estimate does not become a fact merely because it appeared in last year's file.

The sensitivity calculation prevents false precision. Calculate FS at base PVL*, then with PVL* 10% lower, 10% higher, 20% lower, and 20% higher. If the +/-10% range crosses a funded-state boundary, use the lower state. Hysteresis then slows favorable changes but not adverse ones: promotion requires two consecutive annual confirmations, while verified deterioration such as death, incapacity, income loss, leverage, litigation, custody impairment, or newly inflexible spending can demote the state immediately. The asymmetry prevents a favorable estimate from granting extra risk before it proves durable.

Hypothetical example - not a forecast or allocation instruction. Suppose A* is 1.50 million and base PVL* is 1.00 million, so base FS is 1.50. If liabilities 10% higher produce PVL* = 1.10 million, FS falls to about 1.36. Both results remain in the 1.25-1.74 band. If instead base FS is 1.27 and a 10% liability increase reduces it below 1.25, the lower 1.00-1.24 state governs. A later favorable recalculation cannot promote the investor after one year; it needs two annual confirmations.

Stop if assets cannot be reconciled, liabilities lack dates or currencies, tax effects are material but unknown, decision authority is unclear, or the +/-10% sensitivity changes the band without the lower state being applied. The handoff from item 1 is a dated investor-facts pack and a provisional funded-state classification. It goes to the survival review. It does not go to security selection.

2. Survival

> **Survival:** *verify ILF floor/target, D36 or elected D24, fixed commitments, title, beneficiary/incapacity access, Tier-A second custodian and transfer test.*

The investor liquidity floor, or ILF, comes before return design because near-term liabilities cannot depend on selling risky assets at a favorable time. Its target is:

ILF target = D36 + Tax12 + Commit + H x max(0, essential spending beyond D36 - dependable income covering that later interval)

D36 is the amount needed for dated liabilities within 36 months. Tax12 covers known taxes due within 12 months. Commit covers fixed commitments, including known obligations that cannot safely be deferred. H is the applicable income-interruption horizon, but it applies only to essential spending beyond what D36 already funds; this avoids double counting. The hard floor is six months of uncovered essential spending plus twelve months of fixed commitments plus known taxes. The floor is not a source of risk capital: the rule is never to buy risk from the emergency floor.

D36 may become D24 only by an annual documented election when FS is at least 1.75 and spending flexibility is high. That is a narrow conditional permission, not an automatic surplus benefit. Flexibility means essential obligations can genuinely be reduced or deferred without impairing survival. A wish to reduce cash drag is not evidence of flexibility. The election also does not enlarge issuer loss budgets.

Verification covers both quantity and usability. ILF assets should be in the home-liability currency, short duration, and held with high title and custody integrity. Legal title must be demonstrable. Beneficiary and incapacity access must work in practice, not merely appear in an old form. Tier A requires a same-jurisdiction second custodian and an annual transfer test. The second custodian is intended to reduce operational dependence without introducing a second legal system. Tier B plain instruments in another legal system require tax, reporting, title, transfer, sanctions, and estate-situs review; Tier C entities, trusts, and complex wrappers default to zero until the full legal gauntlet passes. Those deeper tiers are not waived by completing the Tier-A line.

Evidence includes current statements, account registrations, custody and deposit-protection terms where applicable, beneficiary designations, powers or mandates for incapacity, a fixed-commitment schedule, tax

amounts, and a recorded transfer test showing that cash or a plain security could move through the intended path. Outsourcing setup or legal review is permitted, but the operator must retain the evidence and outcome.

Hypothetical example - not a live instruction. An investor has 180,000 of liabilities due within 36 months, 30,000 of taxes due within 12 months, 20,000 of fixed commitments, and 90,000 of essential spending beyond D36 that is offset by 50,000 of dependable income over the applicable horizon. With H represented as one applicable horizon unit, the additional uncovered amount is 40,000, producing a target of 270,000. The example illustrates the de-overlap: spending already included in D36 is not counted again. The separately calculated hard floor still applies if it is higher.

Stop if the ILF is below its hard floor, if title or access is uncertain, if the D24 election lacks both required conditions, or if the Tier-A second custodian and transfer test are absent. A failed transfer test is not paperwork noise; it is survival evidence. The handoff is a verified ILF and access certificate for strategic design. No active admission should proceed while survival is unresolved.

3. Strategic design

> **Strategic design:** *name comparator and matched nominal/real liability sleeves; record bands, MDB, EB and illiquid ceiling. No placeholder comparator after 12 months.*

Once facts and survival are established, the policy can define the portfolio that active work must beat. The comparator is a named, low-cost, implementable alternative, not a neutral abstraction. It carries issuers, factors, currencies, taxes, costs, and concentration of its own. Naming it prevents an active decision from being compared opportunistically with whichever benchmark looks easiest after the fact. A diagnostic proxy may organize preliminary work, but no placeholder comparator survives beyond 12 months.

Liability matching is separated from directional market views. Nominal liabilities are matched with appropriate nominal duration. Verified real liabilities may be matched with inflation-linked duration only when index basis, law, market integrity, title, tax, and currency fit are adequate. Liability-matched duration is reported separately from directional rates factor-equivalent weight, with basis risk made visible. The lower-gross-risk portfolio remains the crash null; verified inflation-linked matching is the specific instrument for real grind. Neither is a generic return forecast.

The recorded design must include approved bands and four distinct constraints. The mark-to-market drawdown budget, or MDB, is the maximum peak-to-trough loss the investor can endure without violating trait tolerance, ILF, or liability coverage. The error budget, or EB, limits the aggregate consequences of discretionary mistakes. The illiquid ceiling limits assets that cannot be converted to usable cash on the required schedule. The active tracking-error ceiling limits divergence from the comparator. Under the funded-state table, EB ranges are 0-3%, 3-5%, 5-8%, and 8-10% with a hard 12% cap across the four FS bands; illiquid ceilings are 0-5%, 0-8%, 5-12%, and 10-15% with a hard 20% cap; active tracking-error ceilings are 0-2%, 0-3%, 0-4%, and 0-5%. These are ceilings, not desired exposures.

The second axis still applies. At the same FS, near or inflexible liabilities, or labor and business income highly correlated with market losses, shorten MDB and EB by one state. This is a reminder that a ratio alone cannot represent capacity. Required evidence includes the exact comparator instrument or rules, its costs and tax treatment, a dated liability-matching schedule, approved allocation bands, scenario coverage, trait and liquidity analysis supporting MDB, the EB, the illiquidity inventory, and a tracking-error budget with any abandonment point.

Hypothetical example - not a portfolio recommendation. Two investors both have FS 1.50. One has flexible spending and stable income unrelated to markets. The other owns a cyclical business and faces a fixed liability in two years. Both begin in the same FS band, but the second investor applies the more conservative MDB and EB state. The example shows why equal wealth ratios do not imply equal risk capacity.

Common mistakes are treating a broad index as risk free, using the policy's diagnostic 60/40 proxy as if it were the investor's ratified comparator, counting liability-matched bonds as a speculative rates view, treating ceilings as targets, or allowing the comparator placeholder to roll indefinitely. Stop if no actual comparator can be named by month 12, if liability matching lacks currency/index/duration evidence, or if MDB, EB, and illiquidity are unverified. The handoff is the strategic baseline and its risk budgets, which become inputs to look-through.

4. Look-through

> **Look-through:** *archive top 20 constituents; bucket the rest; calculate issuer, sector, sovereign, currency, FEW/FSL and joint stresses; resolve four-quarter breaches by amendment.*

Look-through asks what the investor really owns after funds, wrappers, and labels are removed. The top 20 constituents of the actual comparator are scored individually because a small number of large holdings can dominate risk. The remainder is grouped by defensible sector proxies. The analysis then aggregates issuers, sectors, sovereigns, currencies, financing chains, counterparties, and common factors across both direct and embedded holdings.

Factor-equivalent weight, FEW_f = sum(weight_i x exposure score_if), estimates how much portfolio weight is economically sensitive to factor f. Factor stress loss, FSL_f = sum(weight_i x loss_if under the common scenario), estimates portfolio loss under the specified shared shock. Joint stresses matter because exposures that appear moderate alone can interact, such as AI capital spending and power availability. Scores are analyst assumptions unless tied to evidence; they need archived holdings, definitions, source dates, and reproducible calculations. The top-20 study is refreshed annually and on a material event, and the first year requires an independent second scoring pass after six months. Sensitivity perturbations and the capped or blended alternative belong in the research file.

Hypothetical example - not a statement about the current portfolio. A fund is 10% invested in Company X, and the archived scoring rubric gives X a 0.8 exposure score to a common factor. X contributes 8 percentage points to FEW. If the common scenario assigns a 50% loss to that exposed portion, X contributes 4 percentage points to FSL. Other holdings are added on the same basis. This is why two different funds can duplicate the same economic bet even when they hold different tickers.

A persistent hard-cap breach cannot become a renewable exception. Four consecutive quarters require a written choice: adopt an implementable capped or blended comparator, or raise the cap through a formal amendment. The amendment accepts and discloses risk; it does not make the risk disappear. The policy's current comparator work remains a certification blocker: the inherited proxy's AI FSL of 8.2% and joint AI-plus-power FSL of 9.0% exceeded the original 5% and 8% caps, but the proxy was not reproducible enough to ratify a cap. Therefore the original caps govern and active additions with material AI/power sensitivity remain zero until the actual comparator study is complete. This is an example of a stop condition produced by look-through, not a view on whether AI-related businesses are attractive.

Required evidence is the dated top-20 constituent file, sector buckets for the remainder, exposure and loss rubrics, issuer and cross-fund aggregation, sovereign and currency maps, individual FEW/FSL worksheets, joint scenarios, independent-review records, and a breach log in currency as well as percentages. Common mistakes include counting the same matched liability exposure as directional risk, assigning missing holdings a zero sensitivity, allowing index concentration to float the cap upward, or changing scenarios after seeing the result.

Stop active additions that would breach a hard cap. If a breach persists for four quarters, stop using exceptions and initiate the amendment choice. If scoring cannot be reproduced or maintained, the eventual default in item 10 is comparator plus ILF plus matching. The handoff is a certified headroom map. Without it, the active-admission formula has no valid factor multiplier.

5. Active admission

> **Active admission:** *all A-E>=2, EC>=60, no hard gate, reverse DCF/hurdle, S and AllowedLoss, derating budget, tax/cost, two disconfirmers for >3% diagnostic size.*

Only now does a specific active idea enter the sequence. The five non-compensating pillars are A, demand and lock-in; B, industry structure, cycle, and payer quality; C, capital economics; D, governance and per-share capture; and E, valuation and falsifiability. Every pillar must score at least 2. A strong demand story cannot offset a zero in governance, and a cheap-looking multiple cannot offset value-destructive capital economics. Class-specific overlays still apply, including reserve, float-cost, and capital evidence for insurers; deposit beta, full-cycle credit, non-performing loans, and capital for banks; and the relevant batteries for other models.

Evidence Coverage is EC = verified required fields/applicable required fields, weighted toward Tier 1 filings, contracts, and regulators. Below 60 is research-only. EC of 60-74 gives an epistemic multiplier of 0.50, 75-84 gives 0.75, and at least 85 gives 1.0. Missing evidence cannot be averaged away, and a conflict between Tier 1 and lower-tier material scores zero until resolved. EC constrains tolerable loss; it is never itself a buy signal.

Valuation is an admission hurdle. A reverse discounted-cash-flow analysis asks what operating and capital assumptions the current price already requires. The conservative expected return must clear the comparator under the policy's relevant hurdle. S is bear-case permanent-impairment severity, subject to anti-gaming floors: at least 40% for ordinary single equity; at least 60% for a levered financial, cyclical, or roll-up; and 80-100% for binary-license, pre-cash-flow, or flagged emerging-market exposure. A lower severity needs written evidence and independent challenge.

After every gate passes, loss capacity is calculated as AllowedLoss = BaseIssuerLoss(FS) x max(0.50, ECmult) x Lmult x Fmult. Base issuer loss is 0.25%, 0.50%, 0.75%, or 1.00% across the four FS bands. The liquidity multiplier is 1.0 for five days or less, 0.75 for 6-30 days, 0.50 for 31-365 days, and 0.25 beyond one year. The factor multiplier is 1.0 with full headroom, 0.75 when factor use exceeds 75%, 0.50 above 90%, and zero if adding would breach a hard cap. Maximum diagnostic weight is AllowedLoss/S. The 0.50 floor applies only to the epistemic multiplier. It never rescues illiquidity or a hard factor breach.

Hypothetical example - not a sizing recommendation. At FS at least 1.75, suppose a fully liquid ordinary equity has EC 80, full factor headroom, and severity 50%. BaseIssuerLoss is 1.00%; ECmult is 0.75; both other multipliers are 1.0. AllowedLoss is 0.75%, and the diagnostic maximum is 0.75%/50% = 1.5%. If the factor multiplier were zero, both AllowedLoss and diagnostic maximum would be zero regardless of the other numbers.

The derating budget is separate. Archive the current and conservative multiples or reverse-DCF outputs and calculate w x (1 - Mconservative/Mcurrent). Across active holdings, derating loss must remain no more than 50% of MDB after strategic-sleeve stress. Up to 75% requires a written amendment and independent challenge. An ex-post overshoot prompts a sizing post-mortem, not an automatic sale of an intact thesis. Tax and cost must be included in the comparison. For a diagnostic size above 3%, two documented attempts at disconfirmation are mandatory; this preserves the checklist's exact decision even though ordinary severity and loss ceilings often make such a size unreachable.

Stop on any pillar below 2, EC below 60, a hard gate, failure of the reverse-DCF hurdle, unresolved tax/cost effects, insufficient loss or derating budget, or zero factor headroom. Common mistakes are turning Wmax into a target, lowering S to manufacture room, treating quality as price, or confusing admission with approval. The handoff is a completed underwrite marked "non-buy until approved" for the separate decision and implementation authority.

6. Options

> **Options:** *total-loss assumption; per-name 0.25-0.75%; sleeve loss cap; no averaging down; 24-month harvest expiry.*

The options item follows ordinary active admission because it is a separate loss-bounded category, not a shortcut around the five pillars. "Option" here means speculative research optionality whose entire market value is assumed lost. Each name is normally limited to 0.25-0.75% of total portfolio market value, and the applicable speculative sleeve loss cap also binds. Across the four funded-state bands, speculative loss is capped at 0-2%, 2-3%, 3-5%, or preferred 5% with a hard 7% ceiling. Loss, not an optimistic notional, is the governed quantity.

No averaging down means a falling price does not itself justify increasing the exposure. New evidence may support a full re-underwrite, but that is a new governed decision, not the continuation of a failed thesis. Speculative and pre-cash-flow positions use $S = 100\%$; the per-name speculative cap binds before any larger diagnostic output from loss arithmetic.

The 24-month harvest expiry applies to two-stage capital-cycle candidates. Candidate status requires observable financing-side conditions such as issuance shutdown, industry capex below depreciation, a major announced exit or bankruptcy, and no sustaining state backstop. It permits only bounded speculative treatment. If exit and shareholder cash-return evidence do not arrive within 24 months, candidate status expires. It cannot be rolled merely because the narrative still sounds plausible. Normal sizing would require full harvest confirmation and ordinary active admission.

Hypothetical example - not an authorization. A research candidate priced at 0.50% of portfolio value is assumed to lose the full 0.50%. If the remaining sleeve loss capacity is only 0.30%, the idea does not fit even though it is below the 0.75% per-name ceiling. If the candidate reaches month 24 without the preregistered exit and cash-return evidence, its harvest status expires; a lower market price does not renew it.

Required evidence includes the option thesis, maximum-loss statement, sleeve loss ledger, financing and covenant map where relevant, preregistered milestones, the start and expiry dates, and the evidence needed for any future re-underwrite. Stop when full loss would breach the sleeve or EB, when ownership introduces unbounded or collateral-call exposure, when averaging down is the rationale, or when the 24-month harvest clock expires. The handoff is to the same decision record as any other idea, still marked non-buy until approved.

7. Decision record

> **Decision record:** *evidence timestamp/tier, causal thesis, countertheme, falsifier, event-anchored pacing, tax lots, action clock and explicit "non-buy until approved."*

The record is placed after admission and options because it captures the decision as it was actually framed before outcomes rewrite memory. Each material item needs a timestamp and evidence tier. The causal thesis should describe the mechanism by which value reaches each share, not merely a theme or addressable market. The countertheme states a credible alternative mechanism. A falsifier is an observable that would disconfirm the thesis; it is stronger than a vague promise to "monitor execution."

Event-anchored pacing, if later authorized, attaches tranches to filings or preregistered milestones and expires by day 180. If all gates and the conservative case clear and no information event is pending, full deployment is the statistical default. Behavioral pacing must record its expected opportunity cost. This rule does not authorize either full deployment or tranching. Tax lots are recorded because deferral is friction, not thesis. Tax deferral uses the applicable statutory rate, a fixed ten-year horizon, and the comparator replacement return. Two annual holds chiefly justified by deferral require an independent challenger. The action clock records the applicable response window: immediate freeze and resolution within 0-5 trading days for title, fraud, or de-licensing threats; a 30-day full re-underwrite and no add for a pillar zero, hard-cap issue, two missed milestones, or EC below 60; and the governed tax/cost clock for a valuation or expected-return decision.

Hypothetical example - not a live decision. A thesis says that installed-base service revenue will grow as equipment ages. The countertheme says independent substitutes will reduce service capture. The falsifier is a sustained six-quarter divergence between installed base and service revenue. The next filing is the first pacing event and day 180 is the outside limit. A price decline is neither an event milestone nor evidence.

Required evidence is the dated source list, fact/inference labels, approved thesis and countertheme, measurable falsifiers, tax-lot inventory, pacing events and expiry, and decision owner. Common mistakes include backdating a thesis, using management guidance as a falsifier, treating a calendar tranche as an evidence event, omitting tax lots, or writing "buy" before authority exists. Stop if the record cannot be completed prospectively. The handoff is a frozen decision packet for approval, monitoring, and later attribution.

8. Monthly/quarterly

> **Monthly/quarterly:** *monthly survival alerts; quarterly holdings/look-through and pillar-breaking evidence only. Do not rescore everything.*

Recurring maintenance is intentionally narrow. Monthly work covers ILF, custody and title alerts, and hard-gate events. Quarterly work reconciles holdings and look-through, checks issuer and factor headroom, reviews evidence capable of breaking a pillar, checks tax lots, and updates the exception log. The instruction not to rescore everything is a governance control: constant full rescoring would consume the hours budget, invite noise-driven trading, and make the one-person system unmaintainable.

This does not mean ignoring adverse evidence. A title defect, fraud allegation supported by credible evidence, de-licensing event, financing-covenant threat, pillar zero, hard-cap breach, two missed milestones, or EC below 60 triggers the relevant faster clock. Conversely, mark-to-market decline alone never forces sale, and valuation deterioration alone does not create a mechanical trim. Expected return below admission forbids an add and requires re-underwriting. Strategic sleeves can rebalance mechanically when they move more than 5 percentage points or more than 20% relative from an approved band, subject to the ILF floor and taxes, but a broken theme is not mechanically refilled.

Hypothetical example - not a live monitoring instruction. A holding falls 18% during a quarter, while its filings confirm the evidence fields and no hard gate changes. The quarterly process records the price and relevant valuation effect but does not rescore all five pillars solely because price moved. If the same filing reveals a license suspension, the hard-gate clock overrides the ordinary cadence immediately.

Evidence is a concise monthly survival log and quarterly reconciliation showing holdings, exposures, exceptions, tax lots, and only the evidence deltas that could change a pillar or gate. Common mistakes are turning monitoring into continuous research, overlooking an event because the quarter has not ended, or treating market movement as thesis evidence. The handoff is a clean exception and evidence-delta file for the annual governance review.

9. Annual governance

> **Annual governance:** *run the prospective 2x active-viability threshold and print the hours budget (about 40 governance hours plus 15-25 per live underwrite); attribute comparator-design value separately from company-selection value; audit ILF drag, exceptions, tax deferrals/challengers, diversifier start-date sensitivity and policy amendments.*

The annual review asks whether the apparatus earns its money and attention before another year of active capital is permitted. Expected annual active-program contribution is portfolio value x active sleeve x evidenced premium x evidenced hit rate. It must be at least twice fully loaded annual cost. Fully loaded cost includes research hours valued at the operator's opportunity cost, data, tax, and other implementation costs. Below the 2x threshold, the program runs in paper-journal mode with zero active capital. Paper mode preserves discipline and the ability to observe decisions without pretending that an uneconomic program should deploy money.

The hours budget must be printed, not hidden. The policy estimates about 40 recurring governance hours plus 15-25 hours for each live full underwrite. Eight names therefore imply roughly 160-240 hours a year before one-off comparator scoring and custody work. Eight is a maximum, not a quota. If the operator can fund only 100 hours, the live set must be reduced prospectively, potentially to zero, before research starts.

Attribution must split the value of comparator and liability-matching design from company selection. Otherwise good benchmark engineering can disguise poor security selection, or vice versa. The annual review also prices the ILF's real drag as insurance cost; audits all exceptions and their expiry; applies the tax-deferral challenger rule; tests diversifiers across rolling start dates, including inflation, deflation, crash, and calm; and records formal policy amendments. A company classification not refreshed within 12 months reverts to Unclassified Research Note. The annual map refresh also reports counter-canon and non-US screens and at least three named examined-and-rejected candidates for any claimed model-class absence.

Hypothetical example - not a forecast. A 5 million portfolio contemplates a 10% active sleeve, a 2% evidenced premium, and a 50% evidenced hit rate. Expected contribution is 5,000,000 x 10% x 2% x 50% = 5,000. If fully loaded annual cost is 4,000, the required 2x hurdle is 8,000, so capital remains in paper mode. The example is deliberately simple; changing unsupported assumptions to force a pass would defeat the gate.

Required evidence is the viability worksheet, time log, cost schedule, decision-level attribution, comparator-design attribution, ILF opportunity-cost audit, exception history, deferral calculations and challenger reviews, rolling-start diversifier results, map refresh, and amendment ledger. Stop active deployment if the prospective threshold fails or if required hours cannot be precommitted. Two consecutive negative realized-viability years after deployment trigger the stronger suspension in item 10. The handoff is an annual governance certificate of process status, not a portfolio or security certificate.

