The Structural Risk Model

Mathematical Foundations

Companion to A Risk Manifesto

Richard Bookstaber

August 31, 2026

1 Overview

This is a self-contained introduction to the core of the mathematical framework underlying *Crisis Without History* (Bookstaber, forthcoming). The model measures financial risk from the observable architecture of business relationships—similarity between firms and directed supply-chain dependencies—rather than from historical return data. The result is a propagation operator whose spectral decomposition reveals how shocks travel through an economic network, a three-way structural decomposition of every portfolio's risk into diffusion, tension, and bottleneck components, and a demonstration of why curvature of the economic surface makes conventional factor models structurally incomplete.

Everyone knows that airlines depend on fuel, semiconductors bottleneck everything, similar businesses tend to move together, and dependencies matter asymmetrically. That intuition is the model. The presentation that follows makes it explicit, operational, and inspectable.

Traditional risk models begin with returns and infer structure from observed comovement. The framework developed here begins with economic structure—the network of business relationships—and determines how shocks propagate through that structure.

The organizing idea is simple: economic shocks move through networks of dependencies and similarities. The model specifies the rules of that motion, in four stages.

1. **Primitives.** A symmetric similarity matrix S and a directed adjacency matrix A, constructed from regulatory filings, encode operational likeness and supply-chain dependency. These are not fitted parameters but structural observables.
2. **Operator.** The propagation operator $M = \alpha L_S + \beta L_A$, built from the two Laplacians, defines the economic surface—the terrain on which disturbances propagate.
3. **Modes.** The spectral decomposition of M reveals $\{\phi_k, \mu_k\}$: propagation modes (coherent firm configurations through which a shock travels intact) and dissipation rates (how

far each mode carries before dissolving). These are traversal patterns, not statistical factors.

4. **Decomposition.** Every portfolio's structural exposure splits into three orthogonal mechanisms—*diffusion*, *tension*, and *bottleneck*—turning risk from a scalar into a point in mechanism-space. Vulnerability is determined by direction, not magnitude alone.

Table 1 summarizes the six analytical objects and their roles; the sections that follow develop each in turn.

Primitive	**Analytical Status**	**Role**
S, A	Matrices	Structural Primitives Operational similarity and supply adjacency, observable from business disclosures.
L_S, L_A	Generators	Local Flow Rules Discrete Laplacians governing how a shock moves from a node to its immediate neighbors.
M	Operator	Global Terrain Eigenvectors: natural propagation routes. Eigenvalues: dissipation rates along each route.
$[L_S, L_A]$	Commutator	Geometric Curvature Algebraic source of path dependence; nonzero iff the manifold curves back on itself.
Π_{diff}, Π_{ten}, Π_{bot}	Projectors	Orthogonal Mechanisms Exact decomposition of every response into Diffusion, Tension, and Bottleneck components.
ϕ_k, μ_k	Spectra	Emergent Geometry Traversal patterns the structure reinforces or disperses, independent of return history.

Table 1: Summary of analytical primitives and their roles in the framework.

2 The Structural Primitives: S and A

Consider a universe of n firms, indexed by $i = 1, \dots, n$. Firms are connected by two distinct kinds of relationships:

1. *Similarity*: firms operating in comparable economic environments—similar products, customers, technologies, regulations, or geographies.

2. *Adjacency*: firms connected by directed dependencies—supplier–customer links, platform reliance, transportation networks, shared infrastructure, or contractual obligations.

Let $S \in [0,1]^{n \times n}$, $S = S^\top$ be a symmetric similarity matrix. The entry S_{ij} quantifies operational proximity: similar products or technologies; similar customer bases or market segments; similar geographic or regulatory exposures; similar cost structures; or similar competitive dynamics. Symmetry reflects mutual exposure: if firm i is similar to firm j, firm j is equally similar to firm i. Normalization: $S_{ii} = 1$ and $0 \leq S_{ij} \leq 1$.

Let $A \in \mathbb{R}_+^{n \times n}$ be a directed adjacency matrix, where

$$A_{ij} = \frac{\text{critical inputs firm } j \text{ obtains from firm } i}{\text{total critical inputs of firm } j}. \tag{1}$$

Asymmetry reflects directed causality: $A_{ij} > 0$ with $A_{ji} = 0$ means firm j depends on firm i, not conversely. Shocks flow from i (supplier) to j (dependent firm). Column normalization: $\sum_i A_{ij} \leq 1$.

2.1 From Filings to Matrices

Both matrices are observable, not fitted.

