

Risk architecture Part 2. Scoring Logic for Archetype, Patterns, and Accelerators

June 2026

Series introduction

The risk architecture is built in three layers.

- 1) Layer 1. Risk points are calculated and assigned for each measure (maximum risk points per measure and assignment logic are defined in Risk Architecture Part 1. Scoring Logic for Individual Measures).
- 2) Layer 2. Risk points are calculated and assigned for measure combinations that are associated with financial trajectories, and that show up in an institution's data. Measure Combinations are subtly different and include:
 - a. **archetypes** - structural patterns that quietly shape an institution's long-term health. They're not fully captured by a single ratio, but rather by combinations of factors, such as revenue dependence, program mix, debt loads, or enrollment trends. So far, the research has identified nine archetypes that signal risk and three that signal resilience;
 - b. **patterns** - recur with some frequency across higher education segments, but not enough to be considered as archetypes. Think of them as emerging archetypes; and
 - c. **accelerators** - patterns that, if present in an organization's data profile, accelerate its financial decline. They include protracted enrollment decline, softness in two or more financial health ratios (e.g., primary reserves, viability, operating margins, etc.), and expenditure patterns suggestive of inefficient operations.

Max risk points per measure combination are defined in Risk Architecture Part 2. Scoring Logic for Archetypes, Patterns, and Accelerators

- 3) Layer 3. Overall risk level assignment (low, medium, high, critical) is based on a percent value calculated as the total risk points assigned to an institution / total available (maximum) risk points. Thresholds are defined in the risk scoring engine.js (analyzer), and are currently set as follows:

Assignment	Range
Low	0-25%
Moderate	26-48%
High	48-68%
Critical	>68%

This file provides details about Layer 2.

If Part 2 has an introductory section that taxonomizes archetypes into families (vulnerability / leverage / cost / brand / etc.), add A10 to whichever family description fits. I'd suggest a small framing addition:

Several archetypes (A1, A2, A5, A9) characterize institutions whose risk derives from a primary financial weakness — small scale, leverage, cost structure, squeeze. A10 belongs to a small but distinct family of archetypes that characterize institutions whose risk derives from a *masking dynamic* — where headline indicators look acceptable but the underlying structure is being supported by an unsustainable mechanism. As the framework matures, additional masking-pattern archetypes may be added (e.g., enrollment-masked operational stress for institutions where graduate enrollment growth conceals undergraduate decline).

I. Archetypes

A1 — High Vulnerability (Small Scale, High Dependency)

What it is: Small school, very tuition/appropriations-dependent, thin cushion.

- Base points: +4
- Trigger (all must be true):

1. Enrollment < 1,000
 2. Dependency (sector-aware):
 - Private: Tuition dependency $\geq 85\%$
 - Public: Appropriations dependency $\geq 70\%$
 3. Thin cushion (any one): $\text{UNA/FTE} < \$20,000$ OR Primary Reserve risk ≥ 5 OR Viability risk ≥ 5
- Inputs referenced: Current enrollment; tuition/appropriations dependency (sector-aware resolver); UNA per FTE (direct, total+enroll, or deep scan fallback); Primary Reserve risk points; Viability risk points.
 - Notes: Structural fragility flag combining small scale + high dependency + thin cushion. Note: the add() path uses PR/Viability ≥ 5 ; the checkA1 bridge uses ≥ 6 . The add() path is the primary detection route.

A2 — Leverage-Led Growth Bet

What it is: High leverage with weak/flat demand.

- Base points: +4
- Trigger:
 - Debt condition (either path): Debt/Revenue risk points ≥ 6 OR Debt ratio (alt) $\geq 20\%$
 - AND Demand condition (either window): 10-yr enrollment change $\leq 0\%$ OR 5-yr enrollment change $\leq 0\%$
- Inputs referenced: Debt-to-Revenue analyzer risk points (and/or alternative debt % field); Enrollment total change (10y) and UG change (5y).
- Notes: "Stagnation" = $\leq 0\%$ in at least one window. The alternative debt-percentage path exists in the add('A2', ...) block; the checkA2_... helper currently implements only the risk-points path. Documenting both paths here avoids code changes; if desired later, mirror the alt-percent path in checkA2_... for perfect symmetry.

A3 — Academic Sprawl Trap

What it is: Too many low-demand programs with no growth.

