Investment Policy 2026-2036

Expanded Prose Edition

A readable explanation of the principles, decade themes, business models, company research map, portfolio architecture, risks, and falsifiers developed through five rounds of Codex/Fable debate.

Evidence cutoff	22 August 2026
Expansion design	One dedicated Codex worker for each of 15 chapters
Manuscript SHA-256	efb9e9291a9a43ea79ebd42f13f085aaee70c7dcb83643fb85beadb6143da55d

Policy boundary. This is a governance instrument and research map, not a portfolio or recommendation. Examples are explanatory only. No order, execution, hedge, outbound message, account opening, or portfolio change is authorized. The concise final policy remains the controlling compact reference.

Contents

Chapter 1 - How the Final Rules Were Chosen 9

F4-1: loss-budget arithmetic and the explicit rejection of a conviction tier . . . 9
F4-2: a persistent comparator breach must produce a choice . . . 10
F4-3: funded status needs a second axis and state-change discipline . . . 10
F4-4: separate crash protection from inflation-grind matching . . . 11
F4-5: budget valuation derating before purchase, not after it . . . 11
F4-6: match custody safeguards to structural complexity . . . 12
F4-7: make look-through factor caps both measurable and honest . . . 12
F4-8: tie entry pacing to information, not comfort . . . 13
F4-9: standardize tax deferral and close the perpetual-hold loophole . . . 13
F4-10: apply financial complexity to the right target . . . 13
F4-11: active management must earn its research cost . . . 13
F4-12: remove overlap inside the liquidity floor . . . 14
F4-13: keep orchestration records out of the constitution . . . 14
F4-14: preserve uncertainty through adjudicating observables . . . 14
What this chapter requires in practice . . . 15

Chapter 2 - The Investment Constitution and Its Durable Principles 16

The constitution begins with an order of purposes . . . 16
Material claims must carry epistemic labels . . . 17
Risk is governed through three separate budgets . . . 17
Measurement starts with liabilities, not the return league table . . . 18
The comparator bears risk and sets the burden of proof . . . 19
Shortfall Parity prevents safety from becoming silent failure . . . 19
Evidence governs tolerable loss; valuation governs admission . . . 20
Recourse leverage is prohibited by default . . . 20
Every exposure must be viewed on a look-through basis . . . 21
Broad strategic rules may move faster than active judgments . . . 22
Tax is friction, not an investment thesis . . . 22
Persistent breaches require amendment, not excuses . . . 23
Governance must be viable enough to endure . . . 23
Research never authorizes implementation . . . 24
Why the principles must coexist . . . 24
What this chapter requires in practice . . . 25

Chapter 3 - From Principles to a Funded-State Operating System 26

Funded status is the selector, not a return forecast . . . 26
Hypothetical worked example - calculating FS and its sensitivity . . . 26
The four funded states . . . 27
State 1 - FS below 1.00 . . . 27

State 2 - FS from 1.00 to 1.24 . . . 28
State 3 - FS from 1.25 to 1.74 . . . 28
State 4 - FS at or above 1.75 . . . 28
What each table column controls . . . 28
The second axis - capacity beyond the FS decimal . . . 29
Hypothetical worked example - the second axis . . . 29
Hysteresis and state transitions . . . 30
Hypothetical worked example - promotion and immediate demotion . . . 30
Speculative options are nested and charged at total loss . . . 30
Hypothetical worked example - market value versus loss charge . . . 31
The Investor Liquidity Floor . . . 31
Hypothetical worked example - ILF target and hard floor . . . 32
Hypothetical worked example - the D24 election . . . 32
Rebalancing architecture - fast for broad sleeves, slow for active risk . . . 33
Hypothetical worked example - the two rebalancing thresholds . . . 33
Decision clocks and fail-closed responses . . . 33
Immediate freeze and a 0-5-trading-day resolution window . . . 33
Thirty-day full re-underwrite and no add . . . 33
Valuation-only deterioration . . . 34
Hypothetical worked example - different evidence, different clocks . . . 34
Entry pacing is an information rule, not calendar theater . . . 34
Policy Exception Events and the freeze on discretion . . . 34
Hypothetical worked example - the rolling exception rule . . . 35
How the pieces fit together . . . 35
What this chapter requires in practice . . . 35

Chapter 4 - How to Judge a Business Model for a Ten-Year Holding Period 37

The five-pillar profile, not an additive quality score . . . 37
Pillar A - Demand and lock-in . . . 38
Pillar B - Industry structure, cycle, and payer . . . 39
Pillar C - Capital economics . . . 40
Pillar D - Governance and per-share capture . . . 41
Pillar E - Valuation and falsifiability . . . 42
Evidence coverage and the standard of proof . . . 42
Class overlays: asking the question appropriate to the model . . . 43
From an admitted thesis to an allowed loss . . . 44
Reading the implied maximum-weight table . . . 46
How the gates interact in practice . . . 46
Boundaries and failure modes of the scorecard itself . . . 47
What this chapter requires in practice . . . 47

Chapter 5 - The Decade Ahead: Scenarios, Themes, and Counterthemes 48

What a decade theme can and cannot do . . . 48

Reading scenarios without pretending to forecast . . . 48
S1 productive diffusion . . . 49
S2 inflation/repression grind . . . 50
S3 fragmented blocs . . . 50
S4 crash/sequence . . . 51
S5 adaptation/physical constraint . . . 51
S6 plateau/muddle . . . 52
The named research themes . . . 52
Verification/standards . . . 52
Installed-base aftermarket . . . 53
Market infrastructure . . . 53
Reliability distribution . . . 53
Semicap and power-control bottlenecks . . . 53
Funded health innovation . . . 54
Regulated resilience . . . 54
Balance-sheet franchises . . . 54
Selective brands . . . 55
EM platforms with audited per-share capture . . . 55
Counterthemes and the discipline of reversal . . . 55
Correlated risks and cross-theme failure . . . 56
Evidence, falsification, and capital-cycle discipline . . . 57
Why no theme authorizes a security purchase . . . 57
What this chapter requires in practice . . . 58

Chapter 6 - The Grind Test: What Must Survive an Unfriendly Decade 59

Why a grind test is necessary . . . 59
The four default-admissible buckets . . . 60
Equity . . . 60
Matched nominal duration . . . 60
Verified duration-matched inflation-linked bonds . . . 60
ILF cash . . . 61
How to read the table . . . 61
Clearly labeled hypothetical application examples . . . 62
Hypothetical example 1 - applying the inherited mix . . . 62
Hypothetical example 2 - applying the boundary mix . . . 62
Hypothetical example 3 - a mislabeled ILB bucket . . . 62
Hypothetical example 4 - an omitted withdrawal reverses the pass . . . 62
Hypothetical example 5 - the liability estimate changes . . . 63
What the test rules out . . . 63
What the test cannot prove . . . 63
What this chapter requires in practice . . . 64

Chapter 7 - Comparator Scoring, Forced Choices, and Certification 65

The comparator is a real position, not a neutral backdrop . . . 65
What proxy scoring is trying to measure . . . 65
How the top-20-plus-buckets method works . . . 66
What the inherited proxy found - and what it did not prove . . . 67
The forced-choice rule . . . 68
Why a temporary diagnostic ceiling is tightly bounded . . . 68
What must be studied before either permanent branch is credible . . . 69
Certification is a blocker, not a blessing . . . 69
Limitations and failure modes . . . 70
What this chapter requires in practice . . . 71

Chapter 8 - The Global Company Research Map 72

A map of questions, not a portfolio . . . 72
How to read the common fields . . . 73
Core Research Candidates - quality plausible, valuation and full evidence still open . . . 73
1. S&P; Global - SPGI, NYSE; United States . . . 74
2. CME Group - CME, Nasdaq; United States . . . 74
3. Visa - V, NYSE; United States . . . 74
4. Fastenal - FAST, Nasdaq; United States . . . 75
5. Berkshire Hathaway - BRK.B, NYSE; United States . . . 75
6. Chubb - CB, NYSE; Switzerland . . . 75
7. DBS Group - D05, SGX; Singapore . . . 76
8. Taiwan Semiconductor Manufacturing - TSM/2330, NYSE/TWSE; Taiwan . . . 76
9. ASML - ASML, Euronext Amsterdam/Nasdaq; Netherlands . . . 76
10. Schneider Electric - SU, Euronext Paris; France . . . 77
11. RELX - REL, LSE/NYSE; United Kingdom . . . 77
Watchlist/Price Sensitive - model plausible, gate or expectation unresolved . . . 77
12. GE Aerospace - GE, NYSE; United States . . . 78
13. Microsoft - MSFT, Nasdaq; United States . . . 78
14. Alphabet - GOOGL, Nasdaq; United States . . . 78
15. NVIDIA - NVDA, Nasdaq; United States . . . 79
16. Eaton - ETN, NYSE; Ireland . . . 79
17. London Stock Exchange Group - LSEG, LSE; United Kingdom . . . 79
18. Moody's - MCO, NYSE; United States . . . 80
19. Constellation Software - CSU, TSX; Canada . . . 80
20. Eli Lilly - LLY, NYSE; United States . . . 80
21. Novo Nordisk - NOVO-B, Nasdaq Copenhagen; Denmark . . . 81
22. MercadoLibre - MELI, Nasdaq; United States legal domicile, Latin America economic exposure . . . 81
23. Ferrari - RACE, NYSE/Euronext Milan; Netherlands/Italy . . . 81
24. Rheinmetall - RHM, Xetra; Germany . . . 82
Speculative Options - research optionality only, with total loss assumed . . . 82
25. CoreWeave - CRWV, Nasdaq; United States . . . 82

26. Samsara - IOT, NYSE; United States . . . 82
27. Symbotic - SYM, Nasdaq; United States . . . 83
28. Nu Holdings - NU, NYSE; Cayman Islands/Brazil . . . 83
29. Sea Limited - SE, NYSE; Singapore/Southeast Asia . . . 83
30. BYD - 1211, HKEX / 002594, Shenzhen; China . . . 84
Avoid/Counterexamples - appealing narratives that fail the gates . . . 84
31. Palantir Technologies - PLTR, Nasdaq; United States . . . 84
32. Intel - INTC, Nasdaq; United States . . . 85
33. Boeing - BA, NYSE; United States . . . 85
34. Coinbase - COIN, Nasdaq; United States . . . 85
35. PG&E; - PCG, NYSE; United States . . . 86
36. Plug Power - PLUG, Nasdaq; United States . . . 86
37. Nike - NKE, NYSE; United States . . . 86
38. Teladoc Health - TDOC, NYSE; United States . . . 87
39. Flagstar Financial - FLG, NYSE; United States . . . 87
40. Bausch Health - BHC, NYSE/TSX; Canada . . . 87
41. First Quantum Minerals - FM, TSX; Canada/Zambia-Panama exposure . . . 88
42. Adani Enterprises - ADANIENT, NSE; India . . . 88
Reading across business-model classes . . . 88
How the map should be used . . . 90
How the map must not be used . . . 91
What this chapter requires in practice . . . 91

Chapter 9 - Monitoring, Falsification, Sell Discipline, and Research Sunset 92

Monitoring is a hierarchy, not a stream of alerts . . . 92
Evidence has an expiry date . . . 93
Falsification separates an investment thesis from a story . . . 94
Price volatility is not thesis breakage . . . 95
The exit hierarchy and its decision clocks . . . 95
The valuation-drawdown budget is set before purchase . . . 97
Research sunset is an attention rule . . . 97
Failure modes the discipline is meant to prevent . . . 98
What this chapter requires in practice . . . 99

Chapter 10 - Risks, Unresolved Disputes, and Portfolio-Level Falsifiers 100

Uncertified liability closure . . . 100
A comparator born in breach . . . 101
The active program may be uneconomic . . . 102
Concentration, quality crowding, and the danger of hidden sameness . . . 103
Procyclical diversifier evidence and the unresolved insurance question . . . 103
Narrative and capital-expenditure risk . . . 104
Political, legal-title, custody, and currency risk . . . 105
Obsolescence and dishonest economic lives . . . 105

Other unresolved allocation disputes . . . 106
How risks combine into correlated failure . . . 106
Company-level and portfolio-level falsification . . . 107
What evidence would change the overall view . . . 108
What this chapter requires in practice . . . 108

Chapter 11 - Twelve Hard Questions and the Policy's Direct Answers 109

Question 1: Were all fourteen Round 4 amendments adjudicated, and how was the concentration dispute resolved? . . . 109
Question 2: What weights does the loss-budget formula actually permit at each severity? . . . 111
Question 3: Did the grind rerun pass using only default-admissible assets? . . . 111
Question 4: Is the comparator already in breach, and what forced choice follows? . . . 112
Question 5: Did the company map include balance-sheet financials and genuine harvest candidates? . . . 113
Question 6: Is every Core row complete enough to be valid? . . . 113
Question 7: Does the Avoid class cover every major model class and reject admired companies when the arithmetic fails? 114
Question 8: Do the financial batteries constrain favorable as well as unfavorable cases? . . . 114
Question 9: Did the map comply with every one of the ten company-research prohibitions? . . . 115
Question 10: Is the debate traceable from proposal to disposition and surviving disagreement? . . . 115
Question 11: Does the assembled policy survive attacks on bias, obsolescence, and implementation complexity? . . . 116
Question 12: What, if anything, is certified or authorized? . . . 117
What this chapter requires in practice . . . 117

Chapter 12 - The Operating Checklist, Explained Step by Step 118

1. Annual facts . . . 118
2. Survival . . . 119
3. Strategic design . . . 120
4. Look-through . . . 121
5. Active admission . . . 122
6. Options . . . 122
7. Decision record . . . 123
8. Monthly/quarterly . . . 124
9. Annual governance . . . 124
10. Certification blockers and stop rules . . . 125
What this chapter requires in practice . . . 126

Chapter 13 - The Prohibition Self-Audit 127

1. No buy call, target weight, entry instruction, order, or portfolio mutation . . . 127
2. No narrative shortcut to expected return . . . 128
3. No return to the old additive Business-Model Score . . . 128
4. No short-history challenger as a Core Research Candidate . . . 129
5. Do not ignore exposure already embedded in the comparator . . . 129
6. No inflation of weak, conflicting, or missing evidence . . . 130
7. Harvest status is neither normal sizing nor renewable by inertia . . . 130
8. State budgets are not collectible revenue without the Payer Quality Gate . . . 131

