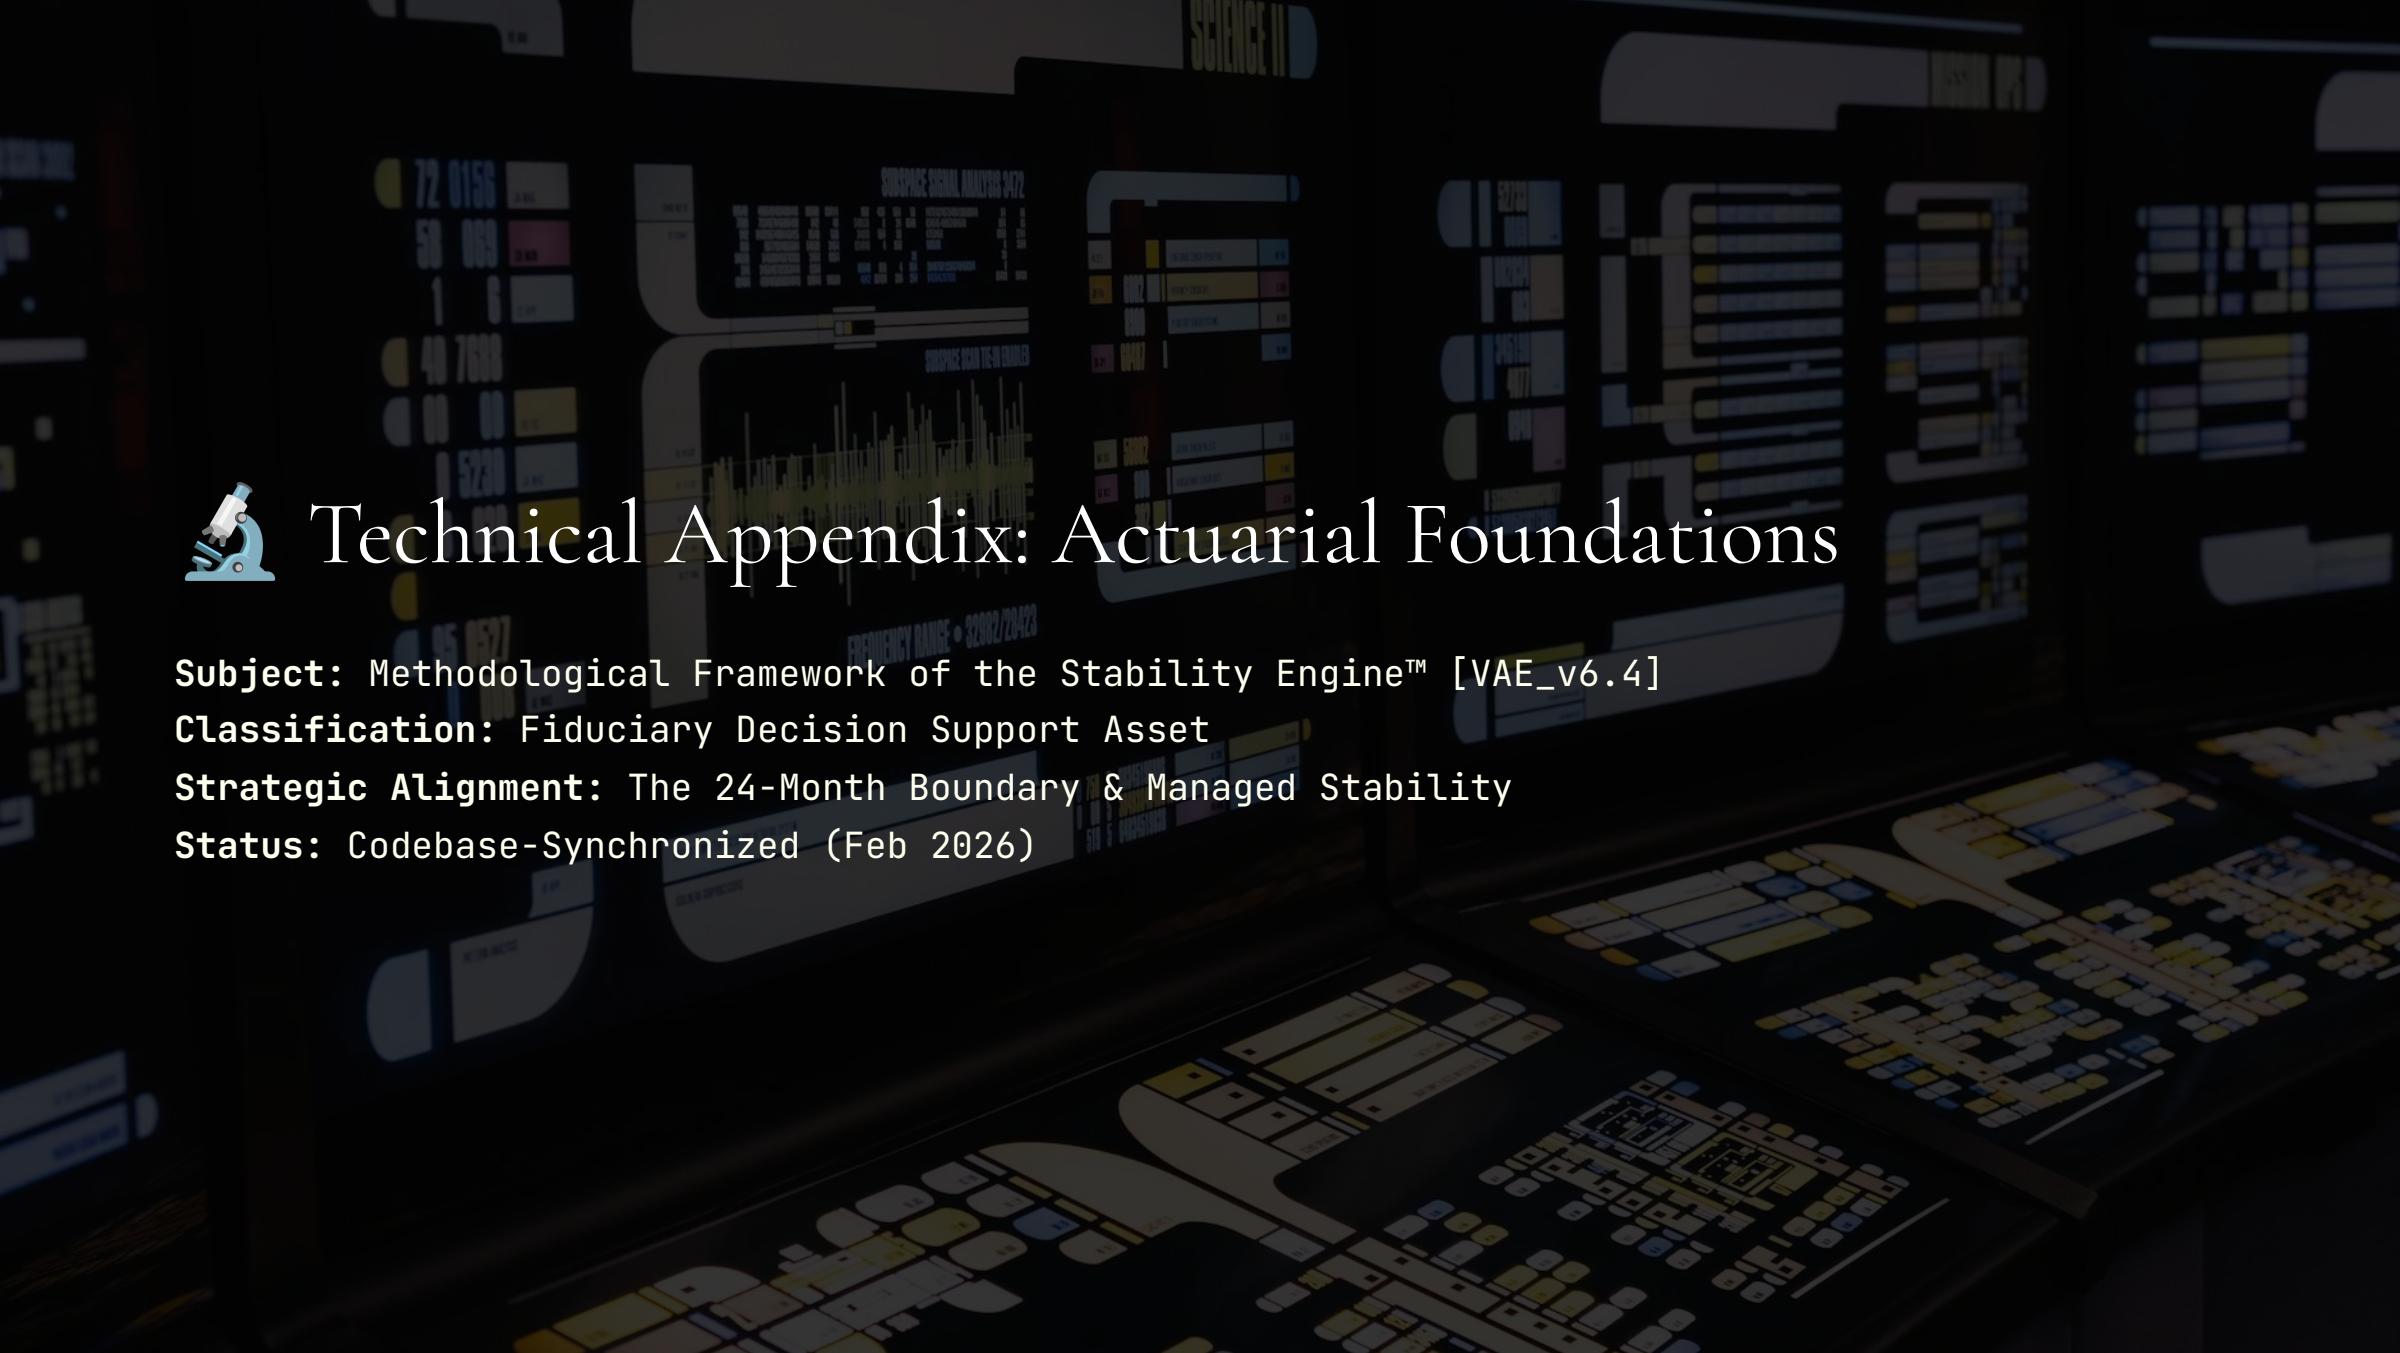

🔬 Technical Appendix: Actuarial Foundations

Subject: Methodological Framework of the Stability Engine™ [VAE_v6.4]
Classification: Fiduciary Decision Support Asset
Strategic Alignment: The 24-Month Boundary & Managed Stability
Status: Codebase-Synchronized (Feb 2026)

Executive Summary: The 24-Month Boundary

The Stability Engine™ uses multi-dimensional survival analysis and information theory to quantify **Intrinsic Flight Risk** in leadership trajectories. Our framework identifies a definitive temporal boundary in the predictive validity of career history:

- **Phase I: Intrinsic Risk (0–12 Months):** Audits of "Career Physics" (Entropy, Velocity, Contagion) yield **90% precision** in identifying historical "Quick-Quits" within a 12-month horizon.
- **Phase II: Managed Stability (12–36+ Months):** While the engine generates 36-month survival curves, accuracy beyond Month 12 is context-dependent. Long-term retention is reclassified as an **Organizational Outcome**, where the engine identifies the "Management Premium" required to extend tenure beyond the inherent 24-month trajectory.