10. Certification blockers and stop rules

> **Certification blockers and stop rules:** *withdrawal-path/grind budget, liability closure, actual-comparator scoring and company EC audits must all be complete before certification. Two exception events freeze discretion; failure of the prospective viability threshold means paper mode; two negative deployed years suspend active work; unmaintainable scoring defaults to comparator+ILF+matching.*

The last item prevents partial completion from being presented as certification. The withdrawal-path and grind budget must show how assets and liabilities interact through time, including spending; an immediate stress snapshot cannot establish the duration of a slow failure when withdrawals are unknown. Liability closure requires a verified cash-flow schedule, currencies, index basis, title, tax, and adequate coverage. Actual-comparator scoring requires the archived top 20, sector buckets, independent scoring, sensitivity ranges, and an implementable capped or blended alternative. Company EC audits require verified evidence, not provisional analyst scores or links that merely point to a disclosure. All four are joint blockers. Completing three does not waive the fourth.

The policy's published grind example shows why. With assumptions of 1.50 million of assets and 1.10 million of stressed real liabilities, the default-admissible 45% equity, 25% nominal duration, 20% inflation-linked bond, and 10% ILF mix ends at 1.0875 million and fails immediately by 12,500. A 40/20/30/10 diagnostic ends at 1.1400 million and passes by 40,000 only if the 30% real-liability match is genuine. Neither result certifies a real investor because the liability path and withdrawals remain missing.

Two material Policy Exception Events within 12 months freeze new discretion pending review. One event can be a narrow, risk-monotone response, recorded within 24 hours, independently reviewed within seven days, and expired within 30 days; it cannot waive unaffected rules. The second event indicates that exception handling is becoming the operating model. The freeze prevents repeated emergencies from silently rewriting policy.

Failure of the prospective 2x viability threshold means paper mode. Two consecutive negative deployed years suspend new active work until the program is re-justified. If factor scoring or evidence maintenance cannot be performed reproducibly within the available hours, the default is the named comparator plus ILF plus verified nominal and real liability matching. This default is not an improvised substitute and does not authorize implementation. It is the policy state that minimizes discretionary complexity while retaining the liability-first structure.

Common mistakes are calling the document certified because a single stress passed, treating an exception as precedent, waiting two losing years even when the prospective threshold already fails, or continuing active work because research has already consumed time. Sunk research cost has no authority. The final handoff is either a complete certification packet for a separate approval process, paper mode, a freeze, suspension, or the simple comparator-plus-ILF-plus-matching state. None authorizes a transaction by itself.

The final line of the original one-page section is also a boundary around this expanded explanation. The checklist omits full 22-measure evidence files, legal memoranda, reserve and credit cohorts, reverse-DCF schedules, top-20 worksheets, and dispute exhibits. Those files are opened only for at most eight active full underwrites, relevant events, or the annual audit. Omission is safe only while every item that needs the missing work remains at zero admissible weight. The checklist compresses the sequence; it never compresses away a gate.

Annex C - Debate Adjudication and Hard Questions

C.1 Status and purpose

Historical / provenance only. This annex is the compact audit trail for the five-stage Codex/Fable policy debate. It replaces the old chapter-length retelling of the amendments and the twelve hard-question answers. It records how the governing positions changed, which objections were adopted, which alternatives were rejected, and which disputes remain open. It does not restate every operative rule; those rules belong in their canonical main-body chapters.

The evidence cutoff is 22 August 2026. The concise investment_policy_2026_2036.md remains controlling unless a later, explicit, dated, owner-approved amendment says otherwise. This revised edition is a structural restatement and is not yet an owner-approved successor instrument. Debate records explain provenance and rationale; they are not investment evidence and cannot waive a main-body rule.

Nothing in this annex certifies liability closure, funded status, trait tolerance, the mark-to-market drawdown budget, tax basis, domicile, custody, the actual comparator, factor look-through, company evidence, valuation, or implementation readiness. Every named company has current admissible weight zero. No amendment, minority view, reopening test, diagnostic ceiling, hard-question answer, example, or research label authorizes a trade, order, transfer, hedge, message, account opening, target allocation, entry price, or portfolio change.

C.2 Controlling hierarchy

When debate history and operative prose appear inconsistent, apply the following order. A lower source can explain a higher source but cannot override it.

Priority	Authority	Effect
1	Later explicit, dated, owner-approved amendment	Controls only the provisions it expressly changes. None was supplied for this revision pass.
2	Controlling concise policy	Governs final dispositions, numeric boundaries, uncertainty labels, certification status, and non-authorization.
3	F5-1 through F5-14	Controls inconsistent legacy language, including F4 implementation defects and the stale D-1 reopening wording.
4	Surviving F4 dispositions and the F1-F4 ledger	Controls earlier proposals unless a named F5 correction supersedes it.
5	Expanded prose and chapter files	Supplies explanation, examples, and failure modes only when consistent with priorities 1-4.
6	Debate records, receipts, examples, proxy results, and analyst classifications	Provenance or illustration only; never independent authority or investment evidence.

Within this revised edition, standing certification and non-authorization language in Chapters 1 and 15 comes first, followed by owner-approved amendments, durable constitutional rules, the canonical parameter or trigger in its owner chapter, operating interpretations, dated worldview material, and annexed history. No annex creates permission. Repeated numbers must be checked against their canonical owner chapter.

C.3 Five-stage debate sequence

The exchange was cumulative. Each Codex position integrated the prior challenge; each Fable response attacked contradictions, missing countercases, or implementation weaknesses. Adoption means the criticism earned a binding policy response, not that every proposed sentence survived unchanged.

Stage	Codex position	Fable challenge and resulting direction
1	P1-P17 established the purpose hierarchy, epistemic labels, liability-first measurement, a named comparator, liquidity floor, leverage and look-through discipline, tax awareness, evidence-based sizing, behavior controls, and amendment governance.	F1-1 through F1-15 added Shortfall Parity, treated the comparator as a risky position, separated trait from state, tightened liquidity and emergency rules, formalized missing-fact treatment, and made lump-sum deployment a prior rather than a command.
2	P18-P27 integrated those objections and added common-factor thinking, a symmetric capital-cycle lens, fast broad rules versus slow active judgment, minimum viable governance, theme non-authorization, research sunset, and the principle that an AI tool is not investment edge.	F2-1 through F2-15 forced factor look-through, distinct AI layers, a payer-quality gate, a shortfall remedy ladder, custody questions, harder falsifiers, base-rate attachments, explicit disputes, and broader consumer and financial counterweights to producer bias.
3	P28-P37 replaced narrative scoring with a five-pillar business-model profile, continuous Evidence Coverage, payer and capital-cycle gates, economic depreciation, acquisition and dilution analysis, per-share capture, and validation requirements.	F3-1 through F3-13 retired the old additive score, decomposed financial and emerging-market risks, created a two-stage harvest test, separated assurance liabilities, constrained short-history challengers, strengthened roll-up triggers, and treated small samples as audits rather than statistical validation.
4	P38-P52 converted the framework into funded-state bands, sleeves, loss arithmetic, factor caps, reverse-DCF gates, custody structure, rebalancing rules, and explicit classification boundaries.	F4-1 through F4-14 exposed arithmetic contradictions, a born-in-breach comparator, double counting, mismatched crash and grind defenses, generic treatment of financials, unpriced governance work, and unresolved one-person capacity. The final Codex position adopted all fourteen, with some later corrected.
5	The assembled policy applied F4 to the company map, comparator proxy, grind test, twelve hard questions, prohibition audit, and one-person stress test.	F5-1 through F5-14 performed the closing consistency pass. They corrected the F4-1 floor defect, imposed severity floors, hardened comparator certification, made favorable financial scores provisional, added map decay and factor bars, converted active viability into a prospective gate, repaired D-1 reopening, and printed the true hours burden.

The traceability ledger retains P1-P52 and F1-F4 by origin and disposition. The provenance record retains four technical Fable receipts for rounds 1-4. Those receipts prove which debate artifacts were recorded; they do not validate a portfolio conclusion. The final F5 dispositions derive their authority from their explicit incorporation in the controlling concise policy, not from provenance metadata.

C.4 Round 4 adjudication

All fourteen F4 amendments were accepted. F4-1 was accepted in principle but corrected by F5-1 and F5-3. F4-2 was executed only provisionally and was then hardened by F5-4. The compact disposition record is:

Amendment	Adopted conclusion	Surviving qualification
F4-1	Remove valuation as a second loss multiplier; size from allowed loss divided by permanent-impairment severity; close D-1 without a generic conviction tier.	The originally proposed floor was defective. F5-1 limits the floor to evidence, and F5-3 imposes severity floors.
F4-2	Four consecutive quarters of an embedded hard-cap breach force a written choice between a capped/blended comparator and a cap amendment.	No renewable exception. The proxy branch is unratified pending the actual comparator study.
F4-3	Funded status needs liability flexibility/duration and income-market correlation as a second axis, sensitivity tests, promotion hysteresis, and immediate demotion after verified deterioration.	Numeric bands remain conditional on verified investor facts.
F4-4	Lower gross risk is the crash default; verified inflation-linked duration is the real-grind matching tool when index, basis, title, law, and market integrity pass.	Trend, commodities, and tail-premium strategies remain conditional rather than default-admissible.
F4-5	Budget valuation derating at admission; active derating loss uses 50% of the drawdown budget by default.	Up to 75% requires amendment and independent challenge; later overshoot triggers a process review, not an automatic sale.
F4-6	Tier custody safeguards: same-jurisdiction redundancy is mandatory; a second legal system needs focused review; complex wrappers default to zero.	Tier B and C sizing remains dependent on domicile, tax, title, sanctions, transfer, and estate facts.

Amendment	Adopted conclusion	Surviving qualification
F4-7	Use top-20-plus-buckets comparator look-through, annual/event refresh, a first-year independent repeat, and explicit risk-versus-feasibility reporting.	Proxy error remains material; a feasibility cap is not proof of safety.
F4-8	Tie any staged entry to a filing or preregistered information milestone and expire it by day 180.	Full deployment remains the prior after all gates clear and no event is pending; the rule itself authorizes no entry.
F4-9	Standardize tax deferral and challenger escalation; define the next liquid window and price any proposed tax delay at day zero.	Tax can support hold-not-add but cannot become the investment thesis.
F4-10	Apply the generic finance-complexity flag only to ancillary finance; use class batteries for insurers, banks, and asset managers.	Missing or failed battery evidence blocks as strongly as the generic flag.
F4-11	Test active work against its fully loaded hours, data, tax, and opportunity cost; two negative deployed years suspend new active work.	The zero-active minority and the dislocation-option counterview both survive.
F4-12	Remove overlap between the dated-liability term and the later income-interruption buffer in the liquidity floor.	D24 may replace D36 only through an annual documented election at high funded status and flexibility.
F4-13	Move orchestration hashes and run metadata out of the constitution and into provenance records.	Provenance remains auditable but is not investment evidence.
F4-14	Register one-person executability and procyclical diversifier evidence with adjudicating observables.	U4-1 was later resolved adversely for the full apparatus; U4-2 remains open.

C.5 Final Round 5 amendments

F5 is the closing correction layer and controls inconsistent F4 or expanded-prose wording.

Amendment	Final adopted effect
F5-1	Floor only ECmult. Liquidity and factor penalties retain full force, and a hard factor breach produces zero allowed loss.
F5-2	Enforce gate-state consistency: an E0 row prints diagnostic maximum weight zero; an E1 row remains blocked.
F5-3	Floor severity at 40% for ordinary equity, 60% for levered financials/cyclicals/roll-ups, and 80%-100% for binary, pre-cash-flow, or flagged emerging-market exposures. Legacy generic 5% and 8% rows are void.
F5-4	Do not ratify a comparator feasibility cap without the actual worksheet, a second independent scoring pass, sensitivity ranges, and a capped/blended alternative study. Original caps govern meanwhile.
F5-5	Exclude genuine liability-matched duration from directional rates exposure, while continuing to report basis risk.
F5-6	Treat favorable financial-battery scores as provisional analyst assumptions capped at 2 until the specified reserve, deposit, credit, float-cost, and stressed-capital exhibits attach.
F5-7	A claimed class absence needs at least three named rejected candidates plus annual non-US and counter-canon screens.
F5-8	Enable active capital only when prospective expected value is at least twice fully loaded annual cost; otherwise use paper-journal mode with zero active capital.
F5-9	Expire company classifications after 12 months without refresh; they revert to an unclassified research note.
F5-10	Put the current factor bar on affected company rows: while the unratified comparator proxy exceeds original AI/power caps, active additions are zero.
F5-11	Correct evidence-pointer language, domicile and naming hygiene, and speculative-cap priority. These are consistency repairs, not new conclusions.
F5-12	Label withdrawal-path/grind closure, verified liabilities, actual-comparator scoring, and company Evidence Coverage audits as certification blockers.
F5-13	Replace the stale at-least-10-year D-1 prerequisite with the achievable but demanding reopening test stated below.
F5-14	Print the hours arithmetic: roughly 40 governance hours plus 15-25 hours per live underwrite, or about 160-240 hours for eight names before one-off work; reduce scope prospectively if capacity is absent.

C.5.1 Controlling D-1 reopening rule

F5-13 supersedes every stale reference to an at-least-10-year D-1 record, including the older dispute-register wording and any expanded-prose summary. The controlling test is conjunctive: a preregistered five-year record; at least 20 fully underwritten decisions; independently audited after-tax attribution; verified liability surplus; independent-challenger sign-off; and demonstrated operator capacity to govern the proposed tier. All elements must be present.

Passing that test would reopen D-1 for formal reconsideration and amendment. It would not approve a conviction tier, authorize a position, establish a target weight, validate claimed edge, or waive the then-current investor, evidence, valuation, liquidity, factor, tax, legal, custody, loss-budget, and owner-approval gates. Until a later explicit owner-approved amendment is adopted, D-1 remains closed for this policy and no conviction tier exists.

C.6 Rejected alternatives, minority views, and unresolved disputes

The debate rejected several tempting shortcuts: valuation cannot be charged twice through both admission and a loss multiplier; an evidence floor cannot rescue illiquidity or a hard factor breach; a persistent comparator breach cannot live forever as an exception; lower gross risk is not by itself real-liability matching; a proxy worksheet cannot ratify a higher factor cap; state budgets, backlogs, market size, reputation, or theme fit are not expected-return proof; state support or a depressed share price is not harvest confirmation; and tax deferral cannot justify perpetual ownership of a failed thesis. The full 42-name map was also rejected as a one-person live-underwriting workload. It remains a dated research inventory, not a live portfolio.

Dispute or minority	Current status	Evidence that could adjudicate it
D-1 concentrated value	Closed without a conviction tier; minority argues rare control or demonstrable-edge cases may justify 10%-40%.	Only the F5-13 reopening test can permit formal reconsideration; reopening is not approval.
D-2 capped/blended versus uncapped comparator	Procedure resolved; actual design open.	Actual holdings and tax basis, reproducible top-20 scoring, capped/blended tracking error, cost, abandonment rules, and historical comparison.
D-3 tail insurance and diversifiers	Split: lower gross risk for crash; verified inflation-linked matching for real grind; other diversifiers conditional.	Rolling-start, full-cycle net carry, capacity, crisis correlation, and improvement in liability coverage.
D-4 private and illiquid assets	Open.	Vehicle cash flows, governance, valuation, access, and net return sufficient to beat the liquid alternative by the required margin.
D-5 long-horizon currency hedging	Open within conditional bounds.	Verified liability currencies, carry, basis, tax, and crisis covariance; dated liabilities are matched first.
Zero active versus dislocation option	Structured, not settled.	Prospective two-times-cost gate, paper-mode records, and separate attribution of comparator design versus single-name selection.
AI overbuild, asset lives, and defensive capex	Open by layer.	Utilization, rental and resale prices, cancellations, customer cash returns, retirement curves, capex/depreciation, and free cash flow.
State-capex shareholder capture	Open by contract and jurisdiction.	Binding appropriation or rate recovery, cancellation terms, state take, financing, dilution, and real per-share cash.
Custody calibration and speculative-sleeve behavioral value	Structure or range set; optimal size remains open.	Verified personal legal facts, annual transfer tests, total-loss simulations, exception rates, goal probability, hours, and abstention comparison.
Deterrence, assurance rents, and quality-duration crowding	Open.	Threat-reversal experiments, real price and liability data, new entry, no-multiple-recovery stress, aggregate derating budget, and diluted cash per share.
U4-1 one-person executability	Resolved adversely for the full apparatus; bounded version accepted.	At most eight live full underwrites, subject to actual hours, cost, decision value, and prospective capacity. Scope may fall to zero.
U4-2 procyclical diversifier evidence	Open.	Rolling start dates spanning inflation, deflation, crash, and calm, net costs and capacity, and crisis-conditional correlation.

No open dispute is a latent permission. Where evidence is missing, the relevant admission, allocation, or certification state remains blocked.

C.7 The twelve final hard questions - closure map

The twelve questions closed a policy-design interrogation, not an investment decision. Their full answers are not repeated here.

Question	Compact answer and controlling destination
1	All F4 amendments were adjudicated. D-1 closed on the no-conviction-tier branch, subject only to the F5-13 reopening route. See Chapters 11 and 14 and Sections C.4-C.6.
2	The loss arithmetic was made visible, severity floors voided the decorative 5%/8% rows, and hard gates remain zero. The canonical formula and ceiling table belong in Chapter 11.
3	The inherited grind example failed; the illustrative boundary passed only if real-liability matching is genuine. Withdrawal-path and liability closure remain blockers in Chapters 5 and 15.
4	The proxy comparator was born in breach, but higher diagnostic ceilings were not ratified. Original caps and zero active AI/power additions govern pending Chapter 6's actual-comparator work.
5	Balance-sheet financials were included under class batteries; no full harvest candidate qualified. Company-specific records belong in Annex A.
6	All eleven former Core rows contained the required fields, but completeness of fields did not establish evidence quality, valuation, portfolio role, or authorization. All remain weight zero.
7	The rejection set covered the major model classes and rejected admired businesses when valuation, governance, capital, payer, or evidence gates failed. A rejection is not a short recommendation.
8	Financial batteries constrain favorable and unfavorable cases alike; F5-6 keeps favorable scores provisional pending named exhibits. Chapter 10 owns the batteries.
9	The ten company-research prohibitions were checked against the highest-risk rows. The detailed self-audit belongs in Annex D, not here.
10	P1-P52, F1-F5, disputes, receipts, and surviving minorities are traceable. Technical receipts and hashes belong in Annex E and are not investment evidence.
11	The assembled policy remains vulnerable to consensus, recency, survivorship, narrative, obsolescence, politics, concentration, and complexity. The full apparatus failed the one-person test; Chapter 13 and Annex D own the controls.
12	Nothing is certified or authorized. Chapters 1 and 15 control the standing boundary.

C.8 Practical boundary

Read this annex backward from authority, not forward from debate drama. F5 controls inconsistent F4 language; surviving F4 controls earlier proposals; canonical main-body rules control this historical summary. Rejected views remain visible so future amendments cannot manufacture consensus, and unresolved views remain tied to observables so they cannot become excuses.

The current operating result is unchanged: liability closure and implementation readiness are uncertified; numeric bands await verified investor facts; actual comparator scoring, factor look-through, withdrawal-path grind analysis, and company evidence and valuation audits are incomplete; every named company has current admissible weight zero; and no transaction or portfolio change is authorized.

Annex D - Bias, Prohibition, and Implementation Self-Audit

The company map is a research inventory, not a portfolio. That distinction is easy to state and surprisingly hard to preserve once named companies, scores, valuation language, and percentage ceilings appear on the same page. The purpose of this self-audit is therefore not to repeat a general disclaimer. It is to test the map against ten specific prohibitions imposed on the final classification round by Round 4 Section 15. Each prohibition closes a different route by which incomplete research could be mistaken for a recommendation or a completed underwrite.

The audit result is compliance, but it is qualified compliance. The map contains 42 named companies: 11 Core Research Candidates, 13 Watchlist/Price Sensitive names, 6 Speculative Options, and 12 Avoid/Counterexamples. Every one has a current admissible weight of zero. The map's scores, Evidence Coverage percentages, severity assumptions, and diagnostic maximum weights are provisional analyst assumptions, not independently corroborated conclusions. A link to an investor-relations page or filing is an evidence pointer and earns no Evidence Coverage credit until the relevant disclosure is actually extracted and checked. The investor's liabilities, funded status, tax position, account structure, domicile, custody arrangements, actual comparator, factor headroom, and current valuations are not verified. The audit therefore establishes that the document did not cross its own research boundary; it does not certify any company, portfolio design, or implementation decision.

This difference between compliance and certification matters throughout the chapter. A row can comply because it openly records missing work and assigns zero current admissible weight. That is not the same as proving the row's preliminary score correct. Likewise, a diagnostic ceiling can comply because it is expressly identified as a ceiling that could bind only after every gate passes. It must not be read as a target, a suggestion, or even evidence that the company is presently eligible for a nonzero position.

1. No buy call, target weight, entry instruction, order, or portfolio mutation

The first prohibition is the most direct: the map may not call any company a buy, assign a target weight, give an entry price or order instruction, or change the current portfolio. The highest-risk points were all 11 Core rows, especially Taiwan Semiconductor Manufacturing and ASML. The word "Core" can sound like a portfolio status, while their diagnostic maximum weights - 1.25% for TSMC and 1.43% for ASML - can look actionable when separated from their qualifiers.

The compliance evidence is explicit. Every company row includes non-buy language and gives current admissible weight as zero. The map's reading rules state that Wmax is a diagnostic policy ceiling, never a recommendation. TSMC and ASML go further: while the unratified comparator proxy remains above the original AI and power risk caps, admissible active addition is zero regardless of any later underwriting outcome. Similar row-level factor bars appear for other AI- or power-sensitive names. The policy as a whole also says that it authorizes no trade, order, hedge, message, account opening, execution, or portfolio change.

This prohibition matters because a maximum permitted by arithmetic is not an optimal allocation and is not current eligibility. Wmax answers a hypothetical question: if the evidence, valuation, liquidity, factor, investor-fact, and approval gates were all satisfied, what upper bound might the loss-budget formula produce? TSMC's and ASML's current answer remains zero because those conditions are not satisfied. Treating a ceiling as a target would reverse the policy's logic: it would let a downstream sizing calculation substitute for the upstream admission work.

Residual ambiguity remains in the persuasive effect of labels and numbers. "Core Research Candidate" may still be read casually as "core holding," and a percentage printed beside a famous company can anchor attention even when surrounded by caveats. The audit controls that risk through repeated non-buy language, not by pretending the risk disappears. Failure would include any sentence such as "buy," "initiate," "enter below," "target X%," or "place an order"; any silent portfolio-file change; any use of Wmax as a recommended allocation; or any omission of the current zero-weight and non-authorization qualifications when a row is reused.