9. No personalized tax, account, domicile, currency, or custody conclusion . . . 131
10. No silent resolution of disputes or professional-review gates . . . 132
Reading the audit result correctly . . . 132
What this chapter requires in practice . . . 133

Chapter 14 - Final Bias and Implementation Stress Test 134

What the final bias actually is . . . 134
Hidden consensus: when a skeptical process still assumes activity . . . 134
Recency bias: a current map is not a decade-long base rate . . . 135
Survivorship bias: visible winners are not a validated sample . . . 136
Narrative fallacy: causal coherence is not cash capture . . . 137
Concentration: the comparator can violate the risk the active sleeve is forbidden to add . . . 138
Obsolescence: correct evidence may arrive after the loss . . . 138
Political and regulatory dependence: law can dominate business quality . . . 139
Implementation complexity: the institutional process fails the one-person test . . . 140
The final one-person verdict and the approval boundary . . . 141
What this chapter requires in practice . . . 141

Chapter 15 - Certification, Ownership, and the Policy Boundary 142

Final as a document, not certified as an investment plan . . . 142
What the edition establishes . . . 142
What remains uncertified . . . 143
The controlling document hierarchy . . . 144
The evidence cutoff and the meaning of a source snapshot . . . 145
Ownership, maintenance, challenge, and approval . . . 145
Prohibited uses of the policy . . . 146
How final amendments and future amendments should be read . . . 147
Boundary failure modes . . . 147
What this chapter requires in practice . . . 148

Appendices 149

Appendix A - Full debate traceability . . . 149
A1. Policy articles P1-P52 . . . 149
A2. Amendment ledger F1-F4 . . . 151
Appendix B - Open-dispute register . . . 153
Appendix C - Fable provenance receipts . . . 153
Appendix D - Final Fable amendments F5-1 through F5-14 . . . 154
Final unresolved minority views and adjudicating observables . . . 154

Chapter 1 - How the Final Rules Were Chosen

The final policy did not emerge from a single act of portfolio design. It emerged from an adversarial process in which every attractive rule had to survive a harder question: does the rule still work when its arithmetic, evidence burden, implementation cost, and failure mode are made explicit? Amendments F4-1 through F4-14 were the last broad challenge to the policy before the closing corrections. They addressed places where a sound principle had produced a contradictory or unusable operating rule, where a debate had been listed as open even though the mathematics had already decided it, or where a safety mechanism could itself create a different form of risk.

The resulting chapter is an adjudication, not a fresh portfolio proposal. All fourteen F4 amendments were accepted. Some were accepted exactly as proposed; others were later corrected or made more conservative when the fifth-round review found a defect. Acceptance therefore means that the underlying criticism was valid and that the final policy contains a binding response. It does not mean that every original sentence survived unchanged.

The policy also remains a governance instrument rather than an authorization to invest. Its bands depend on investor facts that have not been verified. Funded status, liability closure, tax position, domicile and custody structure, actual comparator, factor look-through, market-drawdown budget, and company valuations remain unverified. The rules below explain how a decision would be governed after those facts are supplied. They do not create a trade, target allocation, account structure, or current recommendation.

F4-1: loss-budget arithmetic and the explicit rejection of a conviction tier

F4-1 corrected the most important hidden decision in the Round 4 draft. That draft claimed to preserve the dispute over concentrated investing, called D-1, while its sizing formula made ordinary concentration arithmetically unreachable. Even under the strongest funded state, full evidence, good liquidity, and complete factor headroom, a 1.00% issuer loss allowance divided by a plausible 50% permanent-impairment severity produced only a 2.0% position. References elsewhere to 5% concentration gates and an 8% hard cap were therefore decorative. The formula had closed the debate without saying so.

The first correction was conceptual. Valuation ceased to be an additional multiplier in the loss budget. The former Vmult was removed because valuation already determines whether a security is admitted and whether its expected return clears the required hurdle. Using valuation again to reduce the permitted loss charged the same uncertainty twice. The final formula is:

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

BaseIssuerLoss(FS) is the base loss allowance set by funded state, or FS. It is 0.25%, 0.50%, 0.75%, and 1.00% across the four funded-state bands. ECmult is the Evidence Coverage multiplier. The floor applies only to that epistemic, or knowledge-related, multiplier. Lmult is the liquidity multiplier, and Fmult is the factor-headroom multiplier. Those two measure physical portfolio constraints rather than uncertainty about the thesis, so their penalties retain full force. If an addition would breach a hard factor cap, Fmult is zero and allowed loss is zero. No floor can rescue the position.

This last point is a Round 5 correction to the original F4-1 wording. The initial proposal floored the product of evidence, liquidity, and factor multipliers at 0.50. That would have turned ECmult x Lmult x 0 into 0.50 and granted loss capacity precisely when the hard factor limit said there was none. It would also have softened the penalty for illiquidity. The final policy records the error rather than concealing it: only the evidence multiplier is floored.

Position size is then MaxWeight = AllowedLoss / S, where S is bear-case permanent-impairment severity. The fifth-round review found that an analyst could game this last input by declaring an implausibly mild loss. The final policy therefore floors S at 40% for ordinary single common equity, 60% for a levered financial, cyclical, or roll-up, and 80% to 100% for binary-license, pre-cash-flow, or emerging-market exposure with governance flags. A severity below its class floor requires written evidence and independent challenge.

Consider a hypothetical ordinary equity at the highest funded state. Assume the base issuer loss is 1.00%, Evidence Coverage is high enough for an ECmult of 1.0, liquidity is within five days so Lmult is 1.0, and there is full factor headroom so Fmult is 1.0. Allowed loss is 1.00%. The ordinary-equity severity floor is 40%, so the best-case ceiling is 1.00% / 40% = 2.5% of the total portfolio. If a hard factor breach makes Fmult zero, the ceiling is zero regardless of the company analysis.

The policy consequently chose branch (a) of D-1 openly. It intentionally rejects a generic 5% or 8% conviction tier. The reason is not that concentrated investing can never work. It is that investor edge, staffing, operational capacity, and liability closure are unverified, while a one-person process cannot honestly govern the 10% to 40% single-name positions associated with the concentrated-value minority view. That minority view survives as rejected with reasons, not 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 those conditions are met, apparent conviction cannot substitute for verified capacity and evidence.

F4-2: a persistent comparator breach must produce a choice

The policy treats its low-cost comparator as a real position, not a neutral background. That creates an uncomfortable but necessary consequence: an index can contain issuer or common-factor exposures that breach the same risk caps applied to active holdings. Round 4 recognized this possibility but allowed the system to settle into an indefinite documented exception. It also proposed an expected-return proof that the debate's own epistemology could not reliably supply. The result was a deadlock in which the largest embedded exposure could neither be reduced nor honestly accepted.

F4-2 replaced that deadlock with a forced choice. If an embedded issuer or factor hard-cap breach persists for four consecutive quarters, the policy must either adopt a capped or blended comparator, or raise the cap by formal amendment. A capped or blended comparator must be treated as a deliberate portfolio design, including its implementation consequences. Raising the cap must state that the embedded risk is being accepted and re-underwrite it. A renewable exception is not an available third path.

The diagnostic proxy used in the policy breached the old artificial-intelligence, or AI, factor stress-loss cap and the joint AI-plus-power factor stress-loss cap. The final policy therefore executed the cap-amendment branch provisionally. That was not a claim that the higher caps are proven correct. The actual investor comparator, current holdings, tax lots, and executable capped alternative were unknown. The provisional amendment makes the accepted risk visible while requiring the real comparator analysis before certification. Its trade-off is deliberate: a cap that is impossible to comply with from inception creates false discipline, but a raised cap does not remove the underlying risk. The policy prefers an explicit, amendable acceptance over permanent silent breach.

F4-3: funded status needs a second axis and state-change discipline

Funded status is FS = A* / PVL*. A* means investable assets after known taxes and transaction costs. PVL* means the probability-weighted present value of liabilities plus an explicit safety margin. The ratio is useful, but F4-3 rejected the idea that a single decimal can fully describe risk capacity.

Two investors can have the same FS while facing very different hazards. Near, essential, and inflexible liabilities make a given ratio less forgiving than distant or deferrable liabilities. Income that falls when markets fall also weakens capacity because the household or business loses both portfolio value and outside cash flow at the same time. The final selector therefore modifies the funded-state reading for liability duration and flexibility and for the correlation between income and market losses.

The rule also addresses estimation error and cliff effects. PVL* must be reported at its base value and at plus or minus 10% and plus or minus 20%. If ambiguity changes the applicable state, the lower state is used. A promotion

requires two consecutive annual reviews; it cannot be earned by one favorable mark. Verified capacity deterioration, by contrast, causes immediate demotion.

In a hypothetical example, suppose investable assets are 1.50 million and base PVL* is 1.00 million, producing FS of 1.50. If a 10% higher liability estimate raises PVL* to 1.10 million, FS falls to about 1.36. The policy displays both results and uses the lower applicable state if the sensitivity crosses a boundary. It also asks whether the liabilities are near and inflexible and whether income is exposed to the same market shock. This approach sacrifices some simplicity, but it prevents spurious precision from authorizing more risk than the underlying life circumstances support.

F4-4: separate crash protection from inflation-grind matching

The earlier policy used lower gross risk as the default answer to both sudden crashes and long inflationary grind. F4-4 found that the same answer does not work for both problems. Cash and nominal duration can reduce crash exposure, yet still lose purchasing power during a prolonged inflationary period. A lower equity allocation is therefore a plausible null for crash and sequence risk, but not necessarily a hedge for real liabilities.

The final policy treats inflation-linked duration as structural liability matching when three conditions pass: the liability is genuinely real, the relevant index is an appropriate match, and law and market integrity are credible. In that setting, the instrument is not admitted because it is a fashionable diversifier. It is admitted because its contractual behavior is intended to match the liability. Trend and commodity exposures remain conditional diversifiers; they are not smuggled into the default portfolio to make a stress test pass.

The grind test was rerun using only default-admissible assets and was required to print the boundary honestly. Under its stated diagnostic assumptions, the inherited 45% equity, 25% nominal duration, 20% inflation-linked bond, and 10% liquidity-floor mix ended with 1.0875 million against 1.10 million of stressed real liabilities. Coverage was 0.989, a failure of 12,500 immediately before any spending. A boundary mix of 40% equity, 20% nominal duration, 30% verified inflation-linked matching, and 10% liquidity floor ended at 1.14 million, or coverage of 1.036, a 40,000 pass. That pass exists only if roughly 30% can genuinely match the real liability and its index, basis, title, and market integrity hold. It is a diagnostic boundary, not a recommendation or certification.

F4-4 matters because it prevents a stress test from borrowing success from an asset the policy would otherwise disallow. It also prevents the label "lower risk" from hiding a mismatch between nominally safe assets and real obligations. The trade-off is basis risk: an inflation-linked security is useful only to the degree that its index and legal claim actually correspond to the liability being matched.

F4-5: budget valuation derating before purchase, not after it

The policy distinguishes permanent business impairment from a mark-to-market decline caused by valuation compression. A high-quality business can retain its cash economics while its market multiple falls sharply. Selling automatically after that fall can convert ordinary repricing into a governance error. Yet refusing to sell after the event is credible only if the potential derating was budgeted before purchase.

F4-5 supplies that missing rule. For each active position, the estimated derating loss at purchase is:

w x (1 - conservative-case multiple / current multiple)

Here w is the proposed portfolio weight. The two multiples come from the current valuation and the conservative reverse-discounted-cash-flow case. The losses are summed across active positions. By default, total active derating loss must fit within 50% of the market-drawdown budget, or MDB, after accounting for the strategic sleeve's own stress. Raising the allowance as high as 75% requires a written amendment and independent challenge.

For a clearly labeled hypothetical, a 2% position bought at 25 times the relevant earnings or cash-flow measure with a conservative-case multiple of 15 times has estimated derating loss of 2% x (1 - 15/25) = 0.8% of the portfolio. That 0.8% consumes the active derating budget even if the business thesis remains intact. If the position cannot fit at purchase, its size must be reduced then.

An ex-post breach does not mandate an automatic sale of an intact thesis. It triggers a sizing-process post-mortem. This preserves the distinction between bad business evidence and market repricing while making the cost of duration-like equity exposure visible beforehand. The trade-off is that conservative inputs can reduce entry size and cause opportunity cost, but the alternative is to discover portfolio-level valuation concentration only after prices fall.

F4-6: match custody safeguards to structural complexity

The earlier second-jurisdiction proposal bundled simple operational redundancy together with cross-border legal structures and subjected both to the same professional gauntlet. That could leave the investor at zero protection indefinitely because an exhaustive review was too costly for the simplest remedy. F4-6 separated the problem into three tiers.

Tier A is a second custodian in the same jurisdiction. It is mandatory. Its purpose is to reduce a single operational point of failure without adding a new legal system or complex wrapper. Tier B places plain instruments under a second legal system. It uses a focused review of tax, reporting, title, transfer, sanctions, and estate-situs consequences. Tier C covers entities, trusts, and complex wrappers. It defaults to 0% until the full gauntlet passes.

The rule matters because "diversification" can be false if a second account creates tax, succession, title, or sanctions risk larger than the custody risk it removes. Tiering takes the low-complexity protection immediately while refusing to treat legal complexity as free. It also preserves a hard boundary: nothing about Tier B or Tier C can be personalized until the relevant domicile and professional facts are verified.

F4-7: make look-through factor caps both measurable and honest

Factor Exposure Weight, or FEW, estimates how much of the portfolio is exposed to a common causal factor. Factor Stress Loss, or FSL, estimates portfolio loss under the specified factor shock. The policy applies these measures to the comparator as well as to active holdings. The original ambition to score every comparator constituent repeatedly was not executable by one person, and a risk cap already breached by the comparator could not function as a real limit.

F4-7 adopted a reproducible proxy protocol. The top 20 comparator constituents are scored individually. The remainder is grouped into sector-proxy buckets. The analysis is refreshed annually and on a material event. In year one, an independent repeat after six months measures scoring drift. This is a deliberate balance between false precision and neglect: the largest contributors receive name-level work, while the long tail is represented rather than ignored.

