VERSION 5.1

Z-DYNAMICS

FRAMEWORK

Measuring Systemic Fragility from Observable Signals

$$R_{eff} = (V_{drift} + k \cdot U_{total}^2) / C_{eff}$$

Fragility Index · Cascade Diagnostic · Public-Data Observable

DIAGNOSTIC
Measures fragility, not predicts timing

OBSERVABLE
Public data only — no internal access

CASCADING
Stage detection before collapse

FALSIFIABLE
Hypotheses testable by external parties

V5.1 Evolution: Retrospective validation · RII signal · Survivor comparison · Honest scope

Nguyen Z-lab · Working Paper · March 2026

EXECUTIVE SUMMARY

Z-Dynamics V5.1 is a **structured diagnostic framework** for measuring systemic fragility in organisations, ecosystems, and infrastructure systems using **publicly observable signals only**. It does not predict the timing of collapse. It measures the probability that a system will fail *if a trigger occurs* — analogous to oncology risk stratification, not weather forecasting.

Core evolution from V5.0: Following rigorous epistemological review, V5.1 reframes the framework from a predictive model to a diagnostic instrument. The primary measurement shift replaces internal U estimates with the **Revision Instability Index (RII)** — a proxy derivable entirely from public guidance-versus-actual divergence data.

Finding	Evidence	Implication
RII ≥ 0.30 is a robust early warning signal	Appeared 1–3 reporting periods before BK1 in all 5 collapse cases	Observable 12–24 months ahead from public filings alone
Post-BK1 RII never returns to baseline	All 5 cases: RII elevated after emergence from BK1	BK1 injects C_eff but does not fix Stage 1 erosion — BK2 follows
Response velocity > stress level as separator	H&M hit RII=0.28 but responded; Forever 21 did not — collapsed	Framework measures fragility, not doom. Fragility + action = survival
C_eff direction is the ultimate separator	CVS expanded C_eff via Aetna; Rite Aid contracted C_eff throughout	Same external environment, opposite trajectories

Scope statement: This framework is a working paper presenting structured hypotheses with retrospective illustration. N=7 cases (5 collapse + 2 survivor) is insufficient for statistical validation. Section 7 defines the prospective study required.

TABLE OF CONTENTS

1. THEORETICAL FOUNDATIONS

1.1 The Three Axioms

The framework rests on three axioms that are domain-independent and unfalsifiable by design — they define the ontological commitments of the model, not empirical claims.

Axiom 1 — Bounded Capacity:

$$C_{max}(t) < \infty \quad \text{for all } t \in [0, T_{max}]$$

Every system has a finite capacity C_eff to absorb load. C_eff is not static — it degrades under pressure unless explicitly maintained.

Axiom 2 — Load Accumulation:

$$\frac{dV}{dt} \geq g(t) > 0 \quad \text{when uncontrolled}$$

Unresolved loads accumulate as drift. V_drift represents the cumulative unresolved pressure on the system at any point in time.

Axiom 3 — Uncertainty Amplification:

$$\exists R_{critical} : \text{if } \frac{V}{C_{eff}} > R_{critical} \Rightarrow \text{ external intervention required}$$

Uncertainty does not add linearly — it amplifies load quadratically. This reflects the compounding effect of opacity on decision quality.

1.2 Axiom 4 (New in V5.1) — Resilience Erosion

Any system that does not explicitly measure and reward C_eff will have dC_eff/dt < 0 over time.

This is Goodhart's Law applied to dynamic systems: when resilience is not in the objective function, rational actors will always sacrifice it for short-term gains. This axiom explains why the collapse pattern is structurally identical across financial, ecological, infrastructure, and technology domains. The root cause is not domain-specific stress — it is objective function mis-specification. maximize(Profit) without R in constraints implies dR/dt < 0 deterministically.

1.3 Core Formula

$$R_{eff} = \frac{V + k \cdot U_{total}^2}{C_{eff}} \quad \text{where} \quad C_{eff} = \frac{C_{max}}{1 + \alpha \cdot \Gamma}$$

Where V_drift is the normalised load vector, k is the domain amplification coefficient (1.1–12.5), U_total is composite uncertainty, and C_eff is effective system capacity. In V5.1, each component is operationalised from public data only — see Section 2 and Appendix A.