2. No narrative shortcut to expected return

The second prohibition bars the use of a theme, market size, model-class prior, historical reputation, old Business-Model Score composite, price decline, or current index membership as evidence of expected return. NVIDIA, Intel, and Rheinmetall were the highest-risk entries because each comes with an intuitively powerful story: AI infrastructure demand, semiconductor capacity rationalization, or European rearmament. Those stories may motivate questions, but they do not establish shareholder cash capture at the current price.

The map demonstrates compliance by separating each narrative from valuation and cash economics. NVIDIA reports extraordinary operating facts yet receives Pillar E, the valuation and falsifiability pillar, of 0 and Wmax of 0 because E0 fails admission. Its fiscal 2026 revenue growth, Data Center growth, 71.1% gross margin, ecosystem reputation, and index prominence do not rescue the missing conservative return case. Intel is an Avoid/Counterexample with A2/B1/C0/D2/E1, Evidence Coverage of 72%, severity of 100%, and Wmax of 0. A lower share price, fab assets, or public support does not prove harvest economics. Rheinmetall remains Watchlist/Price Sensitive with B2 and E1; backlog and rearmament budgets are not treated as cash until funding, cancellation, working-capital, margin, and conversion evidence is established.

This rule exists because an economically important theme can still be a poor equity investment. Demand may accrue to customers, labor, suppliers, lenders, or the state. It may be offset by competition, overcapacity, depreciation, dilution, regulation, or an already demanding valuation. The Round 3 critique of the old additive score is central here: a composite dominated by correlated measures of quality could allow a company to compensate for a weak capital-cycle or valuation result with repeated versions of the same attractive moat story. The final five-pillar structure refuses that compensation.

The residual ambiguity is that causal theses still occupy prominent space and can influence judgment before the hard evidence arrives. The final map also materially correlates with the 2010-2025 consensus-quality factor, so its composition reveals the screen's lens rather than future return. Compliance would fail if theme growth, addressable market, reputation, index status, a past high, or a price decline were cited as the reason expected return clears the comparator; if an E0 or E1 name were promoted because the business is admired; or if market-size language replaced a reverse discounted-cash-flow analysis. A theme may open a research file. It may not close the valuation case.

3. No return to the old additive Business-Model Score

The third prohibition applies to all 42 companies. The old additive Business-Model Score, or BMS, may not be reported without the five pillars, weakest-pillar logic, Evidence Coverage, hard gates, reverse-DCF range, and dated evidence. In the final map, no additive total appears. Every row instead reports an A-E profile, an Evidence Coverage percentage, dated 2026 evidence or an evidence pointer, valuation sensitivity, missing evidence, a failure mode, a falsifier, and non-buy language.

The five pillars are non-compensating. A covers demand and lock-in; B covers industry structure, cycle, and payer quality; C covers capital economics; D covers governance and per-share capture; and E covers valuation and falsifiability. Ordinary admission requires every pillar to be at least 2. Evidence Coverage below 60 is research-only; the multiplier is 0.50 at 60-74, 0.75 at 75-84, and 1.0 at 85 or more, but coverage can only reduce an admissible loss budget. It cannot waive a severity floor or hard gate. The Core classification has the stricter expectation of mature evidence, Evidence Coverage of at least 75%, no hard-gate failure, all pillars at least 2, A-D normally at least 3, completed bear impairment and reverse DCF, and mapped common factors. The present Core rows are still not completed underwrites: the map expressly records their missing reverse-DCF and other work.

This prohibition matters because averages hide fatal weaknesses. A strong demand franchise cannot compensate for value-destructive reinvestment; sound reported accounting cannot compensate for owner-cash leakage; and a wonderful business cannot compensate for a price that makes conservative expected return inadequate. The

Round 3 adversarial review found that the old 22-dimension composite quadruple-counted customer lock-in and gave valuation only about 6.5% of the total. Weakest-pillar logic is the repair.

Residual ambiguity lies in the provisional pillar scores themselves. They are analyst assumptions, no independent adversarial-corroboration pass has been completed, and an investor-relations link is not verified evidence. The self-audit proves that the required fields are visible; it does not prove that A4, EC78%, or any other input is right. Failure would be publication of a standalone BMS total; omission of the weakest pillar, date, reverse-DCF status, or hard gates; averaging across a zero pillar; presenting Evidence Coverage as a buy signal; or treating an evidence pointer as extracted and verified proof.

4. No short-history challenger as a Core Research Candidate

The six highest-risk companies under the fourth prohibition are CoreWeave, Samsara, Symbotic, Nu Holdings, Sea Limited, and BYD. All six are classified as Speculative Options, not Core Research Candidates. They are assumed to suffer total loss for sizing purposes, carry severity of 100%, and show a 0.75% diagnostic maximum. Their current admissible weights are nevertheless zero.

The numerical record reinforces the classification. CoreWeave has EC55 and E0; Samsara EC58 and E0; Symbotic EC52 and E0; Nu EC57 and E0; Sea EC58 and E1; and BYD EC59 and E1. None meets the Core evidence threshold. The 0.75% number is the per-name speculative cap and binds before any larger diagnostic result; it is not permission to establish a position. The broader speculative sleeve remains conditional on verified funded status and is itself governed as a total-loss budget.

The rule matters because a short operating history cannot manufacture a full cycle. An emerging company may be genuinely exceptional, but it may lack multi-period evidence, recession or credit-cycle behavior, reliable cohort economics, mature controls, or proof that value reaches each diluted share. The Round 3 debate concluded that the scorecard is principally an instrument for mature companies. Speculative classification preserves the option to research a challenger without laundering that uncertainty into Core status.

The residual ambiguity is that calendar age alone does not define maturity, and Evidence Coverage is partly judgmental. A young company could accumulate unusually strong evidence; an old company could remain opaque. Promotion therefore depends on maturity and evidence gates, not the passage of time by itself. Failure would be moving any of the six into Core on growth, excitement, or model-class fit; describing 0.75% as a recommended starter position; omitting the total-loss assumption; or allowing a speculative label to imply concentration eligibility.

5. Do not ignore exposure already embedded in the comparator

The fifth prohibition is most at risk in Microsoft, NVIDIA, Schneider Electric, and Eaton, because all could appear to diversify an active list while duplicating AI-capex or power exposure already present in the strategic comparator. The policy first looks through the comparator rather than treating it as neutral. Its diagnostic proxy - 60% MSCI ACWI and 40% liability-matched high-grade fixed income - produced AI factor stress loss of 8.2% and joint AI-plus-power factor stress loss of 9.0%. Those exceed the original 5% and 8% risk caps.

The compliance response is deliberately restrictive. The figures are described as plausible order-of-magnitude assumptions, not reproducible enough to ratify a cap, because the archived top-20 worksheet and required second independent scoring pass do not yet exist. Until the actual named comparator is studied, the proxy decision is unratified, the original caps govern, actual-comparator scoring is a certification blocker, and active additions with material AI or power sensitivity are zero. The affected rows repeat that current factor bar.

This matters because apparent diversification by ticker can be concentration by economic cause. Several issuers can depend on the same hyperscaler capital budgets, grid buildout, power scarcity, or equipment replacement cycle. Adding a different company name may increase rather than reduce the portfolio's loss under the shared

shock. Comparator exposure consumes headroom first because the investor already owns that risk through the strategic position.

Residual ambiguity is substantial and explicitly retained. The proxy sensitivities are not a completed actual-comparator study, and raising a feasibility cap would accept risk rather than remove it. The required work includes the actual comparator, archived top-20 weights, sector buckets for the remainder, two independent scoring passes, plus-or-minus 0.25 sensitivity perturbations, a capped or blended alternative, a tracking-error budget and abandonment point, and historical comparison. Failure would be calling an AI or power name a diversifier without look-through; treating missing comparator exposure as zero; authorizing an active addition while the original caps are breached; or converting the proxy's 8.2% and 9.0% into ratified limits.

6. No inflation of weak, conflicting, or missing evidence

The sixth prohibition has three parts: a management-only claim may not receive a score above 1; a conflict among Tier 1, Tier 2, and Tier 3 evidence may not be averaged; and missing look-through may not be treated as zero. Symbotic, Adani Enterprises, and Rheinmetall were the most exposed rows. The audit evidence is their constrained scores, explicit missing-evidence lists, and zero current admissible weights. Symbotic carries B1/C1/D1 and EC52; Adani carries B1/C1/D0/E0 and EC55; Rheinmetall carries B2, E1, and EC68. Symbotic's backlog is not counted as revenue or cash. Adani still requires an independent ownership and guarantee map, cash waterfalls, payer contracts, and auditor evidence. Rheinmetall's public budgets and backlog still require contract-funding and conversion proof.

This rule matters because management controls much of the initial narrative and much of the reported artifact. The debate used General Electric as the warning: a captive balance sheet and opaque segment structure can make audited, management-produced numbers look clean while corrupting cash conversion. Accordingly, a material Tier 1 conflict with Tier 2 or Tier 3 evidence scores zero until resolved, rather than being averaged into a comfortable middle. Missing look-through is an unknown exposure, not no exposure.

Residual ambiguity remains because the evidence tiers do not make any source infallible. Tier 1 can still be wrong or incomplete; lower-tier evidence can be noisy or interested. "Material conflict" also requires judgment. The conservative response is to preserve and investigate the conflict, not select the friendliest source. Failure would include assigning 2-4 from management commentary alone; averaging contradictory filing, regulator, customer, supplier, or channel evidence; counting backlog or announced budgets as cash without enforceability; giving missing factor or ownership exposure a zero; or allowing a low-coverage row to borrow confidence from its theme.

7. Harvest status is neither normal sizing nor renewable by inertia

Intel is the highest-risk map point for the seventh prohibition. It was the proposed harvest candidate, but it is classified as Avoid/Counterexample with A2/B1/C0/D2/E1, EC72, severity 100%, and Wmax 0. It explicitly fails the F3-5 harvest test: confirmed capacity exit and owner cash returns are absent, while continuing investment and a state backstop remain. No full harvest candidate is claimed anywhere in the map.

The rule says harvest-candidate status cannot confer normal-size eligibility and cannot roll beyond 24 months without exit and cash evidence. The debate rationale is historical. Shipping and European banks repeatedly appeared ready for harvest, yet state support and easy refinancing kept marginal capacity alive for years. Price weakness and consolidation rhetoric were not enough. The abuse-resistant indicators are financing-side and observable: closed issuance windows, capex below depreciation, actual exits or bankruptcies, and absence of a state or covenant-lite mechanism keeping zombie capacity alive.

This matters because a cyclical value trap can become more optically attractive while its economics deteriorate. A 24-month expiry prevents a temporary research hypothesis from becoming a permanent exception through repeated renewal. Even a valid experimental candidate belongs in the speculative, total-loss framework until the full confirmation set clears.

Residual ambiguity concerns when exit is truly sufficient and how state involvement should be interpreted. Public support can sometimes accelerate rational restructuring, but it can also preserve uneconomic supply. The policy does not resolve that question by story. Failure would be calling Intel a normal-size candidate because it fell in price, because fabs have replacement value, or because subsidies exist; carrying any harvest candidate beyond 24 months without new exit and cash evidence; or quietly renewing the same thesis under a new label.

8. State budgets are not collectible revenue without the Payer Quality Gate

Rheinmetall, PG&E, and Intel are the highest-risk entries under the eighth prohibition. Each is linked to public spending or regulated support, but the map does not translate an announced budget, allowed return, appropriation, or subsidy into owner cash. Rheinmetall remains constrained by B2 and requires funding, cancellation, cash-conversion, and program-return evidence. PG&E is Avoid with B1/C1/D1/E1 and a ceiling of zero because an allowed regulatory return is not realized cash after wildfire mitigation, financing, affordability pressure, and political disallowance. Intel is Avoid because state support does not establish foundry yield, external demand, or after-sustaining-capex return.

The Payer Quality Gate separates payment certainty from shareholder capture. A sovereign may be highly capable of paying while retaining strong renegotiation, price-control, windfall-tax, or monopsony power. The relevant object is realized supplier return on invested capital through a cycle, not the direction of a headline budget. Similarly, a utility's approved return is not automatically collectible if financing, disallowances, liability, or bill affordability prevent conversion to per-share cash.

Residual ambiguity persists because state contracts differ by jurisdiction, legal form, cancellation rights, indexation, and funding status. There is no universal discount that solves payer risk. Failure would be counting a proposed defense budget as funded backlog, a subsidy announcement as equity value, an allowed return as received cash, or a government policy objective as proof of margin. Contract, appropriation, rate-case, collection, and per-share capture evidence must remain separate.

9. No personalized tax, account, domicile, currency, or custody conclusion

The ninth prohibition is most exposed in DBS, TSMC, and Adani, as well as in the policy's custody architecture. Those names carry jurisdiction, listing, currency, sanctions, estate-situs, transfer, or related-party considerations that can differ radically by investor. The audit finds compliance because all domicile and funded-state conclusions remain conditional and the required investor and liability facts are expressly missing. DBS is not a completed domicile-fit review; TSMC's geopolitical and legal work is incomplete; Adani still lacks an independently verified ownership, guarantee, and cash-flow map. All remain at current admissible weight zero.

The broader custody rules do not change that result. Tier A requires a same-jurisdiction second custodian. Tier B, involving a second legal system and plain instruments, requires the tax, reporting, title, transfer, sanctions, and estate-situs subset review. Tier C entities, trusts, and complex wrappers default to 0% until the full professional gauntlet passes. These are policy structures, not a personalized conclusion that any named country, broker, wrapper, account, or currency is suitable for this investor.

This prohibition matters because business quality cannot cure defective title, an unlawful structure, a tax burden, an estate problem, a currency-liability mismatch, or inability to transfer assets during stress. The same security can be sensible in one legal and personal context and unacceptable in another.

Residual ambiguity is not a drafting flaw here; it is the consequence of missing facts. Domicile, citizenship, tax residence, beneficiaries, succession needs, liability currencies, account types, and custody access must be verified before a conclusion is possible. Failure would be stating that Singapore, Taiwan, India, the United States, or any other jurisdiction is suitable or unsuitable for this investor without that work; recommending an account or wrapper; assuming tax rates or treaty treatment; or treating legal domicile and economic exposure as interchangeable.

10. No silent resolution of disputes or professional-review gates

The tenth prohibition prevents the company map from settling constitutional disagreements merely by the way it classifies companies. It originally protected D-1 through D-5; U2-1, U2-2, U2-3, U2-4, U2-5, U2-7, U2-8, and U2-9; U3-1 and U3-2; and the second-jurisdiction professional-review gate. D-1, D-3, and U4-1 were the points most at risk in the assembled policy. Appendix B records status, surviving minority, and adjudicating observable rather than deleting the contrary view.

D-1, concentration, is resolved explicitly for this policy: there is no generic conviction tier, and ordinary equity is intentionally capped near 2-2.5% in the best funded state because severity cannot ordinarily be below 40%. The concentrated-value view survives as a rejected-with-reasons minority. The adopted F5-13 reopening test is a preregistered five-year record covering at least 20 fully underwritten decisions, independently audited after-tax attribution, verified liability surplus, and challenger sign-off. Two inherited source summaries once said at least ten years; they are preserved only in the production history and do not appear as operative rules in this edition. The adopted amendment is the five-year, at-least-20-decision test.

D-3 is resolved as a split, not as a universal diversifier endorsement: lower gross risk is the crash null, verified inflation-linked liability matching is the real-grind instrument, and trend, tail-premium, or commodity approaches remain dependent on full-cycle net carry and coverage improvement. U4-1 is resolved adversely for the full apparatus: one person cannot maintain 42 live underwrites. The bounded version permits at most eight live full underwrites and still must pass the hours and active-program viability tests. Other disputes remain open or conditional, including comparator design, private assets, long-horizon currency hedging, AI overbuild, economic lives, quality-duration crowding, and procyclical diversifier evidence. The Tier B and Tier C professional gates are not waived by a favorable company classification.

This prohibition matters because an apparent implementation detail can smuggle in a constitutional decision. A map overweight in quality companies might quietly settle the quality-crowding dispute; a nonzero foreign allocation might imply legal approval; or restrictive sizing arithmetic might eliminate the concentration minority without acknowledging it. Explicit adjudication keeps disagreement inspectable and makes future reopening depend on evidence rather than memory.

Failure would include omitting a live minority, marking an open issue resolved without its observable, treating the map's quality-heavy composition as proof that U3-2 is settled, bypassing professional review for cross-border structures, or applying the eight-underwrite ceiling as evidence that an active program is economically worthwhile. The ceiling is a capacity limit, not a mandate to use all eight slots.

Reading the audit result correctly

Across the ten tests, the map's strongest compliance evidence is not that its candidates look attractive. It is that the gates bind against attractive stories. NVIDIA and Ferrari have E0 and Wmax 0. Palantir is an admired Avoid. Intel receives no harvest promotion. The six short-history challengers remain Speculative Options with a 0.75% diagnostic cap and zero current admissible weight. State-backed narratives receive no automatic cash credit. AI and power additions remain at zero while the original comparator risk caps are exceeded. Every named company remains non-buy.

The remaining caveats are equally important. All classifications expire after 12 months without refresh and revert to Unclassified Research Note. The comparator study and company Evidence Coverage audits remain certification blockers. Liability closure and actual funded status remain uncertified. The 8.2% AI and 9.0% joint AI-plus-power proxy stresses are unratified assumptions, not final risk measurements. The audit does not validate preliminary facts simply because they are dated, and it does not convert a research taxonomy into an approved portfolio.

What the final bias actually is

This policy does not emerge from five rounds of debate as neutral machinery. It has a definite bias, and the purpose of the final stress test is to make that bias visible before it is mistaken for objective truth. The policy favors a liability-first strategic structure built around a named low-cost comparator, an investor liquidity floor, nominal and real liability matching, and only a bounded amount of active research. Within active research, its evidence gates naturally favor businesses with recurring demand, defensible industry structure, strong cash economics, transparent per-share capture, and a valuation that survives a conservative reverse discounted-cash-flow test. In ordinary language, it tends to surface established "quality" franchises with good disclosure. That preference may be sensible, but it is still a preference produced by the screen.

The final bias is therefore cautious in two different ways. First, it is skeptical about individual-company conviction. Even a fully liquid ordinary equity at the highest funded-state band is generally constrained to a diagnostic ceiling near 2% to 2.5% under the permanent-impairment arithmetic, and that ceiling is not a recommended weight. Second, it is skeptical about the value of research itself. The policy does not assume that a carefully governed active program must exist. It gives the zero-active view an operational route to prevail whenever the expected benefit of active work does not justify its fully loaded cost. The result is not a conventional stock-picking policy with some risk controls attached. It is a strategic policy that permits a small active layer only if that layer earns the right to consume capital, time, attention, data expense, tax work, and error capacity.

That distinction matters because a process can be internally disciplined and still be wrong about what deserves attention. Forty-two company dashboards, five non-compensating business-model pillars, factor look-through, tax analysis, custody analysis, valuation work, and evidence audits can create an impressive structure. They can also create research theater: a large body of work that rationalizes activity without improving liability coverage or after-tax results. The stress test asks whether the apparatus protects the investor from that possibility. Its answer is qualified. There are meaningful controls, but each control leaves residual risk, and the full apparatus is not executable by one person.

The section is a stress test of governance, not a trading plan. It does not select securities, set target allocations, approve entries, authorize portfolio changes, or convert the company map into holdings. Every named company remains at a current admissible weight of zero because verified investor facts, Evidence Coverage audits, current valuation, factor headroom, and implementation approval are incomplete. Any deep company exhibit, legal memorandum, tax analysis, or custody analysis omitted from the simplified operating process remains a blocking item for the relevant decision. Omission saves time only because the affected item stays at zero; omission is never deemed equivalent to passing the work.

Hidden consensus: when a skeptical process still assumes activity

The first attack is hidden consensus. Across the policy, the visible debate is often about how to govern a core-plus-bounded-active structure. That framing can quietly assume the answer to a prior question: whether an active sleeve is worth building at all. The company map reinforces the assumption because it devotes extensive attention to candidates and counterexamples. The five-pillar scorecard reinforces it because it offers a detailed way to assess a name. The loss-budget formula reinforces it because it explains how an admitted name would be bounded. None of those components proves that the total active program has positive value after all costs.

There is a second hidden consensus inside the first. A quality-gated process will tend to select established franchises with recurring demand, high returns on incremental capital, strong disclosure, and durable market positions. In the final map, many Core Research Candidates resemble the 2010-2025 consensus-quality canon. This does not make them bad businesses, and it does not invalidate the screen. It does mean the screen's composition is evidence of its lens, not evidence of forward returns. The open quality-duration dispute remains real: expensive, long-duration franchises can suffer severe derating even when their operating businesses remain intact.

The policy controls this bias in several ways. Financial businesses and harvest-phase possibilities remain in the counter-canon, rather than allowing the quality list to define the whole investable world. The sixth decade scenario, S6, explicitly contemplates a plateau or muddle-through environment in which a low-cost comparator, ordinary compounders, and disciplined capital return beat dramatic thematic research. A theme that produces neither a decision nor an explicit abstention threshold after 36 months moves to annual-check status and records its hours and cost. Annual map refreshes must include counter-canon and non-US screens, and classifications expire after 12 months without refresh. These controls make it harder for the original list to become a permanent canon through mere repetition.

The strongest control is the active-program viability test. Expected annual active contribution is calculated as:

portfolio value x active sleeve x evidenced premium x evidenced hit rate.

"Portfolio value" is the relevant capital base. "Active sleeve" is the fraction exposed to the active program. "Evidenced premium" is the after-cost return advantage the program can credibly support over the named comparator, rather than an aspirational excess return. "Evidenced hit rate" recognizes that not every underwritten decision will deliver that premium. The product is compared with fully loaded annual cost: research hours valued at the operator's opportunity cost, plus data, tax, and other implementation costs. Active capital may be deployed only when expected contribution is at least 2 times that fully loaded annual cost at the annual review. Below that prospective threshold, the program is confined to paper-journal mode with zero active capital. If it clears the prospective threshold and is deployed, two consecutive negative realized-viability years still suspend new active work and return the operating default to the comparator plus the investor liquidity floor and verified liability matching until the program is re-justified.