Each feasibility cap is calculated as max(risk cap, comparator FSL + 1 percentage point). In a hypothetical example, if the risk cap is 5% and comparator FSL is 6.4%, the feasibility cap is the greater of 5% and 7.4%, so it is 7.4%. The 2.4 percentage-point gap from the original risk cap is not declared safe. It must be reported in currency as accepted embedded risk. If portfolio value were 2 million, that gap would correspond to 48,000 of stress loss. This example is only arithmetic, not a statement of an actual comparator or portfolio.

The rule prevents a born-in-breach limit from becoming meaningless, but it introduces a serious trade-off: feasibility can follow an already concentrated comparator. The final policy therefore treats the diagnostic result as provisional because the actual comparator and worksheet are missing. F4-7 provides the measurement procedure; it does not certify the embedded exposure or turn it into a recommendation. Where the original risk cap is already consumed, factor headroom can remain zero for an active addition even when the company underwrite is otherwise attractive.

F4-8: tie entry pacing to information, not comfort

Calendar tranching can look disciplined while serving only to reduce the emotional discomfort of acting. F4-8 keeps staged entry available for single-name thesis risk, but it attaches each tranche to evidence. Tranches must be linked to a filing or a preregistered milestone and expire by day 180.

The default is full deployment when every gate and the conservative case clear and no information event is pending. If the operator nevertheless uses behavioral pacing, the journal must record the expected opportunity cost. The rule therefore allows time to reveal thesis error when a genuine information event is near, but does not allow "waiting" to become an unpriced habit.

The 180-day limit is an outside boundary, not a requirement to wait six months. A filing or milestone can arrive earlier, and if no relevant event is pending, the default points back to full deployment. Most importantly, this is a rule only. It is not authorization to buy any security, and no tranche exists unless the separate investment and approval gates have already been satisfied.

F4-9: standardize tax deferral and close the perpetual-hold loophole

Tax is a real friction, but an embedded gain can make deferral appear more valuable each year and thereby turn "after tax" into a perpetual excuse to hold a deteriorating opportunity. F4-9 standardizes the comparison. Deferral value uses the applicable statutory rate, a fixed ten-year horizon, and the return of the comparator that would replace the position. Those inputs must not be changed opportunistically to reach a preferred answer.

If two consecutive annual hold decisions are justified chiefly by deferral, an independent challenger must review the position. The phrase "next liquid window" is defined as five trading days unless a security-specific market-integrity event prevents normal execution. A proposed 90-day wait for tax reasons requires its cost to be calculated at day 0, not rationalized near the end of the delay.

This rule does not make tax irrelevant. It preserves the value of deferral when the business case remains intact and the standardized comparison supports holding. What it prevents is tax friction from becoming the thesis itself. The policy's broader boundary remains: tax may support hold-not-add, but it cannot indefinitely rescue a failed business case.

F4-10: apply financial complexity to the right target

The former complexity flag was designed for a non-financial company whose financing activities, receivables structures, or off-balance-sheet arrangements obscure the operating business. Applied literally, however, it also blocked every bank and insurer simply because their principal business is a balance sheet. That reversed an earlier decision to include balance-sheet financials in the research taxonomy.

F4-10 confines the generic complexity flag to ancillary finance inside a non-financial business. A principal balance-sheet business instead faces a class-specific battery. Insurers are tested on reserve development, float cost, and capital. Banks are tested on deposit beta, full-cycle credit, non-performing loans, or NPLs, and capital. Asset managers are tested on net flows, fee compression, client concentration, and capital.

These batteries are not an easier path to Core classification. Missing or failed evidence blocks Core with the same force as the complexity flag. The amendment changes the question from "does this firm have a balance sheet?" to "is the balance-sheet business understood with the evidence appropriate to its class?" Its trade-off is greater analytical specialization, but categorical exclusion would be simpler only by being wrong.

F4-11: active management must earn its research cost

The policy asks active complexity to beat a low-cost comparator after tax, cost, hours, and error. F4-11 made that principle calculable at the program level. Annual active-program value is:

active sleeve x evidenced premium x evidenced hit rate

That expected portfolio contribution is compared with research hours valued at the operator's opportunity cost, plus data and tax costs. A hypothetical 10% active sleeve, a 2 percentage-point evidenced premium, and a 50% evidenced hit rate imply 10% x 2% x 50% = 0.10%, or 10 basis points of expected portfolio value per year before fully loaded research costs. For a portfolio of 1 million, that is 1,000. The example illustrates why a demanding research apparatus can be uneconomic even when individual ideas appear promising.

Two consecutive negative years suspend new active work and return the operating state to comparator plus the Investor Liquidity Floor, or ILF, until active work is re-justified. The minority view that the correct active allocation may be zero remains live. So does the counterview that maintaining a bounded research capability has option value during dislocations. The observable is not enthusiasm for either view, but demonstrated value relative to hours and costs.

This amendment matters because governance itself consumes scarce capital and attention. A complex policy can fail by asking one person to maintain more dashboards than can be kept current. The suspension rule makes simplification a valid outcome rather than an admission of failure.

F4-12: remove overlap inside the liquidity floor

The ILF protects near-term essential liabilities, fixed commitments, taxes, and income interruption. In the earlier formula, the dated spending ladder D36 and the extra interruption buffer H could count the same spending twice. F4-12 removes that overlap: H covers only essential spending beyond the 36 months already funded through D36.

The final formula is expressed as ILF target = D36 + Tax12 + Commit + H x max(0, essential spending beyond D36 - dependable income covering that later interval). Tax12 is the next twelve months of known taxes, and Commit covers fixed commitments. The policy also permits an annual documented election to replace D36 with D24, but only when FS is at least 1.75 and spending flexibility is high.

This does not turn a high funded state into permission to shrink liquidity automatically. The D24 substitution must be elected and documented each year. Nor does H vanish; it applies to the later interval rather than duplicating the first 36 months. The trade-off is transparent. More liquidity can reduce forced-sale risk but impose real return drag. The amended formula makes that insurance cost deliberate rather than an accidental product of double counting.

F4-13: keep orchestration records out of the constitution

Debate provenance matters. The process retained hashes, run metadata, authentication and cache details, receipts, and the record that no trade execution was authorized. But those facts describe how the policy debate was orchestrated, not how the portfolio should be governed.

F4-13 therefore places orchestration hashes and run metadata only in Appendix C and the traceability record, not in the governing constitution. This is an editorial amendment with substantive value: it keeps the durable rules readable while preserving an auditable chain for anyone who needs to verify the debate artifacts. Removing metadata from the constitution does not erase it. It locates it where it can serve provenance without being mistaken for investment evidence.

F4-14: preserve uncertainty through adjudicating observables

The final amendment registered two disputes rather than pretending to resolve them by assertion. U4-1 asks whether the full apparatus is executable by its intended one-person operator. Its observable is actual annual hours, cost, and decision value. The final policy resolves the question adversely for the full apparatus but accepts a bounded version: no more than eight live full underwrites; otherwise the operating default is comparator plus ILF plus nominal and real liability matching.

U4-2 asks whether diversifier evidence bars are themselves procyclical. A rule that admits a diversifier after a strong trailing decade and rejects it after a weak one may be a disguised regime bet. The dispute remains open. Its observables are rolling-start-date results across inflation, deflation, crash, and calm; net implementation cost and capacity; and crisis correlation. No single decade decides the question.

F4-14 matters because a dispute register should contain testable uncertainty, not a museum of unresolved opinions. U4-1 has an operational boundary because the implementation evidence was already adverse. U4-2 remains open because the evidence depends on cycle coverage and start-date sensitivity. In both cases, the policy states what evidence could change the conclusion.

What this chapter requires in practice

The fourteen adjudications establish a common operating posture. Arithmetic must reveal rather than conceal a policy choice. A hard breach cannot be rescued by a floor or renewed exception. Funded status must be tested against liability uncertainty and life circumstances. Crash, permanent impairment, valuation drawdown, and inflation grind must not be treated as the same risk. Custody, tax, entry pacing, and financial-company research must be governed at the level of the actual problem rather than by a generic label. Active management and governance must earn their full cost.

In practical terms, an operator must use the corrected loss formula and severity floors; treat a factor breach as zero headroom; resolve any four-quarter comparator breach through a real written choice; report PVL* sensitivities and apply promotion hysteresis; verify real-liability matching before relying on inflation-linked duration; budget derating loss at purchase; maintain Tier A custody redundancy while subjecting cross-border complexity to the appropriate review; score the comparator through the top-20-plus-buckets protocol; anchor any entry tranche to information and day 180; standardize tax-deferral inputs and escalation clocks; use class-specific financial batteries; test whether active work pays for its hours; avoid double counting D36 and H; keep technical provenance in Appendix C; and monitor U4-1 and U4-2 through their stated observables.

None of these steps certifies liability closure or authorizes a transaction. Their purpose is narrower and more demanding: to ensure that when a future decision is considered, the rule governing it is explicit, arithmetically coherent, operationally feasible, and honest about the evidence it does not yet possess.

Chapter 2 - The Investment Constitution and Its Durable Principles

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.

What this chapter requires in practice

In practice, begin each material decision by stating its place in the purpose hierarchy and labeling every material claim as fact, inference, assumption, forecast, classification, or unresolved. Report FS = A* / PVL*, stressed Coverage, real after-tax progress, currency mismatch, and liquidity before celebrating portfolio return. Test permanent impairment, mark-to-market drawdown, and decade-long real grind separately.

Name the low-cost implementable comparator and look through both it and every active exposure to issuers, sovereigns, factors, currencies, financing chains, and counterparties. Require active complexity to beat the comparator after tax, cost, hours, and error. Treat modeled shortfall as seriously as excess risk, while never exceeding capacity or trait tolerance. Let evidence govern tolerable loss and valuation govern admission; let neither confidence nor narrative grant weight.

Maintain the no-recourse-leverage default and analyze mortgages, paid options, and collateralized structures on the full balance sheet and through notional, path, expiry, basis, and repeat-premium effects. Rebalance broad strategic sleeves through approved mechanical rules, but require fresh evidence for single-name action except narrow solvency or integrity protection. Re-underwrite holdings as if unowned, allowing tax deferral to support hold-not-add but never an indefinitely failed thesis.

Turn persistent breaches into amendments or compliant design changes. Keep any emergency exception risk-monotone, record it within 24 hours, obtain independent review within seven days, and let it expire within 30 days. Preserve the mandatory ILF check, look-through dashboard, decision journal, annual funded-status review, exception log, and active-program viability test; require all additional complexity to earn its hours. Finally, keep research and implementation separate: no theme, classification, example, ceiling, or conclusion in this chapter authorizes any order, message, account action, or portfolio change.

Chapter 3 - From Principles to a Funded-State Operating System

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.

What this chapter requires in practice

Maintain a verified annual schedule of adjusted investable assets, liabilities by time and currency, known taxes and commitments, dependable income, spending flexibility, income-market correlation, custody access, and decision authority. Calculate FS = A* / PVL* at base PVL* and at +/-10% and +/-20%; use the lower state if +/-10% changes the band. Require two consecutive annual confirmations before promotion, but recalculate and demote immediately after a verified capacity event.

Record one coherent strategic/opportunity design and the associated EB, MDB, illiquid, speculative-loss, and tracking-error limits. Apply the second-axis one-state reduction to MDB and EB when liabilities are near or inflexible or income is highly market-correlated. Calculate and protect the ILF target and hard floor in liability currency, document any annual D36 to D24 election, and report safe-liquidity drag as insurance cost.

Use the strategic rebalancing thresholds, do not mechanically refill a broken theme, and obey the event clocks: immediate freeze and 0-5 trading days for hard gate, title, fraud, or de-licensing concerns; 30-day full re-underwrite and no add for Pillar 0, hard-cap problems, two missed milestones, or EC <60; no mechanical trim for valuation alone, but no add when E fails admission. If entry is authorized elsewhere, anchor it to information and complete it

within 180 days. Log exceptions, and after two material Policy Exception Events in 12 months, freeze new discretion pending review. Until these facts and controls are complete, the chapter remains a conditional operating map and authorizes no transaction.

Chapter 4 - How to Judge a Business Model for a Ten-Year Holding Period

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.

What this chapter requires in practice

For any company considered for a ten-year holding period, build a current evidence file appropriate to its business class; score A through E without adding them; require every pillar to be at least 2; calculate EC separately and keep EC below 60 research-only; resolve material evidence conflicts at zero; run every legal, title, valuation, factor, and implementation gate; choose a bear-case permanent-impairment severity no lower than the applicable floor; apply the funded-state, evidence, stress-liquidity, and factor multipliers exactly; and divide allowed loss by severity to obtain a maximum diagnostic weight. Record the hard disconfirmers before acting. Treat any zero, hard-cap breach, or missing required class battery as a stop, and treat every resulting ceiling as a limit rather than a recommendation.

Chapter 5 - The Decade Ahead: Scenarios, Themes, and Counterthemes

What a decade theme can and cannot do

A decade theme is a conditional causal map. It begins with a need, constraint, or change in behavior; identifies the paid solution that might address it; traces the economic value created; asks who has bargaining power over that value; and ends with the cash, after reinvestment and financing, that could reach each share. It is conditional because every link can break. It is a map because its proper use is to direct research toward bottlenecks, payers, business models, and observable disconfirming evidence. It is not a forecast that one future will occur, a label for an industry, a promise that social value will become shareholder value, or an argument that any security is attractively priced.

The governing chain is:

> *Need or constraint -> paid solution -> volume times unit economics -> allocation by bargaining power -> reinvestment and financing -> per-share owner cash -> price paid -> portfolio role.*

This chain prevents several common category errors. Rapid adoption does not prove economic value: customers may experiment without changing production. Economic value does not prove corporate profit: lower prices, higher wages, taxes, or procurement savings may absorb it. Corporate profit does not prove per-share capture: dilution, acquisitions, executive compensation, related-party leakage, state take, or debt service may intervene. Per-share cash does not prove a good investment: an optimistic price can discount more success than the business ultimately delivers. Finally, an attractive security in isolation does not establish an appropriate portfolio role, because the same underlying factor may already be embedded in the comparator or other holdings.