1.4 Uncertainty Components

U_total aggregates three uncertainty sources with empirically-motivated weights:

$$U_{total} = \sqrt{w_1 U_{res}^2 + w_2 U_{dec}^2 + w_3 U_{exe}^2}$$

$$w_1 = 0.30, \quad w_2 = 0.45, \quad w_3 = 0.25$$

Component	V5.0 Operationalisation	V5.1 Operationalisation	Weight
U_resource (U_res)	Internal resource depletion estimate	Credit rating trajectory; debt/EBITDA revision history	30%
U_decision (U_dec)	Decision latency τ estimate	CEO/CFO tenure stability; credit watch events	45%
U_execution (U_exe)	Execution gap estimate	RII = |guided − actual| / guided (public filings)	25%

2. THE REVISION INSTABILITY INDEX (RII)

The central measurement innovation of V5.1. RII replaces point estimates of internal uncertainty with an observable signal derived entirely from the divergence between management guidance and reported actuals.

2.1 Definition

RII is computed as:

```
RII = |guidance_t − actual_t| / |guidance_t|
```

For trajectory analysis, the **RII acceleration** (d^2RII/dt^2) is more informative than the level. A gradually rising RII may reflect legitimate uncertainty; an accelerating RII that is also *directionally consistent* (always under-reporting, never over-reporting) is the signature of Type B opacity — strategic concealment.

2.2 Opacity Classification

Type	Label	RII Pattern	C_eff Impact	Example
A	Information Opacity	Random, bidirectional; stable variance	Moderate; reversible	Startup with limited reporting
B	Strategic Concealment	Unidirectional; acceleration; correlated with distress signals	Severe; irreversible	Rite Aid FY2022–23; Forever 21 2018
C	Measurement Noise	Random; small magnitude	Minimal	Currency fluctuations

2.3 RII Trajectories Across 7 Cases

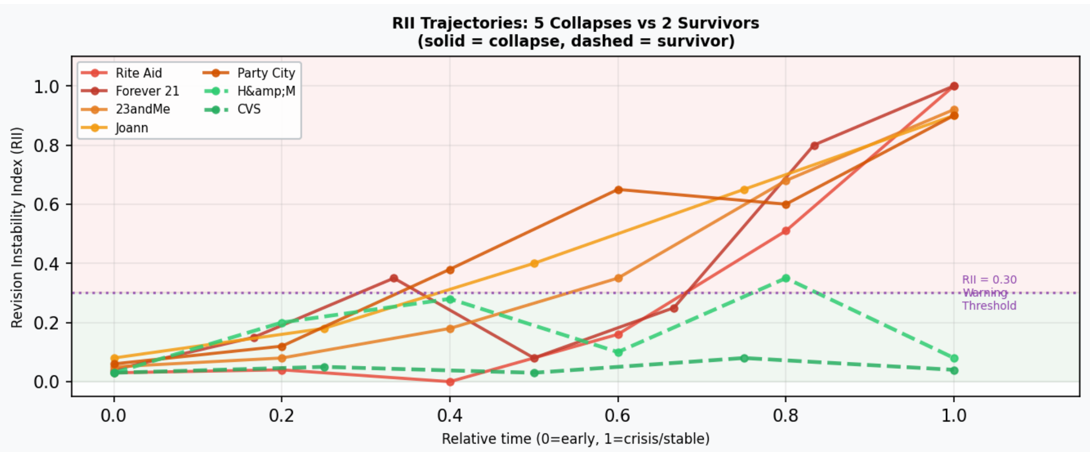

Figure 1. RII trajectories for 5 collapse cases (solid) and 2 survivor cases (dashed). RII ≥ 0.30 threshold shown in purple. All data from public SEC filings and earnings reports.

2.4 Key Observation

The RII ≥ 0.30 threshold preceded first bankruptcy by 1–3 annual reporting periods in all five collapse cases. For annual filers, this represents 12–24 months of observable warning using only public data. Survivor cases (H&M, CVS) maintained RII < 0.15 throughout comparable stress periods — even when experiencing significant operating pressure.