Similarity is the easy half, because it already exists. Text-based pairwise similarity between firms—constructed from business descriptions in 10-K filings—is a maintained research dataset: Hoberg and Phillips (2016) built precisely this object, validated across empirical finance for over a decade with history to the mid-1990s. The framework generalizes scope (adding geographic, regulatory, cost-structure, and infrastructure overlap to product-market language) but not kind: modern language models score operational overlap from the same filings along any specified dimension. The similarity half of the input is a solved measurement problem being continuously re-solved at higher resolution.

Adjacency requires a distinction. The *support* of A—who supplies whom—is reasonably observable: firms disclose major customers (concentrations above ten percent of revenue under segment-reporting rules), supply relationships surface in filings and trade data, and commercial datasets assemble these links at scale. The *weights*—the fraction of firm j's critical inputs sourced from firm i—are estimates layered on an observed skeleton, and they carry error.

Crucially, measurement error is a perturbation, and this framework prices perturbations. An error in S or A is a ΔS, ΔA—exactly the object the perturbation machinery controls. The stability of every spectral quantity is governed by computable bounds (eigenvalues move by at most the perturbation's norm; eigenvector blocks rotate by at most the perturbation over the spectral gap). Empirically: perturb every nonzero link by independent random errors of up to $\pm 20\%$; across five hundred such corrupted networks, the slow eigenvalue stays within $\pm 2\%$ of its true value, the systemic mode's direction is essentially unmoved (alignment 0.9999), and portfolio exposure moves by about $\pm 10\%$. Twenty percent error in, two percent error out, for the quantities that drive decisions.

Parameters. Beyond the matrices, the model has four scalars: the channel weights α, β, the propagation intensity η, and the shock scale γ. They are set once—by economic argument or by matching structural covariance to realized covariance over a calm window—not refitted to improve return-based performance. A conventional factor model for 500 stocks fits tens of thousands of parameters to returns. This model takes two matrices of observations plus four scalars.

The interaction of similarity and adjacency produces four distinct propagation regimes, illustrated in Table 2.

	Low Similarity	**High Similarity**
High Adjacency	**Cascade Zone** Directed transmission *Example: TSMC → chip customers*	**Clustered Fragility** Reinforcement through both channels *Example: Regional banks, 2008*
Low Adjacency	**Isolation Zone** Shocks remain localized *Example: Unrelated firms*	**Diffusion Zone** Broad spread without cascade *Example: Airlines + fuel shock*

Table 2: Four propagation regimes defined by the interaction of S and A.

3 The Generators and the Propagation Operator

A shock is represented as a vector $x \in \mathbb{R}^n$, where x_i is shock intensity at firm i. At this stage x carries no probabilistic interpretation; it records where a disturbance enters the system.

3.1 The Similarity Laplacian L_S

The normalized similarity Laplacian governs diffusion:

$$L_S = D_S - S, \qquad (D_S)_{ii} = \sum_j S_{ij}. \tag{2}$$

For a shock x, the one-step update $x \mapsto x - \delta L_S x$ moves shock intensity from firm i toward similar firms j in proportion to S_{ij}, while attenuating at i in proportion to its total similarity exposure $\sum_j S_{ij}$. This is heat diffusion on the similarity graph: a fuel price shock spreads to all airlines simultaneously, proportional to their operational overlap.

L_S is symmetric positive semidefinite. Its smallest eigenvalue is zero (corresponding to the uniform vector), and all eigenvalues are non-negative.

3.2 The Adjacency Generator L_A

The adjacency generator governs directional transmission:

$$L_A = I - A^\top. \tag{3}$$

The one-step update $x \mapsto x - \delta L_A x$ transmits from suppliers to their dependent customers along directed edges, attenuating at the source and amplifying at the destination in proportion to the dependency weight A_{ij}. A disruption at a mining firm flows downstream to refiners, then to battery manufacturers, then to electric vehicle producers.

L_A is generally asymmetric, reflecting the directed causality of supply chains.

3.3 The Combined Operator M

The propagation operator combines both channels:

$$M = \alpha L_S + \beta L_A, \qquad \alpha, \beta \geq 0. \tag{4}$$

This single operator governs all propagation dynamics. The parameter ratio α/β controls the relative speed of diffusion (through similarity) versus transmission (through dependency chains).