Audit I: Information Theory & Trajectory Entropy

We apply Shannon Information Theory to career paths to distinguish coherent specialization from chaotic mobility.

Shannon Entropy (H)

The engine calculates the informational uncertainty of a candidate's role transitions. A path characterized by frequent, non-linear shifts generates high entropy (H>2.0), signaling a high probability of early-stage misalignment.

$$H(X) = -\sum_{i=1}^{n} P(x_i) \log_2 P(x_i)$$

Industry-Contextualized Benchmarks

Unlike static models, the Stability Engine utilizes industry-specific entropy thresholds to account for varying sector norms. These thresholds are periodically recalibrated to reflect current labor market volatility.

Sector-Specific Entropy Anchors

Thresholds are recalibrated monthly to reflect sector volatility:

High-Velocity Tech (FAANG+)

Specialist signals are anchored below 0.7

Healthcare & Life Sciences

Sector stability is historically higher; specialist signals are anchored below 0.3

Finance & Private Equity

Progressive mobility signals are expected between 0.5 and 1.2

Audit II: Survival Analysis (Cox Proportional Hazards)

We utilize the Cox Proportional Hazards Model to quantify retention risk over a 36-month horizon. While the model visualizes a long-term decay curve, its primary fiduciary utility is identifying the **Hazard Inversion Points** occurring in the first 12 months.

Time-to-Event Modeling

The engine identifies "Jarring Events" and calculates specific Hazard Ratios based on historical data. While base survival is modeled via Cox, environmental hazards are applied as Retention Multipliers to calibrate the probability curve.

$$\lambda(t|Z) = \lambda_0(t) \exp(\sum_{j=1}^{p} \beta_j Z_j)$$

Audit III: Environmental Friction (Policy Shock)

We quantify non-linear "Shock" turnover risks driven by policy-to-lifestyle mismatch.

RTO Shock & Autonomy Loss

- **Remote → Office:** Candidates transitioning from fully remote to >45m commutes receive a 0.825x Retention Multiplier (~18% penalty).
- **Executive Calibration:** For C-Suite roles, the model dampens RTO friction by 50% to account for high-agency negotiation power (e.g., dedicated drivers, flexible hours).

Strategic AI Mismatch

High-skill "AI Superusers" joining legacy organizations experience high strategic friction. These assets receive a 0.875x Retention Multiplier to account for efficiency frustration. This penalty is not dampened for executives, as strategic misalignment is a primary churn driver at the leadership level.

Audit IV: Social Dynamics & Network Contagion

Retention is social. Peer departures are the #1 leading predictor of individual turnover.

The 135-Day Contagion Window

Based on Visier (2023) research, the probability of a "Viral Resignation" peaks 4.5 months post-departure.