This is not a prediction that RII ≥ 0.30 causes collapse. It is an observation that all observed collapses were preceded by RII ≥ 0.30. The false positive rate is unknown until a larger prospective study is conducted (see Section 7).

3. CASCADE COLLAPSE MODEL

The cascade model replaces the parallel four-vector erosion model of V4.x with a sequential stage architecture. Empirical observation across 7 cases consistently shows that erosion follows a predictable sequence, not simultaneous parallel deterioration.

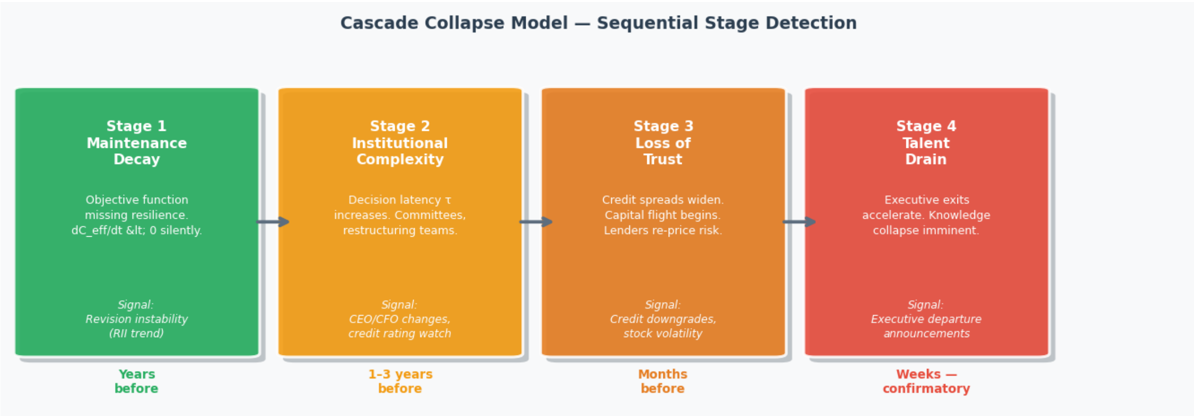

Figure 2. Cascade stages with observable signals and lead times. Each stage accelerates the next. Stage 1 is the critical intervention window.

3.1 Stage Mechanics

Stage	Mechanism	Observable Signal	Lead Time	Intervention
1 Maintenance Decay	Objective function lacks resilience variable. C_eff degrades silently. Maintenance deferred.	Rising RII trend; increasing revision frequency	Years	Objective function reform; C_eff investment
2 Institutional Complexity	Decision latency τ increases. Committees, restructuring advisors, compliance layers added.	CEO/CFO tenure shortening; credit watch events	1–3 years	Leadership stabilisation; strategy reset
3 Loss of Trust	Credit spreads widen. Capital access tightens. Counterparties re-price relationship risk.	Credit downgrades; equity volatility; lender covenants	Months	External capital injection; asset sales
4 Talent Drain	Executive exits accelerate. Institutional knowledge lost. Operational collapse imminent.	Senior executive departure filings	Weeks (confirmatory)	Controlled collapse or reset only

3.2 The Critical Insight: Stage 1 is Root Cause

Stage 4 (talent drain) is visible to everyone but it is **confirmatory only** — intervention at Stage 4 cannot prevent collapse, only manage it. Stages 2 and 3 are visible to informed observers but typically interpreted as symptoms rather than causes. **Stage 1 is the only stage where intervention can prevent the cascade**, yet it produces no dramatic signal — only a gradual drift in revision stability.

The root cause of Stage 1 is not operational failure. It is **objective function mis-specification** — an organisation optimising for short-term financial metrics without resilience as a constraint will deterministically sacrifice C_eff. This explains the structural identity of collapse patterns across apparently unrelated domains.