Consider a clearly hypothetical example, not a forecast or authorization. Suppose a portfolio value were 2,000,000, the contemplated active sleeve were 10%, the evidenced premium were 2%, and the evidenced hit rate were 50%. The modeled annual contribution would be 2,000,000 x 0.10 x 0.02 x 0.50 = 2,000. If fully loaded annual cost were 1,500, the required 2 times cost threshold would be 3,000, so the program would fail prospectively and remain in paper-journal mode. The conclusion would not be that a particular security should be sold or avoided. It would be that capital has not been authorized for the active apparatus because its expected program-level contribution does not yet pay enough for its burden.

The residual risk is behavioral. Research can rationalize itself. An operator who enjoys company analysis may overstate the evidenced premium, underprice personal time, treat a hoped-for hit rate as evidence, or describe optional learning as return. Paper mode can also become a large unpaid hobby if its hours are not counted honestly. The zero-active minority therefore remains live, alongside its counter-position: standing research capability may have option value in a dislocation, harvest phase, or sector bust when an index-only investor lacks an evaluative framework. The policy does not force consensus between these views. It structures the disagreement through the prospective 2 times cost threshold, paper mode below it, retrospective suspension above it, and separate annual attribution between value created by comparator and matching design and value created by single-name work.

Recency bias: a current map is not a decade-long base rate

The second attack is recency bias. The evidence date is 22 August 2026, and the named map is visibly shaped by the subjects commanding attention at that date: artificial intelligence, power infrastructure, obesity treatment, and defense spending. Current filings and current conditions are necessary for a present research map, but they can cause the most salient recent trends to occupy too much of the long-horizon imagination. A decade policy built from a current snapshot can mistake the latest investment cycle for the structure of the whole decade.

The six scenarios are the first defense. They include productive diffusion, inflation or financial repression, fragmented geopolitical blocs, crash and sequence risk, adaptation to physical constraints, and the S6 plateau or muddle-through case. These are conditional causal maps, not forecasts with assigned certainty. The point of

carrying all six is to prevent one fashionable mechanism from becoming the unspoken base case. Old-cycle normalization performs a related function at the company level: margins, capital intensity, utilization, credit performance, payer behavior, and demand must be assessed across a relevant cycle rather than extrapolated from the latest year.

Company falsifiers and the S6 research sunset are further controls. A memorable 2026 thesis must be tied to observable failure conditions, normally with a 36-month horizon in the map. A research theme that cannot produce a decision or a defined abstention threshold within 36 months loses attention and moves to annual-check status. This does not declare the theme false. It prevents open-ended fascination from consuming recurring governance capacity.

A hypothetical example clarifies the distinction. Imagine that current data-center spending makes a power-equipment thesis appear unusually strong. Recency bias would treat current orders as proof of decade-long shareholder returns. The policy instead asks whether orders convert to cash, whether new capacity floods the market, whether customers earn enough from the spending to sustain demand, whether economic depreciation is understated, and what evidence would falsify the thesis. The example does not imply that the theme will succeed or fail. It shows how a current fact is prevented from becoming an unbounded forecast.

The residual risk cannot be eliminated by the scenario framework. The source set is intentionally current and is not a systematic ten-year base-rate study. The policy has not assembled every comparable cycle, nor does it claim that its current examples are statistically representative. A six-scenario table can broaden attention, but it can also provide a narrative slot into which almost any outcome is placed after the fact. The final implementation implication is modest but strict: current classifications are dated research priorities, not timeless findings, and no amount of repeated reference upgrades them into evidence.

The same recency problem applies to diversifier evidence. A requirement for a long historical record can still be procyclical if the decision depends on the date at which the record is measured. A trend strategy assessed after a weak trend decade may be rejected just before a favorable regime, while the same rule assessed after a strong run may admit it when it is crowded. Dispute U4-2 therefore remains open. Its adjudicating evidence is rolling-start, full-cycle net performance across inflation, deflation, crash, and calm, including implementation cost, capacity, and crisis correlation. No single decade decides the question, and this chapter neither admits nor rejects a diversifier.

Survivorship bias: visible winners are not a validated sample

The third attack is survivorship bias. The Core Research Candidate list is populated mainly by firms that have survived, won, and produced relatively clean disclosures. Failed predecessors are less visible because they disappear from indices, lose research coverage, delist, merge, restructure, or no longer present tidy data. Studying current winners can therefore teach the characteristics of survivors while concealing how often the same apparent characteristics existed before failure.

The policy's twelve Avoid or Counterexample cases and its model-class counterexamples are a deliberate control. They demonstrate that attractive themes can fail the gates: an installed-base or backlog story can be undermined by safety, leverage, or governance; a regulated-capex story can fail on liability, payer collectibility, or financing; a scarcity or growth narrative can fail on price, per-share capture, or economic depreciation. The map also rejects admired companies when the valuation pillar fails. These are useful adversarial examples because they show that the scorecard is capable of saying no.

However, an Avoid roster is not the same as a systematic failure cohort. The policy has no comprehensive delisted-company sample and no matched historical study demonstrating that its five-pillar classifications predict superior outcomes. The preliminary company scores are analyst judgments, and no name has completed the independent adversarial-corroboration pass required for a live underwrite. Accordingly, the map cannot validate its own predictive power.

Consider a hypothetical historical screen that selects ten surviving firms with excellent current disclosure and then identifies two obvious failures as warnings. Even if the warnings are well chosen, the exercise says nothing about how many now-failed firms would have scored equally well before their problems emerged. Without that denominator, apparent discrimination may be hindsight. The policy's final conclusion is therefore deliberately narrow: counterexamples are a control against complacency, but they do not establish a base rate, and the map remains a research-priority tool rather than evidence of investment skill.

This is also why classifications decay after 12 months. A stale Core label can become a miniature survivorship mechanism, preserving the favorable classification while adverse or contradictory evidence accumulates elsewhere. Reversion to an Unclassified Research Note prevents an unmaintained label from masquerading as a current conclusion. It does not mean a company has deteriorated; it means the operator no longer has the evidence freshness required to assert the classification.

Narrative fallacy: causal coherence is not cash capture

The fourth attack is narrative fallacy. The decade themes are memorable because they connect technologies, political choices, physical constraints, and business models into causal stories. A strong story can feel explanatory before the shareholder economics are demonstrated. Artificial intelligence may increase productivity without allowing an equipment supplier to retain the value. Defense budgets may rise without converting promptly into funded, collectible, profitable contracts. Healthcare demand may grow while net pricing, payer access, competition, or pipeline concentration erodes per-share value. Infrastructure spending may benefit labor, customers, lenders, or the state rather than shareholders.

The five non-compensating pillars are the main control. Demand and lock-in, structure and payer quality, capital economics, governance and per-share capture, and valuation and falsifiability must each meet the admission minimum. Strength in one pillar cannot average away a zero in another. A large market does not repair poor capital economics. A strong balance sheet does not repair a failed valuation. An attractive social purpose does not establish collectibility. A compelling founder story does not establish that value reaches each diluted share.

The payer and cycle gates make this control concrete. Announced budgets are not contracts, backlog is not cash, and public need is not a funded rate case. Per-share capture requires attention to dilution, stock-based compensation, acquisitions, related parties, state leakage, debt, leases, and maintenance capital. Reverse discounted-cash-flow analysis asks what operating and valuation assumptions are already embedded in the current price and whether a conservative return clears the comparator. Hard falsifiers are preregistered so that later disappointment cannot be explained away by rewriting the thesis.

A hypothetical example makes the sequence visible. Suppose a public plan announces 10 billion of spending in an industry served by a research candidate. Narrative reasoning might multiply the budget by an assumed market share and margin. The policy requires additional evidence: appropriation, contract award, cancellation rights, timing, working-capital needs, cost escalation, cash collection, sustaining investment, and diluted per-share retention. Until those links are evidenced, the budget is context rather than shareholder cash. This example is neither a forecast nor an endorsement of any company.

The residual risk is that the pillar scores themselves are preliminary analyst judgment. A five-pillar framework can discipline a story, but an analyst can still score the story. Evidence Coverage, abbreviated EC, helps show how much required evidence is verified, yet EC is not a buy signal and cannot validate the judgment used to interpret the evidence. Management-produced data may be incomplete or framed favorably, conflicts across evidence tiers score zero until resolved, and evidence may arrive too late. The policy therefore preserves the distinction between a well-structured hypothesis and a completed underwrite.

Concentration: the comparator can violate the risk the active sleeve is forbidden to add

The fifth attack is concentration. The diagnostic comparator proxy contains mega-cap exposures that breach the original artificial-intelligence factor stress-loss limit and the joint artificial-intelligence-plus-power limit, while the active loss arithmetic caps ordinary single-name exposure near 2% in a normal 50% impairment case and near 2.5% in the best ordinary 40% severity case. This creates an uncomfortable asymmetry. The policy may be extremely strict about a deliberate active addition while inheriting a larger related exposure through the supposedly neutral comparator.

The first conclusion is that the comparator is not neutral. Its constituents, sectors, currencies, sovereigns, and common factors must be looked through before active exposure is considered. The top 20 constituents are scored individually and the remainder is bucketed by sector proxies. The proxy exercise reported an artificial-intelligence factor stress loss of 8.2% and a joint artificial-intelligence-plus-power factor stress loss of 9.0%, both above the original limits. Those figures are diagnostic and unratified because the actual investor comparator, the archived worksheet, the independent second pass, taxes, and the capped or blended implementation study remain incomplete.

The control is a forced choice rather than a renewable exception. A persistent embedded issuer or factor hard-cap breach cannot remain a piece of permanent paperwork. After four consecutive quarters it must lead to a written comparator modification, such as a compliant capped or blended design, or to a formal policy amendment that raises the cap and explicitly accepts the risk. Meanwhile, the original caps govern, and active additions to the affected artificial-intelligence and power exposures remain inadmissible. The annual report must state the gap between the risk-based cap and any feasibility cap in currency units. The policy also requires the actual top-20 work, an independent second scoring pass, and study of an executable capped alternative before a feasibility cap can be ratified.

The most important residual risk is semantic. Raising a cap does not diversify or remove concentration; it accepts concentration. A formal amendment may improve honesty without improving the exposure. Conversely, modifying a broad comparator can create tracking error, tax cost, complexity, and abandonment risk, and history contains periods in which capped or tilted alternatives underperformed an uncapped benchmark for many years. The comparator-design dispute therefore remains open. Its adjudicating work must compare capped or blended implementation with the broad design rather than presuming that either one is safe.

A hypothetical example illustrates the governance issue without giving a trade instruction. Suppose an actual comparator's look-through factor stress loss remains above a stated hard limit for four quarterly reviews. Continuing to call the breach temporary would make the limit fictional. The policy requires the owner-level governance process to document which design risk is being accepted: the embedded concentration of the broad comparator, or the tracking-error and implementation risks of a modified comparator. This chapter does not choose the product, raise the cap, or authorize a portfolio change. It explains why an explicit amendment is required and why amendment is acknowledgement, not risk removal.

Obsolescence: correct evidence may arrive after the loss

The sixth attack is obsolescence. Economic lives can change much faster than accounting schedules or historical franchise descriptions. Artificial-intelligence hardware can be replaced sooner than assumed. Software can be bundled, automated, or substituted. A healthcare franchise can lose relevance through a clinical result, safety finding, payer decision, or pipeline failure. A consumer brand can lose cultural relevance before trailing financial statements fully reveal the shift.

Economic-depreciation tests are the first control. Reported depreciation is not automatically accepted as the economic cost of keeping a business competitive. The research process looks for retirement curves, impairment, resale or rental prices, utilization, vintage margins, maintenance capital, and customer cash return. The policy also

refuses a defensive-capex add-back while the underlying dispute remains open. Spending described as necessary to defend a franchise is still a cost unless evidence shows that it preserves or expands owner cash returns.

The 36-month company falsifiers force each thesis to identify evidence that would disconfirm it. They are not promises that obsolescence will reveal itself neatly within three years. They are controls against a thesis becoming unfalsifiable. A hypothetical example is an accelerator fleet whose accounting life is five years but whose rental price and utilization collapse after two years because a newer generation is far more efficient. Reported earnings may initially look stronger than economic earnings because depreciation follows the slower schedule. The policy would examine actual retirement, impairment, resale, rental, utilization, and customer-return data rather than treating the accounting life as proof. This example is not a claim about any named company or current asset.

The residual risk is timing. Evidence can arrive after value loss. A safety event, regulatory decision, technical leap, or consumer shift may be discontinuous. Quarterly filings may lag operating reality, and even well-defined falsifiers can confirm only after the market has repriced. The policy reduces the size of this error through severity floors, loss budgets, factor limits, and hard-gate review clocks, but it cannot guarantee that analysis precedes price discovery. This is one reason no story quality, evidence score, or classification creates an entitlement to a large weight.

Political and regulatory dependence: law can dominate business quality

The seventh attack is political and regulatory dependence. Ratings, payments, semiconductors, defense, utilities, healthcare, banks, and emerging-market platforms all depend in different ways on law, licensing, reimbursement, sanctions, capital rules, procurement, export controls, rate cases, or state tolerance. A firm may have strong demand and attractive operating economics while a legal or political discontinuity overwhelms both.

The payer gate controls one part of this problem by distinguishing announced intent from funded and collectible cash. Binary-license severity recognizes that the loss distribution can change abruptly when permission to operate is withdrawn. Jurisdiction and currency look-through prevent the company label from hiding where cash is earned, which sovereign can intervene, which currency funds liabilities, and how capital can move. The custody tiers separately address legal title and operational access: a same-jurisdiction second custodian is the basic Tier A control; a second legal system holding plain instruments is Tier B and requires the relevant tax, reporting, title, transfer, sanctions, and estate-situs work; complex entities, trusts, or wrappers are Tier C and default to zero until the full professional review passes.

A hypothetical example would be a highly profitable regulated service whose allowed economics depend on periodic rate approval. Demand may be necessary and recurring, but shareholder cash still depends on the regulator allowing cost recovery, the customer base being able to pay, and the legal framework remaining credible. The policy therefore does not let necessity substitute for the appropriation or rate-case-to-cash bridge. Again, this is an explanatory example, not a recommendation.

The residual risk is irreducible discontinuity. Diversification across issuers may not protect against a state action that affects an entire sector, currency, legal system, or transfer channel at once. A sanctions change, capital control, licensing decision, reimbursement reset, or expropriation can occur faster than a normal evidence cadence. Custody diversification can reduce a single operational point of failure, but it cannot make every legal system independent or eliminate tax and estate interactions. Where required professional evidence is missing, the relevant structure or position remains at zero. This policy does not authorize creation of a custody structure, account, hedge, transfer, or transaction.

Implementation complexity: the institutional process fails the one-person test

The eighth attack is implementation complexity, and it produces the chapter's clearest verdict. Forty-two dashboards multiplied by five pillars already imply 210 pillar assessments before factor scoring, reverse discounted-cash-flow work, evidence grading, tax lots, custody and title checks, jurisdiction mapping, and exception monitoring. The apparent simplicity of a one-page checklist can hide the work needed to populate it. A process that cannot be maintained is not conservative. It creates stale labels, missing evidence, selective monitoring, and false confidence.

The full apparatus is not executable by one person. That conclusion resolves dispute U4-1 adversely for the full version. The bounded version is accepted only conditionally. The full research queue is not a live portfolio, and one operator may maintain no more than eight live full underwrites. Everything else remains annual-check research and loses its classification after 12 months if not refreshed. Eight is a ceiling, not a target and not an authorization. Prospective infeasibility reduces the live count before work begins, possibly to four or zero; it is not something to discover after the operator has already taken active risk.

The hours arithmetic makes the limit concrete. Recurring governance is estimated at roughly 40 hours per year. Each live full underwrite requires approximately 15 to 25 hours per year. Eight names therefore imply about 160 to 240 hours annually in total, before one-off comparator work, the year-one independent factor-scoring pass, or Tier B custody work. This is a substantial part-time workload. Even eight may fail the active-viability and operator-capacity tests. If the hours cannot be explicitly budgeted, the active work is suspended or the live set is reduced in advance.

Cadence is deliberately narrowed. Survival matters are checked monthly: the investor liquidity floor, custody or title alerts, and hard-gate events. Exceptions and relevant holdings or factor headroom are reviewed quarterly rather than every company being fully rescored. The actual comparator's top 20 constituents and sector buckets are refreshed annually and on a material event, with the independent second pass required in year one. The actual comparator must be scoreable reproducibly within 16 operator hours per year; inability to do so is a portfolio-level falsifier that freezes new active work. Annual review also covers funded status, liability sensitivities, the viability calculation, hours, research sunset, map decay, and any required amendment.

The simplified process has seven linked parts. It begins with verified liabilities and funded status, where funded status is investable assets after known taxes and transaction costs divided by the probability-weighted present value of liabilities plus the safety margin. It then establishes the investor liquidity floor and separates nominal matching for nominal liabilities from verified inflation-linked matching for real liabilities. It names one implementable comparator and performs annual top-20 look-through. It applies issuer and factor hard limits with forced amendment for persistent breaches. It caps live active full underwrites at eight or fewer. It uses one decision journal and one exception log rather than parallel paperwork. Finally, it runs the annual viability audit, including the prospective 2 times cost threshold and the retrospective two-negative-year rule.

These seven parts are not a lighter route around the substantive gates. They determine which work is live. Deep company exhibits, legal and tax memoranda, reserve triangles, deposit cohorts, stressed capital, and current reverse discounted-cash-flow schedules are omitted from the recurring one-page layer only because a decision that requires them stays at zero until they are completed. The simplification is therefore risk-monotone: it removes active scope rather than lowering evidence standards.

The remaining implementation failure cases are straightforward. An operator may try to keep all 42 names mentally current despite the eight-name rule. A classification may remain visible after its evidence has expired. The top-20 proxy may consume more than 16 hours or prove irreproducible. Personal opportunity cost may be valued at zero, making any research program appear viable. Exceptions may be recorded without changing behavior. Legal and tax work may be postponed while an economic thesis is treated as complete. Each failure undermines the control

at its weakest point. The policy's response is not an improvised substitute or a relaxed gate. It is less active scope, paper mode, or the named comparator plus the investor liquidity floor and verified nominal and real matching.

The final one-person verdict and the approval boundary

The final bias is intentionally asymmetric about complexity. Strategic design receives priority because it applies to most of the portfolio: liabilities, liquidity, matching, custody integrity, the comparator, and look-through concentration. Single-name research is secondary because it applies to a bounded sleeve and offers only a small modeled contribution at the permitted position sizes. The comparator-engineering-first minority makes this resource claim explicit: a small improvement or avoided embedded loss across 75% to 95% of a strategic sleeve can matter more than the entire gross contribution of a small active program. Annual attribution must therefore distinguish value associated with comparator and matching design from value associated with company selection.

At the same time, the policy does not declare all active research useless. Paper-journal mode preserves a record and some evaluative readiness without committing capital below the 2 times cost threshold. Above that threshold, the live program is still bounded, monitored, and subject to suspension after two negative years. The zero-active view and the dislocation-option counterview remain visible; neither may be converted into a hidden permanent default without the stated evidence.

Most importantly, this chapter is not implementation approval. It authorizes no trade, order, target allocation, entry price, hedge, account opening, transfer, outbound message, or portfolio change. The policy remains uncertified because liability closure, actual funded status, trait tolerance, the mark-to-market drawdown budget, tax basis, domicile and custody structure, the actual comparator, factor look-through, the withdrawal-path grind budget, company Evidence Coverage audits, and current valuations remain incomplete or unverified. A diagnostic maximum weight is a ceiling produced by governance arithmetic, never a recommendation. A research classification is a queue position, never ownership approval. Any portfolio action requires the separate implementation process and the relevant owner approval after all applicable evidence, valuation, legal, tax, portfolio, and policy gates pass.

Annex E - Glossary, Formula Registry, and Parameter Register

E.1 Purpose, scope, and interpretive rules

This annex is the controlled dictionary and numerical register for the Investment Policy 2026-2036. It consolidates the material acronyms, variables, formulas, bands, floors, caps, decision clocks, review horizons, expiry rules, and breach consequences found in the controlling concise policy and explained in the expanded manuscript. It does not create an allocation, certify an investor, or authorize implementation. Every company named in the policy has a current admissible weight of zero. A numerical ceiling is a limit after all gates pass, never a target or permission to trade.

The governing hierarchy is: (1) the final debate-integrated concise policy dated 22 August 2026; (2) its Appendix D, which adopts F5-1 through F5-14 and supersedes inconsistent earlier formulations; (3) the remainder of the integrated concise policy; (4) the expanded chapters, which explain but cannot amend the concise policy; and (5) earlier debate rounds and traceability records, which preserve rationale and disagreement but cannot override an adopted final rule. Company research files, worksheets, legal and tax memoranda, and later evidence supply inputs; they do not silently amend policy.

The status labels in this annex mean:

Status	Meaning	Treatment
Durable	A constitutional or operating principle intended to survive the dated worldview.	Change only by explicit written policy amendment through the owner-approved governance process. An analyst, operator, or explanatory chapter cannot waive it.
Calibrated	A numerical operating rule selected by the debate but acknowledged to depend on verified investor facts or later experience.	Apply as written once its inputs are verified; change by written amendment, with independent challenge where the rule expressly requires it.
Provisional	A temporary, first-pass, proxy-based, or unaudited value.	It cannot certify implementation. Refresh or ratify on the stated clock; otherwise it expires, reverts, or fails closed.
Illustrative	A worked example used to expose arithmetic or a boundary, not a forecast or live input.	Never substitute it for investor facts or convert it into a portfolio target.
Unresolved	A disputed value or design for which the policy preserves competing views and adjudicating observables.	Do not choose silently. Maintain the conservative or zero default specified by the policy until evidence and the formal authority resolve it.

Unless a row states otherwise, **amendment authority** means the policy owner acting through a documented formal amendment and the separate implementation-governance process. The source does not name a broker, adviser, committee, researcher, or model as the holder of trading authority. Independent challenge tests an amendment or input but is not approval to transact. A breach consequence stated as “freeze,” “no add,” “close/reduce,” or “act” remains subject to the separate implementation process; this annex itself grants no execution authority.

Several types of number must not be conflated. A portfolio band is not a target. A hard cap cannot be exceeded; a preferred cap identifies the normal bound below the hard cap. A diagnostic ceiling is the most an arithmetic test could admit after every other gate passes. A falsifier is a preregistered research condition, not necessarily an automatic sell trigger. A worked-example number has no authority outside the example. A dated proxy output can identify a governance problem but cannot ratify a new cap.