M decomposes into symmetric and antisymmetric parts:

$$M_s = \tfrac{1}{2}(M + M^\top) = \alpha L_S + \tfrac{\beta}{2}(L_A + L_A^\top) \tag{5}$$

$$M_a = \tfrac{1}{2}(M - M^\top) = \tfrac{\beta}{2}(L_A - L_A^\top). \tag{6}$$

The symmetric core M_s drives the spectral decomposition; the antisymmetric part M_a encodes the directionality of supply dependencies.

3.4 The Propagation Response Operator

The full propagation response, incorporating all multi-hop cascades, is encoded in the operator

$$E = P_\eta(\gamma S)P_\eta^\top, \qquad P_\eta = (I - \eta A^\top)^{-1}, \tag{7}$$

where η is propagation intensity and γ scales the similarity kernel. The resolvent $P_\eta = I + \eta A^\top + \eta^2(A^\top)^2 + \cdots$ sums all multi-hop cascades, attenuating by η per hop. Entry E_{ij} answers: if firm j experiences a unit disturbance, how much does firm i feel after all similarity diffusion and adjacency cascades have run?

In the weakly-asymmetric regime ($\varepsilon = \|M_a\|_2/\lambda_{\min}(M_s) < 1$), the response admits the mode-filter representation

$$E \approx \sum_k \frac{\gamma}{\mu_k + \eta}\, \phi_k \phi_k^\top, \tag{8}$$

with controlled error. The weight $\gamma/(\mu_k + \eta)$ is transmission efficiency: slow modes (small μ_k) transmit strongly; fast modes dissipate.

4 The Commutator and Geometric Curvature

The commutator

$$[L_S, L_A] = L_S L_A - L_A L_S \tag{9}$$

measures whether the order in which the two channels act matters. When $[L_S, L_A] \neq 0$, a shock arriving through similarity-then-adjacency produces different exposure than adjacency-then-similarity.

For the S&P 500, $\|[L_S, L_A]\|_F \approx 0.34$: the economic surface curves substantially back on itself.

Geometric meaning. When $[L_S, L_A] \neq 0$, the manifold on which firms live is curved. The sequence of shocks matters: the state a firm reaches after experiencing two shocks depends on their order. This is not a modeling assumption; it is a measurable property of the economy, arising because similarity and adjacency are generically non-aligned.

A concrete illustration. A smartphone firm is similar to other technology companies and dependent on TSMC for fabrication. First a tariff (similarity channel) hits all technology firms, then a TSMC disruption (adjacency channel) hits downstream customers. Versus the reverse order: the TSMC shock first, then the tariff. The firm's accumulated position differs. In both orderings, the first event changed the environment in which the second arrived. Two firms that experienced the same two shocks in opposite sequences are genuinely different companies today, even if their balance sheets look identical. The commutator is the algebraic representation of this fact.

When $[L_S, L_A] = 0$, the manifold becomes locally flat, recovering the standard Euclidean geometry assumed by linear models. Commutativity is a special case.

5 The Propagation Modes: ϕ_k and μ_k

The spectral decomposition of M_s yields eigenpairs $\{(\phi_k, \mu_k)\}_{k=1}^n$:

$$M_s \phi_k = \mu_k \phi_k, \qquad 0 \leq \mu_1 \leq \mu_2 \leq \cdots \leq \mu_n, \tag{10}$$

with $\phi_k^\top \phi_j = \delta_{kj}$.

Interpretation. These eigenvectors are *traversal patterns*, not statistical factors. If a shock aligns with ϕ_k, it does not spread randomly—it travels along the supply-chain paths and similarity neighborhoods that define that pattern, maintaining spatial structure as it propagates:

$$x(t) = e^{-\mu_k t} \phi_k. \tag{11}$$

The eigenvalue μ_k is a *dissipation rate*, not a variance share. Small μ_k: the pattern propagates far before dissipating. Large μ_k: the pattern dissolves within a few hops.

Contrast with factor models. In PCA, eigenvectors are directions of maximum variance in return space—statistical abstractions summarizing historical co-movement. Here, eigenvectors are configurations of firms connected by supply dependencies and operational similarity through which a shock, once entering the network, will travel. They are derived from the structure of business relationships, not from return data.

Latent modes. Because propagation modes derive from network geometry (S, A) rather than return correlations, they exist whether currently active or dormant. A mode ϕ_k can be present structurally—as a supported propagation pathway—before any shock has activated it in observable returns. This is the mechanism by which the framework identifies vulnerabilities invisible in historical data.