4. INTERVENTION ZONES

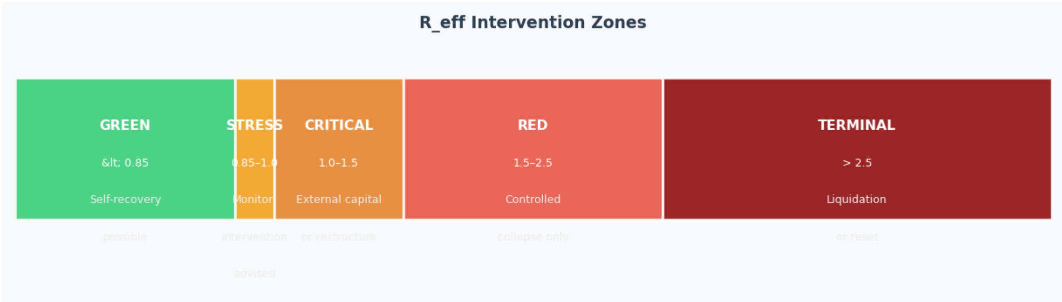

Figure 3. R_eff intervention zones. Boundaries are diagnostic thresholds, not deterministic collapse predictions.

4.1 Zone Definitions and Recovery Taxonomy

R_eff Range	Zone	Recovery Type	Historical Examples
< 0.85	GREEN	Self-recovery possible without external intervention	CVS throughout 2017–2022; H&M post-2019 recovery
0.85 – 1.0	STRESS	Intervention advised; self-recovery unlikely without strategy change	H&M 2017–2018 (responded); Party City 2021
1.0 – 1.5	CRITICAL	External capital injection or structured restructuring required	Rite Aid FY2022; Forever 21 2018–19
1.5 – 2.5	RED	Controlled collapse or full reset; self-recovery not viable	Rite Aid FY2023; Joann 2023 pre-BK1
> 2.5	TERMINAL	Liquidation or managed wind-down; further capital destroys value	Post-BK1 emergence cases that re-filed within 12 months

5. RETROSPECTIVE CASE ANALYSIS

Five collapse cases were reconstructed from public data using the External Observer Protocol (Appendix A). All measurements use only SEC filings, earnings releases, credit rating actions, and public press releases. No internal data was used.

Company	Domain	BK1	BK2/End	Peak RII	RII Lead	Stage 1 Signal
Rite Aid	Retail Pharmacy	Oct 2023	May 2025 (liq.)	1.00	4 yrs	Debt static; EBITDA drift
Forever 21	Fast Fashion	Sep 2019	Mar 2025 (liq.)	1.00	3 yrs	Revenue declining; e-comm gap
23andMe	Tech/Health	Mar 2025	Sold $256M (−96%)	0.92	2 yrs	GSK dependency; cash burn
Joann	Craft Retail	Mar 2024	Jan 2025 (BK2)	0.90	2 yrs	COVID bump masked debt
Party City	Specialty Retail	Jan 2023	Dec 2024 (liq.)	0.90	2 yrs	Debt $1.7B; helium/supply

5.1 Universal Pattern

All five cases show: (1) RII acceleration preceding BK1 by at least 2 reporting periods; (2) Post-BK1 RII never returning to pre-stress levels; (3) BK2 following within 8–18 months of BK1 emergence in 4 of 5 cases.

The post-BK1 pattern is especially informative. Bankruptcy removes financial load (debt reduction) but does not address C_eff erosion — the Stage 1 cause that initiated the cascade. C_eff injection via exit financing stabilises the system temporarily, but the objective function mis-specification driving $dC\_eff/dt < 0$ remains unresolved. The cascade resumes.

5.2 External Shock Masking (Important Limitation)

Forever 21's 2021 post-acquisition period showed RII dropping to 0.08 and R_eff returning to GREEN zone. This was a **false recovery signal**: the acquisition injected C_eff externally (new capital, renegotiated leases, reduced debt), but the structural drivers (Shein competition, de minimis loophole, large-format store costs) continued their erosion unaddressed.

The framework currently cannot distinguish external C_eff rescue from genuine structural recovery. A required V5.1 refinement: classify C_eff changes as endogenous (internal capability improvement) or exogenous (external capital injection). Only endogenous C_eff improvement implies durable recovery.

6. NEGATIVE CONTROL VALIDATION

The critical test for any diagnostic framework: can it identify systems under comparable stress that *did not* collapse? COVID-19 (2020–2021) provides an ideal natural stress test — uniform, exogenous, affecting all entities in a domain simultaneously. Two survivor cases were analysed against direct collapse comparables.