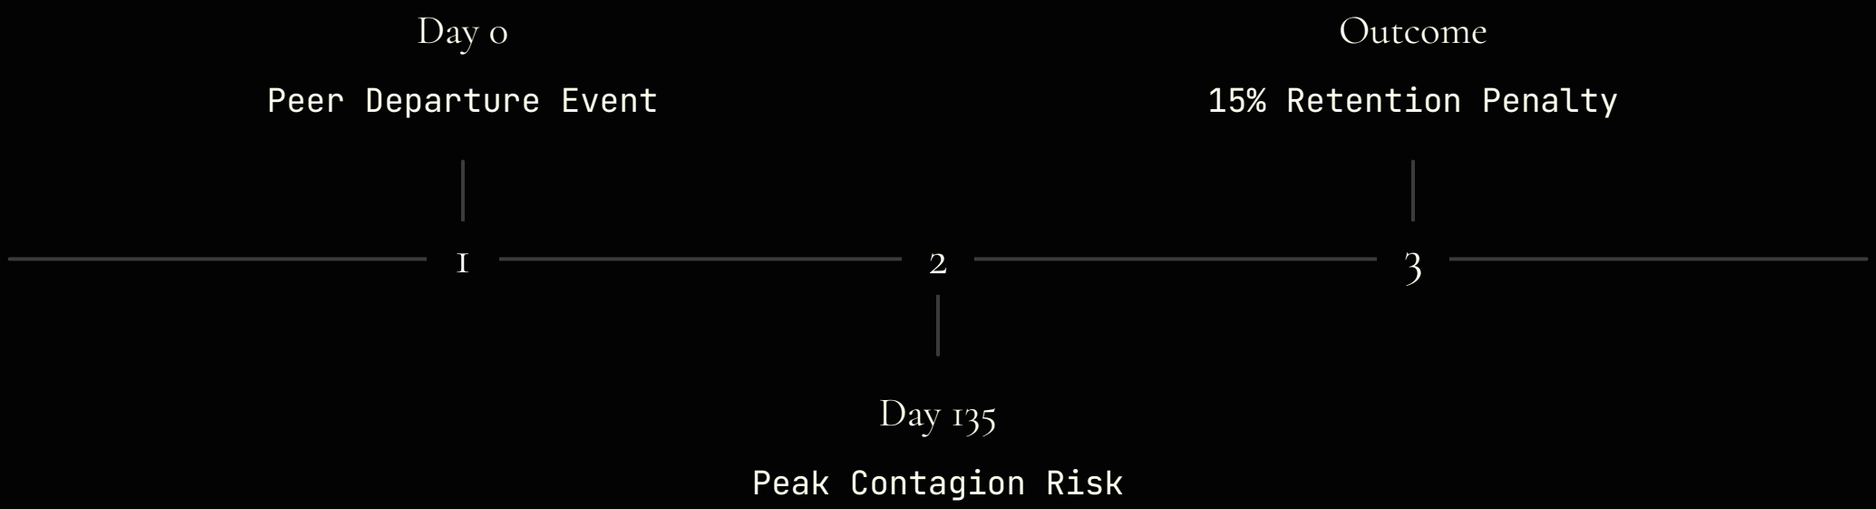

- **Exodus Clusters:** Candidates hired out of an active peer exodus cluster receive a 0.85x Retention Multiplier (15% penalty) to account for the social "pull" risk of their former network.

The Boomerang Premium

Candidates returning to a former employer ("Boomerangs") receive a 1.44x Retention Multiplier, reflecting the 44% stay-premium documented in Shipp et al. (2014).

Calibration & Statistical Significance

To prevent AI-driven "grade inflation," the engine utilizes **Nuclear Calibration**:

Nuclear Calibration (Mean 50)

Forced distribution calibration anchors scores to a true mean of 50 to prevent AI-driven "grade inflation."

Statistical Rarity

In current validation datasets (internal and expanding), a Stability Index of 80+ represents a >2-sigma outlier. Such scores indicate candidates with exceptional historical resistance to common turnover catalysts.

Compliance & Defensibility (NYC LL144.1)

The Stability Engine is a "Glass-Box" system.

Audit Trail

Every score adjustment is logged with its specific research citation (e.g., Griffeth et al., 2000) and logic trace.

Decision Support

The engine is categorized as a Decision Support Layer, ensuring the final hiring liability remains with the human fiduciary owner.

📚 Selected Bibliography & Methodological Note

Selected Bibliography

- Griffeth, R. W., et al. (2000): Meta-analysis of antecedents of turnover.
- Cox, D. R. (1972): Regression models and life-tables.
- Felps, W., et al. (2009): Turnover contagion and job embeddedness.
- Shipp, A. J., et al. (2014): Boomerang employee retention premiums.
- Holtom, B. C., et al. (2008): The Unfolding Model of Voluntary Turnover.

> **Methodological Note:** Internal validation based on $n = 42$ leadership resumes (Feb 2026). Continuous production validation is in progress. All Hazard Ratios and Multipliers are subject to monthly actuarial recalibration.