- Base points: +1
- The tool evaluates both the breadth of academic programs and the strength of demand signals such as enrollment, applications, and net tuition revenue. A prerequisite gate applies: the Program Distribution analyzer's report narrative must indicate "sprawl" (i.e., the institution's portfolio actually exhibits breadth/fragmentation, not concentration). If the narrative does not contain "sprawl," A3 will not fire regardless of risk points. If the gate passes, it fires if either of the following paths is satisfied:
 - A. Full (tiered) rule – Program breadth and weak demand:
 - Breadth indicator: Program Distribution Risk (sprawlRisk) from the program distribution analyzer.
 - Demand-softness signals:

- Five-year enrollment change $\leq 0\%$
 - Applications trend since 2014 $\leq 0\%$
 - Net tuition revenue (NTR) trend $\leq 0\%$
- Trigger:
 - If sprawlRisk ≥ 9 and any one of the demand-softness signals is true, or
 - If sprawlRisk ≥ 8 and two or more demand-softness signals are true.
- B. Simplified (check function) rule – Basic fallback:
 - sprawlRisk ≥ 7
 - Five-year enrollment change $\leq 0\%$
- Interpretation: An institution with broad program offerings but no evidence of growth or market strength is at risk of academic sprawl — a condition where fixed instructional costs rise while student demand fragments across too many low-volume programs.
- Notes: This documentation intentionally captures both paths now present in code. There is no enforced mutual exclusivity with A4 in the current logic, but the narrative-based sprawl gate provides de facto separation: A3 requires "sprawl" in the report narrative while A4 requires "monoculture," so both cannot fire on the same program distribution result.

A4 — Monoculture Dependency

What it is: One field dominates academic mix.

Scoring: + base points

How it triggers:

The analyzer looks for either a general concentration risk from the *Program Distribution Analyzer* or a large program-share imbalance combined with financial or enrollment strain. It fires if **either** of these pathways is met:

A. Program Distribution risk path:

Prerequisite: The Program Distribution analyzer's report narrative must indicate "monoculture" (i.e., the institution's portfolio actually exhibits concentration, not sprawl). If the narrative does not contain "monoculture," this path will not fire.

Program Distribution (pdRisk) ≥ 6 AND monoculture narrative gate passes

→ Fires on high concentration risk confirmed by narrative.

B. Concentration share path:

Largest program share (mShare) $\geq 50\%$

→ Fires on high concentration share alone (no additional fragility factors required).

Interpretation:

A “monoculture” institution depends too heavily on one dominant field — nursing, education, business, etc. — without diversification to buffer against shifts in labor-market demand or enrollment. The risk rises if that concentration coincides with weak enrollment trends or limited reserves, as downturns in a single field can quickly destabilize operations.

- Inputs referenced: Program Distribution riskPoints; PD concentration/“largest program share” field; enrollment 5-yr change; UNA per FTE; sector-aware dependency %.
- Notes: Path A requires the monoculture narrative gate plus pdRisk ≥ 6 . Path B fires on concentration share $\geq 50\%$ without additional fragility factors. The monoculture narrative gate on Path A provides de facto mutual exclusion with A3 (which requires a “sprawl” narrative).

A5 — High-Cost Architecture, Stalling Demand

What it is: Cost base too heavy for the demand profile.

- Base points: +3
- Trigger (all must be true):
Demand (either):
10-yr enrollment change $< +5\%$ OR 5-yr enrollment change $\leq 0\%$

Margins/Cost:

Path A — add('A5'): AOM risk ≥ 6 OR Operating Margin risk ≥ 6 (both sectors, no NTR gate)

Path B — checkA5 (sector-aware, stricter for publics):

Private: AOM risk ≥ 6 OR Operating Margin risk ≥ 6

Public: AOM risk ≥ 7 OR Operating Margin risk ≥ 7 AND Net tuition per FTE trend ≤ 0

- Inputs referenced: Enrollment 10-yr total change, 5-yr UG (or GD) change; Adjusted Operating Margin risk points; Operating Margin risk points; Net tuition per FTE trend (publics gate).
- Notes: The add() path applies ≥ 6 uniformly without sector distinction. The checkA5 bridge applies stricter publics logic: ≥ 7 plus a non-positive net-tuition-trend gate.

A6 — Brand-Weakening Channel

What it is: Selectivity deteriorating without offsets.

- Base points: +2
- Trigger:
 - Core signal: Acceptance-Yield (AY) risk ≥ 4
 - Additional behavior in add('A6'): will not fire if a “strong now” pipeline is present (accept rate $\leq 60\%$ or yield $\geq 30\%$) or if A8 would fire (i.e., a combination of strong brand signals and thin margins per its stricter test).
- What actually happens (union of code paths):

- Because the checkA6_... helper fires on AY risk ≥ 4 without the exclusions, A6 can still trigger even when the “strong now” or “A8 would fire” conditions are present.
- Inputs referenced: AY analyzer risk points; acceptance rate; yield rate; AOM/OM risk, NTR trend, 3-yr AOM (used only in the add('A6') mutual-exclusion logic).
- Notes: A6 and A8 are intended to be mutually exclusive in add('A6'), but the helper’s simpler rule means A6 may still appear alongside A8 in practice.

A7 — High-Pell “Mission Pressure Cooker”

What it is: High need, retention & cushion stress.