E.2 Glossary

E.2.1 Portfolio, liability, and governance terms

Term or acronym	Meaning and units	Governing interpretation	Primary source
A*	Adjusted investable assets, in portfolio currency. Gross accessible assets less known taxes, transaction/exit costs, settlement costs, and other known implementation friction.	Inaccessible assets are not freely investable merely because they appear on a statement. It is the numerator of funded status.	Concise Sections 2-3; Expanded Chapters 3 and 15.
PVL*	Probability-weighted present value of liabilities plus an explicit safety margin, in portfolio/liability currency.	Includes future claims the portfolio must fund: essential spending, taxes, health, education, insurance gaps, property, fixed and contingent obligations. Timing, inflation, currency, discount rates, and probabilities must be explicit.	Concise Section 3; Expanded Chapter 3.
FS	Funded status, dimensionless ratio A* / PVL*.	Solvency selector, not a return forecast or risk-tolerance score. FS=1.00 means adjusted assets equal stressed liabilities.	Concise Sections 2-3; Expanded Chapter 3.
Coverage	Assets divided by the relevant liability measure after a stated stress; dimensionless ratio.	Coverage >= 1 is the conceptual closure condition, but actual liability closure remains uncertified. In the grind example, Coverage is end assets divided by stressed real liabilities.	Concise Sections 6 and 15; Round 5 verdict.
Strategic sleeve	Percentage of total investable assets assigned to the low-complexity architecture.	Contains the named comparator, nominal and verified real-liability matching, ILF, and only validated diversifiers. Broad does not mean riskless.	Concise Section 3; Expanded Chapter 3.
Opportunity sleeve	Percentage of total assets available for fully underwritten active names and nested speculative options.	It is capacity, not a required allocation. If nothing qualifies, unused capacity remains strategic. Strategic and opportunity maxima may not be added independently.	Concise Section 3; Expanded Chapter 3.
EB	Error budget: maximum portfolio budget for active permanent impairment, in percent of total portfolio.	Outer capacity, not expected loss and not an amount to spend. Tightened by funded state, capacity, traits, currency, liability timing, and the second axis.	Concise Section 3; Expanded Chapter 3.
MDB	Mark-to-market drawdown budget, in percent or currency loss.	Maximum peak-to-trough quoted-value loss consistent with liabilities, liquidity, and behavior. Separate from permanent-impairment EB and real-grind risk. Its personal value is unresolved/unverified.	Concise Sections 1, 3, 9, 15; Expanded Chapters 3 and 9.
Three budgets	Permanent impairment, mark-to-market drawdown, and decade-long real grind.	Passing one does not pass either of the others. Price decline alone is not permanent impairment; a crash test without withdrawals is not a full grind test.	Concise Section 2; Expanded Chapters 2, 6, and 9.
ILF	Investor Liquidity Floor, in liability currency.	Survival reserve for near liabilities and forced-sale prevention. Target and hard floor are distinct; the hard floor can never fund risk.	Concise Section 3.1; Expanded Chapter 3.
D36 / D24	Dated liabilities within the next 36 or 24 months, in liability currency.	D36 is the default ILF term. D24 may replace it only through the documented annual election at strong funded status and high spending flexibility.	Concise Section 3.1; Expanded Chapter 3.
Tax12	Known tax due within the next 12 months, in liability currency.	Not discretionary capital. Included in ILF target and known taxes also enter the hard floor.	Concise Section 3.1; Expanded Chapter 3.
Commit	Fixed or already undertaken commitments, in liability currency.	Includes verified legal, medical, contractual, and similar obligations. The hard floor protects twelve months of fixed commitments.	Concise Section 3.1; Expanded Chapter 3.
H	Income-interruption horizon, in years or other units consistent with the spending input.	Applies only to uncovered essential spending beyond D36/D24. No universal value is certified; derive from income stability, concentration, correlation, and liabilities.	Concise Sections 1 and 3.1; Expanded Chapter 3.
Second axis	Capacity overlay based on liability nearness/flexibility and labor/business-income correlation with market losses.	Can shorten MDB and EB one state. Flexibility can permit D24 but does not enlarge issuer loss budgets.	Concise Section 3; Expanded Chapter 3.

Term or acronym	Meaning and units	Governing interpretation	Primary source
Hysteresis	Asymmetric funded-state transition rule.	Promotion needs persistence; verified capacity deterioration demotes immediately. Prevents noisy threshold flipping and procyclical risk expansion.	Concise Section 3; Expanded Chapter 3.
Policy Exception Event (PEE)	A material deviation from the policy, such as using the ILF floor for risk, an unjournaled panic action, an undefended missed mandatory review, or deviation without amendment.	Logged as a governance failure, not as an exception-renewal mechanism. Two in a rolling year freeze new discretion.	Concise Sections 2-3 and 10; Expanded Chapters 3, 9, 10, and 15.
Risk-monotone	A response that reduces or contains the affected risk.	Emergency authority cannot be used to increase gross risk or buy a dislocation. Only the minimum action needed to protect title, solvency, or compliance is contemplated.	Concise Sections 2 and 9; Expanded Chapter 9.
Comparator	Named, low-cost, implementable baseline against which active complexity is judged.	It is a real position with issuer, factor, currency, sovereign, fee, tax, and path risk; it is not neutral. The actual investor comparator is missing.	Concise Sections 2, 7, 10, and 15; Expanded Chapter 7.
Tracking error	Relative return/path deviation from the named comparator, normally measured as a percentage.	A ceiling on path risk, not a return promise. Comparator modification requires a tracking-error budget and predeclared abandonment point.	Concise Section 3; Expanded Chapters 3 and 7.
Trait tolerance	The investor's durable behavioral ability to bear loss and uncertainty.	Different from financial capacity and currently unverified. It cannot be inferred from confidence, a questionnaire alone, or a bull-market experience.	Concise Section 15; Expanded Chapters 2 and 15.
Shortfall Parity	Principle that excessive caution producing a modeled liability shortfall is governed as seriously as excess risk.	Does not authorize risk beyond verified capacity or trait tolerance. Remedy order is contributions, spending/timing changes, objective changes, then admissible risk.	Concise Section 2; Expanded Chapters 2-3.

E.2.2 Research, evidence, and security terms

Term or acronym	Meaning and units	Governing interpretation	Primary source
A-E pillars	Five non-compensating scores, each on a 0/2/4 anchor scale: A Demand/lock-in; B Structure/cycle/payer; C Capital economics; D Governance/per-share capture; E Valuation/falsifiability.	No additive total. Every pillar must be at least 2 for ordinary admission; any current pillar at 0 makes diagnostic maximum weight zero. E below 2 means no add and re-underwrite, not an automatic price-based sale.	Concise Section 4; Expanded Chapter 4; F5-2.
Pillar 0	A score of zero on any pillar.	Foundational gate failure. It triggers no add and full re-underwrite; A-D zero can lead to close/reduce under the governed exit clock.	Concise Sections 3.2 and 9; Expanded Chapters 4 and 9.
EC	Evidence Coverage: verified applicable required fields divided by applicable required fields, expressed as a percentage.	Coverage, not favorability. A complete adverse file can have high EC; a sparse positive file can have low EC. Weighted toward primary filings, contracts, and regulators.	Concise Section 4; Expanded Chapters 4 and 9.
ECmult	Epistemic multiplier determined by EC; dimensionless.	The only multiplier protected by the 0.50 floor after admission. EC below 60 remains research-only and is not rescued by the floor.	Concise Section 4.1; F5-1; Expanded Chapter 4.
T1 / T2 / T3 evidence	Evidence hierarchy: T1 denotes primary filings/contracts/regulators; T2 and T3 are secondary or lower tiers.	A material T1-versus-T2/T3 conflict scores zero until resolved; it is not averaged. A link or “evidence pointer” receives no EC credit until extracted and assessed.	Concise Sections 4, 8, and 13; Expanded Chapters 4, 8, and 15.
[A]	Assumption label, especially for provisional analyst scoring.	Current company-map EC, pillar, severity, and favorable financial-battery scores are unaudited analyst assumptions. The label prevents provisional work from masquerading as verified evidence.	Concise Sections 4, 8, and Appendix D; F5-6.
Reverse DCF	Reverse discounted-cash-flow analysis starting from price to infer required operating and cash-flow assumptions.	Valuation admission test. It must be current; quality or a low share price does not substitute for it.	Concise Sections 4 and 9; Expanded Chapters 4 and 9.
S	Bear-case permanent-impairment severity, expressed as a fraction or percentage of the position.	Not volatility and not probability. Subject to class floors to prevent weight inflation through optimistic severity.	Concise Section 4.1; F5-3; Expanded Chapter 4.
BaseIssuerLoss(FS)	Starting portfolio permanent-loss allowance for one issuer, as a percent of total portfolio, selected by funded state.	Applies only after all admission gates pass. It is then reduced by evidence, stress-liquidity, and factor-headroom multipliers.	Concise Section 4.1; Expanded Chapter 4.
AllowedLoss	Maximum total-portfolio percentage permitted to be permanently impaired by the issuer under the diagnostic assumptions.	A loss budget, not a proposed weight. Hard factor breach makes it zero.	Concise Section 4.1; Expanded Chapter 4.
Lmult	Stress-liquidity multiplier, dimensionless.	Based on stressed exit time, not normal-market impressions. Applies unfloored under F5-1.	Concise Section 4.1; Expanded Chapter 4.
Fmult	Common-factor-headroom multiplier, dimensionless.	Measures portfolio context after comparator look-through. Applies unfloored; if an addition breaches a hard cap, Fmult=0.	Concise Section 4.1; F5-1; Expanded Chapters 4 and 7.
MaxWeight / Wmax	AllowedLoss / S, expressed as percent of total portfolio. Wmax is the company-map shorthand.	Optimistic diagnostic ceiling only. It never supplies approval, timing, account suitability, or an order. A failed gate overrides arithmetic and prints zero.	Concise Sections 4.1 and 8; F5-2/F5-11.
Speculative option	Short-history, pre-cash-flow, paid-option, or research-optional exposure treated as capable of total loss.	Nested inside the opportunity sleeve, not a third sleeve. Per-name speculative cap binds before a larger formula result. No averaging down merely because price fell.	Concise Sections 3 and 4.1; Expanded Chapters 3-4.
Core / Watch / Speculative / Avoid	Research-priority classifications, not portfolio roles or buy ratings.	Every label is provisional and expires after 12 months without refresh. “Core” has no predictive claim. Current admissible weight remains zero for every class.	Concise Sections 8-9 and 15; F5-9; Expanded Chapters 8-9.
Unclassified Research Note	Status after a company classification expires unrefreshed.	Preserves old evidence as history while removing any claim that the classification is current. It is not a sell signal.	Concise Sections 8-9; F5-9; Expanded Chapter 9.
Falsifier	Preregistered observable that would materially disconfirm a causal thesis.	Must concern business reality and per-share economics, not price alone. A stated horizon organizes monitoring; it is not automatic execution authority.	Concise Sections 8-9; Expanded Chapter 9.
Research sunset	Attention rule that downgrades a theme with no decision or explicit abstention threshold by the review horizon.	Sunsets attention, not evidence. The archive remains; the line moves to annual-check status and records time and cost.	Concise Section 9; Expanded Chapter 9.

Term or acronym	Meaning and units	Governing interpretation	Primary source
Financial battery	Class-specific evidence gates for principal balance-sheet businesses.	Insurers: reserves, float cost, capital. Banks: deposit beta, full-cycle credit/NPL, capital. Managers: net flows, fee compression, client concentration, capital. Favorable scores are provisional and capped at 2 until named exhibits attach.	Concise Sections 1 and 4; F5-6; Expanded Chapter 4.
NPL	Non-performing loan ratio or behavior.	Part of the bank battery and a company-specific falsifier; the exact measurement definition belongs in the current bank underwrite.	Concise Sections 1, 4, and company map.
CET1	Common Equity Tier 1 regulatory capital measure.	Stressed CET1 evidence is required for bank capital claims. A favorable score is capped until the exhibit is attached.	Concise Sections 4 and 8; F5-6.
FCF / FCF per share	Free cash flow, in currency or per diluted share.	Must be normalized for maintenance capex, leases, dilution, financing, and business-class economics. Adjusted EBITDA or reported growth cannot substitute for owner cash.	Concise company map; Expanded Chapters 4 and 8.
ROE / ROIC / ROTCE	Return on equity, return on invested capital, and return on tangible common equity, expressed as percentages.	Research evidence only when normalized through a cycle and reconciled to capital, dilution, and risk. Current company-specific thresholds are provisional falsifiers, not portfolio parameters.	Concise company map; Expanded Chapters 4 and 8.
ILB	Inflation-linked bond.	Structural match only for verified real liabilities where index, law, currency, market, title, and tax integrity pass. It is not granted an assumed speculative gain.	Concise Sections 1 and 6; Expanded Chapter 6.
AI / EM	Artificial intelligence / emerging markets.	Theme or exposure descriptors, never expected-return proof. Flagged EM and certain AI-linked exposures carry severity/factor constraints.	Concise Sections 4-5 and 7-8.

E.2.3 Factor, scenario, provenance, and notation terms

Term or acronym	Meaning and units	Governing interpretation	Primary source
FEW_f	Factor-equivalent weight for factor f, percent of total portfolio.	Sum of each holding or bucket's total-portfolio weight times its 0-to-1 factor sensitivity. Liability-matched duration is excluded from directional rates FEW and reported separately.	Concise Section 7; F5-5; Expanded Chapter 7.
FSL_f	Factor stress loss for factor f, percent or currency loss of total portfolio.	Sum of stressed losses across holdings/buckets. It is a scenario diagnostic, not a forecast. Hard FSL caps consume headroom before active additions.	Concise Section 7; Expanded Chapter 7.
w_i	Weight of holding or bucket i in the total portfolio.	Not merely weight inside a fund. Used in FEW/FSL.	Expanded Chapter 7.
x_if	Estimated sensitivity of holding/bucket i to factor f, dimensionless from 0 to 1.	Standard grid: none, indirect, material, high, dominant. Archived rationale and evidence are required.	Expanded Chapter 7.
s_if	Stress loss rate applied to the factor-exposed portion of holding i, as a fraction.	Used in the decomposed FSL formula. Must be scenario-specific and reproducible.	Expanded Chapter 7.
f	A named common factor.	Mandatory families: AI-capex sustainability, state payer, China/export control, directional real rates/refinancing/quality duration, and power availability/cost; joint shocks are also required.	Expanded Chapter 7.
S1-S6	Six decade scenarios: productive diffusion; inflation/repression grind; fragmented blocs; crash/sequence; adaptation/physical constraint; plateau/muddle.	Conditional causal maps, not probabilities or forecasts. No scenario authorizes a security purchase.	Concise Section 5; Expanded Chapter 5.
D-1...D-5 / U2...U4	Numbered disputes and unresolved questions carried through the debate.	A closed dispute can reopen only on its stated evidence test; an open dispute cannot be settled by convenience.	Concise Appendix B; Expanded Chapter 10.
F1-F5	Sequential Fable amendment rounds.	F5-1 through F5-14 are the final corrections. Earlier wording remains provenance only where F5 supersedes it.	Concise Sections 1 and Appendix D; Round 5.
P1-P52	Policy article identifiers in the traceability appendix.	Provenance and cross-reference aids, not separate implementation authority.	Concise Appendix A.
pp / bps	Percentage points / basis points; 100 bps equals 1 percentage point.	“More than 5 percentage points” is an absolute allocation deviation; “more than 20% relative” is proportional to the approved boundary.	Concise Section 3.2; Expanded Chapter 3.
Stress day	A trading day under stressed exit assumptions.	Used for Lmult rather than ordinary-volume estimates.	Concise Section 4.1; Expanded Chapter 4.
Evidence cutoff	22 August 2026.	Later facts are outside the edition until incorporated through a current review. A source snapshot proves only that a disclosure existed, not that its claims were independently verified.	Concise opening and Section 15; Expanded Chapter 15.

E.3 Formula registry

ID and formula	Variables, units, and current calibration	Status and source	Amendment authority, review, and breach consequence
FR-1 Funded status: FS = A* / PVL*	A* and PVL* use consistent currency; FS is a ratio. Report base and liability sensitivities.	**Durable formula; calibrated bands.** Concise Sections 2-3; Expanded Chapter 3.	Policy owner may amend the definition only formally. Recalculate annually and immediately after a verified capacity event. If +/-10% PVL changes the band, use the lower state; if liabilities remain unmeasured at the next annual review, freeze new active work.
FR-2 Adjusted assets: A* = gross accessible assets - known taxes - known transaction/exit costs - other known implementation friction	Currency amount. The expanded worked example uses 1,500,000 - 100,000 - 50,000 = 1,350,000.	**Durable concept; example illustrative.** Expanded Chapter 3.	Inputs are verified investor facts, not amendment choices. Missing access/title facts exclude the amount or fail closed. Review annually and on tax, title, or access changes.
FR-3 ILF target: D36 + Tax12 + Commit + H x max(0, later essential spending - dependable income covering that later interval)	All cash-flow terms use liability currency; H uses time units consistent with spending. D24 may replace D36 only under FR-5.	**Durable structure; H unresolved; numerical inputs investor-specific.** Concise Section 3.1; Expanded Chapter 3.	Formula changes require amendment. Calculate at least annually and monitor ILF monthly. Falling below target requires restoration; it does not make the target-floor gap available for risk.
FR-4 ILF hard floor: six months uncovered essential spending + twelve months fixed commitments + known taxes	Currency amount.	**Calibrated minimum.** Concise Section 3.1; Expanded Chapter 3.	Formal amendment only. Check monthly and on liability events. Breach is a survival event: no funding of risk, restore liquidity, log any use for risk as a material Policy Exception Event.
FR-5 D24 election: replace D36 with D24	Requires verified FS >= 1.75, survival of the sensitivity test, and documented high spending flexibility.	**Calibrated annual election.** Concise Sections 1 and 3.1; F4-12; Expanded Chapter 3.	Annual documented owner-governance election; not an analyst default. It lapses when either condition fails. Consequence is reversion to D36 and ILF recalculation, not larger issuer budgets.
FR-6 Evidence Coverage: EC = verified required fields / applicable required fields	Percentage, weighted toward T1.	**Durable formula; calibrated thresholds.** Concise Section 4; Expanded Chapter 4.	Evidence methodology changes require amendment. Maintain continuously enough for quarterly breakage and annual audit. EC below 60 makes research-only, triggers no add and a 30-day full re-underwrite.
FR-7 Allowed loss: AllowedLoss = BaseIssuerLoss(FS) x max(0.50, ECmult) x Lmult x Fmult	Percent of total portfolio. Base loss, evidence, stress-liquidity, and factor headroom are separate inputs.	**Durable corrected formula.** Concise Sections 1 and 4.1; Appendix D F5-1; Expanded Chapter 4.	Only a formal amendment can change it. Recompute before any proposed admission and on material evidence, liquidity, factor, or funded-state change. Failed gate prevents use; Fmult=0 makes AllowedLoss zero.
FR-8 Maximum diagnostic weight: MaxWeight = AllowedLoss / S	AllowedLoss and S must use compatible fraction/percentage notation; result is percent of total portfolio.	**Durable formula; inputs calibrated/provisional.** Concise Section 4.1; Expanded Chapter 4.	Formal amendment required. Recompute with the current underwrite and portfolio. Apply the lowest of this result and every other cap. Any current pillar at 0, hard gate, or implementation failure makes Wmax/admissible addition zero.
FR-9 Factor-equivalent weight: FEW_f = sum_i(w_i x x_if)	Percent of total portfolio. x_if uses 0, .25, .50, .75, 1.00 grid.	**Durable method; scores provisional.** Concise Section 7; Expanded Chapter 7.	Method/caps change formally. Refresh annually, on material event, and independently after six months in year one. Unreproducible scoring is a certification blocker; inability to maintain the actual comparator within 16 hours/year freezes new active work.
FR-10 Factor stress loss: FSL_f = sum_i(w_i x loss_if) or sum_i(w_i x x_if x s_if)	Percent or currency loss of total portfolio under a named common stress.	**Durable method; scenario assumptions provisional.** Concise Section 7; Expanded Chapter 7.	As for FR-9. A proposed active addition that breaches a hard factor cap sets Fmult=0; an embedded comparator breach persisting four quarters forces the comparator/cap choice.
FR-11 Former feasibility-cap rule: max(risk cap, comparator FSL + 1 pp)	Would have produced a floating cap.	**Superseded and not controlling.** Concise Section 1 records F4-7; F5-4 rejects ratification based on the incomplete proxy and forbids floating with index concentration.	Do not use. The original risk caps govern until the actual study and formal ratification. Any temporary diagnostic ceiling must be fixed, no more than 2x the original cap, and auto-revert at month 12.

ID and formula	Variables, units, and current calibration	Status and source	Amendment authority, review, and breach consequence
FR-12 Derating loss: w x (1 - Mconservative / Mcurrent)	w is total-portfolio weight; both multiples use the same operating denominator; result is portfolio-level loss fraction.	**Durable formula; calibration at 50%/75% MDB.** Concise Sections 1 and 9; Expanded Chapter 9.	Current/conservative multiples and reverse DCF must be archived before purchase. Aggregate across active holdings after strategic stress. Above 50% and up to 75% of MDB requires written amendment and independent challenge; above 75% is not admitted. Ex-post overshoot with intact thesis causes a sizing post-mortem, not automatic sale.
FR-13 Prospective active-program value: portfolio value x active sleeve x evidenced premium x evidenced hit rate	Currency per year; premium and hit rate as fractions. Compare with fully loaded annual research/implementation cost.	**Calibrated/provisional operating test.** Concise Sections 1, 10, and 14; Appendix D F5-8; Expanded Chapters 10 and 14.	Run annually before deploying active capital. Active capital requires expected contribution at least 2 x fully loaded annual cost. Below threshold: paper-journal mode and zero active capital. Above threshold: two consecutive negative realized-viability years still suspend new active work.
FR-14 Tax-deferral comparison	Applicable statutory tax rate, fixed 10-year horizon, and comparator replacement return; the manuscript does not provide a single closed-form equation.	**Calibrated method, investor-specific inputs.** Concise Sections 1 and 9; F4-9; Expanded Chapter 9.	Tax assumptions require current professional/investor facts; changing the 10-year standard requires amendment. Two annual holds chiefly justified by deferral require independent challenge. Price every proposed delay on day 0.
FR-15 Grind end-assets: A_end = A_start x sum_j(weight_j x (1 + real_return_j)), equivalently the printed weighted-loss arithmetic	Currency; compare A_end/PVL* for Coverage.	**Durable diagnostic structure; current inputs illustrative and incomplete.** Concise Section 6; Expanded Chapter 6; F5-12.	Re-run when the actual liability/withdrawal schedule and admissible vehicles exist. A static immediate pass does not certify a withdrawal-path grind. Missing withdrawal path remains a certification blocker.