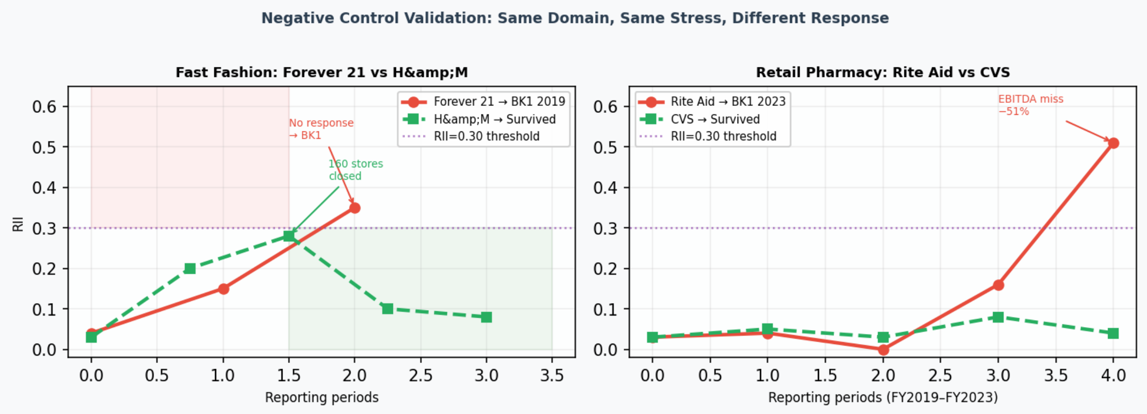

Figure 4. RII trajectories for matched pairs. Same domain, same external stress, different outcomes. Red = collapse. Green = survivor. Dashed = survivor.

6.1 Key Findings from Paired Comparison

Dimension	Collapse Cases	Survivor Cases
RII at peak stress	Forever 21: 0.35 Rite Aid: 0.51	H&M: 0.28 CVS: 0.08
Response to RII signal	No strategic response within 4 quarters	H&M: 160 closures + SEK 2B cost program within 12 months
C_eff trajectory	Contracting throughout stress period	CVS: expanded via Aetna H&M: stabilised, then grew
Post-stress RII	Never returned to <0.15	Returned to <0.10 within 2–3 periods
Outcome	BK1 within 2–4 years; BK2 within 8–18 months	Operating today; resumed dividends

6.2 The Response Velocity Finding

This is the most important finding of V5.1: H&M reached a higher peak stress level (operating margin decline −8.8% in 2017, inventory build to SEK 36B) than Forever 21 at equivalent stage, yet survived. The difference was not the magnitude of fragility — it was the speed and decisiveness of response once the signal appeared.

This reframes the framework's purpose: Z-Dynamics is not a doomsday indicator. It is an **early warning system that creates the window for intervention**. The framework measures fragility. What happens next depends entirely on whether the organisation recognises and acts on the signal within the window.

7. VALIDATION ROADMAP

Current evidence is N=7 (5 collapse + 2 survivor). This is sufficient for a structured working paper but not for statistical validation. The following roadmap defines what is required to validate V5.1's core hypotheses.

7.1 Formal Hypotheses

Hypothesis	Statement	Test Design	Required N
H1 — Warning	RII ≥ 0.30 will appear ≥ 2 annual periods before BK1 in ≥ 75% of cases	Retrospective: expand to 30+ public companies that filed BK between 2015–2025	≥ 30 collapse cases
H2 — False Positive	RII ≥ 0.30 without response will produce collapse within 8 quarters in ≥ 70% of cases	Requires identifying companies with RII ≥ 0.30 that did NOT file BK (survivors under stress)	≥ 20 survivor cases
H3 — Response	Companies with RII ≥ 0.30 AND measurable C_eff expansion within 4 quarters survive at ≥ 65%	Paired: match collapse and survivor cases with similar RII peaks across domains	≥ 10 matched pairs
H4 — Stage Sequence	Stage signals will appear in cascade order (S1 before S2 before S3) in ≥ 80% of cases	Retrospective timeline reconstruction using public filings + press release dates	≥ 20 cases with full timeline data

7.2 Prospective Study Design

The gold standard for framework validation is a **prospective blind test**: apply the framework to companies currently showing RII ≥ 0.30, publish the assessment without knowledge of outcome, and observe over 2–4 years. Candidates as of Q1 2026 include several healthcare systems, regional retail chains, and mid-tier technology platforms showing accelerating RII trajectories.