- Base points: +3
- Trigger (fires if any tier below is satisfied):
 - Tier 1 (extreme Pell override):
 - Pell share $\geq 60\%$ AND Retention $< 78\%$
 - Tier 2 (high Pell + medium cushion):
 - Pell share $\geq 50\%$ AND Retention $< 78\%$ AND (any of: UNA/FTE $< \$35,000$ OR Primary Reserve risk ≥ 3 OR Viability risk ≥ 2)
 - Tier 3 (original/strong cushion test):
 - Pell share $\geq 45\%$ AND Retention $< 80\%$ AND (any of: UNA/FTE $< \$25,000$ OR Primary Reserve risk ≥ 4)
 - Additional broader path from add('A7'):
 - Pell share $\geq 45\%$ AND Retention $< 82\%$ AND (UNA/FTE $< \$25,000$ OR Primary Reserve risk ≥ 5)
- Inputs referenced: Pell share; first-year retention; UNA per FTE; Primary Reserve risk points; Viability risk points (as an alternate cushion in Tier 2).
- Notes: The helper adds tiering and a $<80\%$ retention cut for common cases; the add('A7') path is slightly broader (allows $<82\%$ retention but expects a tighter cushion: PR ≥ 5 or UNA/FTE $< 25k$). Documenting all tiers keeps the spec faithful to current behavior without changing code.

A8 — Brand-Strong but Margin-Thin

What it is: Strong pipeline, weak economics.

- Base points: +2
- Trigger (fires if BOTH brand-strong AND margin-thin, with A6 exclusion in add):
 - Path A — add('A8') (segment-aware, stricter):
 - Brand-strong:
 - Determine *research segment* = Carnegie {15,16,17,18}.
 - Signals: (1) AY risk ≤ 1 (selective), (2) accept rate $\leq 60\%$ ($\leq 68\%$ if non-research), (3) yield $\geq 28\%$.
 - Fire brand side if research: ≥ 2 of 3; non-research: ≥ 1 of 3.
 - Margin-thin:

- Either one very high margin risk: $AOM \geq 7$ or $OM \geq 7$, OR
- Any two of: $AOM \geq 6$, $OM \geq 6$, $NTR \text{ trend} \leq 0$, $AOM \text{ 3-yr avg} \leq 1.5\%$.
- Mutual exclusion with A6: block A8 if A6 would fire (i.e., $AY \text{ risk} \geq 4$ and not (accept $\leq 60\%$ or yield $\geq 30\%$)).
Path B — checkA8_... (looser helper):
- Brand-strong: $AY \text{ risk} \leq 1$.
- Margin-thin: $AOM \geq 5$ or $OM \geq 5$ or $NTR \text{ (5y/recent) change} \leq 0$.
- Inputs: AY risk; acceptance %; yield %; AOM/OM risk; NTR trend or 5y/recent change; AOM 3-yr %.
- Notes: Because the helper is looser and doesn't enforce A6 exclusion, A8 can co-appear with A6 in practice. Documentation reflects both paths.

A9 — The Squeezed Middle

What it is: Middling margins, some leverage, small decline.

- Base points: +2
- Trigger (fires if ALL families below pass):
Demand softness (either window):
 - 10-yr enrollment change in $[-20\%, +2\%]$ OR recent/5-yr change in $[-10\%, +2\%]$.
Margins (modest risk):
 - Take $\max(AOM \text{ risk}, OM \text{ risk})$ and require it in $[0, 4]$.
Leverage/solvency (any one qualifies):
 - Debt/Revenue risk in $[1, 6]$ OR Primary Reserve risk in $[2, 6]$ OR Viability risk in $[1, 4]$.
(Narrower add('A9') variant also exists):
 - $AOM \text{ risk} \leq 2$ AND Debt/Revenue risk $\in [3,5]$ AND 10-yr enrollment $\in (-15\%, 0\%)$.
- Inputs: Enrollment 10-yr and recent/5-yr change; AOM/OM risk; Debt/Revenue risk; Primary Reserve risk; Viability risk.
- Notes: The helper broadens the original "moderate pressure" idea to accept multiple leverage indicators, while keeping margins only *modestly* stressed.

A10 — Endowment-Masked Operating Deficit

What it captures

A structurally distinct failure mode for endowment-supported institutions: operations are running a meaningful deficit (visible only after stripping investment-return revenue), but the deficit is being absorbed by an unsustainable endowment draw against an asset base where most net assets are donor-restricted and therefore cannot support operations or service debt indefinitely.

The pattern is dangerous precisely because the component signals can each look explicable in isolation. A 5–6% endowment draw is "high but defensible." A negative structural margin is "explained by one-time investments." High donor restrictions are "evidence of fundraising success." It is the *compounding* of all four — operating deficit AND elevated draw AND high dependency AND restricted asset base — that signals a structural problem the institution cannot easily exit.

Without A10, the framework correctly flags each component but does not synthesize them into the systemic finding. With A10, the synthesis happens explicitly and produces an interpretable narrative for the case investigator.