The complete spectral basis satisfies:

$$M_s = \sum_{k=1}^{n} \mu_k \, \phi_k \phi_k^\top. \tag{12}$$

We retain the K smallest eigenvalues as dominant propagation modes, typically $K \approx 20$–40 for $n = 500$ stocks. These modes change only when S or A change—quarterly at most, not daily like statistical factors.

5.1 Spectral Risk Measures

The spectrum provides risk measures capturing propagation structure beyond variance.

Integrated spectral exposure. For portfolio weights w and horizon T:

$$I(w) = w^\top M_s^{-1} w = \sum_{k=1}^{n} \frac{(w^\top \phi_k)^2}{\mu_k}. \tag{13}$$

Slow modes (small μ_k) receive the largest weight, because disturbances entering those channels propagate far before dissipating. A portfolio concentrated in slow modes has high $I(w)$ relative to recent volatility—structurally vulnerable even if historical variance appears moderate.

Mode concentration. For each mode k:

$$C_k(w) = \frac{(w^\top \phi_k)^2}{\|w\|^2}, \tag{14}$$

the fraction of portfolio norm aligned with mode k. Slow-mode dominance indicates vulnerability. Aggregate slow-mode concentration:

$$C_{\text{slow}}(w) = \sum_{k:\, \mu_k < 0.2} \frac{(w^\top \phi_k)^2}{\|w\|^2}. \tag{15}$$

Effective dimension.

$$D_{\text{spec}} = \frac{\left(\sum_k \mu_k\right)^2}{\sum_k \mu_k^2}, \qquad H(M_s) = -\sum_{k=1}^{n} p_k \log p_k, \quad p_k = \frac{\mu_k}{\sum_j \mu_j}. \tag{16}$$

High entropy $H \to \log n$: flat spectrum, many modes contributing equally. Low entropy $H \to 0$: one or few modes dominating.

Worst-case scenario. Given portfolio w, the worst persistent shock is:

$$k^* = \arg\max_k \frac{|w^\top \phi_k|^2}{\mu_k}, \tag{17}$$

$$x^* = \phi_{k^*}, \qquad E_{\text{worst}} = \frac{|w^\top \phi_{k^*}|}{\sqrt{\mu_{k^*}}}. \tag{18}$$

This is structural vulnerability: a scenario determined by network geometry, not historical frequency. The worst-case shock is not a statistical abstraction—it is a specific configuration of firms, readable directly from the eigenvector.

6 The Structural Decomposition

Traditional risk models give you a single number: portfolio volatility $\sigma_w = 18\%$. This collapses a geometric reality into a scalar. Risk is not a number—it is a position in three-dimensional space, and vulnerability is a direction.

The structural decomposition separates every propagated portfolio response into three orthogonal mechanisms. Applied to the exposure $X(w;x) = w^\top Ex$:

$$X(w;x) = X_{\text{diff}} + X_{\text{ten}} + X_{\text{bot}}, \tag{19}$$

where the three components are mutually orthogonal and are constructed from the framework's own operators—the spectrum of M_s and the commutator $[L_S, L_A]$. The construction is inspired by the Hodge decomposition of network flows, but operates on firm-level response vectors with two non-commuting operators, a setting the classical machinery has no hold on. Their orthogonality follows from that construction, as developed in *Crisis Without History*.

Diffusion. The component of propagation whose outcome does not depend on the order in which the channels act: shared-exposure co-movement, the fuel shock reaching every airline simultaneously. This is the component conventional models approximate best, because it is the component a flat model can represent. Empirically, diffusion accounts for approximately 21% of propagation in the S&P 500.

Tension. Propagation that circulates: it arises exactly where the two channels conflict, concentrates in identifiable directions (computed from the commutator $[L_S, L_A]$), and is the geometric home of feedback—shocks that return by a different channel than they left on. A shock arriving via similarity-then-adjacency produces different exposure than adjacency-then-similarity; the tension component is precisely this difference. Tension accounts for approximately 71% of propagation, and is the dominant channel for systemic risk.

Bottleneck. Propagation that persists: the network's slowest patterns, concentrated where many paths converge on few nodes. TSMC's fabrication capacity serves automotive, consumer electronics, healthcare, and industrial equipment with no alternative source at scale. These exposures are topologically determined, independent of portfolio construction—

no rebalancing within an affected ecosystem can remove them. Empirically, bottleneck accounts for approximately 8% of propagation.