The scenarios in this chapter are therefore coherent alternatives, not assigned probabilities and not switches for tactical trading. The themes are research priorities, not allocation buckets. The counterthemes are not automatic hedges or short cases; they are mechanisms through which the failure, commoditization, or second-order consequence of a popular thesis may create a different value pool. All are dated to the policy evidence cutoff of 22 August 2026 and remain subordinate to liability funding, the Investor Liquidity Floor (ILF), the named low-cost comparator, the five-pillar business-model scorecard, security-level valuation, loss budgets, and explicit owner approval.

Reading scenarios without pretending to forecast

The six scenarios are deliberately broad enough to expose different weaknesses in a portfolio and a research process. More than one can occur at once. Productive artificial intelligence (AI) diffusion can coexist with trade fragmentation. A physical power constraint can coexist with an economy-wide credit bust. Financial repression can be the policy response to both inflation and high public debt. A decade may also move through several regimes rather than remain in one. The scenarios should therefore be used to ask whether liabilities can still be met, whether legal title and liquidity remain sound, whether supposedly independent exposures share a failure mechanism, and which claims need to be re-underwritten.

The authoritative scenario map is preserved below.

Scenario	Mechanism	Winners to research	Failure/countertheme
S1 productive diffusion	AI/electrification raise output after complementary capital	mission-critical data, semicap, automation, power management, workflow	capex glut, weak customer monetization, shorter asset lives
S2 inflation/repression grind	nominal assets and cash lose real value; states steer capital	indexed liabilities, pricing power, short-cycle distributors, sound balance sheets	controls, windfall taxes, index basis/integrity failure
S3 fragmented blocs	export controls, localization, defense and duplicate supply chains	testing, trusted foundry, grid, sovereign procurement bottlenecks	monopsony, stranded localization, retaliation

Scenario	Mechanism	Winners to research	Failure/countertheme
S4 crash/sequence	valuation and credit unwind together	ILF, matched duration, liquid quality after re-underwrite	false safety, forced selling, comparator concentration
S5 adaptation/physical constraint	power, water, climate and insurability become gating inputs	efficiency, grid control, underwriting discipline	unfunded public plans, rate suppression, litigation
S6 plateau/muddle	nominal growth, mediocre productivity, range-bound themes	low-cost comparator, boring compounders, disciplined capital return	research theater and overtrading

S1 productive diffusion

S1 is the favorable technology case, but it is not a simple continuation of infrastructure spending. Its mechanism is that AI and electrification raise output only after complementary capital arrives. Complementary capital includes redesigned workflows, clean and permissioned data, employee training, liability controls, electrical connections, automation integration, and changes in customer behavior. The delay matters. The historical analogy is factory electrification: replacing a steam engine with an electric motor did not by itself create the full productivity gain; factories had to be reorganized around the new capability. In the same way, purchasing compute or software does not establish productive diffusion.

The observable path is broader adoption in governed production, lower task or unit cost, improved throughput and quality, and customer cash savings that exceed implementation expense. As of the debate evidence, US surveys showed any generative-AI use at about 41% in November 2025 but daily use at about 12%, while Federal Reserve work cautioned that investment and adoption precede measurable aggregate productivity (Federal Reserve adoption note, Federal Reserve buildout dashboard). Those observations support a diffusion question, not a conclusion.

The winner classes to research are mission-critical data, semiconductor capital equipment, or semicap, automation, power management, and workflow. These classes may benefit because reliable data, scarce fabrication, integration know-how, electrical control, and embedded workflows become necessary complements. Yet the transmission to business economics differs by layer. A semicap supplier may capture scarcity rent through qualification and a concentrated equipment market; a workflow company must prove retention and pricing after AI changes seat counts and customer value metrics; an automation vendor must show customer payback and an installed service base; a power-management supplier must convert backlog into cash without being displaced when lead times normalize.

The failure case is equally central: capex glut, weak customer monetization, and shorter asset lives. The debate documented extraordinary capital commitment, including roughly $495 billion to $515 billion of combined 2026 capex described by Amazon, Alphabet, and Meta, before other buyers, and $95.2 billion of NVIDIA manufacturing, supply, and capacity commitments. These dated figures show a heavy buildout, not that a security is overvalued. The adverse telecom/fiber analogy is that demand can explode while capacity owners lose money because unit prices collapse and capital is financed ahead of monetization. Asset-life error compounds the danger: if economically useful lives are shorter than accounting lives, reported depreciation understates the real cost of serving customers.

Within 36 months, the relevant evidence includes recurring production deployment, customer savings, utilization, rental or token prices, cancellations, inventory charges, and incremental gross profit after depreciation and financing. The theme weakens if pilot activity does not become recurring governed use, if prices fall without enough volume response, or if customer AI gross profit does not cover incremental depreciation and capital cost. A minority view must remain open: hyperscaler spending may defend existing search, cloud, advertising, or software rents even when direct returns on AI capex look low. Under that interpretation, persistence of the core franchise matters alongside utilization. Neither interpretation grants ownership; it changes which facts the underwrite must gather.

S2 inflation/repression grind

S2 combines two related but distinct stresses. In an inflationary grind, nominal revenues, nominal asset values, and nominal cash balances can rise while purchasing power and real liability coverage deteriorate. In financial repression, states steer capital or cap financing costs so that savers and minority owners absorb part of the adjustment. The mechanism may include administered rates below inflation, captive demand for government securities, directed bank or retirement capital, capital controls, windfall taxes, or regulation that favors nominal solvency over owners' real return.

The research candidates are indexed liabilities, genuine pricing power, short-cycle distributors, and sound balance sheets. "Indexed liabilities" refers here to matching verified real obligations with claims whose index, law, duration, and market integrity genuinely track those obligations; it does not mean that any security marketed as inflation-linked is safe. Pricing power must be measured after volume, mix, wage cost, input cost, and working capital, rather than inferred from nominal revenue growth. Short-cycle distributors may reprice inventory and serve necessary small-ticket demand more quickly than long-cycle project businesses, but thin margins, substitution, and working-capital needs remain. Sound balance sheets matter because inflation and policy steering can raise refinancing cost even for firms with nominal growth.

The countermechanisms are controls, windfall taxes, and index basis/integrity failure. A politically visible surplus may attract regulation or taxation. A regulated utility may experience rising nominal rate base yet fail to receive timely cost recovery. A resource owner may see higher prices while royalties, taxes, local-content obligations, and capex absorb the gain. An inflation-linked bond may fail to match the actual consumption basket, may carry the wrong duration, or may depend on an index or legal promise whose integrity weakens. Home-currency matching can itself become the channel of repression if convertibility, title, or real purchasing power is impaired.

Evidence should distinguish true short repricing from accounting optics. Useful observables include realized price versus unit volume, wage and input lags, interest coverage, refinancing schedules, real free cash flow per share, rate-case recovery, effective tax and state take, index basis, convertibility, and custody integrity. The historical record offers no universal asset that always preserves value under repression. Productive equities, property, commodities, gold, inflation-linked claims, and foreign assets have each failed in some episodes through valuation, leverage, controls, taxes, confiscation, or custody. The policy response is therefore liability matching and jurisdictional integrity analysis, not a mechanical "real assets" trade.

S3 fragmented blocs

S3 begins with export controls, localization, defense requirements, and duplicate supply chains. Trade or technology restrictions cause firms and states to replicate fabs, inventories, logistics routes, grids, payment systems, and strategic production. This can increase nominal orders and improve resilience while reducing global utilization and productivity. Enterprise quality cannot be assessed separately from supplier access, export licenses, customer nationality, currency, legal system, and the location of critical physical nodes.

The winner classes to research are testing, trusted foundry, grid, and sovereign procurement bottlenecks. Testing can matter because products and systems must satisfy more jurisdictions, standards, and security requirements. A trusted foundry may become indispensable within a bloc, but its owner still faces capex intensity, geographic concentration, export law, and state bargaining. Local grid and control infrastructure may receive accelerated demand as production is duplicated. Sovereign procurement bottlenecks may possess qualifications, security clearances, installed bases, or scarce capacity.

The path from policy demand to owner cash is especially fragile. Public announcements and budget aspirations are not funded contracts. Funded contracts are not necessarily attractive contracts. Backlog must be tested for appropriation, cancellation, fixed-price risk, escalation clauses, delivery capacity, cash conversion, and clawbacks. The principal countertheme is monopsony: a dominant state buyer can demand local capacity while capping supplier margins, renegotiating terms, or taxing visible profit. Stranded localization occurs when subsidized

capacity arrives after the strategic need changes or when parallel plants operate below economic utilization. Retaliation can remove export markets, raise input costs, or jeopardize legal title.

Observable support includes widening control lists, funded multiyear contracts, increasing local-content requirements, durable capacity utilization, and cash conversion after compliance and duplication expense. Falsification includes cancellations, low utilization, recurring cost overruns, margin compression by state buyers, and retaliation that eliminates the assumed market. The policy looks through regional revenue labels to critical nodes such as foundry, shipping lane, payments access, custodian, cloud, and export license. It does not assume that a domestic champion is an attractive security merely because its government wants the capacity.

S4 crash/sequence

S4 is a joint valuation and credit unwind. High expectations reverse, collateral values fall, refinancing tightens, and forced sales reinforce price declines. Sequence matters because liabilities due early in the drawdown can force an investor to sell assets before recovery, making a temporary market loss permanent at the household level. A technology or capital-spending bust can also remain concentrated in one complex while the broader economy muddles through; that acute variant is particularly important when the low-cost comparator itself is concentrated in the affected issuers.

The relevant "winners" are not securities selected in advance as crisis trades. They are the ILF, matched duration, and liquid quality after re-underwrite. The ILF is the policy's emergency reserve for essential spending, known taxes, fixed commitments, and the defined later-liability component. Matched duration means claims aligned to the timing and type of verified liabilities, not a directional rates bet. Liquid quality means a business whose demand, capital economics, governance, balance sheet, and valuation survive fresh analysis - not a prestigious label purchased mechanically after a price decline.

The failure modes are false safety, forced selling, and comparator concentration. Cash can be false safety if custody, convertibility, inflation, or bank exposure is impaired. Long-duration government bonds can fail if the inflation or issuer-integrity assumption is wrong. A private asset can appear stable because marks are stale while its underlying economics have deteriorated. A broad comparator can contain large shared exposures to mega-cap technology, country, currency, or duration. None of these risks is cured by a low volatility label.

The operational sequence is liability and access first, re-underwriting second, and already-approved mechanical action only within the surviving capacity and liquidity constraints. A price decline alone is not evidence that expected return has improved, because earnings, credit, title, or dilution may have worsened. Emergency authority restores access or reduces risk; it does not authorize opportunistic buying. If income impairment or early claims break liability coverage, restoring the liability plan outranks "buying low."

S5 adaptation/physical constraint

S5 treats power, water, climate, and insurability as gating inputs rather than background costs. Rising load can encounter slow transmission and equipment delivery. Heat, drought, flood, fire, and storm can impair facilities, labor productivity, municipalities, insurance availability, and supply chains. The IEA reported more than 2,500 gigawatts of generation, storage, and large-load projects in connection queues and estimated that annual grid investment would need to rise roughly 50% from about $400 billion by 2030, while grid projects can require 5 to 15 years (IEA Electricity 2026). Yet queues include speculative and duplicate applications, so they are evidence of both constraint and possible phantom demand.

The research classes are efficiency, grid control, and underwriting discipline. Efficiency can avoid the need for scarce generation or network capacity. Grid-control equipment, sensors, software, and power management can improve the use of existing assets. Underwriting discipline applies to insurers, lenders, property owners, and project developers: the central skill is to price physical exposure, obtain rate or premium adequacy, limit correlated loss, and retain the option to withdraw or reprice.

The payer problem is decisive. UNEP estimated developing-country adaptation needs of $310 billion to $365 billion annually by 2035 against $26 billion of international public flows in 2023. That roughly 12-to-1 aspiration-to-payment gap demonstrates large need and weak funded demand at the same time. Unfunded public plans, rate suppression, and litigation are the stated failure cases. A municipality may need water infrastructure but lack tariff capacity. A regulator may prevent an insurer from repricing while courts expand liability. A utility may be required to harden its network but be denied timely cost recovery. A contractor may record orders yet consume cash through working capital and fixed-price exposure.

The theme requires funded appropriation, rate recovery, contract-to-cash conversion, and real per-share free cash after maintenance. Within 36 months, evidence should include paid orders rather than plans, falling loss or outage intensity per unit of exposure, adequate premium or tariff repricing, and cash conversion after working capital. Failure appears when public plans remain unfunded, cost recovery is suppressed, insurance exposure grows without adequate pricing, or litigation transfers the benefit away from shareholders.

S6 plateau/muddle

S6 is the antidote to dramatic-scenario bias. Nominal growth continues, productivity is mediocre, and themes remain range-bound. AI becomes a useful but ordinary tool; major capability gains slow; capital spending grows closer to revenue; inflation and term premia stabilize; no single crisis or boom dominates. The sector-confined bust may occur as an acute episode without producing an economy-wide deflationary collapse.

The winners to research are the low-cost comparator, boring compounders, and disciplined capital return. The comparator wins by delivering broad exposure with low cost and little decision burden. Boring compounders win only if recurring demand, incremental returns, governance, and price are real; "boring" is not a waiver of the scorecard. Disciplined capital return matters where managers lack attractive reinvestment and return surplus cash per share rather than pursue empire building.

S6's failure cases are research theater and overtrading. The investor may manufacture narratives, frequent monitoring, and unnecessary decisions because the process was designed around dramatic change. The costs include research hours, data, taxes, spreads, abandoned comparator exposure, and errors. A theme that generates no decision value should not be kept alive simply because it sounds important. The debate's research-sunset rule moves an inactive theme to annual-check status after 36 months without a journaled decision, including a documented abstention with stated thresholds. At the program level, two consecutive negative annual active-viability results suspend new active work to comparator-plus-ILF status until re-justified. S6 is not an argument for complacency; it is a requirement that complexity earn its cost.

The named research themes

The final policy names ten research themes. They are business-economics lenses that can operate across several scenarios, not conventional sector baskets. Their boundaries matter because a broad narrative can conceal radically different payer, capital, and governance structures.