Trigger conditions

Fires when **all four** of the following hold:

Input	Threshold Captures
-------	--------------------

Input	Threshold	Captures
Structural Operating Margin riskPts	≥ 4	Meaningful operating deficit after stripping investment returns
Endowment Draw Sustainability riskPts	≥ 4	Draw at or above NACUBO sustainable ceiling
Endowment Dependency Ratio riskPts	≥ 2	Operations meaningfully dependent on endowment payout (≥ 8%)
Net Asset Restriction Profile riskPts	≥ 3	Donor-restricted share elevated, cushion partially illusory

Single-signal noise cannot trigger A10. The four-input AND-gate enforces that the systemic pattern is genuinely present, not the artifact of any one analyzer's threshold sensitivity.

Points awarded

+6 risk points when triggered. This sits in the upper half of the archetype point distribution (A1/A2 at 4, A3/A4 at 1, A5/A7 at 3, A6/A8/A9 at 2) — reflecting that A10 is a high-specificity, low-false-positive synthesis finding. When it fires, the supporting case is unambiguous and the implications are significant.

What it tells the analyst

A10 firing means: *"The institution's headline balance sheet strength is masking a structural operating problem the endowment cannot indefinitely absorb."* This is a different finding than any of the input analyzers produce in isolation, and it carries a different remediation conversation:

- Component analyzers each suggest local interventions (raise draw, reduce expenses, fundraise more).
- A10 reframes the problem as: *the local interventions have already been deployed, and the combination is unsustainable.*

The implied conversation for institutions firing A10 is structural — spending policy reform, restricted-fund repurposing efforts, fundamental cost-base reset, or strategic combinations — rather than tactical adjustment.

Interaction with other archetypes

A10 is designed to be complementary, not redundant, with existing archetypes:

- **A1 (Small-Scale High Dependency)** captures scale and tuition concentration; A10 captures endowment-side fragility. The two can co-fire on small-to-mid liberal arts institutions where both patterns exist.
- **A2 (Leverage-Led Growth)** captures debt-financed expansion; A10 captures operating coverage of *existing* debt. Different time horizon, different remediation.
- **A9 (Squeezed Middle)** is broader and less specific; when A10 fires on an institution that would otherwise trigger A9, A10 is the more informative finding.

Anchor example — Brandeis University FY24

All four gates clear cleanly:

- Structural Operating Margin: 6/8 (Critical) — current -13.3%, 3-yr avg -11.0%
- Endowment Draw Sustainability: 4/10 (Moderate) — 5.23% current, 5.20% three-year avg
- Endowment Dependency Ratio: 4/6 (High) — 13.77% of operating revenue
- Net Asset Restriction Profile: 5/6 (Critical) — 86.1% donor-restricted

A10 fires with +6 archetype points. The combined story aligns with the FY24 audit's own characterization of the institution's position: operating cash burn worsened materially year-over-year, the spending policy draw rate rose toward 6%+, and the unrestricted asset base remained thin relative to long-term obligations. No single component analyzer alone would compel a structural conversation; A10's synthesis does.

R10 — Resilient by Scale and Diversification

What it is: Big scale, diversified revenue, solid cushion.

- Base points: -3
- Trigger (ALL must be true):
Path A — add('R10') (stricter):
 - Enrollment $\geq 12,000$
 - Revenue dependency $\leq 55\%$ (*sector-aware: tuition for privates; appropriations for publics via depAnyPct*)

- Cushion: $\text{UNA/FTE} \geq \$45,000$ OR (Primary Reserve risk ≤ 2 AND Viability risk ≤ 2)
Path B — checkR10_... (looser helper):
- Enrollment $\geq 10,000$
- Revenue dependency $\leq 50\%$ (*sector-aware: pulls tuition/appropriations separately*)
- Cushion: $\text{UNA/FTE} \geq \$40,000$ OR (Primary Reserve risk ≤ 2 AND Viability risk ≤ 2)
- Inputs: Current enrollment; sector-aware revenue dependency; UNA per FTE; Primary Reserve risk; Viability risk.
- Notes: Expect minor threshold discrepancies across paths (12k/55%/45k vs 10k/50%/40k). Both require scale + diversification + cushion.

R11 — Research-Strong Anchors

What it is: Large, research-anchored stability.

- Base points: -4
- Trigger (ALL required):
Path A — add('R11') (Carnegie or “soft” research profile):
 - Research orientation:
 - Either Carnegie in research set {public: 15,17,18; private: 15,16,17,18}, or research analyzer applicable and
 - Scale: Enrollment $\geq 12,000$ (public) / $\geq 15,000$ (private), and
 - Brand: (AY risk ≤ 2 or accept rate $\leq 60\%$ or yield $\geq 30\%$)
 - Research performance: any research metric risk
(research_grants/research_grants_expenses/grantfunding) ≤ 2
Path B — checkR11_... (helper):
 - isResearchOriented() true (*uses your internal helper*), and
 - At least one research metric present, with risk ≤ 2
- Inputs: Enrollment; AY risk / accept % / yield % (brand check, add-only); Carnegie/segment id or isResearchOriented(); research risk from grants/expenses/funding analyzers.
- Notes: The helper doesn't re-check scale/brand; the add path does. Both require low research risk (≤ 2).