Two portfolios with identical total volatility can occupy completely different positions in mechanism-space, for example diffusion-dominant $(0.85, 0.10, 0.05)$ versus tension-dominant $(0.20, 0.65, 0.15)$. They face entirely different vulnerabilities. There exists, for each position, a worst-case direction: the specific pattern of disruptions that most efficiently excites the mechanisms the portfolio is exposed to.

7 Portfolio Exposure and Diagnostics

For portfolio weights w and shock scenario x, the structural exposure is:

$$X(w; x) = w^\top E\, x. \tag{20}$$

This scalar measures how much of shock x loads onto portfolio w through economic structure—through all similarity diffusion and adjacency cascades. Decomposing by dominant propagation modes:

$$X(w; x) = \sum_{k=1}^{K} \underbrace{(w^\top \phi_k)}_{\text{mode loading}} \cdot \underbrace{(\phi_k^\top E\, x)}_{\text{shock activation}} + w^\top E_{\text{residual}}\, x. \tag{21}$$

The exposure measure supports scenario comparison (cobalt disruption vs. semiconductor shortage vs. energy shock), stress testing, bottleneck identification, and hedging decisions.

7.1 Directional Diagnostics

Normalize shocks to the unit sphere $\mathbb{S}^{n-1}$. The portfolio's directional exposure is $X_w(\omega) = w^\top E\, \omega$. The relevant object is the mode-exposure distribution:

$$\ell_k(w) = w^\top E\, \phi_k, \qquad \pi_k(w) = \frac{\ell_k(w)^2}{\sum_j \ell_j(w)^2}. \tag{22}$$

Three diagnostics follow.

Effective structural directions.

$$N_{\text{eff}}(w) = \exp\!\left(-\textstyle\sum_k \pi_k \log \pi_k\right) \in [1, n]. \tag{23}$$

A portfolio with $N_{\text{eff}} = 2$ behaves, under propagation, as if only two independent shock types matter—however many names it holds.

Fragility. The concentration multiple $\text{Fragility}(w) = n / N_{\text{eff}}(w) \geq 1$.

Worst-case direction.

$$x^*(w) = \frac{E^\top w}{\|E^\top w\|}. \tag{24}$$

The largest components *name the firms*: the worst case is a list of nodes, readable off the vector. For a supply-chain-concentrated book, x^* points at the book's own suppliers.

7.2 Additional Diagnostics

Exposure amplification from tension:

$$A_p(w;x) = \frac{X_{\text{ten}}(w;x)}{X_{\text{diff}}(w;x)}. \tag{25}$$

Quantifies how much of the portfolio's response travels through the order-dependent tension channel relative to simple diffusion.

Bottleneck dependency:

$$B_s(w) = \sum_j A_{sj}\, w_j. \tag{26}$$

Portfolio-weighted reliance of holdings on supplier s: identifies single points of failure in the supply network.

Eigenvalue perturbation (structural stability):

$$\dot{\mu}_k = \phi_k^\top \dot{M}_s\, \phi_k. \tag{27}$$

The Hellmann–Feynman relation from quantum mechanics applied to the symmetric core. Eigenvector rotation is bounded by the Davis–Kahan theorem:

$$\sin\theta_k \lesssim \frac{\|\Delta M_s\|}{\Delta\mu_k}, \tag{28}$$

where $\Delta\mu_k = \min_{j\neq k} |\mu_k - \mu_j|$ is the spectral gap. For the S&P 500, eigenvalue changes are approximately 0.01–0.02 for slow modes per quarter—about 10–20% relative change. Well-separated modes are structurally stable.

8 Why Factor Models Fail: A Structural Account

Traditional multi-factor models take the form:

$$r = Bf + \varepsilon, \qquad \Sigma = B\,\mathrm{Cov}(f)\,B^\top + D, \tag{29}$$

where B is a loading matrix, f are factors, and D is a diagonal residual. When the economic surface is curved—when $[L_S, L_A] \neq 0$—this structure fails in three computable, non-calibration ways.

1. Structural covariance conflation. The covariance observed by a factor model conflates the similarity and adjacency channels:

$$\Sigma_{\text{structural}} = \alpha^2\, G_S + \beta^2\, G_A + \alpha\beta\,(G_{SA} + G_{SA}^\top), \tag{30}$$

where $G_S = \Phi_S\Phi_S^\top$ and $G_A = \Phi_A\Phi_A^\top$ are the separate channel covariances and G_{SA} is the cross-term. Factor models cannot separate these; they fit a single B to what is in fact a structured sum.