Verification/standards

Verification/standards covers testing, inspection, certification, audit, identity, model monitoring, and other activities that establish whether an asset, process, output, or claim meets a required standard. The causal mechanism is a rising cost of failure combined with regulation, liability, or customer trust. Capture is strongest when accreditation is scarce, the purchase is mandatory, the service is a small share of the customer's cost, and failure would be catastrophic. This model may benefit in S1 as AI deployment requires audit trails, in S3 as blocs impose separate standards, and in S5 as physical systems require inspection.

It is not a single homogeneous industry. Mandated assurance differs from discretionary consulting, and an accredited testing network differs from a labor-heavy field contractor. AI may reduce labor hours or commoditize parts of the work; regulators may impose liability; conflicts of interest can damage the franchise; and customers

can consolidate vendors. Evidence must include the actual mandate, accreditation barriers, renewal and retention, real price, labor hours per job, error and liability history, entrants, and how much customer savings the provider retains. The open dispute over AI and assurance rents is adjudicated by real price, labor input, errors, liability, entry, and client-retained savings - not by the mere growth of regulation.

Installed-base aftermarket

Installed-base aftermarket refers to proprietary or highly advantaged service, spares, consumables, and upgrades attached to equipment already in use. The installed base creates recurring need; downtime cost and qualification can make switching difficult; and service may require less capital than manufacturing the original equipment. This can support durable cash economics in aircraft, industrial systems, healthcare equipment, grids, and automation.

The theme is not equivalent to all maintenance. Independent alternatives, customer insourcing, right-to-repair rules, standardization, or predictive diagnostics can erode the original equipment manufacturer's control. New equipment may be reliable enough to reduce service intensity, while an aging base may be retired rather than maintained. The necessary evidence includes installed-base age and growth, utilization, attach rate, service retention, independent-alternative share, aftermarket gross profit, parts availability, and maintenance capex. Organic service cash must be separated from acquisition-led roll-up accounting.

Market infrastructure

Market infrastructure includes exchanges, clearing, settlement, payment rails, custody, financial data, identity, and compliance systems. Its causal mechanism is that participants value liquidity, trusted records, regulatory permission, low error, and network access. Two-sided density can lower customer acquisition cost and improve the product as participation grows. Revenue may arise from transactions, data, clearing, float, software, or reduced fraud.

The countercase is that public instant-payment rails, regulation, fee compression, cyber losses, or internal bank systems commoditize private tolls. Tokenization may remain in pilots or fragment liquidity into walled gardens. Crypto-asset price exposure is not the same as market-infrastructure economics. Underwriting must measure take rate, volume quality, customer concentration, fraud and credit loss, switching, regulatory capital, liquidity, and per-share cash after investment. Network size alone does not establish owner capture.

Reliability distribution

Reliability distribution is the provision of small-ticket but mission-critical parts, supplies, field support, and availability through a dense network. Customer value comes from avoiding downtime and reducing procurement complexity; the distributor may capture through breadth, local inventory, response speed, and low share of customer cost. It can serve the reliability needs in S1, S3, and S5 without depending on a single large project.

The label does not imply high margins or a moat. Distribution can be price competitive, working-capital intensive, vulnerable to manufacturer direct sales, and dependent on technician or warehouse labor. Evidence should show organic growth, retention, inventory turns, receivables, service levels, supplier concentration, substitution, local density, and incremental cash return. Acquisition growth and leverage require separate scrutiny. A roll-up that buys distributors is a financing structure, not proof of reliability economics.

Semicap and power-control bottlenecks

Semicap and power-control bottlenecks combine scarce enabling nodes in semiconductor production with electrical management, switching, cooling, and control. The mechanism is that productive diffusion, electrification, localization, and data-center growth require specialized equipment and reliable power before downstream capacity can operate. Qualification barriers, technical complexity, and long lead times can create bargaining power.

This theme contains one of the policy's largest correlated risks. AI/compute and energy/grids share the data-center-load variable; both depend to varying degrees on hyperscaler capex, East Asian fabrication,

equipment delivery, and state policy. A slowdown can hit semiconductor orders, grid backlogs, automation, and data-center services together. Scarcity also attracts supply. The evidence must include book-to-bill, lead times, cancellations, utilization at customers, economic asset lives, replacement and service revenue, customer concentration, power reservations that become operational load, and free cash after capex. It is not enough to cite a queue or backlog.

Funded health innovation

Funded health innovation means therapies, diagnostics, tools, and care-productivity systems that deliver material clinical or economic benefit to a payer able and willing to pay. "Funded" is the critical qualifier. Demography, disease burden, or a scientific breakthrough establishes need, not collectible revenue. Shareholder capture depends on evidence, reimbursement, net price, access, manufacturing, exclusivity, replacement pipeline, and total cost of care.

The theme includes a built-in counterforce. Prevention can reduce procedures, admissions, devices, and labor-intensive chronic care; payers or employers may capture more of the benefit than the therapeutic or provider. Patents and exclusivity expire, me-too pipelines invite price competition, and longer life can increase other care costs. Evidence must therefore track real-world adherence and outcomes, payer access, net price, hours or cost saved per episode, exclusivity cliffs, replacement pipeline, research and development productivity, acquisition spend, and cash conversion per share. No terminal value may bridge a patent cliff by assumption.

Regulated resilience

Regulated resilience covers utilities, grids, water, adaptation, defense supply, and other businesses whose demand or economics depend materially on public rules, appropriations, or rate cases. The need may be essential and the technical position scarce, but the payer can be a monopsony or a regulator. The Payer Quality Gate therefore asks whether money has been appropriated, whether a contract or rate case converts to cash, whether allowed returns compensate the capital, and whether cancellation, clawback, state take, or dilution changes per-share economics.

The qualifier "regulated" is not a synonym for safe. Rate suppression, affordability politics, fixed-price procurement, disallowed costs, windfall taxes, litigation, and forced investment can transfer resilience benefits away from owners. Backlog and public targets receive no cash credit until funding, margin, cancellation, and conversion are verified. The same state-payer factor can correlate grid, defense, adaptation, and localization holdings, especially during S2 and S3.

Balance-sheet franchises

Balance-sheet franchises include insurers, banks, and asset managers whose product and economics reside principally in the balance sheet. They are a counterweight to the producer-goods bias of many decade narratives, but no automatic tilt under S2 or S3 follows. For insurers, the evidence battery includes reserve development, float cost, and capital. For banks it includes deposit beta, full-cycle credit, non-performing loans, and capital. For managers it includes net flows, fee compression, client concentration, and capital.

Transmission occurs through underwriting or credit discipline, low-cost and durable funding, risk selection, reinvestment, and per-share capital allocation. The failure mechanism is leverage hidden inside apparently stable accounting: reserve deficiency, correlated catastrophe loss, deposit flight, credit deterioration, duration mismatch, liquidity dependence, or fee compression. Nominal earnings in inflation are not enough. A sound balance sheet must survive stress without diluting owners or requiring state support on punitive terms.

Selective brands

Selective brands acknowledges that consumer desire, trust, identity, and scarcity can produce pricing power and attractive capital economics outside the industrial-capex research agenda. "Selective" is essential. A familiar logo, premium history, or high gross margin is not proof that the next cohort retains the same desire. Pricing can conceal unit decline, channel stuffing, promotion, or loss of relevance.

The evidence is cohort desire, unit volumes, price and mix separately, repeat purchase, resale or scarcity where relevant, distribution control, marketing efficiency, geographic concentration, and incremental cash return. AI and digital distribution can either strengthen customer reach or reduce the information and creative tolls charged by incumbents. S6 may favor a genuinely durable brand with disciplined capital return, but crowding and valuation duration can still make an excellent business a poor security.

EM platforms with audited per-share capture

Emerging-market, or EM, platforms may benefit from formalization, urbanization, digital payments, logistics, financial access, and productivity convergence. The final qualifier - audited per-share capture - prevents country growth from being mistaken for investor return. Gross domestic product can grow while currency depreciation, dilution, promoter transactions, state leakage, related parties, foreign-currency debt, or poor minority rights consume the result.

The evidence must reconcile audited cash, ownership, guarantees, funding currency, related parties, local regulation, capital allocation, and real per-share earnings or free cash in a relevant hard-currency and liability context. User growth, payment volume, or nominal local-currency revenue is insufficient. There is no single EM capital cycle: some countries or platforms may be crowded, some manufacturing chains overbuilt, and other markets starved of capital. The theme is a selective governance-and-capture test, not a geographic allocation instruction.

Counterthemes and the discipline of reversal

Counterthemes make the map two-sided. They ask who benefits if the popular mechanism works differently from expected, if abundance replaces scarcity, or if political and capital-cycle responses transfer value. They remain conditional sketches unless developed to the same payer, capture, cycle, valuation, and falsification standard as a primary theme.

In the final policy's compressed formulation, the counterthemes are: energy abundance erodes scarcity rents; AI replaces seat/value metrics and some information tolls; procurement duplication wastes capital; prevention shifts payer economics; peace dividend reverses defense; financial repression punishes visible surplus; capex harvest may reward survivors but not builders. The paragraphs below unpack those exact claims without promoting them into forecasts.

Energy abundance erodes scarcity rents. Solar, storage, efficiency, demand response, transmission reform, or new supply can reduce marginal power cost faster than load grows. Consumers, flexible-load operators, and electrified industries may gain while merchant generators and high-cost fuels lose. The falsifier is that all-in firmed power cost, including storage and networks, does not fall or reliability deteriorates. Even if abundance occurs, downstream businesses must prove that competition does not pass all savings to their own customers.

AI replaces seat/value metrics and some information tolls. If AI performs more work per employee, software priced per seat may face shrinking units or pressure to move toward usage, outcome, or value pricing. Routine information intermediaries, undifferentiated applications, and some proprietary-looking data tolls may be absorbed by foundation models or customer workflows. Beneficiaries may include low-cost adopters, verifiers, identity and security providers, liability transfer, and human escalation. Evidence includes normalized retention after seat changes, price per outcome, customer savings, model switching, proprietary rights, error rates, and assurance spend. This is not a blanket short thesis against software or infrastructure.

Procurement duplication wastes capital. Localization and strategic autonomy can create nominal orders for factories, inventories, logistics, and defense while lowering utilization and raising taxes. Security value may be real even when shareholder return is poor. The falsifier is that diversified supply chains achieve utilization, productivity, and free cash returns comparable to the prior global system. Until then, public support must be treated as a payer and governance question, not as free demand.

Prevention shifts payer economics. Effective obesity, cardiovascular, cancer-screening, vaccine, or other preventive interventions may shrink downstream procedure and chronic-care volumes. The value can accrue to payers, employers, and patients rather than to providers or drug owners, particularly if reimbursement negotiation captures the savings. The countertheme fails if real-world adherence and outcomes disappoint or if longer life raises total downstream costs more than morbidity falls.

Peace dividend reverses defense. Ceasefires, political turnover, burden-sharing fatigue, or sovereign financing stress can slow procurement before new capacity earns its cost of capital. Households and civil investment may benefit; late-cycle defense capacity may be stranded. The falsifier is multiyear funded procurement and production contracts that remain intact through political turnover and debt stress. A backlog is not enough without margin, cancellation, delivery, and appropriation evidence.

Financial repression punishes visible surplus. Businesses perceived to benefit excessively from inflation, scarcity, public need, or state support may face price controls, windfall taxes, directed investment, reserve requirements, or capital restrictions. The mechanism can harm resource owners, utilities, banks, defense suppliers, and even foreign-revenue firms. The proper evidence is realized state take, real price after controls, capital mobility, convertibility, index integrity, and cash per share. This countertheme explains why a correct inflation or scarcity call can still produce poor owner returns.

Capex harvest may reward survivors but not builders. This is the required counterweight to a uniformly downstream view. Capital cycles move from scarcity to buildout to competitive flood and, sometimes, consolidation and harvest. Telecom around 2000 warns that builders can be impaired while users win; consolidated railroads after 1893 and memory semiconductors after 2012 warn that surviving infrastructure can become attractive after capital exits and industry structure improves. The harvest thesis requires confirmed competitor exit, disciplined survivor capex, improved utilization, balance-sheet survival, and cash return. It does not rehabilitate the original builders or permit buying early merely because consolidation might later occur.

Correlated risks and cross-theme failure

Theme labels can create an illusion of diversification. AI, semicap, power equipment, grid construction, automation, and maintenance may all be partial derivatives of the same capex cycle. Defense, regulated grids, adaptation, and localization can share the same state payer and fiscal constraint. Critical materials, electrification, and data centers may share power, China policy, shipping, and commodity factors. Market infrastructure, balance-sheet franchises, and EM platforms may share regulation, funding liquidity, cyber risk, and currency. Selective brands and software can share long-duration valuation and consumer demand even when their products differ.

The look-through task is to aggregate the causes, not count the names. At minimum, research should identify sensitivity to AI-capex sustainability, state-payer share, East Asian fabrication, power availability, interest-rate duration, commodity inputs, currency, legal system, and comparator-embedded exposure. A portfolio containing a chip-equipment supplier, a transformer producer, an automation vendor, and a data-center service company may still be one large buildout bet. A broad-market comparator may already hold substantial exposure to the same factor. Adding a thematically different company can increase rather than diversify the common risk.

Crowdedness is also multi-dimensional. Capital commitments, new entrants, issuance, capacity, inventory, depreciation, valuation, and index weight all contain information, but none alone proves overvaluation. A queue can show scarcity and phantom demand. Backlog can show contracted demand and cancellations waiting to happen. A high multiple can reflect durable economics or unrealistic duration. A low multiple can reflect an impending collapse

in cash flow. Security-level reverse discounted cash flow analysis must therefore infer the growth, margins, capital needs, and duration embedded in price.

Evidence, falsification, and capital-cycle discipline

Every decade thesis must produce intermediate causal evidence within 36 months. That requirement does not demand that a general-purpose technology transform aggregate productivity within three years; complementary capital may take longer. It does demand timely evidence at the link actually underwritten. An AI infrastructure thesis can be tested through utilization, pricing, depreciation, customer economics, cancellations, and financing. An installed-base thesis can be tested through retention, attach rate, independent alternatives, and service cash. A regulated-resilience thesis can be tested through appropriation, rate recovery, margins, and contract-to-cash conversion. An EM platform thesis can be tested through audited per-share cash, dilution, related parties, funding currency, and minority rights.