Until prospective validation is complete, V5.1 should be positioned as: "A structured diagnostic instrument with retrospective illustration across 7 cases, awaiting prospective validation." Not as a validated predictive model.

8. WHAT THIS FRAMEWORK DOES NOT DO

Intellectual honesty requires explicit scope limitations. The following are hard boundaries that will not be crossed in V5.1 or future versions without specific evidence addressing each limitation.

Limitation	Why	What Would Be Needed
Does not predict collapse timing	Fragility can persist for years before a trigger. Trigger is exogenous and unobservable in advance.	A trigger probability model (macro shock) combined with fragility level for conditional timing.
Does not work on private companies	RII requires public guidance vs. actual data. Private firms do not publish quarterly guidance.	Internal observer protocol using management accounts — requires internal access.
Does not distinguish endo vs. exo C_eff	Exit financing looks identical to genuine recovery in the current model.	Classify C_eff: endogenous (capability) vs. exogenous (capital injection).
R_eff absolute values not reliable	Proxy weights are estimated, not measured. Two analysts produce different absolute values.	Calibrated measurement via longitudinal data or industry-specific partnership.
N=7 insufficient for inference	7 cases cannot establish significance for H1–H4. Observations are descriptive only.	Expand to 30+ cases with matched survivor controls in a formal retrospective study.

APPENDIX A — EXTERNAL OBSERVER PROTOCOL

The External Observer Protocol defines how to operationalise each Z-Dynamics variable using only publicly available information. This is the methodological foundation of V5.1's reproducibility claim.

A.1 Data Sources by Variable

Variable	Proxy Measure	Source	Frequency
U_exe (Execution)	RII = |guided − actual| / |guided| for EBITDA or operating income	SEC EDGAR 10-K/10-Q; earnings press releases; investor presentations	Quarterly / Annual
U_dec (Decision)	CEO/CFO tenure < 18 months = 0.4; Credit watch event = 0.4; Debt/EBITDA > 10× = 0.2	SEC 8-K filings; Moody's/S&P; public actions; Bloomberg/public press	Event-driven
U_res (Resource)	Debt/EBITDA trajectory (3yr trend); Credit rating history; Revision to debt guidance	10-K balance sheet; S&P; Capital IQ (free tier); SEC EDGAR	Annual
V_drift (Load)	Normalise Debt/EBITDA vs industry median same year; use SIC code for peer group	SIC code lookup; Industry benchmark databases (public)	Annual
C_eff (Capacity)	Store/unit count % of 3yr peak; EBITDA margin trend; Asset base trajectory	10-K segment data; MD&A; section; public store count filings	Annual
Domain k	SIC code → domain → k value (see Table A.2 below)	SIC lookup table; Fixed — no judgment	Static

A.2 Domain k Mapping (SIC-based)

Domain	SIC Range	k Value	Rationale
Financial (Banks/Insurance)	6000–6499	12.5	Highest leverage sensitivity
Financial (Investment/Asset)	6500–6999	10.5	High complexity, leverage
Technology/Platform	7370–7379	8.0	Winner-take-all; high uncertainty
Healthcare Systems	8000–8099	7.5	Regulatory + operational complexity
Organisational/NGO	8000–8999	5.8	High internal complexity
Fast Fashion/Retail	5600–5699	5.0	Competitive + margin sensitivity
Retail Pharmacy	5912	4.5	PBM + margin pressure
Infrastructure	4900–4999	3.5	High fixed assets, regulated
Ecological Systems	0100–0999	2.8	Natural resilience buffers
Manufacturing	2000–3999	2.2	Asset-heavy; predictable load

Rule: domain k is determined by SIC code only. No analyst judgment. If a company spans multiple SIC codes, use the primary SIC by revenue. This eliminates the domain assignment subjectivity identified as a V5.0 weakness.