2. Loading drift. Factor loadings B drift at a rate proportional to curvature:

$$\dot{\beta} \;\sim\; \|[L_S, L_A]\| \cdot v_{\text{structural}}, \tag{31}$$

where $v_{\text{structural}}$ is the portfolio's velocity through structural space. This is geometric necessity, not parameter instability. Loading drift is predictable in magnitude from observable structure; it is not estimable from return history, because return history sees only the projection of the drift onto market prices.

3. Causal misspecification. Adjacency creates confounders and colliders that factor models systematically misread. A common upstream supplier (confounder) induces spurious correlation between two otherwise independent downstream firms. A common downstream customer (collider) induces conditional independence where factor models predict dependence. These are structural features of directed networks; they are invisible to symmetric correlation matrices.

Non-additivity. Factor models assume $\sigma^2(w) = w^\top B\,\mathrm{Cov}(f)\,B^\top w + w^\top D\,w$. Variance decomposes additively into factor risk plus diversifiable residual. But propagation makes total risk larger than the no-cascade baseline, and the interaction terms generated by the commutators $[L_S, L_A]$ do not separate into orthogonal channel components. Risk is fundamentally non-additive on curved manifolds.

These failures are not calibration problems or estimation failures. They are consequences of projecting a non-integrable manifold onto a flat factor space. The framework makes the geometry explicit, which means the failures become *predictable*: holonomy is detectable in closed structural loops, loading drift is bounded by tension strength, and the gap between factor coverage and actual risk is measurable in the structural inner product.

A Appendix

Here are the core theorems underlying the framework, developed in full in *Crisis Without History* (Bookstaber, forthcoming). Proofs rely on standard results in spectral graph theory (Chung 1997), differential geometry (Lee 2018), causal inference (Pearl 2009), and matrix perturbation theory (Kato 1995).

A.1 Geometric Obstruction

Lemma A.1 (Macro Closure Characterization)**.** *Let $L = L_S + L_A$ and let P be a macro-projector. Macro closure holds if and only if $[P, L] = 0$.*

Theorem A.2 (Geometric Obstruction)**.** *Let $(\mathcal{M}, g)$ have generic nonzero curvature. Let P be a macro-projector and define the horizontal distribution $\mathcal{H}$ by $T_x\mathcal{H} = \ker(dP_x)$. Assume $\mathcal{H}$ is non-integrable. Then:*

1. ***Path Dependence.*** *Parallel transport along closed loops is non-trivial: the structural state*

after a portfolio's trajectory depends on the path taken, not just the endpoint. Holonomy provides a geometric mechanism for non-ergodicity. (Ambrose–Singer theorem.)

2. ***Emergence.*** *The macro-projector P does not commute with L: a self-contained macro-dynamics exists iff $[P, L] = 0$, and non-integrability of $\mathcal{H}$ guarantees this fails. The same macro-state can evolve differently depending on micro-detail.* (Frobenius theorem.)

3. ***Local Irreducibility.*** *Nonzero curvature destroys the flat separation law. Where curvature is negative ($K \leq -\kappa^2 < 0$), nearby trajectories diverge exponentially:*

$$d(\gamma(t), \tilde{\gamma}(t)) \gtrsim \frac{e^{\kappa t}}{\kappa} |\dot{\gamma}(0) - \dot{\tilde{\gamma}}(0)|,$$

and prediction degrades exponentially in the horizon. Where curvature is positive, trajectories focus at conjugate points: distinct initial conditions become locally indistinguishable. (Jacobi field estimates; Cartan–Hadamard comparison.)

Corollary A.3 (Necessity of Adaptation)**.** *Under the conditions of Theorem A.2, no static policy $\pi(x)$ remains optimal as $T \to \infty$. Optimization must be replaced by adaptive navigation $\pi(x, t)$ responsive to local curvature and structural change.*

Remark A.4 (Radical Uncertainty). A fourth consequence—radical uncertainty—follows from the open-endedness of $\mathcal{M}$ rather than its curvature. As new firms enter and dependency patterns shift, no fixed finite-dimensional model captures all admissible future deformations. This is the geometric form of Knightian uncertainty: futures that cannot be enumerated, not merely probability-weighted.

A.2 The Structural Filter Theorem

Definition A.5 (Regime and Spectral Parameters)**.** The *asymmetry ratio* is $\varepsilon = \|M_a\|_2 / \lambda_{\min}(M_s)$. The operator M is in the *weakly asymmetric regime* when $\varepsilon < 1$. Write $\mu_{\min} = \mu_1 = \lambda_{\min}(M_s)$ and $\mu_{\tau^+} = \min\{\mu_k : \mu_k > \tau\}$ for the slowest mode outside the bottleneck block.