E.4 Parameter register

E.4.1 Document, certification, and standing boundaries

Parameter	Meaning and units	Current value/range	Status	Source section	Amendment authority and review	Breach consequence
Policy horizon	Stated investment-policy period	2026-2036	**Dated**	Title; Expanded Chapter 15	Policy owner; revisit in a successor edition or material policy review.	Do not extrapolate decade beliefs indefinitely.
Evidence cutoff	Latest date included in the integrated synthesis	22 Aug 2026	**Durable boundary for this edition**	Concise opening and Section 15	Cannot be altered retroactively; later facts enter a refreshed edition/current underwrite.	Facts after cutoff cannot be represented as part of the dated synthesis.
Current security authorization	Admissible weight and transaction authority for every named company	0%; no trade, order, hedge, message, account opening, target, entry, or portfolio change	**Durable standing boundary**	Concise opening, Sections 8 and 15; Expanded Chapter 15	Only a separate current implementation process after all gates, investor facts, tax/legal checks, and approvals.	Any action inferred from a theme, label, ceiling, or example is unauthorized.
Company-map population	Research-map breadth, not live-underwrite count	42 names: 11 Core, 13 Watch, 6 Speculative, 12 Avoid	**Provisional snapshot**	Concise Section 8	Annual map refresh; classifications expire independently.	Do not treat map as a 42-name portfolio or 42 maintained underwrites.
Live full-underwrite maximum	One-person maintenance capacity	At most 8; prospectively fewer, possibly 0	**Calibrated/provisional**	Concise Sections 12 and 14; F5-14; Expanded Chapters 9, 10, 14, 15	Annual viability/hours review; operator reduces count before work. Formal amendment for a higher standing maximum.	Unmaintainable scope freezes/suspends active work; stale names revert to research notes.
Certification blockers	Items that must close before policy-level certification	Withdrawal-path grind, liability closure, actual-comparator scoring, company EC audits; also actual FS, trait tolerance, MDB, tax basis, domicile/custody, factor look-through, valuation	**Unresolved**	Concise Sections 12 and 15; F5-12; Expanded Chapter 15	Verified evidence and separate approval; no silent waiver. Review at least annually and when an item closes.	Continued non-certification and non-authorization.

E.4.2 Funded-state bands and portfolio architecture

All percentages below are percentages of total investable assets unless noted. They are coherent ranges, not targets. The strategic and opportunity maxima cannot be added independently. The actual selected point can be lower or zero because of the second axis, traits, liquidity, evidence, implementation, or another hard gate.

FS state	Strategic sleeve	Opportunity sleeve	EB	Illiquid ceiling	Active tracking-error ceiling	Speculative-loss sleeve	BaseIssuerLoss	Status/source	Amendment, review, and breach consequence
FS < 1.00 shortfall	95-100%	0-5%	0-3%	0-5%	0-2%	0-2%	0.25%	**Calibrated, conditional.** Concise Sections 3 and 4.1; Expanded Chapter 3.	Formal amendment for bands. Verify annually and after capacity events. Zero is meaningful; a shortfall does not authorize gambling for return.

FS state	Strategic sleeve	Opportunity sleeve	EB	Illiquid ceiling	Active tracking-error ceiling	Speculative-loss sleeve	BaseIssuerLoss	Status/source	Amendment, review, and breach consequence
FS 1.00-1.24 marginally funded	90-95%	5-10%	3-5%	0-8%	0-3%	2-3%	0.50%	**Calibrated, conditional.** Same sources.	Same. Near the lower boundary, liability error may require the conservative end or zero.
FS 1.25-1.74 funded with surplus	80-90%	10-20%	5-8%	5-12%	0-4%	3-5%	0.75%	**Calibrated, conditional.** Same sources.	Same. Illiquid lower bound does not compel private ownership; vehicle-specific gates still apply.
FS >= 1.75 strongly funded	75-90%	10-25%	normal 8-10%; hard 12%	normal 10-15%; hard 20%	0-5%	preferred 5%; hard 7%	1.00%	**Calibrated, conditional.** Same sources.	Same. Normal ranges require evidence; hard caps may not be exceeded. Strong funding does not enlarge issuer severity or factor headroom.

Additional funded-state parameter	Units/current value	Status	Source	Authority/review	Breach consequence
PVL sensitivity set	Base, -10%, +10%, -20%, +20%	**Durable method**	Concise Sections 1, 3, and 12; Expanded Chapter 3	Annual FS review and material liability change; formal amendment to change grid.	If the +/-10% cases change the state, the lower state governs; all +/-20% cases still print.
Promotion persistence	2 consecutive annual confirmations	**Calibrated hysteresis**	Concise Sections 1 and 3; Expanded Chapter 3	Annual owner-governance confirmation.	No higher-state discretion after only one favorable annual crossing.
Demotion clock	Immediate after verified death, incapacity, income loss, leverage, litigation, custody impairment, or newly inflexible spending	**Durable risk-monotone rule**	Concise Section 3; Expanded Chapter 3	Recalculate immediately; no amendment needed to tighten.	Apply lower state/budgets before new active decisions; recalculate PVL, ILF, and second axis.
Second-axis reduction	Shorten MDB and EB by one funded state	**Calibrated**	Concise Section 3; Expanded Chapter 3	Apply when liabilities are near/inflexible or income is highly market-correlated; review annually and on capacity events.	More conservative state's MDB/EB governs; at lowest state, most restrictive treatment already applies.
D24 eligibility	FS >=1.75 plus documented high spending flexibility, surviving sensitivity test	**Calibrated annual election**	Concise Sections 1 and 3.1; Expanded Chapter 3	Document every year.	Loss of either condition reverts ILF term to D36 immediately/at recalculation.
Per-option speculative market value	Normally 0.25-0.75% and assumed 100% loss	**Calibrated**	Concise Sections 3 and 4.1; Expanded Chapter 3	Formal amendment for range; current full underwrite/approval still required.	Lower of per-name cap, sleeve-loss cap, formula, factor/legal gates applies. No probability discount or price-only averaging down.

E.4.3 Liquidity, rebalancing, emergency, and exception parameters

Parameter	Meaning and units	Current value/range	Status	Source section	Amendment authority and review	Breach consequence
ILF home	Currency/duration/title standard	Home-liability currency; short duration; high title/custody integrity	**Durable**	Concise Section 3.1; Expanded Chapter 3	Policy owner for principle; implementation uses verified liabilities and instruments. Monthly review.	A reserve that fails currency, duration, or title integrity does not satisfy ILF.
ILF hard floor	Emergency minimum	6 months uncovered essential spending + 12 months fixed commitments + known taxes	**Calibrated**	Concise Section 3.1	Formal amendment only; monthly check and event-driven recalculation.	Never buy risk from it. Any use for risk is a material PEE; restore reserve/freeze as applicable.
ILF target restoration	Response when target, but not floor, is breached	Restoration required; no exact calendar supplied	**Durable consequence; timing unresolved**	Expanded Chapter 3	Owner-governance restoration plan.	Target-floor gap is not spare risk capital. Residual ambiguity: source states restoration but no fixed deadline.
Strategic absolute rebalance trigger	Deviation from approved band boundary	More than 5 percentage points	**Calibrated**	Concise Section 3.2; Expanded Chapter 3	Formal amendment for threshold; monitor through holdings reviews.	Governed action is triggered, subject to ILF and tax; do not refill broken theme mechanically.
Strategic relative rebalance trigger	Relative deviation from approved boundary/reference	More than 20% relative	**Calibrated**	Same	Same	Same. Either absolute or relative test can trigger.
Hard-gate freeze	Title, fraud, de-licensing, financing-covenant, or other integrity threat	Immediate freeze; resolve/minimum risk-monotone action in 0-5 trading days	**Durable clock**	Concise Sections 3.2 and 9; Expanded Chapters 3 and 9	Separate authorized implementation process; analyst cannot trade.	No add immediately; contain/reduce risk within clock if authorized. Five days is a ceiling, not required waiting.
Pillar/evidence re-underwrite	Pillar 0, hard-cap problem, 2 missed milestones, or EC <60	Full re-underwrite within 30 calendar days; no add immediately	**Durable/calibrated clock**	Concise Sections 3.2 and 9; Expanded Chapters 3, 4, and 9	Formal amendment for clock; current research owner performs review.	Fully underwritten status is not restored unless evidence/pillars/gates pass.
Fundamental close/reduce review	Any A-D pillar becomes 0, a hard cap persists, or full-underwrite gate fails	5-30 days	**Durable exit hierarchy**	Concise Section 9; Expanded Chapter 9	Separate implementation approval; review cannot silently amend a cap.	Close or reduce under governed process; no renewable exception.
Valuation-only deterioration	E below admission while A-D thesis intact	No add; re-underwrite; no mechanical trim	**Durable distinction**	Concise Sections 3.2 and 9; Expanded Chapter 9	Current underwrite and separate implementation authority.	Hold may persist only if standardized after-tax comparison clears; price movement alone never authorizes add or forces sale.
Entry-pacing outer clock	Event-anchored entry, if separately authorized	Complete by day 180	**Calibrated, non-authorizing**	Concise Sections 1 and 3.2; F4-8; Expanded Chapter 3	Formal amendment for horizon; entry itself needs separate approval.	Tranche expires; absent pending information after all gates clear, full deployment is statistical default, not authorization.
Exception documentation	Emergency exception record	Within 24 hours	**Durable**	Concise Section 2; Expanded Chapters 9 and 15	No analyst waiver.	Unrecorded exception is a governance breach/PEE.
Exception independent review	Challenger review	Within 7 days	**Durable**	Same	Independent reviewer; not execution authority.	Exception cannot become precedent or remain unreviewed.
Emergency exception expiry	Maximum life	Within 30 days	**Durable**	Same	Formal amendment needed for enduring rule change.	Exception ends; persistent condition requires compliant design or amendment.

Parameter	Meaning and units	Current value/range	Status	Source section	Amendment authority and review	Breach consequence
Rolling PEE limit	Behavioral/process breach threshold	2 material events in any rolling 12 months	**Calibrated**	Concise Sections 3.2 and 10; Expanded Chapters 3 and 9	Formal amendment for threshold; review quarterly via exception log.	Freeze new discretion pending governance review; valid mechanical actions may continue if other gates permit.
Post-decision “next liquid window”	Execution clock after governed tax/valuation decision	5 trading days, absent security-specific integrity event	**Calibrated**	Concise Sections 1 and 9; F4-9; Expanded Chapter 9	Separate implementation authority.	Do not use “next liquid window” as indefinite delay. Integrity events use faster hard-gate logic.

E.4.4 Scorecard, evidence, loss, and sizing parameters

Parameter	Meaning and units	Current value/range	Status	Source section	Amendment authority and review	Breach consequence
Pillar score anchors	Ordinal research score	0, 2, 4	**Durable architecture**	Concise Section 4; Expanded Chapter 4	Formal amendment only; current underwrite refresh.	No additive total; intermediate judgments may be explained, but admission uses the non-compensating rule.
Ordinary pillar minimum	Minimum each of A-E	At least 2	**Durable hard gate**	Same	Formal amendment only.	Any pillar below 2 blocks ordinary admission; any 0 prints Wmax 0 under F5-2.
Pillar E 4 hurdle	Conservative expected return versus comparator	Conservative case clears comparator by +1 pp	**Calibrated**	Concise Section 4	Formal amendment for hurdle; current reverse DCF at each underwrite/annual refresh.	Without this, E cannot score 4; strength elsewhere cannot compensate.
EC research-only boundary	Verified applicable evidence	<60%	**Calibrated**	Concise Section 4; Expanded Chapter 4	Formal amendment for boundary; monitor quarterly and annually.	Research-only; no ordinary admission; 30-day re-underwrite/no add if an existing file falls below.
ECmult low band	Evidence multiplier	EC 60-74 -> 0.50	**Calibrated**	Concise Section 4.1; F5-1	Formal amendment only.	Applies after admission; it cannot rescue a failed pillar or EC <60.
ECmult middle band	Evidence multiplier	EC 75-84 -> 0.75	**Calibrated**	Same	Same	Reduces AllowedLoss; current company-map values remain [A].
ECmult full band	Evidence multiplier	EC >=85 -> 1.00	**Calibrated**	Same	Same	Full coverage multiplier is not a buy signal and cannot cure adverse evidence.
Material evidence conflict	T1 versus T2/T3 contradiction	Affected item scores 0 until resolved	**Durable**	Concise Section 4; Expanded Chapter 4	Evidence reviewer resolves; no averaging discretion.	Failed affected gate/coverage until the contradiction closes.
Financial-battery favorable-score cap	Provisional battery evidence	Maximum 2 for sizing until named exhibits attach	**Durable final correction**	Concise Appendix D F5-6; Expanded Chapter 4	Current independent evidence review; policy amendment to remove requirement.	Missing/failed battery blocks Core; apparent 3/4 cannot enlarge sizing.
Insurer exhibit horizon	Reserve evidence	10-year reserve triangles plus float-cost and stressed-capital exhibits	**Calibrated evidence requirement**	Concise Sections 8 and 11; F5-6	Current underwrite/challenger.	Battery remains provisional and capped at 2.
Ordinary-equity S floor	Permanent-impairment severity	>=40%	**Durable anti-gaming floor**	Concise Sections 1 and 4.1; F5-3	Below-floor claim needs written evidence plus independent challenge; changing class schedule requires amendment.	Use at least 40%; no legacy 5%/8% conviction tier.
Levered financial/cyclical/roll-up S floor	Severity	>=60%	**Durable**	Same	Same	More severe applicable class floor governs.
Binary-license, pre-cash-flow, flagged-EM S floor	Severity	80-100%	**Durable**	Same	Same	Use class-appropriate point in range; speculative/pre-cash-flow positions use S=100%.

Parameter	Meaning and units	Current value/range	Status	Source section	Amendment authority and review	Breach consequence
Liquidity multiplier 1	Stressed exit time	<=5 days -> Lmult 1.00	**Calibrated**	Concise Section 4.1	Formal amendment; reassess on liquidity events.	Use current slower band if evidence no longer supports exit speed.
Liquidity multiplier 2	Stressed exit time	6-30 days -> 0.75	**Calibrated**	Same	Same	AllowedLoss reduced without floor rescue.
Liquidity multiplier 3	Stressed exit time	31-365 days -> 0.50	**Calibrated**	Same	Same	Same.
Liquidity multiplier 4	Stressed exit time	>1 year -> 0.25	**Calibrated**	Same	Same	Same; full evidence does not neutralize illiquidity.
Factor multiplier - full	Use of applicable factor cap	Full headroom -> 1.00	**Calibrated**	Concise Section 4.1	Recompute with quarterly look-through and material events.	Apply next lower multiplier as cap use rises.
Factor multiplier - crowded	Factor use >75%	0.75	**Calibrated**	Same	Same	AllowedLoss reduced.
Factor multiplier - near cap	Factor use >90%	0.50	**Calibrated**	Same	Same	AllowedLoss reduced.
Factor multiplier - breach	Addition breaches hard cap	0	**Durable hard gate**	Concise Sections 1, 4.1, 7; F5-1	Only formal policy amendment can change the cap; current comparator study does not waive it.	AllowedLoss and admissible active addition are zero.
Legacy single-name tiers	Prior generic concentration rows	5% and 8%	**Superseded/void**	Concise Sections 1, 4.1, Appendix D F5-3	Cannot be revived without formal amendment and D-1 evidence.	Do not use; class S floors control.
D-1 reopening record	Evidence to reconsider concentration minority	Preregistered 5 years and at least 20 fully underwritten decisions, independently audited after-tax attribution, verified liability surplus, challenger sign-off	**Durable reopening test; currently unmet**	Concise Sections 1 and Appendix D F5-13; Expanded Chapter 15	Formal amendment after all conditions.	Until met, no conviction tier and ordinary best-case ceiling remains about 2-2.5%.
Additional >3% diagnostic-size scrutiny	Research record	2 preregistered disconfirmers	**Calibrated checklist requirement**	Concise Section 12	Current underwrite; does not override S floors or other caps.	No admission at that diagnostic scale absent disconfirmers; practically such weights are generally unreachable for ordinary equities under final floors.

E.4.5 Implied best-case maximum-weight matrix

This matrix applies only when every admission gate passes, stress liquidity is within five days, factor headroom is full, and the stated severity is permitted. Parentheses show the 0.50 EC multiplier. Values are percentages of total portfolio, not recommendations. The speculative per-name cap of 0.25-0.75% binds before a larger S=100% result.

S	FS <1.00, BaseLoss 0.25%	FS 1.00-1.24, BaseLoss 0.50%	FS 1.25-1.74, BaseLoss 0.75%	FS >=1.75, BaseLoss 1.00%	Status and consequence
100%	0.25 (0.13)	0.50 (0.25)	0.75 (0.38)	1.00 (0.50)	Pre-cash-flow/speculative uses S=100%, but 0.75% per-name speculative ceiling binds first.
80%	0.31 (0.16)	0.63 (0.31)	0.94 (0.47)	1.25 (0.63)	Minimum range for binary-license or flagged-EM exposure where 80% is justified.
60%	0.42 (0.21)	0.83 (0.42)	1.25 (0.63)	1.67 (0.83)	Minimum for levered financial, cyclical, or roll-up.

S	FS <1.00, BaseLoss 0.25%	FS 1.00-1.24, BaseLoss 0.50%	FS 1.25-1.74, BaseLoss 0.75%	FS >=1.75, BaseLoss 1.00%	Status and consequence
50%	0.50 (0.25)	1.00 (0.50)	1.50 (0.75)	2.00 (1.00)	Common ordinary-equity diagnostic case; not target.
40%	0.63 (0.31)	1.25 (0.63)	1.88 (0.94)	2.50 (1.25)	Lowest ordinary-equity severity; intentional best-case ceiling.

Status: **calibrated diagnostic matrix**, derived from final F5-1/F5-3 arithmetic in concise Section 4.1 and Expanded Chapter 4. Amendment requires a formal change to the formula, base budgets, evidence bands, or severity floors. Review for every live underwrite and whenever FS, EC, liquidity, severity, or factor headroom changes. Any hard gate replaces the table result with zero.

E.4.6 Common-factor caps and comparator-governance parameters

The original risk caps remain controlling. The proxy's 8.2% AI FSL and 9.0% joint AI-plus-power FSL reveal a possible breach but are not reproducible enough to ratify higher caps. F5-4 supersedes the earlier suggestion that feasibility caps equal comparator FSL plus one percentage point. Any mention of provisional 9.2% AI or 10.0% joint caps is historical/proxy material, not current authority.

Factor/control	Preferred FEW cap	Concentration sub-cap	Hard FSL/joint-loss cap	Status and source	Amendment, review, and breach consequence
AI-capex sustainability	15% total portfolio or 25% of risk assets	None separately stated	5%	**Calibrated original risk cap.** Concise Section 7; Expanded Chapter 7.	Annual/material-event top-20-plus-buckets review. Proposed addition breaching cap gets Fmult=0. Current material AI/power active additions remain zero until actual study.
State-payer exposure	20% total portfolio	One sovereign 15%	5%	**Calibrated.** Same	Same. Budgets/backlog receive no owner-value credit without payer collectibility. Joint state-payer-plus-rates cap also applies.
China/export-control exposure	15%	One irreplaceable node 10%	5%	**Calibrated.** Same	Same. Look through revenue, supply chain, legal restriction, custody, and node dependency.
Directional rates/refinancing/quality duration	25% preferred **net** FEW	None separately stated	6%	**Calibrated; net-of-matching correction durable.** F5-5; Concise Section 7; Expanded Chapter 7.	Exclude genuine liability-matched duration from directional FEW and report it separately with basis risk. A phantom gross-bond breach must not be carried forward.
Power availability/cost	20%	None separately stated	4%	**Calibrated.** Concise Section 7; Expanded Chapter 7.	Current material AI/power active additions zero while original proxy breach remains unresolved.
Joint AI-capex + power	Not stated	Not stated	8% portfolio loss	**Calibrated.** Same	Test jointly; separate factor compliance does not waive joint cap. Proxy output 9.0% is a provisional breach signal, not a new limit.
Joint state payer + rates	Not stated	Not stated	8% portfolio loss	**Calibrated.** Same	Test jointly after rates are netted for true liability matching.