R12 — Brand-Strong Anchors

What it is: Strong brand with healthy cushion/margins.

- Base points: -2
- Trigger (check-only; no add block):
 - Brand: Acceptance-Yield (AY) risk ≤ 2
 - Scale: Enrollment $\geq 10,000$

- Inputs: AY risk; current enrollment.
- Notes: Pure “brand + scale” resilience flag. Since there’s no add('R12'), this helper defines R12’s behavior end-to-end.

LA1 — Aid Arms-Race / Own-Source Pressure (*Limited-audience flag*)

What it is: NTR below peers + weak demand/tuition signal.

- Base points: +2 (*per config*)
- Trigger (peer-aware; falls back if peers missing):
 - Compute NTR ratio = NTR_now / NTR_peer when peers exist.
 - If peers exist:
 - Public: require NTR ratio ≤ 1.00 AND ≥ 2 weak signals.
 - Private: require NTR ratio ≤ 0.95 AND ≥ 1 weak signal.
 - If peers missing: fire on weak signals alone with ≥ 2 weak signals (both sectors).
 - Weak signals:
 - Applications trend since 2014 $\leq 0\%$
 - 5-yr enrollment change $\leq 0\%$
 - NTR trend $\leq 0\%$
- Inputs: Current and peer NTR; apps since 2014; 5-yr enrollment change; NTR trend; sector.
- Notes: Captures price/discount pressure relative to peers plus weakening demand signals; peerless mode is deliberately conservative.

LA2 — Small-Scale Fragility (*Limited-audience flag*)

What it is: Small, thin cushion, high tuition dependence.

- Base points: +2 (*per config*)
- Trigger (all must be true):
 1. Small scale: Enrollment $< 2,000$
 2. Thin cushion (any): $UNA/FTE < \$30,000$ OR Primary Reserve risk ≥ 4 OR Viability risk ≥ 4
 3. High tuition dependence: Tuition dependency $\geq 75\%$ (*sector-aware: depAnyPct/depPct*)
- Inputs: Current enrollment; UNA per FTE; Primary Reserve risk; Viability risk; tuition dependency (resolver handles depAny/depPct).
- Notes: This is the “lighter” cousin of A1—larger size cut, slightly softer cushion threshold, and a fixed high-tuition-dependency gate.

II. Pattern Scoring Maps

Patterns detect *multi-dimensional correlations*—combinations of metrics that, together, amplify institutional fragility beyond what the individual analyzers capture.

Each adds the indicated number of bonus risk points to the total composite (up to 50 possible from all pattern layers combined).

P1 — “Siena” Price-Pressure (Privates only)

What it is: Tuition/discount moves not translating to NTR.

- Base points: varies by variant
- Sector: Privates only (sector = 2)
- Trigger (two variants exist in code; either can fire):
Variant A (stronger price-pressure):
 - Tuition dependency $\geq 75\%$ AND (Retention $< 75\%$ OR AY risk ≥ 4)
 - Points: 10 (+2 kicker if Tuition dependency $> 85\%$)Variant B (demand-decline pairing):
 - Tuition dependency $\geq 70\%$ AND (Applications since 2014 $\leq -10\%$ OR 5-yr enrollment change $\leq -8\%$)
 - Points: 4 (+2 kicker if Apps $\leq -20\%$ OR 5-yr enrollment $\leq -15\%$)
- Inputs: Tuition dependency; Retention; AY risk; Apps since 2014; 5-yr enrollment change
- Notes: Two calculateP1_* functions exist with different thresholds/points. Due to JavaScript class semantics, the second definition (Variant B) overwrites the first (Variant A), so only Variant B is currently active in production. Variant A (tuition dep $\geq 75\%$ + retention/brand test $\rightarrow 10$ pts) is dead code. Consolidation is recommended.

P2 — High Discount + Flat Enrollment (Privates only)

What it is: Buying volume, not gaining demand

- Base points: 8 (+2 if very low NTR)
- Sector: Privates only
- Trigger:
 - Have NTR normalization (ntrNorm) AND 5-yr enrollment
 - $ntrNorm \leq 0.85$ AND $-5\% \leq 5\text{-yr enrollment change} \leq +5\%$
 - Kicker: if $ntrNorm \leq 0.80 \rightarrow +2$ more (total 10)
- Fallback path (no ntrNorm): If NetTuition analyzer riskPoints $\geq 4 \rightarrow 4$ points
- Inputs: ntrNorm (your peer-normalized NTR/FTE); 5-yr enrollment change; NetTuition risk (fallback)

P3 — Publics: Appropriations Dependence + Enrollment Decline

What it is: Public-funding dependency with demand erosion.