Definition A.6 (Portfolio Exposure Parameters)**.** For portfolio w, the *bottleneck concentration* and *tension alignment* are:

$$C_{\text{bot}}(w) = \frac{\|\Pi_{\text{bot}}\, w\|^2}{\|w\|^2}, \qquad f_{\text{ten}}(w) = \frac{\|\Pi_{\text{ten}}\, w\|^2}{\|w\|^2}.$$

Theorem A.7 (Structural Filter Theorem)**.** *Assume the weakly asymmetric regime, $\varepsilon < 1$. Then the integrated structural exposure $I(w) = w^\top M^{-1} w$ satisfies:*

$$I(w) \leq \left(\frac{C_{\text{bot}}(w)}{\mu_{\min}} + \frac{1 - C_{\text{bot}}(w)}{\mu_{\tau^+}} \right) \|w\|^2 + R(\varepsilon)\, \|w\|^2, \qquad R(\varepsilon) = \frac{\varepsilon^2}{(1 - \varepsilon)\, \lambda_{\min}(M_s)}.$$

Consequently, for any target $I^* > 0$, the filter condition

$$\frac{C_{\text{bot}}(w)}{\mu_{\min}} + \frac{1 - C_{\text{bot}}(w)}{\mu_{\tau^+}} \le \frac{I^*}{\|w\|^2}$$

implies $I(w) \le I^* + R(\varepsilon)\|w\|^2$, with $R(\varepsilon) \to 0$ as $\varepsilon \to 0$.

Corollary A.8 (Necessity of the Bottleneck Condition)**.** *The bottleneck term cannot be dropped: along any family of networks with $\mu_1 \to 0$, portfolios $w = \phi_1$ have $C_{\text{bot}} = 1$ and $I(w) \to \infty$. A portfolio concentrated in a slow structural mode has unbounded integrated exposure regardless of its recent return history.*

Remark A.9 (Slow Modes and Structural Risk). The term $C_{\text{bot}}(w)/\mu_{\min}$ is the formal expression of the book's central empirical claim: long-horizon forces—climate, demographics, geopolitical realignment—appear in M_s as slow modes with small eigenvalues, concentrated at structural bottlenecks that Cheeger's inequality ties to near-partitions of the network (Chung 1997). They are present in S and A before any shock activates them in return history. A portfolio with high C_{bot} is exposed regardless of what recent variance suggests.

B A Worked Example: The DRC Critical Minerals Network

To make the framework concrete, we apply it to ten companies connected through the Democratic Republic of Congo's critical minerals supply chains. All matrices are fully visible; every number follows from the arithmetic. Numerical values are illustrative, calibrated to reflect the qualitative structure of these supply relationships; the purpose is to make the arithmetic visible, not to characterize any firm's actual exposure.

The Network

1. Glencore — cobalt mining, DRC
2. Yageo/KEMET — tantalum capacitors (coltan exposure), DRC
3. Umicore — refining and battery chemicals
4. CATL — battery manufacturing
5. Carpenter Technology — specialty alloys
6. Tesla — electric vehicles
7. Medtronic — medical devices
8. Lockheed Martin — defense and aerospace
9. Stryker — orthopedic implants
10. Coca-Cola — *control firm*: no links to the DRC complex

Coca-Cola is the isolation test throughout. It has no supply dependency on and no operational similarity to any firm in the network. It should remain structurally unaffected by any DRC shock—not by assumption, but as a result the framework must verify.

The Adjacency Matrix A

Entry A_{ij} is the fraction of firm j's critical inputs sourced from firm i. Three propagation paths run from DRC sources: the cobalt path (Glencore $\to$ Umicore $\to$ CATL $\to$ Tesla/Medtronic/Lockheed), the tantalum path (Yageo $\to$ Umicore, and Yageo $\to$ Lockheed directly), and the specialty materials path (Umicore $\to$ Carpenter $\to$ Lockheed/Stryker).

	Glen	Yag	Umic	CATL	Carp	Tesl	Medt	Lock	Stry	Coke
Glencore	0	0	0	0	0	0	0	0	0	0
Yageo	0	0	0	0	0	0	0	0	0	0
Umicore	.40	.15	0	0	0	0	0	0	0	0
CATL	0	0	.35	0	0	0	0	0	0	0
Carpenter	0	0	.25	0	0	0	0	0	0	0
Tesla	0	0	0	.18	0	0	0	0	0	0
Medtronic	0	0	0	.10	0	0	0	0	0	0
Lockheed	0	.05	0	.06	.08	0	0	0	0	0
Stryker	0	0	0	0	.12	0	0	0	0	0
Coca-Cola	0	0	0	0	0	0	0	0	0	0

Table 3: Adjacency matrix A (supplier as row, dependent firm as column).