Capital-cycle states should be recorded explicitly. Scarcity appears in high utilization, long lead times, rising prices, and disciplined additions. Buildout appears when orders, capex, financing, and entrants accelerate. Competitive flood appears when capacity outruns demand, prices fall, inventories rise, cancellations spread, and incremental returns decline. Consolidation or harvest appears when weak owners exit, capacity rationalizes, survivors restrain investment, and cash returns improve. The state is not permanent, and the transition matters more than the label. The policy uses the cycle symmetrically: scarcity and buildout warn against extrapolation; consolidation and harvest invite research, but only after the evidence confirms survival and discipline.

Falsification must reach business economics and per-share capture. A socially successful theme may remain an investment failure if its providers earn below their cost of capital. Conversely, a broad theme may disappoint while one licensed bottleneck or disciplined survivor produces durable owner cash. Two missed primary causal milestones require a full re-underwrite rather than an explanation based on sentiment. A failure of legal title, payer collectibility, financing, or per-share capture cannot be excused by continued end demand.

Why no theme authorizes a security purchase

No theme authorizes ownership. A theme cannot establish the security's rights, governance, balance sheet, economic depreciation, valuation, tax treatment, liquidity, custody, factor headroom, or relationship to the investor's liabilities. It also cannot prove that the named vehicle captures the value pool. The final policy therefore treats theme research as non-authorization. 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, applicable tax and legal review, and all policy gates pass.

Even a fully evidenced causal map must pass the five pillars: demand and lock-in; industry structure, cycle, and payer; capital economics; governance and per-share capture; and valuation plus falsifiability. Evidence Coverage below 60 remains research-only, and coverage at or above 60 is only an input to the separate loss budget, not a buy signal. The valuation case must clear the named low-cost comparator after tax, cost, hours, and error. Factor limits can reduce allowed loss to zero even for an attractive security. The ILF and liability-matching rules always precede thematic expression.

This boundary also prevents themes from smuggling in portfolio changes through language such as "winner," "beneficiary," or "hedge." In this chapter those words mean only a class whose economics warrant research under a stated scenario. They do not mean that its securities will outperform, that the price is acceptable, or that an existing holding should be added, sold, or retained. A countertheme similarly authorizes neither a short position nor an offsetting purchase.

What this chapter requires in practice

Use the six scenarios to stress liability coverage, access, and common-factor exposure, not to forecast the next regime. For each live theme, write the causal chain from need to per-share owner cash; identify the payer, capital-cycle state, countertheme, correlated factors, simplest comparator or abstention alternative, and at least one observable 36-month milestone. Re-underwrite when the causal link, payer, financing, legal title, or capture evidence fails. Move inactive research to annual-check status when it does not earn its cost. Above all, keep the policy boundary explicit: no scenario, theme, countertheme, company class, or research conclusion in this chapter authorizes a security purchase or any portfolio change.

Chapter 6 - The Grind Test: What Must Survive an Unfriendly Decade

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.

What this chapter requires in practice

Before relying on a grind result, verify A^* after known taxes and transaction costs and build PVL^* from the actual liability cash-flow schedule with its safety margin. Record each liability's amount, timing, currency, nominal or real character, relevant inflation index, tax treatment, and required legal access. Classify assets by function rather than label: equity, genuinely matched nominal duration, verified duration-matched ILBs, and ILF cash. Confirm index, law, market, basis, title, custody, duration, and currency integrity before giving an ILB the 0% real treatment.

Run the printed stresses without a speculative +15% diversifier gain and show weighted losses, end assets, coverage, and the pass or failure amount in currency. Keep the official conclusions visible: 45/25/20/10 fails immediately by 12,500 at 0.989 coverage; 40/20/30/10 passes by 40,000 at 1.036 coverage only under the genuine real-match condition. Do not call either row a target allocation. Add the missing withdrawal and income path before claiming any survival duration. Until that work is complete, label the withdrawal-path grind budget and liability closure as certification blockers, preserve lower gross risk as the crash null, preserve verified inflation-linked matching as the grind instrument, and authorize no portfolio change from this chapter.

Chapter 7 - Comparator Scoring, Forced Choices, 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 $F_{mult} = 0$, 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.

What this chapter requires in practice

The operator must treat the comparator as a look-through portfolio position. Until the investor's actual named comparator is known and the required study is complete, the 60% MSCI ACWI and 40% liability-matched high-grade fixed-income proxy remains diagnostic only; its 8.2% AI FSL and 9.0% joint AI+power FSL are plausible order-of-magnitude assumptions, not ratified outputs. The original 5% and 8% risk caps govern, and every active addition with material AI/power sensitivity remains at zero admissible addition under this factor rule.

The actual study must archive the top 20 constituent weights, bucket the remainder, score the five factors and joint shocks, complete two independent passes with the year-one second pass after six months, show ±0.25 perturbation ranges, separate liability matching from directional rates FEW while reporting basis risk, and convert any risk-versus-feasibility gap into currency. It must compare an implementable capped or blended alternative, its costs, tax effects, tracking-error budget, and abandonment point, using Japan-1989, US-1992-2025 and US-2013-2025, Nortel/TSE, and Nokia/HEX as opposing exhibits.

If the actual comparator remains in hard-cap breach for four consecutive quarters, the policy must resolve the breach through a written capped/blended-comparator decision or a formal cap amendment. No renewable exception is available. Any temporary diagnostic ceiling is fixed, cannot exceed twice the original cap, cannot float with index concentration or analyst convenience, and auto-reverts after 12 months. Completing this work removes only the actual-comparator certification blocker. It does not certify the broader policy, make any company a buy, establish any target weight, or authorize a trade, order, hedge, message, account opening, or portfolio change.

Chapter 8 - The Global 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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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.

What this chapter requires in practice

Treat the 42 companies as a dated, expiring research map. Preserve each exact classification, but attach no portfolio meaning to it. When a company is selected for one of the limited live-underwrite slots, rebuild the evidence file from the cited pointers, apply the relevant class battery, test all five non-compensating pillars, normalize the cycle and economic depreciation, complete a conservative reverse DCF, state the falsifier, and check legal, tax, liability, liquidity, and factor constraints. Refresh the classification within 12 months or let it revert to Unclassified Research Note. Run the annual counter-canon and non-US searches, document rejected candidates for any claimed absence, and keep every name at admissible weight zero unless a separate current review and approval process is fully completed. Nothing in this chapter authorizes a trade, order, message, target allocation, or portfolio change.

Chapter 9 - 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.

What this chapter requires in practice

Maintain a narrow monthly survival check, a quarterly holdings and thesis-breakage review, and a comprehensive annual reset. Expire an unrefreshed company classification after 12 months. Stop adds and complete a 30-day re-underwrite after two missed milestones or EC below 60. At every annual map refresh, report the counter-canon and non-US screens and name at least three examined rejections behind any claimed class absence. Move a theme with no journaled decision or explicit abstention threshold after 36 months to annual-check status, while archiving its evidence and recording its hours and cost.

Apply the exit hierarchy without substituting price for evidence: integrity threats use the minimum separately authorized risk-monotone response within five trading days; a zero in Pillars A through D, a persistent hard cap, or a failed full-underwrite gate leads to a governed close-or-reduce review over five to 30 days; an intact thesis with sub-comparator expected return means no add, independent challenge after two deferral-led annual holds, and action within five trading days after the decision; price decline alone never compels sale. Never average down without a current thesis and all applicable gates.

Before any purchase authorized elsewhere, archive the current and conservative valuation cases, calculate and aggregate w x (1 - Mconservative / Mcurrent), and keep active derating loss within 50 percent of MDB after strategic-sleeve stress. Reaching as high as 75 percent requires a written amendment and independent challenge. Treat an ex-post overshoot as a sizing post-mortem when the thesis is intact, not as an automatic trim. Price every tax delay on day 0, preserve all owner-approval and implementation boundaries, and remember that nothing in this chapter authorizes a transaction.

Chapter 10 - 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 contains a drafting tension that must not be hidden. Chapter 10 of the final policy states a record of at least 10 years after tax and cost, with preregistered attribution, verified liability surplus, an independent challenger, and operational capacity. The final F5-13 amendment elsewhere records a replacement test of a preregistered five-year journal covering at least 20 fully underwritten decisions, independently audited after-tax attribution, verified liability surplus, and challenger sign-off. Because both appear in the final instrument, no concentration reopening should be treated as procedurally complete until a versioned amendment reconciles the duration. For this chapter, the explicit Chapter 10 observable of at least 10 years is preserved; the five-year, 20-decision amendment is also disclosed rather than silently discarded.

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.

What this chapter requires in practice

At each annual review, first verify liability closure and the nominal and real matching implementation. Then archive and independently score the actual comparator, including its top 20 holdings, sector buckets, factor ranges, and any hard-cap breach. Run the prospective active-program viability test, print the fully loaded hours and cost, and reduce the live underwrite set before work begins if capacity is insufficient. Review the exception log, EC auditability, derating budget, diversifier start-date sensitivity, custody transfer evidence, and the status of every open dispute that could alter the architecture.

Maintain a bright line between a company falsifier and a portfolio falsifier. Company evidence governs the relevant thesis through the scorecard and exit hierarchy. Portfolio evidence governs whether new active work may continue at all. If any of the five portfolio falsifiers is met, freeze new active work and default to the named comparator, the ILF, and verified nominal and real liability matching. Do not improvise a replacement, infer a trade, or treat a higher cap as risk reduction. Record what failed, what evidence would reopen the process, and which disagreement remains unresolved.

Chapter 11 - Twelve Hard Questions and the Policy's Direct Answers

The twelve questions in this chapter were not requests for another set of investment themes. They were tests of whether the policy could survive contact with its own arithmetic, its own research map, and the limits of a one-person process. Each question demanded a direct answer before the final policy could be treated as an operating framework. The answers therefore have a deliberately narrow status. They explain what the policy concluded, what remains provisional, and what work is still missing. They do not turn the research map into a portfolio, certify the investor's liabilities, or authorize a transaction.

Several definitions recur. Funded status, or FS, compares investable assets after known taxes and transaction costs with the probability-weighted present value of liabilities plus an explicit safety margin. Evidence Coverage, or EC, is the verified share of the applicable required evidence, weighted toward primary sources. Bear-case severity, or S, is the assumed permanent impairment if the thesis fails. Wmax is the largest weight that the loss-budget arithmetic could permit under stated assumptions. It is a diagnostic ceiling, not a target or recommendation. The investable-liquidity floor, or ILF, protects near-term liabilities and spending needs. Factor-equivalent weight, or FEW, measures portfolio exposure to a common factor, while factor stress loss, or FSL, translates that exposure into an estimated portfolio loss under the specified shock. None of these quantities has been personalized or certified.

Question 1: Were all fourteen Round 4 amendments adjudicated, and how was the concentration dispute resolved?

The first question asked for an explicit disposition of every item from F4-1 through F4-14: accept it, reject it with reasons, reframe it, or leave it unresolved with a test. It singled out D-1, the long-running disagreement between diversified position sizing and a concentrated-value approach. The required choice was visible and consequential: either acknowledge that the policy intentionally caps ordinary single-name conviction near 2%, or create a separately governed conviction tier that could support larger positions.

The answer is that all fourteen F4 amendments were accepted, although later final-round corrections refined how some of them operate. F4-5 was calibrated to a 50% default, and F4-2 was executed provisionally. D-1 takes branch (a): the policy intentionally accepts an ordinary ceiling of about 2% at S=50% in the highest funded state and creates no conviction tier. The concentration position is not erased. It survives as a rejected-with-reasons minority and may be reconsidered only through the stated reopening test.

The individual adjudications matter because the aggregate answer hides substantial content:

1. F4-1 removes valuation as an additional loss multiplier, retains the evidence, liquidity, and factor constraints, and makes a hard factor breach set allowed loss to zero. The final loss arithmetic also imposes severity floors. Ordinary single equities use at least 40% severity; levered financials, cyclicals, and roll-ups use at least 60%; binary-license, pre-cash-flow, or flagged emerging-market exposures use 80% to 100%. These rules make the old generic 5% and 8% single-name rows inoperative. The final text describes a best-case 2.5% ceiling at S=40% and, consistently with the adjudicated D-1 branch, about 2% at S=50%.
2. F4-2 accepts the forced-choice rule for persistent comparator breaches. Four consecutive quarters above an embedded issuer or factor hard cap require a written choice between a capped or blended comparator and raising the cap by amendment. There is no indefinitely renewable exception.
3. F4-3 makes funded status a two-dimensional selector. Liability duration and flexibility, and the correlation between income and market losses, can make the same numerical FS support different risk capacity. The policy reports liability value at plus or minus 10% and 20%, uses the lower state if a 10% change crosses a band, requires two consecutive annual reviews for promotion, and demotes immediately after verified capacity deterioration.