Comparator process parameter	Current value/range	Status	Source	Amendment authority and review	Breach consequence
Diagnostic proxy mix	60% MSCI ACWI / 40% liability-matched high-grade fixed income	**Illustrative/provisional; not ratified**	Concise Section 7; Expanded Chapter 7	Replace with actual named implementable comparator within the policy clock; not a target allocation.	Cannot certify, ratify caps, or authorize a product.
Proxy AI FSL	8.2% versus 5% original cap; gap 3.2 pp	**Provisional order-of-magnitude assumption**	Concise Sections 7 and 10	Actual worksheet, second pass, sensitivity, and implementation study required.	Original 5% governs; active AI/power additions zero; disclose accepted embedded risk.
Proxy joint AI+power FSL	9.0% versus 8% original cap; gap 1.0 pp	**Provisional**	Same	Same	Original 8% governs; same factor bar.
Top-constituent resolution	Top 20 scored individually; remainder in documented sector buckets	**Durable bounded method**	Concise Sections 1, 7, 12; Expanded Chapter 7	Formal amendment to method. Refresh annually and on material event.	An unarchived/incomplete worksheet cannot ratify a cap and is a certification blocker.
Factor sensitivity grid	0, 0.25, 0.50, 0.75, 1.00 = none, indirect, material, high, dominant	**Calibrated**	Expanded Chapter 7	Formal amendment for grid; each score needs dated evidence/rationale.	Missing rationale reduces reproducibility; do not conceal exposure in “other.”
Sensitivity perturbation	+/-0.25 around factor scores	**Calibrated robustness test**	Concise Section 7; Expanded Chapter 7	Required before cap ratification.	A conclusion that flips under one grid step is not robust enough to ratify the cap.
Independent scoring passes	2 before ratification; second after 6 months in year 1	**Durable/calibrated**	Concise Sections 1 and 7; F4-7/F5-4	Independent reviewer, not cosmetic approval. Annual refresh thereafter and material-event refresh.	No feasibility-cap ratification before completion.
Embedded hard-cap persistence	4 consecutive quarterly observations on actual comparator	**Durable forced-choice clock**	Concise Sections 1, 7, 9; F4-2	Owner-level written choice.	Must adopt implementable capped/blended comparator or raise cap by formal amendment; no renewable exception.
Temporary diagnostic ceiling outer bound	No more than 2x original risk cap; fixed, not floating	**Durable F5 correction**	Concise Section 7; Appendix D F5-4	Formal, documented temporary construct only; not presently adopted.	For illustration, 10% outer bound against 5% AI cap and 16% against 8% joint cap; these are not current caps or permissions.
Temporary ceiling expiry	12 months	**Durable**	Same	Ratify on actual study or let expire.	Auto-revert to original caps; forced choice reruns and capped/blended-comparator study becomes mandatory.
Actual comparator placeholder limit	12 months	**Calibrated**	Concise Section 12; Expanded Chapter 7	Annual governance must name actual implementable comparator.	No placeholder after the limit; actual scoring remains a certification blocker and new active work freezes as stated.
Comparator scoring capacity	Reproducible within 16 operator hours/year	**Calibrated minimum-viable-governance test**	Concise Section 10; Expanded Chapter 10	Annual hours/quality review; reduce complexity rather than degrade method.	Freeze new active work; default to named comparator + ILF + verified nominal/real matching.
Risk/feasibility gap reporting	Percentage points and portfolio currency	**Durable disclosure rule**	Concise Sections 1 and 7; F4-7	Annual report and any cap study.	An undisclosed gap cannot support amendment; raising cap must be described as accepting, not removing, risk.
Capped/blended study history set	Japan 1989; US 1992-2025; US 2013-2025; Nortel/TSE; Nokia/HEX	**Calibrated evidentiary set**	Concise Section 7; F5-4; Expanded Chapter 7	Required for permanent branch study; future amendment may update set transparently.	Do not select only history favoring capping or uncapped concentration. D-2 remains unresolved.

E.4.7 Monitoring, evidence-decay, and research-attention parameters

Parameter	Current cadence/value	Status	Source section	Amendment authority and review	Breach consequence
Monthly survival review	Monthly: ILF, custody/legal-title alerts, hard-gate events only	**Durable cadence**	Concise Section 9; Expanded Chapter 9	Formal amendment for cadence; operator performs narrow check.	Do not rescore whole company universe monthly; discovered survival breach uses fast clock.
Quarterly portfolio/research review	Quarterly: actual holdings, issuer/factor headroom, pillar-breaking evidence, tax lots, exception log	**Durable cadence**	Same	Same	Hard-cap/pillar/evidence breaches activate their separate clocks; exceptions count toward rolling limit.
Annual comprehensive reset	Annual: FS/PVL sensitivity, comparator top 20/buckets, ILF drag, live scorecards/reverse DCFs, transfer test, active viability, amendments, dispute status	**Durable cadence**	Concise Sections 9 and 12; Expanded Chapters 9-10	Owner/operator with independent components where required.	Missing liability closure, scoring, EC auditability, or viability can freeze active work.
Company-classification life	12 months without refresh	**Durable expiry**	Concise Sections 8-9; Appendix D F5-9	Annual/current underwrite can refresh. No automatic renewal.	Revert to Unclassified Research Note; old label cannot support current standing.
Company causal-falsifier horizon	Normally 36 months	**Provisional/calibrated**	Concise company map and Section 9; Expanded Chapter 9	Set specifically in each underwrite; material events can trigger earlier.	Two missed primary milestones cause immediate no add and 30-day re-underwrite.
Theme research sunset	36 months without journaled decision or explicit abstention threshold	**Durable attention rule**	Concise Section 9; Expanded Chapters 9-10	Annual research-governance review.	Move to annual-check status; archive evidence and record hours/external cost.
Annual class-absence evidence	At least 3 named examined-and-rejected candidates for each claimed absence	**Durable F5 rule**	Concise Sections 8.5 and 9; Appendix D F5-7	Annual map refresh.	Otherwise label search incomplete rather than claiming supported absence.
Annual breadth screens	Counter-canon and non-US screen every annual refresh	**Durable**	Same	Annual map owner.	Failure weakens map governance/EC auditability and may contribute to active-work freeze; do not carry a favorable canon by inertia.
Harvest-status outer life	24 months	**Calibrated**	Concise Section 12 and prohibition audit; earlier F3-5/F5 context	Current underwrite and explicit reclassification; no inertia-based renewal.	Cannot roll as normal sizing after expiry; revert/re-underwrite under applicable option/research status.
Portfolio falsifier: EC maintenance	EC audits must remain maintainable	**Durable**	Concise Section 10	Annual and ongoing operating review.	Freeze new active work.
Portfolio falsifier: realized viability	2 consecutive negative annual results	**Durable/calibrated**	Concise Sections 1 and 10; F4-11/F5-8	Annual realized review; re-justification required to reopen.	Suspend new active work; default to comparator + ILF + verified matching.
Portfolio falsifier: exceptions	2 material PEEs in 12 months	**Calibrated**	Concise Sections 3 and 10	Quarterly log/rolling review.	Freeze new discretion.
Portfolio falsifier: comparator burden	Reproducible scoring within 16 operator hours/year	**Calibrated**	Concise Section 10	Annual.	Freeze new active work if not achievable without sacrificing method quality.
Portfolio falsifier: liability closure	Must be measured by next annual review	**Durable stop rule**	Concise Section 10	Annual; verified liability process.	Freeze new active work; no improvised substitute.

E.4.8 Active-program economics and hours

Parameter	Current value/range	Status	Source	Amendment authority and review	Breach consequence
Prospective activation multiple	Expected annual active contribution >=2x fully loaded annual cost	**Calibrated F5 rule**	Concise Sections 10, 12, 14; F5-8	Annual owner-governance test; premium/hit rate must be evidenced, not hoped.	Below threshold: paper-journal mode, zero active capital.
Fully loaded cost	Investor-specific research hours at opportunity cost + data + tax + external/implementation costs	**Unresolved input**	Same	Annual documentation; cannot price time at zero merely because research is enjoyable.	An understated denominator invalidates activation; reduce scope or remain paper mode.
Recurring governance hours	About 40 hours/year	**Illustrative calibration**	Concise Sections 10, 12, 14; F5-14	Annual prospective capacity budget.	If unavailable, reduce active scope before research.
Per-live-underwrite hours	About 15-25 hours/year	**Illustrative calibration**	Same	Same	Same; do not maintain stale live status without work.
Eight-name workload	Roughly 160-240 hours/year, before one-off comparator/custody work	**Illustrative; arithmetic is deliberately approximate**	Same; Expanded Chapters 10 and 14	Print annually; choose fewer names prospectively.	The maximum of eight is not a quota. Infeasibility can reduce the set to zero.
Realized suspension	2 consecutive negative active-viability years	**Durable/calibrated**	Concise Sections 1 and 10; F4-11/F5-8	Annual retrospective test and formal re-justification.	Suspend new active work; comparator + ILF + matching default.
Attribution split	Annual separate attribution of comparator/matching versus company selection	**Durable**	Concise Sections 12 and Appendix D minority views; Expanded Chapters 10 and 14	Annual governance.	Do not allocate scarce research hours based on blended attribution that hides where value arose.

E.4.9 Valuation, tax, and exit parameters

Parameter	Current value/range	Status	Source	Amendment authority and review	Breach consequence
Active derating-loss default	Aggregate <=50% of MDB after strategic-sleeve stress	**Calibrated**	Concise Sections 1 and 9; F4-5	Archive before any separately authorized purchase; refresh with valuation.	Proposed position cannot be admitted if aggregate budget would exceed default without amendment.
Active derating-loss amendment band	>50% up to 75% of MDB	**Calibrated exceptional range**	Same	Written policy amendment plus independent challenge.	No informal analyst discretion; above 75% is inadmissible under current rule.
Ex-post derating overshoot	If aggregate exceeds budget while A-D remains intact: process post-mortem	**Durable consequence**	Same	Review at event/annual audit.	Not an automatic sale or trim; fix sizing process. Hard gates still control independently.
Tax comparison horizon	10 years fixed	**Calibrated**	Concise Sections 1 and 9; F4-9	Formal amendment; apply current statutory rate and comparator return.	A customized horizon cannot be selected merely to justify holding.
Deferral-led hold challenge	Independent challenger after 2 annual deferral-led holds	**Durable/calibrated**	Same	Annual tax-blind re-underwrite.	Hold rationale cannot roll indefinitely without challenge; never supports an add or a hard-gate failure.
Proposed tax-window wait	90 days	**Illustrative named delay requiring day-0 costing**	Concise Sections 1 and 9	Separate implementation decision and current tax facts.	Record expected opportunity cost/risk on day 0; the wait is not free.
All tax delays	Any length, costed on day 0	**Durable**	Concise Section 9; Expanded Chapter 9	Decision journal.	An unpriced delay is process failure; tax remains friction, not thesis.
Intact thesis below comparator	No add; action within 5 trading days after governed decision	**Durable**	Concise Section 9	Separate implementation authority.	Cannot use deferral indefinitely; price-only drawdown never forces sale.
Price-only drawdown	Any decline, including illustrative 30-40% derating, is never itself a sell/add trigger	**Durable**	Concise Section 9; Expanded Chapter 9	Formal amendment to principle only.	Re-underwrite valuation and evidence; no thesis-free averaging down or mechanical sale.

E.4.10 Custody, jurisdiction, illiquidity, FX, and professional-review parameters

Parameter	Current value/range	Status	Source	Amendment authority and review	Breach consequence
Custody Tier A	Mandatory same-jurisdiction second custodian	**Durable**	Concise Section 1; F4-6; Expanded Chapter 10	Current legal/custody implementation; annual transfer test.	Failure of transfer/access triggers re-underwrite; unresolved title keeps affected allocation at zero.
Custody Tier B	Plain instruments in second legal system; size unresolved; focused professional review	**Unresolved size; durable tests**	Same	Verified domicile/citizens hip/tax/estate/liability facts.	No assumed size; failed review/transfer keeps relevant allocation at zero.
Custody Tier C	Entities/trusts/complex wrappers default to 0% until full gauntlet passes	**Durable fail-closed rule**	Same	Full professional review and separate approval.	Remain at zero.
Jurisdiction transfer test	Annual successful movement/access test	**Durable**	Concise Sections 9 and 12; Expanded Chapters 9-10	Annual; outsourcing permitted but evidence retained.	Failure triggers re-underwrite; unresolved title/situs/succession/tax leaves allocation zero.
Private/illiquid hurdle	Vehicle-specific 4-6 pp net advantage over liquid alternative	**Unresolved D-4 conditional hurdle**	Concise Appendix B; Expanded Chapter 10	Formal future decision using vehicle evidence.	No admission from asset-class averages; actual allocation can be zero.
Long-horizon FX hedge	Conditional 0-50%, after dated liabilities matched first	**Unresolved D-5**	Concise Appendix B; Expanded Chapter 10	Verify currency, carry, basis, tax, covariance, and instruments.	Do not infer a personalized hedge percentage.

E.4.11 Grind-test parameters and diagnostic boundary

Parameter	Current value/range	Status	Source	Amendment authority and review	Consequence/limitation
Starting A*	1.50 million currency units	**Illustrative assumption**	Concise Section 6; Expanded Chapter 6	Replace with verified investor facts.	Cannot certify actual FS or allocation.
Stressed real PVL*	1.10 million	**Illustrative assumption**	Same	Same	Actual liability currency/index/duration/withdrawals remain missing.
Equity real stress	-45%	**Illustrative**	Same	Future policy stress review; not a forecast.	Used only in the printed diagnostic.
Matched nominal-duration real stress	-25%	**Illustrative**	Same	Same	Actual duration/currency/basis must be matched; vehicle-specific path stress remains open.
Verified duration-matched ILB stress	0% real; no assumed +15% gain	**Illustrative/conservative**	Same; F4-4	Same	Only admissible if index/law/market/title match is genuine.
ILF cash real stress	-10%	**Illustrative**	Same	Same	Represents real erosion, not nominal loss forecast.
Inherited mix	45% equity / 25% nominal / 20% ILB / 10% ILF	**Illustrative failed mix**	Concise Section 6	None as a target; replace with verified design.	End assets 1.0875m; Coverage 0.989; fails by 12,500 immediately, before withdrawals.
Boundary mix	40% equity / 20% nominal / 30% ILB / 10% ILF	**Illustrative conditional pass**	Same	None as a target.	End assets 1.1400m; Coverage 1.036; passes by 40,000 only if roughly 30% is genuine real/index matching.
Withdrawal path	Missing; required for a true grind duration	**Unresolved certification blocker**	Concise Sections 6, 10, 12, 15; F5-12	Verify liability cash-flow schedule and rerun.	Neither immediate diagnostic proves decade survival; duration of slow failure remains uncomputable.

E.5 Research-falsifier threshold register

The company map's 36-month tests are provisional research parameters, not policy-wide hard caps. They exist to prevent narrative rewriting. They must be refreshed with the classification and current underwrite; a breach normally invokes the evidence/pillar review hierarchy rather than an automatic trade. The table below captures every material numeric threshold or clock stated in the controlling map. Qualitative falsifiers remain in the source rows and are not converted here into invented numbers.

Company	Numeric threshold or horizon	Status/source	Review and breach consequence
S&P Global	36 months; recurring growth <5%; margin erosion >300 bps	**Provisional [A].** Concise Section 8.1	Monitor quarterly for pillar-breaking evidence; threshold breach supports 30-day re-underwrite/no add, while material de-licensing uses the fast hard-gate clock.
CME Group	36 months; >10% fee compression together with core-contract share loss	**Provisional [A].** Section 8.1	Re-underwrite causal/liability case; material clearing impairment uses hard-gate treatment.
Visa	36 months; net-revenue yield down >100 bps with payment-volume share loss	**Provisional [A].** Section 8.1	Re-underwrite structure/regulation/valuation; no automatic price-based execution.
Fastenal	36 months; cash conversion <75% through a full cycle with absent site/share gains	**Provisional [A].** Section 8.1	Re-underwrite A/C and working-capital economics.
Berkshire Hathaway	36 months; persistent adverse reserve development plus per-share value growth below comparator after tax/cost	**Provisional [A].** Section 8.1	Battery evidence remains capped at 2 pending ten-year exhibits; underperformance alone is insufficient without causal evidence.
Chubb	36 months; cumulative adverse development >5% of opening equity or combined ratio structurally >100	**Provisional [A].** Section 8.1	Re-underwrite reserves/capital; faster action if solvency/title/integrity gate breaks.

Company	Numeric threshold or horizon	Status/source	Review and breach consequence
DBS	36 months; NPL >2.5%, material deposit-share loss, and ROE below cost across 6 quarters	**Provisional [A].** Section 8.1	Re-underwrite bank battery; stressed CET1/deposit cohorts remain required.
TSMC	36 months; leading-node share loss >10 points; advanced-node utilization <70% absent recession	**Provisional [A].** Section 8.1	Current AI/power factor bar already makes admissible active addition zero. Trigger adds thesis re-underwrite.
ASML	36 months; service revenue/installed base diverges for 6 quarters, or material leading-edge share loss	**Provisional [A].** Section 8.1	Same current factor bar; re-underwrite A-C and asset lives.
Schneider Electric	36 months; qualitative break in orders/revenue while capex rises and incremental ROIC falls below cost	**Provisional [A].** Section 8.1	Current AI/power addition zero; numeric horizon governs research attention, not a trade.
RELX	36 months; organic growth <4% with renewal/usage decline, or AI costs cause >300 bps margin erosion	**Provisional [A].** Section 8.1	Re-underwrite AI capture, retention, and valuation.
GE Aerospace	36 months; shop-visit conversion misses plan by >15% and cash conversion <70% absent recession	**Provisional [A].** Concise Section 8.2	E<2 already blocks admission; safety/licensing issues use fast clock.
Microsoft	36 months; AI gross profit below incremental depreciation/power for 6 quarters	**Provisional [A].** Section 8.2	E fails/current factor bar makes active addition zero; re-underwrite capex cohort returns.
Alphabet	36 months; Search revenue per query and operating profit both decline for 4 quarters while capex/sales remains elevated	**Provisional [A].** Section 8.2	E/legal work incomplete and factor bar zeroes addition.
NVIDIA	36 months; Data Center growth <10% with inventory correction, or rental/customer-profit decline while supply rises	**Provisional [A].** Section 8.2	Current E0 makes Wmax zero under F5-2; factor bar independently zeroes addition.
Eaton	36 months; backlog reverses and incremental margins negative through 4 quarters	**Provisional [A].** Section 8.2	E<2 and factor bar block addition; re-underwrite capital-cycle thesis.
Moody's / LSEG / Constellation and other watch rows	36-month causal horizon; some tests use 4-6 quarters or >300 bps in their source rows	**Provisional [A].** Concise Section 8.2	The exact current row controls. If two registered primary milestones are missed, no add and 30-day re-underwrite. Residual register limitation: not every watch-row condition is numerical.
Rheinmetall	36 months; funded awards/backlog conversion <70% of plan and working capital absorbs >50% of cumulative EBIT	**Provisional [A].** Section 8.2	Payer and valuation gates remain incomplete; budget is not cash.
CoreWeave	36 months; rental prices fall >30% while utilization <70%, or interest/lease cash exceeds gross profit	**Provisional [A].** Concise Section 8.3	Research-only EC and E0; per-name cap 0.75%, but current admissible weight remains zero.
Samsara	36 months; dollar retention <105%, gross margin down >500 bps, or share count outruns FCF/share	**Provisional [A].** Section 8.3	EC/E/governance gaps keep research-only; re-underwrite rather than average down.
Symbotic	36 months; contractual throughput miss or cumulative project cash remains negative despite growth	**Provisional [A].** Section 8.3	Hard evidence gaps; current weight zero.
Nu Holdings	36 months; loss-adjusted ROE below cost while deposits reprice and USD FCF/share stagnates	**Provisional [A].** Section 8.3	Financial battery and short-cycle evidence remain insufficient; current weight zero.
Sea Limited	36 months; commerce contribution margin reverts negative without share gain or USD FCF/share fails to grow	**Provisional [A].** Section 8.3	Current weight zero; re-underwrite segment/capital allocation.
BYD	36 months; auto gross margin remains below cost-of-capital economics while capacity/utilization worsens	**Provisional [A].** Section 8.3	EM/state/cycle gates incomplete; current weight zero.
Palantir	36 months; Avoid would be challenged only if fully diluted FCF/share compounds >20% and conservative return clears comparator +1 pp	**Provisional [A].** Concise Section 8.4	Current E0 makes Wmax zero; evidence must restore admission, not merely price appreciation.
Intel	36 months; 2 independent external customers scale on leading nodes; capex/depreciation <1; segment cash return exceeds cost	**Provisional [A].** Section 8.4	Current C0/harvest failure keeps ceiling zero; state support is not harvest proof.
Boeing	36 months; positive cumulative FCF excluding advances, quality stabilization, and no new material charge	**Provisional [A].** Section 8.4	Current C0/D0 and safety/certification risks keep zero. Integrity events use fast clock.

Company	Numeric threshold or horizon	Status/source	Review and breach consequence
Coinbase	36 months; non-transaction owner earnings cover fixed costs across a full crypto drawdown and title independently validated	**Provisional [A].** Section 8.4	Cycle/title/E gates keep zero.
PG&E	36 months; no catastrophic loss, no material-dilution funding plan, and realized cash ROIC reaches allowed return	**Provisional [A].** Section 8.4	Payer/liability/capture gates keep zero; allowed ROE alone is insufficient.
Plug Power	36 months; 2 plants sustain positive unsubsidized contribution and operating cash is positive without equity issuance	**Provisional [A].** Section 8.4	C0/E0 and financing gates keep zero.
Nike	36 months; 2 years of full-price unit growth, stable gross margin, and share recovery without rising marketing intensity	**Provisional [A].** Section 8.4	Current E1 fails admission; F5-2 correction means Wmax zero.
Teladoc	36 months; independently verified clinical savings and retention >95% produce positive diluted FCF/share	**Provisional [A].** Section 8.4	Payer/capture gates keep zero.
Flagstar	36 months; ROTCE exceeds cost without dilution for 8 quarters, with NPL/charge-offs normalization and core-deposit persistence	**Provisional [A].** Section 8.4	B0/C0 and failed bank battery keep zero.
Bausch Health	36 months; net leverage <3x from retained cash plus clean separation and positive organic FCF/share	**Provisional [A].** Section 8.4	Roll-up/leverage/governance gates keep zero.
First Quantum	36 months; durable Panama resolution, balance-sheet survival at conservative copper, positive after-sustaining FCF	**Provisional [A].** Section 8.4	Title/cycle gates keep zero.
Adani Enterprises	36 months; positive USD FCF/share across 2 years, simplified audit, lower leverage, no related-party leakage	**Provisional [A].** Section 8.4	D0/E0 and EC<60 keep zero; current hard governance/coverage failure controls.