- Base points: 10 (+2 severity kicker)
- Sector: Publics only
- Trigger:
 - Appropriations dependency $\geq 40\%$ (or $\geq 35\%$ when *severity.hasSeverity* is set) OR Tuition dependency $\leq 20\%$
 - AND 5-yr enrollment change $\leq -15\%$
 - Kicker: +2 if Appropriations dependency $\geq 50\%$ OR 5-yr enrollment $\leq -25\%$
- Inputs: Appropriations dependency (preferred); Tuition dependency (alt); 5-yr enrollment change
- Notes: The “tuition dep $\leq 20\%$ ” branch functions as a proxy for high state support when approps metric is missing.

P4 — Net Tuition Risk (peer-normalized)

What it is: Weak pricing power vs peers.

- Base points: 5 (base) or 7 (severity mode)
- Trigger (primary path with *ntrNorm*):
 - Base hit: $ntrNorm \leq 0.80 \rightarrow 5$ points
 - Severity hit (when *severity.hasSeverity* is true): $ntrNorm \leq 0.90$ AND at least one paired stressor:
 - 5-yr enrollment $\leq -15\%$ OR Apps since 2014 $\leq -10\%$ OR Acceptance $\geq 90\%$ OR Yield $\leq 10\%$
 - $\rightarrow 7$ points
- Fallback (no *ntrNorm*): $NetTuition\ riskPoints \geq 4 \rightarrow 4$ points
- Inputs: *ntrNorm*; 5-yr enrollment; Apps since 2014; Acceptance %; Yield %; *NetTuition risk* (fallback)
- Notes: Use *severity.hasSeverity* to enable the 0.90 “warning” band + stressor.

P5 — UNA Cushion Risk (Liquidity + Earnings/Reserves corroboration)

What it is: Thin liquidity/long-term cushion.

- Base points: 6 (+2 if very low UNA)
- Trigger (preferred with *unaNorm*):
 - $unaNorm \leq 0.60$ AND (AOM 3-yr avg $\leq -2.0\%$ (public) | $\leq -3.0\%$ (private) OR Primary Reserve ratio $< 20\%$ (public) | $< 25\%$ (private))
 - Kicker: if $unaNorm \leq 0.40 \rightarrow +2$ (total 8)
- Fallback (no *unaNorm*): $UNA\ analyzer\ riskPoints \geq 3 \rightarrow 4$ points
- Inputs: *unaNorm*; AOM 3-yr avg; Primary Reserve ratio; *UNA risk* (fallback)

- Notes: Sector-aware thresholds for both AOM-3yr and PR ratio.

P6 — Research Coverage Erosion (scale-gated; research data present)

What it is: IDC coverage falling vs cost growth. Base points:

- Severity 2: 4 (publics) / 3 (privates)
- Severity 1: 2
- Scale gate: Enrollment $\geq 8,000$ (publics) / $\geq 6,000$ (privates)
- Primary signals (any contribute to severity):
 - Coverage trend ≤ -3 pp (5y)
 - Spread (grants trend – expense trend) ≤ -3 pp
 - Low current coverage ($< 80\%$) with rising costs
→ 1 signal = Severity 1, 2+ signals = Severity 2
- Fallback path (if coverage signals absent): Grant stress via analyzer point ratio:
 - Publics: grant risk / max ≥ 0.60 ; Privates: ≥ 0.70 (or infer with small max scales)
 - Boost to Severity 2 if costs rising or grants non-growing
- Private dampener: if private and Enrollment $< 10,000$, reduce computed severity by 1 (min 0)
- Inputs: Enrollment; coverage ratio now & 5y change; grant/expense 5y trends; grant funding risk & max
- Notes: Returns no points if severity falls to 0 after dampener.

P7 — Brand Positioning (pipeline deterioration)

What it is: Weak selectivity/yield signal.

- Base points: 6 (+2 if very adverse)
- Trigger:
 - Applications since 2014 $\leq -20\%$ AND (Acceptance $\geq 85\%$ OR Yield $\leq 12\%$)
 - Kicker: +2 if Acceptance $\geq 90\%$ OR Yield $\leq 10\%$
- Inputs: Apps since 2014; Acceptance %; Yield %
- Notes: Captures a sharp demand drop paired with poor selectivity/yield.

P8 — Market Position (multi-signal deterioration)

What it is: Losing share in served market. Base points: 5 (or 7 if 3+ signals)

- Trigger: At least 2 of the following adverse signals:
 - Acceptance $\geq 85\%$
 - Yield $\leq 12\%$
 - Applications since 2014 $\leq -10\%$

- Market share change $\leq -5\%$ (if available)
- Points: 7 if ≥ 3 signals, else 5
- Inputs: Acceptance %; Yield %; Apps since 2014; Market share change (optional)
- Notes: Lightweight add-on—keeps weights modest by design

P9 – Endowment Drain Risk

What it is: Sustained endowment draws above the NACUBO sustainable ceiling against an endowment base that is thin relative to Carnegie peers, in institutions running structural operating deficits. Identifies institutions that are quietly depleting modest endowments to fund recurrent operating shortfalls they cannot cover from core revenues.