4. F4-4 makes verified inflation-linked duration a structural match for verified real liabilities when the relevant index, law, market, basis, and title integrity pass. Lower gross risk remains the first response to crash risk, but is not treated as a complete answer to an inflationary grind.
5. F4-5 creates a valuation-derating budget. At purchase, derating loss is w x (1 - conservative-case multiple/current multiple). Here w is the proposed portfolio weight. The sum for active positions must fit within 50% of the maximum drawdown budget, or MDB, by default. Raising that share as high as 75% requires a written amendment and independent challenge. A later breach prompts a process post-mortem; it is not an automatic sale when the operating thesis remains intact.
6. F4-6 separates custody into three tiers. A same-jurisdiction second custodian is mandatory. Plain instruments in a second legal system require the limited tax, reporting, title, transfer, sanctions, and estate-situs review. Entities, trusts, and other complex wrappers default to zero until the full professional gauntlet passes.
7. F4-7 makes comparator scoring feasible by examining the top 20 constituents individually and using sector proxies for the rest. It requires annual and event-driven refresh, plus an independent repeat after six months in year one. The risk-versus-feasibility gap must be reported rather than concealed.
8. F4-8 ties entry tranches to filings or preregistered milestones and gives them a 180-day outside limit. If every gate and the conservative case clear and no information event is pending, full deployment is the default. Behavioral pacing must record its expected opportunity cost. This rule describes a method only; it authorizes no entry here.
9. F4-9 standardizes tax deferral using the applicable statutory rate, a fixed ten-year horizon, and the comparator replacement return. Two annual holds justified mainly by deferral trigger an independent challenger. The next liquid window is five trading days unless a security-specific market-integrity event is documented, and a 90-day tax wait needs its cost calculation at day zero.
10. F4-10 confines the financial-complexity flag to ancillary finance inside a non-financial business. Principal balance-sheet businesses instead face class-specific batteries: reserve development, float cost, and capital for insurers; deposit beta, full-cycle credit or non-performing-loan behavior, and capital for banks; and net flows, fee compression, client concentration, and capital for asset managers. Missing or failed battery evidence blocks Core status.
11. F4-11 compares expected active-program value - active sleeve times evidenced premium times evidenced hit rate - with research hours valued at the operator's opportunity cost plus data and tax costs. Two consecutive negative years suspend new active work and return the process to comparator-plus-ILF status until it is re-justified. The minority view that active exposure should be zero remains live.
12. F4-12 prevents double counting inside the ILF. The income-interruption horizon covers only essential spending beyond the period already funded in the dated-liability term. At FS>=1.75, high spending flexibility may support an annually documented election to use 24 rather than 36 months.
13. F4-13 moves run hashes and other orchestration metadata into the traceability appendix. They establish provenance, not investment evidence.
14. F4-14 registers two further disputes: whether one person can execute the apparatus, U4-1, and whether time-stamped evidence bars for diversifiers become procyclical, U4-2. Each remains tied to observable tests.

The concentration conclusion is intentionally conservative because the policy lacks verified investor edge, complete liability facts, and institutional staffing. It does not claim that concentrated investing can never work. Section 1 gives a concrete reopening route: a preregistered five-year record spanning at least 20 fully underwritten decisions, independently audited after-tax attribution, verified liability surplus, and challenger sign-off. The dispute appendix summarizes the hurdle more broadly as a verified record of at least ten years, liability surplus, independent challenge, and operator capacity, while preserving the minority case for rare 10% to 40% control or edge positions. Those two formulations are not silently collapsed here. The controlling conclusion is unchanged: no conviction tier exists now, and reopening requires a formal amendment supported by a demonstrated record rather than confidence alone.

Question 2: What weights does the loss-budget formula actually permit at each severity?

The second question demanded the implied maximum-weight table before any company was classified. Its purpose was anti-theatrical: a label such as "Core Research Candidate" means little unless the reader knows whether the policy could hold 1%, 5%, or 20% if all remaining work cleared.

The governing formula is AllowedLoss = BaseIssuerLoss(FS) x max(0.50, ECmult) x Lmult x Fmult, used only after all admission gates pass. Base issuer loss is 0.25%, 0.50%, 0.75%, or 1.00% across the four funded-status bands. ECmult is the evidence multiplier; the formula prevents epistemic evidence from reducing it below the stated 0.50 floor after admission. Lmult is 1.0 for liquidity within five days, 0.75 for six to 30 days, 0.50 for 31 to 365 days, and 0.25 beyond one year. Fmult is 1.0 with full factor headroom, 0.75 when factor use exceeds 75%, 0.50 when it exceeds 90%, and zero if the addition would breach a hard cap. There is no floor that rescues illiquidity or a factor breach. Maximum weight is AllowedLoss/S.

At the best funded state, with full liquidity, full factor headroom, and the best evidence multiplier, the direct answer prints the following mathematical implications. The parenthetical value is the 0.50 epistemic-floor equivalent:

Bear-case severity S	Best-case Wmax	At 0.50 evidence floor
100%	1.00%	0.50%
80%	1.25%	0.625%, rounded to 0.63% in the formal table
60%	1.67%	0.835%, rounded to 0.83%
50%	2.00%	1.00%
40%	2.50%	1.25%
30%	3.33%	1.67%
20%	5.00%	2.50%
12.5%	8.00%	4.00%

A simple hypothetical shows the logic. Suppose the highest-FS base loss is 1%, the position is fully liquid, full factor headroom remains, evidence is at the 0.50 floor, and severity is 50%. Allowed loss is 1% x 0.50 x 1.0 x 1.0 = 0.50%. Dividing by 50% severity gives a 1% maximum portfolio weight. With the best evidence multiplier, the same position reaches 2%. Neither result says the investor should own it.

The lower-severity rows expose the old table's algebra, but they do not bypass the final anti-gaming floors. An ordinary equity cannot claim 20% or 12.5% permanent-impairment severity merely to reach a 5% or 8% weight. Under the final rule, ordinary equity uses at least 40%, levered or cyclical cases at least 60%, and binary cases 80% to 100%, unless the special written challenge for a lower severity is satisfied. Hard admission gates still make the permitted weight zero. In addition, every named company currently has an admissible weight of zero because investor facts, evidence audits, current valuation, factor headroom, and implementation approval are incomplete. The table makes constraints visible; it never becomes a portfolio proposal.

Question 3: Did the grind rerun pass using only default-admissible assets?

The third question challenged an earlier stress example for relying on a real diversifier that the policy itself did not admit by default. The rerun therefore uses only equity, liability-matched nominal bonds, verified duration-matched inflation-linked bonds, or ILBs, and the ILF. It makes no forecast and hides no withdrawal in the arithmetic that was actually supplied.

The stated hypothetical begins with 1.50 million of investable assets and 1.10 million of stressed real liabilities. Equity falls 45%, matched nominal duration loses 25% in real terms, verified ILBs earn 0% real, and ILF cash loses 10% real. ILBs receive no speculative 15% gain.

For the inherited 45% equity, 25% nominal bonds, 20% ILBs, and 10% ILF mix, end assets are 1.5 x (1 - .2025 - .0625 - 0 - .01) = 1.0875 million. Coverage is 0.989, so the mix fails immediately by 12,500. This is not a small rounding issue: coverage below 1 means the stressed assets do not fully meet the stressed liability estimate.

A diagnostic boundary mix of 40% equity, 20% nominal bonds, 30% ILBs, and 10% ILF ends at 1.5 x (1 - .18 - .05 - 0 - .01) = 1.1400 million. Coverage is 1.036, so it passes by 40,000. That pass is conditional on approximately 30% of the portfolio being genuinely matched to real liabilities and on index basis, legal title, market integrity, currency, and duration all being sound. It is neither a recommended allocation nor certified liability closure.

The question also required the duration of accepted slow failure. The honest answer is unbounded once withdrawals occur, because no withdrawal schedule was supplied. The failure can only worsen, but its maximum duration cannot be calculated without the missing cash-flow path. This is the practical distinction between a drawdown statistic and a liability test: the same percentage market loss can be survivable or terminal depending on when and how much the investor must withdraw. Lower gross risk remains the policy's crash-risk default. Verified inflation-linked matching is the specific grind instrument. Nothing in this example certifies that the investor owns, can access, or can correctly match a 30% ILB allocation.

Question 4: Is the comparator already in breach, and what forced choice follows?

The fourth question required scoring the comparator across five common factors at top-20-plus-sector-proxy resolution. The five are AI-capex sensitivity, state-payer exposure, China and export-control exposure, rates exposure, and power exposure. Each holding receives a 0 to 1 sensitivity estimate; the top 20 comparator constituents are scored individually and the remainder is grouped using sector proxies. FEW aggregates the amount exposed. FSL applies the specified shock and estimates the portfolio loss. Comparator holdings consume the same risk headroom as active holdings because an index is a risky position, not a neutral absence of choice.

The direct answer is that the diagnostic comparator is born in breach of the AI and combined AI-plus-power caps. The inherited proxy produced AI FSL of 8.2% and joint AI-plus-power FSL of 9.0%, above the original 5% and 8% risk caps. The Round 4 direct answer records F4-2 branch (b) as provisionally executed with diagnostic ceilings of 9.2% for AI and 10.0% for the joint factor, with the gap from the original caps reported. It also identifies a rates FEW breach caused by matched bonds, although the final comparator section treats liability-matched duration separately from directional rates FEW because it serves a liability rather than a speculative rate view.

The final certification block is more restrictive than those provisional figures might sound. The actual investor comparator is missing, the archived top-20 worksheet does not exist, and the required second independent scoring pass has not been completed. Therefore the proxy is diagnostic, not ratified. The original caps continue to govern; no active AI or power additions are permitted while the original limits are exceeded. Any temporary diagnostic ceiling may never exceed twice the original cap, cannot float upward with index concentration or analyst convenience, and automatically expires after 12 months. The actual comparator decision must be made within that 12-month window.

Ratification requires the actual named comparator, archived top-20 weights, sector buckets for the remainder, two independent scoring passes, sensitivity ranges using plus or minus 0.25 changes in factor scores, a capped or blended alternative, a tracking-error budget and abandonment point, and the required historical comparison exhibits. The important practical meaning is not that raising a cap cures risk. A cap amendment acknowledges and records accepted embedded risk. Until the work is complete, actual-comparator scoring is a certification blocker, the proxy decision remains unratified, and active AI or power additions remain at zero.

Question 5: Did the company map include balance-sheet financials and genuine harvest candidates?

The fifth question tested whether the research map would quietly reproduce a quality-growth bias by omitting financials and capital-cycle survivors. The answer is mixed by design. Balance-sheet financials are present. A full harvest candidate is not.

Chubb and DBS are Core Research Candidates using the insurer and bank batteries rather than a generic financial-complexity penalty. Berkshire Hathaway is represented as an insurer and capital allocator. Nu is a Speculative Option, reflecting its shorter record and risk profile. Flagstar is an Avoid/Counterexample that demonstrates the bank battery can reject rather than merely bless a financial company. These classifications expand the model map without implying an automatic financial-sector tilt.

No full harvest candidate qualifies. The policy examined the capital-cycle idea but requires more than a distressed share price, state support, or a story about eventual capacity discipline. Full harvest confirmation needs actual financing-window closure, net capacity exit, balance-sheet survivability, and cash returns to shareholders. The map records that Micron or memory, Maersk or container shipping, and Transocean or offshore services were examined but left out because the required evidence was incomplete.

Intel is the explicit rejected case. It has continuing manufacturing investment and government-linked support, and it lacks confirmed net capacity exit and shareholder cash returns. Its screen is A2/B1/C0/D2/E1, EC72%, severity 100%, and Wmax zero. Its possible reversal test requires two independent external customers scaling on leading nodes, capex divided by depreciation below 1, and segment cash return above cost. Until that evidence appears, state support cannot be treated as harvest proof, book or fab value is not owner value, and Intel remains an Avoid/Counterexample rather than a normal-size or speculative harvest candidate.

The limit is explicit: absence of a qualifying harvest name is a documented research conclusion for this edition, not proof that no future company could qualify. The search itself remains incomplete.

Question 6: Is every Core row complete enough to be valid?

The sixth question required every named Core Research Candidate to print five pillar scores A through E, Evidence Coverage, bear-case severity, maximum diagnostic weight, and the standard non-buy caveat. A missing field would make the row void under the map rules.

All 11 Core rows contain the required fields. The names are S&P Global, CME Group, Visa, Fastenal, Berkshire Hathaway, Chubb, DBS Group, Taiwan Semiconductor Manufacturing, ASML, Schneider Electric, and RELX. Each row states its A-to-E profile, EC, S, and Wmax; describes valuation sensitivity, failure conditions, missing evidence, and a 36-month falsifier; and repeats that the classification is not a buy recommendation. This completeness does not mean the underlying evidence is complete. The scores are provisional analyst assumptions, links are evidence pointers rather than credited verification, no independent adversarial-corroboration pass is finished, and any unresolved conflict between higher- and lower-tier evidence must score zero.

Most importantly, every Core name remains at current admissible weight zero. The diagnostic Wmax figures range according to severity - for example, 2.50% for the 40% severity assigned to Berkshire and Chubb, 2.00% for several 50% severity cases, 1.67% for S&P Global at 60%, 1.43% for ASML at 70%, and 1.25% for TSMC at 80%. These are optimistic ceilings under best-case portfolio conditions. Current valuation, factor capacity, investor circumstances, evidence audits, and authorization are all still missing. "Core" means worthy of deeper research, not approved for ownership.

Question 7: Does the Avoid class cover every major model class and reject admired companies when the arithmetic fails?

The seventh question required the map to demonstrate that its gates can reject attractive stories, including at least one admired consensus-quality company. The direct answer identifies 12 Avoid/Counterexample rows: Palantir, Intel, Boeing, Coinbase, PG&E, Plug Power, Nike, Teladoc, Flagstar, Bausch Health, First Quantum Minerals, and Adani Enterprises.

Together they cover or stress the major classes: vertical workflow and AI software through Palantir; semiconductors and harvest claims through Intel; installed-base aftermarket, aviation, and defense through Boeing; exchanges and custody rails through Coinbase; regulated power and grid spending through PG&E; energy abundance through Plug Power; consumer brands through Nike; healthcare platforms through Teladoc; banking through Flagstar; serial acquisition and leverage through Bausch; natural resources through First Quantum; and emerging-market infrastructure, state dependence, and per-share capture through Adani. Where a listed pure-play counterexample is unavailable, the model-class table documents the absence or incomplete search rather than pretending the class was tested.

The policy also separates admiration from admission. NVIDIA and Ferrari are admired consensus-quality Watch names whose E-pillar arithmetic fails. NVIDIA is A4/B3/C4/D4/E0 with EC74%, severity 70%, and Wmax zero because E0 fails the admission gate. Ferrari is likewise kept from ownership by its valuation pillar rather than rescued by brand quality. Palantir is the explicit admired Avoid: A4/B3/C3/D1/E0, EC68%, severity 80%, and Wmax zero. Its market narrative of mission-critical software and ontology-based switching costs does not override per-share dilution, concentration, governance, or conservative-return arithmetic.

This is the practical purpose of the Avoid class. It is not a permanent blacklist or a prediction that each security must fall. Each row includes evidence that could reverse the conclusion. But until that evidence satisfies the applicable hard gates and a current reverse discounted-cash-flow review, popularity, index membership, price decline, market size, and reputation cannot create expected-return proof.