E.6 Superseded hierarchy and residual contradictions

Topic	Earlier wording or conflict	Final controlling rule	Residual issue, if any
Multiplier floor	F4-1 said to floor the combined remaining multiplier, which would rescue Fmult=0 and dilute liquidity penalties.	F5-1: floor only ECmult; AllowedLoss = BaseLoss x max(0.50, ECmult) x Lmult x Fmult; hard factor breach zeroes budget.	No residual rule conflict. Traceability must retain that the original defect came from Fable and was faithfully implemented before correction.
Gate-state Wmax	NVIDIA/Ferrari once printed nonzero Wmax with E0; Palantir printed zero; Nike cited the wrong A gate.	F5-2: any pillar 0 prints Wmax 0; Nike is blocked by E1.	A Watch classification can remain as a research-priority label despite E<2, but it has no current admissible diagnostic weight.
Severity/concentration	Legacy 5%/8% generic single-name rows and an unguarded S input enabled rule shopping.	F5-3 class floors; legacy tiers void; ordinary best case about 2-2.5%.	D-1 is closed but can reopen after the F5-13 record. Earlier Appendix B language requiring >=10 years is superseded by 5 years/>=20 decisions.
Comparator feasibility caps	F4-7 formula max(risk cap, comparator FSL + 1 pp) and concise Section 1's statement that a provisional cap-amendment branch was executed.	F5-4/F5 final integrated Section 7: no ratification on proxy, no floating cap, original caps govern, fixed temporary ceiling no more than 2x and 12-month auto-reversion.	The wording “F4-2 executed provisionally” survives in Section 1, but Section 7 and Appendix D make the practical hierarchy clear: execution is unratified and confers no current higher cap. This is a disclosed semantic tension, not permission to select 9.2%/10.0%.
Rates factor	Earlier proxy rates FEW counted 40% matched bonds and reported a breach.	F5-5: compute directional rates exposure net of matched liabilities; report matching separately with basis risk.	Corrected actual worksheet is still missing, so no replacement rates FEW value is certified.
Financial battery	Early map scores of 3/4 appeared favorable without underlying triangles/cohorts/stress exhibits.	F5-6: all favorable scores >=3 are provisional [A] and capped at 2 for sizing until exhibits attach.	Current concise rows have been relabeled/capped. No completed favorable battery demonstration exists; first full underwrite must supply it.
Speculative S=100 table	Generic matrix can show 1.00% at highest FS, above normal per-name speculative range.	F5-11: 0.25-0.75% per-name speculative cap binds first; speculative sleeve assumes total loss.	No conflict remains if lower-of-all-caps rule is applied.
D-1 reopening	Earlier open-dispute row retained a >=10-year record.	F5-13: preregistered 5 years and >=20 audited decisions, plus after-tax attribution, liability surplus, challenge, capacity.	Some explanatory text may still mention ten years as historic provenance; it is superseded.
Map decay	Original map had no automatic expiry.	F5-9: 12 months without refresh -> Unclassified Research Note.	No current company label may be carried beyond its stamp by inertia.
Active viability	Original rule waited for two negative deployed years.	F5-8 adds prospective >=2x cost activation and paper mode below; retrospective two-negative-year suspension remains.	Evidenced premium and hit rate have no universal numeric default; they are unresolved investor/program inputs and must not be fabricated.
Underwrite capacity	“At most eight” lacked printed arithmetic.	F5-14 prints about 40 governance hours plus 15-25 per live name, roughly 160-240 hours/year for eight.	Arithmetic spans 160-240 only approximately because 40 + 8x(15-25) mathematically yields 160-240; one-off comparator/custody work is excluded and must be separately budgeted.
Grind result	Static test was sometimes described as a grind rerun/pass boundary.	F5-12: withdrawal-path grind is a certification blocker; illustrative mix fails immediately and boundary mix passes only conditionally.	Residual unresolved item is the actual withdrawal schedule and vehicle-specific ILB path/basis stress. Do not call liability closure resolved.
Company factual citations	Some rows called an IR landing-page link a “2026 fact.”	F5-11: call it an evidence pointer and grant no EC credit until extraction.	Row-level EC scores remain unaudited [A]; current full EC audits are certification blockers.
Comparator-clock trigger	Durable forced choice is after four quarters on the actual comparator, while proxy controls are already conservative.	Current original caps and zero active AI/power additions apply pending study; the permanent forced choice waits for actual reproducible measurement.	The policy does not specify a new portfolio trade to cure the existing proxy breach. It freezes affected additions and requires the study.

E.7 Amendment and breach authority matrix

Decision type	Who may do what under the text	Required evidence/challenge	What cannot be inferred
Update a factual input	Operator/research process may update assets, liabilities, tax, holdings, evidence, valuation, liquidity, factor scores, and costs when verified.	Current source extraction, timestamps, appropriate T1 weighting, and independent pass where specified.	Updating an input is not a policy amendment or trading approval.
Tighten after a breach	Policy applies lower state, no add, freeze, re-underwrite, or zero automatically as a governance consequence.	Evidence of the triggering event and decision journal.	Risk-reducing posture does not grant a specific transaction without separate implementation authority.
Change a calibrated band/cap	Policy owner through explicit written amendment.	Rationale, effects, surviving minority views, and independent challenge where required. Comparator cap changes require full actual study.	A persistent breach, proxy output, or analyst convenience is not an amendment.
Use exceptional 50-75% MDB derating allowance	Formal written amendment plus independent challenge.	Archived valuation cases, aggregate stress, and challenger record.	Ex-post overshoot cannot retroactively authorize the larger range.
Set S below class floor	Written evidence plus independent challenge within a current underwrite.	Class-specific loss history and causal impairment evidence; confidence/reputation insufficient.	Does not recreate a conviction tier or override other caps.
Reopen D-1 concentration	Formal amendment after the full five-year/>=20 decision record and other conditions.	Preregistered fully underwritten decisions, independent audited after-tax attribution, verified liability surplus, challenger sign-off, capacity.	Past intuition or one successful concentrated holding is insufficient.
Ratify comparator cap/design	Owner-level formal decision after actual comparator study.	Archived top 20/buckets, two passes, +/-0.25 perturbations, net-rates treatment, capped/blended alternative, tax/cost/TE/abandonment, balanced history.	Proxy 8.2%/9.0% figures do not authorize 9.2%/10.0% caps.
Approve an implementation	Separate process not identified in this policy.	Verified investor facts, liability closure, current underwrite, tax/legal/custody checks, all policy gates, explicit owner approval.	Researcher, model, challenger, table, classification, or annex has no trading authority.

E.8 Completeness checklist and practical boundary

The parameter hierarchy is fail-closed. Where investor facts are missing, use the lower state, zero, or non-certification as specified. Where an earlier formula conflicts with F5, F5 controls. Where the policy preserves a dispute, record it and its observable rather than selecting a favorable branch. Where a clock says action, the policy still requires the separate implementation authority. No table in this annex turns a scenario, company classification, formula result, cap, or diagnostic ceiling into a buy, sell, hedge, order, target allocation, account action, or portfolio mutation.

Practical boundary: the current authorized company weight remains zero; original factor caps govern; actual liability closure, withdrawal-path grind, comparator scoring, EC audits, MDB, investor facts, and implementation approval remain incomplete. The correct response to an incomplete or breached parameter is the stated freeze, re-underwrite, reversion, expiry, lower state, or zero - never an improvised looser interpretation.

Annex F - Traceability, Open Disputes, Provenance, and Final Amendments

The following appendices are reproduced from the concise final policy so that the expanded prose remains auditable back to the five-round debate. They are intentionally not paraphrased: their role is traceability and evidence integrity.

Appendix A - Full debate traceability

A1. Policy articles P1-P52

A = adopted into this integrated draft; A/R = adopted as later reframed. Numeric bands remain contingent. The minority column records the live contrary view rather than manufacturing consensus.

P	Origin	Subject	Final disposition / minority
P1	R1 Codex	goal hierarchy	A; Shortfall Parity later makes excessive caution symmetric
P2	R1 Codex	epistemic labels	A; no minority
P3	R1 Codex	liability segmentation	A; flexibility second axis added
P4	R1 Codex	real after-tax measurement	A
P5	R1 Codex	named low-cost comparator	A/R; comparator is a risky position, not neutral
P6	R1 Codex	active-risk burden	A; zero-active minority strengthened
P7	R1 Codex	circle-of-competence gate	A; unknown remains zero
P8	R1 Codex	capacity/tolerance weakest link	A/R; trait/state split, capacity binds
P9	R1 Codex	ILF	A/R; formula, bands, de-overlap and drag audit
P10	R1 Codex	no recourse leverage default	A/R; narrow operational exceptions only
P11	R1 Codex	look-through concentration	A/R; forced choice for persistent breach
P12	R1 Codex	liability-led currency	A; long-horizon hedge still open
P13	R1 Codex	tax discipline	A/R; standardized deferral/challenger
P14	R1 Codex	valuation ranges	A; reverse DCF and derating budget added
P15	R1 Codex	evidence-to-size	A/R; evidence-to-loss plus bounded total-loss option
P16	R1 Codex	behavioral controls	A/R; observable exceptions, journal
P17	R1 Codex	amendment process	A; persistent breach cannot become exception
P18	R2 Codex	Shortfall Parity	A; never licenses excess capacity risk
P19	R2 Codex	comparator non-neutral	A/R; top-20 proxy and feasibility cap
P20	R2 Codex	loss-budget sizing	A/R; F4 formula without Vmult
P21	R2 Codex	symmetric capital cycle	A; harvest minority preserved
P22	R2 Codex	home-sovereign integrity	A/R; tiered custody
P23	R2 Codex	fast rules/slow discretion	A
P24	R2 Codex	minimum viable governance	A/R; hours/viability added
P25	R2 Codex	theme non-authorization	A
P26	R2 Codex	36-month milestones/sunset	A
P27	R2 Codex	AI tool is not edge	A
P28	R3 Codex	canonical business-model scorecard	A/R; five non-compensating pillars
P29	R3 Codex	evidence caps	A/R; continuous EC-to-loss
P30	R3 Codex	symmetric cycle states	A
P31	R3 Codex	payer gate	A

P	Origin	Subject	Final disposition / minority
P32	R3 Codex	economic depreciation	A
P33	R3 Codex	acquisitions as reinvestment	A
P34	R3 Codex	AI moat burden	A; defensive add-back prohibited while open
P35	R3 Codex	per-share capture	A
P36	R3 Codex	validation protocol	A/R; 200+ needed for statistical claims
P37	R3 Codex	quality is not price	A; central to map rejects
P38	R4 Codex	FS selector	A/R; second axis, sensitivity, hysteresis
P39	R4 Codex	three-layer portfolio	A
P40	R4 Codex	pillar-gated sizing	A
P41	R4 Codex	loss arithmetic	A/R; Vmult removed, floor 0.50, D-1 closed
P42	R4 Codex	error budget	A
P43	R4 Codex	comparator modification	A/R; exposure-persistence forced choice
P44	R4 Codex	ILF bands/glide	A/R; D36/H de-overlap; D24 election
P45	R4 Codex	common-factor caps	A/R; proxy tier and feasibility gap
P46	R4 Codex	reverse-DCF hurdle	A
P47	R4 Codex	fast broad/slow active	A
P48	R4 Codex	sector-bust behavior	A/R; pre-bust comparator breach resolved
P49	R4 Codex	D-3 default	A/R; lower gross for crash, ILB matching for grind
P50	R4 Codex	second jurisdiction	A/R; Tier A/B/C
P51	R4 Codex	minimum viable governance	A/R; active viability
P52	R4 Codex	Round-5 classification gates	A/R; F4-10 financial battery and explicit non-buy

A2. Amendment ledger F1-F4

Amendment	Origin	Final disposition	Surviving minority / qualification
F1-1 Shortfall Parity	R1 Fable	Accepted in R2; retained	capacity still binds
F1-2 comparator de-neutralization	R1 Fable	Accepted; refined F4/F5	capped-index versus broad-market warnings both survive
F1-3 trait/state tolerance	R1 Fable	Accepted with reframe	genuine life events can override blackout
F1-4 leverage carve-outs	R1 Fable	Reframed/tests retained	paid options not treated as harmless
F1-5 ILF formula/cap/glide	R1 Fable	Reframed, bands set R4/F5	numerical calibration awaits facts
F1-6 loss-cap/speculative sleeve	R1 Fable	Accepted	abstention remains valid
F1-7 solvency/integrity exception	R1 Fable	Accepted	minimum scope only
F1-8 home-sovereign integrity	R1 Fable	Accepted, domicile-calibrated	triggers not universal
F1-9 tax-blind re-underwrite	R1 Fable	Reframed; standardized F4-9	deferral can support hold-not-add
F1-10 observable exceptions	R1 Fable	Accepted	self-report can tighten only
F1-11 stress precedence	R1 Fable	Accepted	liquidity real drag accepted
F1-12 one diligence battery	R1 Fable	Accepted	five-pillar battery supersedes duplicates
F1-13 minimum governance	R1 Fable	Accepted	institutional apparatus scales down
F1-14 missing facts	R1 Fable	Accepted	present policy remains conditional
F1-15 lump-sum prior	R1 Fable	Accepted as prior	event risk can justify paced single names
F2-1 common-factor map	R2 Fable	Accepted; operationalized R4/F5	proxy error remains
F2-2 three AI layers	R2 Fable	Accepted	equipment ≠ hyperscaler ≠ merchant cloud
F2-3 symmetric capital cycle	R2 Fable	Accepted	harvest thesis preserved, no current qualifier
F2-4 shortfall remedy ladder	R2 Fable	Accepted	risk increase last and capacity-bound
F2-5 emergency precedence	R2 Fable	Accepted	unaffected rules continue
F2-6 payer-quality gate	R2 Fable	Accepted/broadened	budgets are not contracts
F2-7 second-jurisdiction floor	R2 Fable	Unresolved R3; resolved structurally by F4-6	Tier-B size is domicile-specific
F2-8 split maintenance/verification	R2 Fable	Accepted	no composite class shortcut
F2-9 S6/research sunset	R2 Fable	Accepted	evidence archived, attention reduced
F2-10 patent cliff	R2 Fable	Accepted/generalized	no terminal bridge by assumption
F2-11 countertheme sketches	R2 Fable	Reframed	not promoted to themes
F2-12 hard falsifiers	R2 Fable	Accepted	firm-level, per-share tests
F2-13 detector repair	R2 Fable	Accepted	no mood-based loosening
F2-14 base-rate attachments	R2 Fable	Accepted with caution	priors lower, never settle
F2-15 disputes/producer bias	R2 Fable	Accepted	consumer/financial classes added
F3-1 pillar restructure	R3 Fable	Accepted	additive quality score retired
F3-2 continuous coverage	R3 Fable	Accepted/tightened	EC not a buy signal
F3-3 complexity flag	R3 Fable	Accepted then scoped by F4-10	principal financials use batteries

Amendment	Origin	Final disposition	Surviving minority / qualification
F3-4 balance-sheet financials	R3 Fable	Accepted	no automatic S2/S3 tilt
F3-5 two-stage harvest	R3 Fable	Reframed	option-only until full confirmation
F3-6 deterrence observables	R3 Fable	Accepted	causal attribution remains open
F3-7 option value	R3 Fable	Accepted	external observable price cap
F3-8 split assurance	R3 Fable	Accepted	binary liability severity retained
F3-9 EM decomposition	R3 Fable	Accepted	USD per-share capture required
F3-10 validation redesign	R3 Fable	Accepted with feasibility gate	small cohort is audit, not validation
F3-11 disappointment roster	R3 Fable	Accepted as stress	no class underweight by story
F3-12 challengers	R3 Fable	Accepted	option/research-only
F3-13 roll-up indicators	R3 Fable	Accepted	30-day re-underwrite
F4-1 loss recalibration/D-1	R4 Fable	Accepted; branch (a)	concentration view rejected with reopening test
F4-2 comparator forced choice	R4 Fable	Accepted/executed cap-amendment branch	capped alternative still must be studied
F4-3 FS second axis	R4 Fable	Accepted	numeric bands contingent
F4-4 grind honesty/D-3 split	R4 Fable	Accepted; result printed	ILB integrity conditional
F4-5 derating budget	R4 Fable	Accepted; default 50% MDB	75% requires amendment
F4-6 custody tiering	R4 Fable	Accepted	Tier-B/C size open
F4-7 FEW feasibility	R4 Fable	Accepted	comparator proxy error reported
F4-8 event pacing	R4 Fable	Accepted	full deployment prior remains
F4-9 deferral standard	R4 Fable	Accepted	tax facts required
F4-10 complexity scope	R4 Fable	Accepted	battery equally blocking
F4-11 active viability	R4 Fable	Accepted	zero-active vs option-value views preserved
F4-12 ILF de-overlap	R4 Fable	Accepted	D24 election conditional
F4-13 provenance relocation	R4 Fable	Accepted	metadata Appendix C only
F4-14 U4 disputes	R4 Fable	Accepted	both remain observable disputes

Appendix B - Open-dispute register

Dispute	Current status	Adjudicating observable / surviving view
D-1 concentration	**Closed for this policy:** no conviction tier; ~2% at S50 best case	Reopen only after a preregistered five-year record covering at least 20 fully underwritten decisions, independently audited after-tax attribution, verified liability surplus, challenger sign-off and operator capacity. Reopening is not approval.
D-2 strategic valuation tilt/comparator modification	Procedure resolved; actual design open	Actual comparator/tax basis/top-20 risk; compare capped/blended TE, cost and abandonment with broad-market Japan-1989 and US-1992-2025 exhibits.
D-3 tail insurance/diversification	Split resolution	Crash null is lower gross risk; real grind default is verified ILB matching; trend/tail premium/commodities require full-cycle net carry and Coverage improvement.
D-4 private/illiquid assets	Open; FS bands only	Manager/vehicle cash flows, governance, valuation and net return must beat liquid alternative by 4-6pp; actual investor access unknown.
D-5 long-horizon FX hedge	Open, 0-50% conditional	Liability currencies, carry, basis, tax, crisis covariance; dated liabilities matched first.
U2-1 diversifier carry	Open	Rolling-start full-cycle net results versus equal worst-case-loss reduction in gross risk.
U2-2 AI overbuild	Open by layer	Utilization, rental/token prices, cancellations, customer AI gross profit, capex/depreciation and free cash. Minority: capex defends existing rents.
U2-3 state-capex shareholder return	Open by contract/jurisdiction	Appropriation/rate recovery, cancellation, state take, dilution and real per-share cash.
U2-4 home-sovereign/custody calibration	Structure resolved; size open	Verified domicile/citizenship/tax/estate/title/sanctions and annual transfer test.
U2-5 speculative sleeve	Conditional ranges set; behavioral value open	Total-loss simulations, exception rate, goal probability and hours versus abstention.
U2-6 inherited-dispute umbrella	Superseded by individual D rows	No separate test; kept for provenance.
U2-7 accelerator economic lives	Open	Retirement/impairment/resale, rental price, utilization and vintage margin versus accounting lives.
U2-8 deterrence versus empire-building	Open; no defensive add-back	Ratchet when threat recedes, natural experiments, rival response and spend capped by rents at risk.
U2-9 muddle-through modal decade	Open	S6 comparator attribution and research-sunset cost; minority says dramatic scenarios are overbuilt.
U3-1 AI and assurance rents	Open	Real price, labor hours, error/liability, entrants and client-retained savings.
U3-2 quality-duration crowding	Open but governed at purchase	No-multiple-recovery stress, aggregate derating budget, Coverage and cash per share; no ex-post thesis rewrite.
U4-1 one-person executability	**Resolved adversely for full apparatus; bounded version accepted**	Actual annual hours/cost and decision value. Maximum eight live full underwrites; otherwise comparator+ILF+matching.
U4-2 procyclical diversifier evidence	Open	Rolling start dates across inflation, deflation, crash and calm; net costs/capacity and crisis correlation.

Appendix C - Fable provenance receipts

All four receipts record model claude-fable-5, authentication subscription_oauth_cli, success, one call, prompt cache off, web search off, and trade authorization false. They are provenance, not investment evidence.

Round	Prompt SHA-256	Response SHA-256
1	1d0c64d9700da7622e01fdf10d1191ac573fe5569ac2ef4ed917aed8e7d0fced	b6dded772a3ec075c7281c88b9b0d8ed02cc9b6bc2b157bace6ee61160239ad4
2	d229f2865c113d712920b338b27e490132d203fd720558bcba76526a05ebbd08	2346243d978485c353e92a8ece8ec30a963eb417c655b2fd81693801a10d21ea
3	d92871ccbf02b74351074223540449981070d1ff5afbd4f6913e08aff769ad93	1df1eb6bcc4dff4493a73442214254cc28d4258e4741e308dedde4e6d4226d7e
4	f8876d5b85b87e7952c1faa35a08c32f3a493b385a88e55cb7d849ef41f8e23f	acacbc75f3a0559b6a500ac36d1f761d9b8a973cfb85f2ef131fb0fbe49520df

Appendix D - Final Fable amendments F5-1 through F5-14

Amendment	Final disposition
F5-1 Floor rescope	**Adopted.** The floor applies only to ECmult; liquidity and factor penalties retain full force; hard factor breach means zero.
F5-2 Gate-state consistency	**Adopted.** E0 rows print Wmax 0; Nike is blocked by E1.
F5-3 Severity floors	**Adopted.** Equity 40%; levered/cyclical/roll-up 60%; binary/pre-cash/flagged-EM 80-100%; legacy 5%/8% rows void.
F5-4 Comparator hardening	**Adopted conservatively.** No feasibility-cap ratification without the actual worksheet, second pass and capped-comparator study; original caps govern meanwhile.
F5-5 Rates net of matching	**Adopted.** Liability matching is separated from directional rates exposure and retains basis-risk reporting.
F5-6 Financial batteries	**Adopted.** Favorable scores are provisional [A], capped at 2 pending named exhibits.
F5-7 Documented absence / counter-canon	**Adopted.** Three named rejected candidates and annual non-US/counter-canon screens required.
F5-8 Prospective active viability	**Adopted.** Capital only above a 2× fully loaded cost threshold; otherwise paper mode.
F5-9 Map decay	**Adopted.** Classifications expire after 12 months.
F5-10 Row-level factor bar	**Adopted.** AI/power-sensitive rows state current admissible addition is zero.
F5-11 Hygiene	**Adopted.** MELI domicile, internal codename, evidence-pointer language and speculative-cap priority corrected.
F5-12 Blocking-items label	**Adopted.** Grind path, liability closure, actual comparator and EC audits are checklist blockers.
F5-13 D-1 reopening	**Adopted.** Five years and at least 20 audited decisions replace the impossible ten-year prerequisite.
F5-14 Hours arithmetic	**Adopted.** The policy prints the 160-240-hour eight-name workload and reduces scope prospectively.

The traceability record explicitly notes that the original F4-1 floor defect came from Fable and was faithfully implemented by Codex before Fable corrected it in Round 5. Errors on either side remain visible.

Final unresolved minority views and adjudicating observables

- **D-2 comparator design:** capped/blended versus uncapped remains open; the mandatory actual-comparator study carries evidence in both directions.
- **D-4/D-5:** private assets and currency hedging remain dependent on verified investor facts and implementation evidence.
- **AI overbuild, asset lives and defensive capex:** current factor bar is a dated risk control, not a forecast resolution.
- **Quality-duration crowding:** the map's own composition is evidence of the bias; annual counter-canon/non-US screens adjudicate.
- **Zero-active versus dislocation option:** the prospective 2× cost threshold and paper mode structure the disagreement without forcing consensus.
- **Comparator-engineering-first:** annual attribution must compare value from comparator/matching design with value from single-name work.

Final boundary: liability closure is uncertified; all numeric bands await verified investor facts; the withdrawal-path grind budget is open; actual comparator scoring and EC audits are incomplete; and no company in this map is a buy absent a current, full evidence and valuation review. No transaction is authorized.