The Similarity Matrix S

S_{ij} encodes operational proximity (symmetric, unit diagonal). Three entries deserve attention. **Glencore–Yageo (0.65)** captures DRC infrastructure co-exposure: both mine different minerals but share the same transport network, power infrastructure, and political environment in eastern DRC. **Medtronic–Stryker (0.70)** reflects shared medical device regulation, materials certification, and hospital customer base. **CATL–Tesla (0.40)** reflects the battery-EV ecosystem: shared technology platforms and regulatory exposure. The entire Coca-Cola column is zero.

Results

With $\eta = 0.30$, $\alpha = 0.6$, $\beta = 0.4$, the resolvent $P_\eta = (I - 0.3A^\top)^{-1}$ sums all multi-hop cascades, attenuating by 0.3 per hop. A cobalt shock at Glencore (unit disturbance) propagates: 0.40 to Umicore at one hop, then $0.35 \times 0.40 = 0.14$ to CATL at two hops, then $0.18 \times 0.14 = 0.025$ to Tesla at three hops. Coca-Cola receives exactly zero at every hop.

The dominant slow mode ϕ_1 has large positive weights on the cobalt complex (Glencore, Umicore, CATL, Tesla) and zero on Coca-Cola, confirming this is the systemic DRC channel. A second mode ϕ_2 separates the cobalt and tantalum paths. Coca-Cola appears in neither.

A portfolio equal-weighted across the nine connected firms has structural exposure approximately 1.12 times that of a portfolio equal-weighted only across the cobalt path—the

	Glen	Yag	Umic	CATL	Carp	Tesl	Medt	Lock	Stry	Coke
Glencore	1.00	.65	.20	.10	.15	.05	.05	.10	.05	0
Yageo	.65	1.00	.15	.08	.10	.05	.05	.12	.05	0
Umicore	.20	.15	1.00	.25	.30	.10	.10	.15	.10	0
CATL	.10	.08	.25	1.00	.15	.40	.20	.25	.15	0
Carpenter	.15	.10	.30	.15	1.00	.10	.20	.25	.30	0
Tesla	.05	.05	.10	.40	.10	1.00	.15	.35	.10	0
Medtronic	.05	.05	.10	.20	.20	.15	1.00	.20	.70	0
Lockheed	.10	.12	.15	.25	.25	.35	.20	1.00	.20	0
Stryker	.05	.05	.10	.15	.30	.10	.70	.20	1.00	0
Coca-Cola	0	0	0	0	0	0	0	0	0	1.00

Table 4: Similarity matrix S (symmetric, unit diagonal).

tantalum and specialty materials channels add 12% additional exposure through shared infrastructure similarity (the Glencore–Yageo link), even though they share no direct supply dependency. A portfolio including Coca-Cola at 10% weight shows the exposure reduced exactly in proportion: the control firm contributes zero to every propagation mode.

This last point illustrates the framework's diagnostic value. The Glencore–Yageo similarity link (0.65) captures DRC infrastructure co-exposure that any sector classification would miss: both firms would be assigned different industry codes, yet they share the same transport network, power infrastructure, and political environment. A geopolitical shock to that infrastructure propagates to both through L_S even though they have no direct supply relationship in A. The structural model sees it; a factor model calibrated to return history would not, until the shock activated it in prices.

Note on the Book

Crisis Without History (Bookstaber, forthcoming) develops the tools used in the application and summarized here, with additional chapters on structural decomposition, operator interactions, adaptive dynamics, tracking structural flows, regime transitions, and structural portfolio design.

References

Bookstaber, R. (forthcoming). *Crisis Without History*.

Chung, F. R. K. (1997). *Spectral Graph Theory*. American Mathematical Society.

Hoberg, G. and Phillips, G. (2016). Text-based network industries and endogenous product differentiation. *Journal of Political Economy* 124(5), 1423–1465.

Kato, T. (1995). *Perturbation Theory for Linear Operators*. Springer.

Lee, J. M. (2018). *Introduction to Riemannian Manifolds*. Springer.

Pearl, J. (2009). *Causality*. Cambridge University Press.