Question 8: Do the financial batteries constrain favorable as well as unfavorable cases?

The eighth question asks whether F4-10 merely creates a new route for approving banks and insurers. The answer is no. The same battery that allows a principal balance-sheet business to be analyzed on the right dimensions also blocks Core status when evidence is absent or weak.

Chubb is tested on reserve development, float cost, and capital. DBS is tested on deposit beta, full-cycle credit behavior, and capital. In the final map, all favorable scores are provisional analyst judgments capped at 2: Chubb is reserves 2, float cost 2, capital 2; DBS is deposit beta 2, full-cycle credit 2, capital 2. The cap remains until ten-year reserve triangles, float-cost and stressed-capital exhibits for Chubb, and deposit cohorts, full-cycle loss and stressed common-equity-tier-1 capital exhibits for DBS are attached. Strong current reported results do not substitute for those through-cycle exhibits.

The contrasting cases prove that the battery can constrain. Nu scores 2/1/2 on its applicable bank dimensions and remains a Speculative Option. Flagstar scores deposit beta 1, full-cycle credit 0, and capital 2; its broader pillars are A2/B0/C0/D2/E1, with severity 100% and Wmax zero, so it is an Avoid/Counterexample. The battery is not a blessing direction. It replaces an ill-targeted categorical penalty with harder class-specific evidence, and failure or unavailability has the same blocking force.

Remaining uncertainty is material. Summary scores are not the reserve triangles, deposit cohorts, or stressed capital calculations themselves. Until the exhibits are attached and independently challenged, the favorable battery direction is not demonstrated, Core remains a research label only, and admissible weight remains zero.

Question 9: Did the map comply with every one of the ten company-research prohibitions?

The ninth question required a compliance audit, not a general assurance. Each prohibition had to be linked to the map entry where violation was most tempting and to the mechanism that prevented it.

1. No buy call, target weight, entry order, or portfolio mutation. This was most at risk across every Core row, especially TSMC and ASML. Every row states non-buy and current admissible weight zero. Wmax is labeled diagnostic, never instructional.
2. No theme, total addressable market, model-class prior, reputation, price decline, or index membership may serve as expected-return proof. NVIDIA, Intel, and Rheinmetall are the sharpest tests. Their AI, rearmament, and harvest narratives are separated from cash capture; NVIDIA fails E, Intel is rejected, and Rheinmetall's state-payer gate remains open.
3. No return to the old additive business-model score without the five pillars, weakest-pillar constraint, EC, gates, reverse DCF, and dates. All 42 rows show A through E, EC, dated evidence, valuation sensitivity, and missing reverse-DCF work. No additive total is used.
4. No short-history challenger may enter Core. CoreWeave, Samsara, Symbotic, Nu, Sea, and BYD are all Speculative Options, assumed capable of total loss, with a 0.75% diagnostic per-name maximum rather than concentration eligibility.
5. Comparator-embedded exposure cannot be ignored. Microsoft, NVIDIA, Schneider, and Eaton make this most tempting. Comparator exposure is scored first, and active AI or power additions are barred while the original caps are exceeded.
6. A management-only claim cannot score above 1, conflicting evidence cannot be averaged away, and missing look-through cannot be treated as zero exposure. Symbotic, Adani, and Rheinmetall carry low pillar or EC scores, explicit missing evidence, and zero current weight. Backlogs and budgets receive no automatic cash credit.
7. Harvest-candidate status cannot justify normal size or roll forward after 24 months without exit and cash-return evidence. Intel is the direct test and fails F3-5; no full candidate is claimed.
8. State budgets are not collectible revenue until the payer gate clears. Rheinmetall, PG&E, and Intel remain constrained by missing contract, appropriation, rate-recovery, or state-support economics.
9. The map cannot make personalized tax, account, domicile, currency, or custody conclusions. DBS, TSMC, Adani, and the custody framework remain conditional because investor facts and professional review are missing.
10. Disputes and professional-review gates cannot be silently resolved. D-1, D-3, and U4-1 were especially exposed. The dispute appendix states each status, minority view, and observable. D-1 and U4-1 are expressly resolved for this policy; several others remain open or conditional.

This audit does not prove the company work is investment-grade. It demonstrates that the map obeyed its own research-output boundaries. Every named company is still non-buy, and every missing exhibit continues to block rather than receive the benefit of narrative doubt.

Question 10: Is the debate traceable from proposal to disposition and surviving disagreement?

The tenth question demanded a full traceability appendix. The answer is yes. Appendix A identifies every policy article P1 through P52 and every Fable amendment in the F1 through F4 ledgers by origin, final disposition, and surviving minority or qualification. This matters because later reframing can otherwise make an accepted proposal look as though it always had its final form.

Appendix B lists the dispute families D and U with current status and adjudicating observables. D-1 concentration is closed for this policy without a conviction tier, while the concentration minority and reopening requirements

remain visible. D-2 comparator modification has a resolved procedure but an open actual design. D-3 is split: lower gross risk is the crash default, verified ILB matching addresses real grind, and trend, commodities, or tail-premium strategies still require full-cycle net-carry and Coverage evidence. D-4 private assets and D-5 long-horizon currency hedging remain open. The U2, U3, and U4 rows preserve uncertainties about diversifier carry, AI overbuild, state-capex cash capture, custody size, speculation, economic lives, deterrence, muddle-through, assurance rents, quality-duration crowding, one-person execution, and procyclical diversifier evidence.

Appendix C contains all four Fable receipts. Each records model claude-fable-5, authentication by subscription_oauth_cli, a successful single call, prompt cache off, web search off, and trade authorization false. It also records prompt and response SHA-256 hashes for rounds 1 through 4. Those details demonstrate which artifacts entered the debate record. They are explicitly provenance rather than investment evidence and therefore do not appear as support for a company score.

Traceability does not eliminate disagreement. Its function is to prevent disagreement from disappearing through editing. A reader can see which ideas were adopted, which were later reframed, which minorities survive, and what observable could reopen or resolve an issue. The appendices do not certify that the adopted numbers are correct; all numeric bands remain contingent on the missing investor facts.

Question 11: Does the assembled policy survive attacks on bias, obsolescence, and implementation complexity?

The eleventh question turned the stress test against the policy itself. It covered hidden consensus, recency bias, survivorship bias, narrative fallacy, concentration, obsolescence, political and regulatory change, and implementation complexity. The answer is not that the apparatus is bias-free. It is that each vulnerability is named, paired with a control, and left with a residual risk.

Hidden consensus remains because the framework still assumes a core-plus-bounded-active architecture is worth building and tends to find quality franchises. The annual active-program viability test can reduce active work to zero, and balance-sheet financials, harvest research, and the quiet S6 plateau scenario preserve alternatives. Yet research can still rationalize activity.

Recency bias remains because the 2026 map emphasizes AI, power, obesity drugs, and defense. Six decade scenarios, the research sunset, company falsifiers, and normalization against older cycles mitigate it. They do not turn the current source set into a ten-year base-rate study. Survivorship bias remains because Core candidates are successful, well-disclosed firms. Twelve Avoid cases provide counterexamples, but the policy has no systematic delisted-company cohort and cannot claim predictive validation.

Narrative fallacy remains because decade theses are memorable. Pillar minimums, payer and cycle gates, per-share capture, reverse DCF, and hard falsifiers constrain the stories. Preliminary scores are still analyst judgment. Concentration remains because the comparator's mega-caps can breach AI limits even while direct active positions are capped around 2%. The forced amendment, annual risk-gap report, and no-active-add rule expose this mismatch, but raising a cap accepts rather than removes the embedded risk.

Obsolescence can outrun monitoring in AI hardware lives, software substitution, healthcare pipelines, and consumer relevance. Economic-depreciation tests and 36-month falsifiers help, yet evidence can arrive after value is lost. Political and regulatory discontinuities affect ratings, payments, semiconductors, defense, utilities, healthcare, banks, and emerging markets. Payer gates, binary severities, custody tiers, and jurisdiction or currency look-through reduce preventable errors, but state action cannot be diversified away completely.

The implementation verdict is adverse for the full apparatus. Forty-two dashboards multiplied by five pillars, factor scoring, valuation, tax, and custody work is an institutional workload. One person cannot maintain 42 live underwrites. The executable version allows at most eight active full underwrites, scores the comparator top 20

annually, checks survival monthly, reviews exceptions quarterly, and leaves everything else as annual-check research. Even eight may fail the hours test.

The one-person version therefore contains verified liabilities and FS; ILF plus nominal and real liability matching; one named comparator with an annual top-20 proxy; issuer and factor hard limits with a forced amendment mechanism; no more than eight live active underwrites; one decision journal and exception log; and an annual viability audit. If this cannot fit within roughly 40 recurring governance hours plus separately budgeted research time, active work is suspended. The one-page operating checklist intentionally omits deep company exhibits and legal or tax memoranda. That omission is safe only because anything needing those materials stays at weight zero until the work is completed.

Question 12: What, if anything, is certified or authorized?

Nothing is certified. Liability closure remains uncertified. Actual funded status, trait tolerance, maximum drawdown budget, tax basis, domicile and custody structure, actual comparator, factor look-through, and every company valuation remain unverified. Every numeric band is a dated policy hypothesis awaiting verified investor facts.

The existence of a source link proves only that a disclosure existed as of the evidence date. It does not independently verify a management claim, complete an evidence-coverage audit, or establish a current valuation. Likewise, a Core label, favorable pillar profile, or positive company fact does not constitute an entry judgment. Every named company is non-buy unless and until a current, complete evidence and valuation review clears all relevant gates in the investor's actual circumstances.

The policy authorizes no trade, order, transfer, hedge, account opening, custody change, message, target weight, or other portfolio mutation. The standing non-buy and owner-approval boundaries apply to every theme, model class, and named company. The correct outcome when a required fact is missing is zero admissible weight, not an assumption chosen for convenience.

What this chapter requires in practice

Treat the twelve answers as closure of a policy-design interrogation, not closure of an investment decision. Use the visible loss arithmetic and severity floors; do not infer a conviction tier. Record the failed and conditional grind results without presenting the passing diagnostic mix as an allocation. Complete the actual comparator study within the stated window while original caps and the ban on active AI or power additions remain operative. Keep all 11 Core names, every Watch name, every Speculative Option, and every Avoid case at zero unless a current full review independently changes that status under the governing gates. Attach the missing financial exhibits before treating favorable battery scores as demonstrated. Maintain traceability and the ten prohibitions. Limit live full underwrites to eight, apply the hours test, and suspend active work if the one-person version is not executable. Above all, preserve the final boundary: liability closure is uncertified, numeric bands await verified investor facts, no company is presently a buy, and this policy authorizes no transaction.

Chapter 12 - The Operating Checklist, Explained Step by Step

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.

What this chapter requires in practice

Use the original one-page checklist as the live control sheet and this chapter as its interpretation. Complete items 1 through 10 in order. Preserve dated evidence at every handoff, stop whenever a required fact, gate, budget, or capacity test is missing, and reduce complexity before relaxing a control. Keep no more than eight live full underwrites, and fewer whenever the printed hours budget requires it. Treat paper mode, freezes, suspension, and comparator plus ILF plus matching as valid governed outcomes. Above all, keep the non-authorization boundary visible: research, admission, diagnostic size, checklist completion, and certification status do not themselves authorize any trade, order, hedge, message, account opening, or portfolio change.

Chapter 13 - The Prohibition 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 summaries elsewhere still say at least ten years; that drafting inconsistency is residual ambiguity, not permission to choose the easier version. The final 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 this chapter requires in practice

Whenever a company-map row is read, copied, updated, or presented, preserve its class, A-E profile, Evidence Coverage, hard gates, valuation status, missing evidence, factor constraints, current admissible weight of zero, and non-buy language together. Never detach Wmax from its diagnostic-only meaning. Keep short-history and harvest cases inside their explicit experimental boundaries. Score management-only, conflicting, and missing evidence conservatively. Look through the comparator before calling anything a diversifier. Require contracts and realized cash conversion before crediting state support. Leave personal tax, legal, currency, account, and custody conclusions unresolved until verified facts and professional gates are complete. Finally, keep every dispute, minority, and adjudicating observable visible. Nothing in this audit authorizes an order, message, transaction, or portfolio change.

Chapter 14 - Final Bias and Implementation Stress Test

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.

What this chapter requires in practice

In practice, the owner must treat the final map as a dated and biased research queue, not as a portfolio. The recurring system is limited to verified liabilities and funded status, the investor liquidity floor and liability matching, one named comparator with annual top-20 look-through, issuer and factor limits with explicit amendment for persistent breach, no more than eight live full underwrites, one journal and exception log, and an annual viability audit. Roughly 40 recurring governance hours must be budgeted before active research, and each live underwrite adds about 15 to 25 hours; eight names imply roughly 160 to 240 hours before one-off work. If expected active contribution does not reach 2 times fully loaded annual cost, active capital stays at zero in paper-journal mode. If a deployed program records two consecutive negative viability years, new active work is suspended until re-justified. If the process cannot be maintained, scope is reduced rather than standards relaxed. All named companies and all omitted deep-work items remain at zero until separately reviewed and approved, and nothing in this chapter authorizes a transaction.

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.

What this chapter requires in practice

Read the document as a dated governance policy and research map, not as certification of an investor, portfolio, or security. Preserve the 22 August 2026 evidence cutoff. Treat Section 15's exact non-authorization and Appendix D's final boundary as controlling. Keep every named company at current admissible weight zero unless a separate, current process verifies investor facts, liability closure, comparator and factor look-through, EC and valuation work, relevant tax and legal matters, every policy gate, and implementation approval.

Use final adopted amendments rather than superseded formulations, but keep corrections visible through the existing traceability record. Do not turn research pointers into proof, diagnostic ceilings into targets, exceptions into amendments, challenge into execution authority, or explanation into new governance. If the required evidence, maintenance capacity, or approval is absent, the practical result is not a looser interpretation. It is continued non-authorization.

Appendices

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
P32	R3 Codex	economic depreciation	A

P	Origin	Subject	Final disposition / minority
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 with verified ≥10-year edge, liability surplus, independent challenge and operator capacity. Minority: rare control/edge cases warrant 10-40%.
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.