- **Base points:** +4
- **Sector:** Private only (sector = 2). Endowment draw sustainability data not calibrated for public sector draw patterns; public endowment-dependent institutions are addressed through other pattern and archetype flags.
- **Trigger:** All three gates must hold simultaneously:
 - Gate 1 — Structural Operating Margin = Critical. Confirms the draw is structural (institution cannot cover operations from core revenues without endowment income) rather than discretionary spending policy from a position of strength.
 - Gate 2 — Three-year average endowment draw rate $> 5.5\%$. The NACUBO sustainable spending ceiling is 4.5–5.0%. The 5.5% threshold adds a 0.5 percentage point buffer to focus on institutions with a multi-year sustained pattern above the ceiling, excluding borderline single-year exceedances. Draw rates above 50% are treated as data quality artifacts (near-zero endowment denominator) and excluded from consideration.
 - Gate 3 — Endowment per FTE in the bottom quartile of Carnegie peers. Threshold set at $0.480 \times$ peer median (P25 of the endowment/FTE-to-peer-median ratio distribution across private institutions, empirically derived from sector-wide calibration run). Institutions above this threshold — including well-endowed institutions drawing generously from per-FTE bases above peer median — are excluded. The gate ensures the pattern fires only when the draw is against a genuinely thin base, not as a feature of a sustainable endowment spending policy.
- **Mutual exclusion with A10:** If Archetype A10 (Endowment-Masked Operating Deficit) has already triggered for the institution, P9 returns 0 points and defers entirely to A10. A10 is the more specific, higher-severity version of the underlying condition (+6 points, additionally requires elevated asset restriction profile). When both conditions are technically present, A10 takes precedence. This prevents double-counting of a shared underlying mechanism.
- **Inputs:** Structural Operating Margin risk level (from unified_structural_margin_analyzer); three-year average endowment draw rate (from unified_endowment_draw_sustainability_analyzer, field: three_year_avg, expressed as decimal); endowment per FTE institution value and Carnegie peer median (from unified_endowment_analyzer, fields: rawData.perFTEEndowment.currentValue and rawData.peerContext.peerMedian)
- **Relationship to A10:** P9 and A10 detect related but distinct conditions. A10 requires four inputs including high asset restriction — it targets the "fragile wealth" configuration where the endowment looks large but most assets are donor-locked. P9 requires only three inputs with no restriction gate — it targets the broader population of institutions running structural deficits against thin peer-relative endowment bases regardless of restriction profile. In the sector-wide private calibration run (890 institutions), P9 fires on 34 institutions; 12 of those also trigger A10. The 22 institutions triggering P9 but not A10 represent the net new catch — same underlying behavior without the restriction overlay.
- **Calibration data:** Sector-wide private calibration run conducted June 2026 across 890 private not-for-profit institutions (Carnegie C15–C22). Bottom quartile threshold ($0.480 \times$ peer median) derived empirically from that distribution. Draw rate floor (5.5%) and artifact cap (50%) set based on NACUBO guidance plus sector distribution review. Results validated against known cases: Woodbury University (three-gate pass, deferred to A10 ✓); Thomas Aquinas College (three-gate pass, net new catch, 14.05% draw, \$47,657/FTE vs \$158,058 peer median ✓); St Olaf College

(correctly excluded — 4.08% draw below 5.5% threshold, above-median EFT ✓); Harvard University (correctly excluded — draw below threshold, EFT far above peer median ✓).

- **Notes:** Implementation added to `risk_adjustment_analyzer.js` as `calculateP9_EndowmentDrainRisk()` method, registered in `calculatePatternAddons()` alongside P8. The three-year average draw rate is stored as a decimal in the draw sustainability analyzer (e.g., 0.1931 for 19.31%); all threshold comparisons in the code use decimal values. The 0.480 bottom quartile ratio is a sector-wide empirical value; it should be recalibrated if the tool is extended to additional sectors or if significant changes are made to the endowment/FTE analyzer peer group construction.

III. Accelerators

The presence of accelerators typically speeds the pace of financial decline.

Implementation note: In the current codebase, all three accelerators (Operational Inefficiency, Enrollment Cliff, and Financial Distress) are evaluated inside the archetype detection pipeline (`detectArchetypeFlags`) rather than in a separate accelerators method. The `detectAccelerators()` method exists but contains only placeholder stubs. This is a structural detail; the scoring logic is fully implemented and active.

Operational Inefficiency

What it is: Under-filled sections, low utilization, high admin share.

Base Points: +3

Logic (fires if either condition set is met):

A. Cost-mix signals (need ≥ 2 of 3)

- Low instructional share: below peer mean by a pad (-4 pp general; -7 pp for R1/R2) or, if no peers, $<35\%$ (private) / $<40\%$ (public).
- High administrative share: above peer mean by $+4$ pp ($+6$ pp for R1/R2) or, if no peers, $>26\%$ (private) / $>27\%$ (public).
- High student-services share: above peer mean by $+4$ pp ($+6$ pp for R1/R2) or, if no peers, $>12\%$ (private) / $>13\%$ (public).