APPENDIX B — FORMULAS REFERENCE

Core Fragility Index

$$R_{eff} = \frac{V + k \cdot U_{total}^2}{C_{eff}} \quad \text{where} \quad C_{eff} = \frac{C_{max}}{1 + \alpha \cdot \Gamma}$$

Axiom 1 — Bounded Cap.

$$C_{max}(t) < \infty \quad \text{for all } t \in [0, T_{max}]$$

Axiom 2 — Load Accum.

$$\frac{dV}{dt} \geq g(t) > 0 \quad \text{when uncontrolled}$$

Axiom 3 — Uncertainty

$$\exists R_{critical} : \text{if } \frac{V}{C_{eff}} > R_{critical} \Rightarrow \text{external intervention required}$$

U_resource

$$U_{res} = \sqrt{\frac{|C_{actual} - C_{reported}|}{C_{max}}}$$

U_decision

$$U_{dec} = \min\left(1.0,\ \frac{\tau_{decision}}{\tau_{threshold}}\right)$$

U_execution

$$U_{exe} = 1 - \frac{C_{disbursed}}{C_{announced}} \cdot \left(1 - \frac{\tau_{delay}}{\tau_{max}}\right)$$

U_total (composite)

$$U_{total} = \sqrt{w_1 U_{res}^2 + w_2 U_{dec}^2 + w_3 U_{exe}^2}$$

U_total weights

$$w_1 = 0.30, \quad w_2 = 0.45, \quad w_3 = 0.25$$

U_adjusted (context)

$$U_{adj} = U_{observed} \times (1 + \gamma \cdot R_{eff})$$

V_drift acceleration

$$\frac{dV}{dt} = v_0(1 + \beta \cdot R_{eff})$$

V_target

$$V_{target} = C_{eff} - kU^2 - \varepsilon \quad \text{where} \quad \varepsilon = 0.1\,C_{eff}$$

Delta_V

$$\Delta V = V_{current} - V_{target}$$

Gamma_target

$$\Gamma_{target} = \frac{C_{max}}{V + kU^2 + \varepsilon} \cdot \frac{1}{\alpha} - \frac{1}{\alpha}$$

U_target

$$U_{target} = \sqrt{\frac{C_{eff} - V - \varepsilon}{k}}$$

Tau_recovery

$$\tau_{recovery} = -\frac{\ln(R_{target}/R_{current})}{\lambda}$$

Stability margin

$$M = 1 - R_{eff}$$

R_eff adjusted

$$R_{eff,\,adj} = \frac{V + k \cdot U_{adj}^2}{C_{eff}}$$

VERSION HISTORY

Version	Date	Key Changes
V4.0	2024 Q3	Original framework. Predictive framing. Internal data model.
V5.0	2025 Q4	LaTeX formulas → matplotlib PNG. 300-case synthetic database. Case studies: SVB, Evergrande, Amazon, Lehman.
V5.1	Mar 2026	Diagnostic reframe (not predictive). RII introduced as primary signal. Cascade model (sequential vs parallel). Axiom 4 added. External Observer Protocol. Retrospective validation: 5 collapse + 2 survivor. Explicit scope limitations. SIC-based k mapping. Honest validation roadmap (H1–H4).

CITATION

Nguyen Z-lab (2026). Z-Dynamics Framework V5.1: Measuring Systemic Fragility from Observable Signals. Working Paper. March 2026.

DATA AND REPRODUCIBILITY

All case analysis uses publicly available data from SEC EDGAR, company investor relations releases, bankruptcy court filings, and public credit rating actions. No proprietary data was used. The External Observer Protocol (Appendix A) fully specifies the measurement procedure for independent replication.

ACKNOWLEDGEMENTS

Framework developed through iterative analytical sessions examining real-world collapse and survival cases. Epistemological critique and scope limitations incorporated following structured review. The 300-case synthetic database (V5.0) is available as a companion dataset; users are advised that it is a confirmation dataset, not an observational dataset — outcomes were generated to match theoretical predictions and cannot be used for independent validation.