...with analyzer risk-point fallbacks treated as “strong” (≥ 5 ; ≥ 6 for R1/R2) if the percent isn’t available.

B. Delivery/throughput signals (need ≥ 1)

- Student-faculty ratio risk: S/F RP ≥ 2 (everywhere).
- Low retention: $<78\%$ (general) / $<76\%$ (research segments).

Decision rule: Trigger if (≥ 2 cost-mix signals) OR (1 cost-mix signal AND both delivery signals). Returns +3 with reasons list.

Note: Two code paths exist. The freestanding `detectArchetypes()` function uses the looser rule (≥ 1 cost + ≥ 1 delivery). The class-based `checkOperationalEfficiency()` method uses the tighter rule above (1 cost requires both delivery signals). The class method is the primary detection route. Notes: Peer-aware whenever peer means are present; otherwise sector floors apply.

Enrollment Cliff

What it is: A severe enrollment decline paired with market position loss or extreme contraction.

Base Points: +10

Trigger (all must be true):

5-year enrollment change $\leq -10\%$

AND either: market share change $< -10\%$, OR enrollment change $\leq -30\%$

Where it's also used: As an explicit signal inside Financial Distress, the code uses the fixed -10% 5-year gate as one of six signals ("enrollment $\downarrow 5yr \geq 10\%$ " in the reason string).

Related scoring map (context): The standalone Enrollment analyzer (feeding this tool) scores long- and short-run performance versus peers, with absolute fallbacks (e.g., $\leq -20\% \rightarrow +4$). This is informative to readers of the risk packet even though the "cliff" flag used here is the fixed -10% gate.

Financial Health — “Financial Distress” Pattern

What it is: Multiple ratio risks “lit” simultaneously.

Base Points: +12

Signals considered (any may contribute):

- Primary reserve weakness (low reserve ratio vs rule set).
- Viability weakness (viability ratio below sector/research cut).
- Thin UNA/FTE: UNA index < 0.60 .
- Weak margin (3-yr AOM): $\leq -2.0\%$ public / $\leq -5.0\%$ private.
- Enrollment decline: $\leq -10\%$ over 5 years.
- High revenue dependency: $\geq 50\%$ appropriations (public) or $\geq 70\%$ tuition (private).

Decision rule:

- Research-oriented segments: need ≥ 4 total signals, ≥ 2 “hard” signals (reserve/viability/margin), and either the enrollment-cliff signal or a very weak margin ($\leq -3\%$).
- Non-research segments: ≥ 3 total signals with ≥ 1 “hard” signal.

When triggered, it awards +12 points and enumerates which signals fired (e.g., “low primary reserves; weak viability; ... enrollment $\downarrow 5yr \geq 10\%$ ”).

Notes: This is a composite “crisis” detector rather than a single-metric score; it purposely mixes solvency, margin, demand, and dependency to avoid false positives from any one analyzer. (For broader enrollment context see the Enrollment scoring map referenced above.)

IV. Critical Mass Escalation Rule

When an institution is rated ‘Critical’ on three or more individual measures, and those critical ratings span at least two of the four risk domains evaluated by the tool (demand, revenue, margin, and balance sheet), the overall risk level is escalated by one tier. For example, an institution that would otherwise be rated ‘High’ would be escalated to ‘Critical.’

The rule includes guardrails to prevent false positives: at least one of the critical ratings must come from the demand domain (enrollment, acceptance/yield, or retention) and at least one must come from a financial domain (revenue, margin, or balance sheet). Purely outcome-based metrics (graduation rate, transfer-out rate) and mission-linked metrics (Pell share) are excluded from the escalation calculus.

A measure is considered 'Critical' for escalation purposes if it has an explicit 'Critical' risk level or if its risk points reach 80% or more of the maximum available for that measure.

The four risk domains and their constituent analyzers are:

Demand: Enrollment, Acceptance/Yield, Retention

Revenue: Tuition Dependency, Net Tuition per FTE, Appropriations Dependency, Grant Funding

Margin: Operating Margin, Adjusted Operating Margin, Net Income Ratio

Balance Sheet: Primary Reserve, Viability, Unrestricted Net Assets, CFI, Capitalization Ratio, Debt/Revenue, Return on Net Assets

Excluded from the escalation calculus: 150% Graduation Rate, Transfer-Out Rate, Pell Share, Student-Facing Expenditure, Instructional Expenditure, Institutional Expenditure, Student Services Expenditure, Salaries & Benefits, Student-Faculty Ratio, Research Grants, Program Distribution, Programs Sustainability, Federal Policy.

This rule does not add risk points. It adjusts the final risk level assignment to reflect the compounding effect of simultaneous critical weaknesses across multiple institutional dimensions.”