
Quantum Fragility Index (QFI) ¹

Bob Prieto

This paper focuses on Fragility in Large Complex Projects (LCPs) and builds on the body of work on Quantum Project Management (QPM)², that encompasses not only its theoretical framework but also the various properties and associated metrics that enable its implementation. Existing project controls excel at measuring variance, probability, and exposure, but they do not measure how close a project is to a state transition. Traditional risk registers, Monte-Carlo simulations, and performance dashboards quantify dispersion around expected outcomes, yet they remain blind to the system’s susceptibility to crossing governance boundaries when coherence degrades. What they measure is risk; what they miss is fragility. The Quantum Fragility Index (QFI) fills this gap by providing a coherence-sensitivity metric—one that quantifies how small routine perturbations can trigger a shift into a higher-risk or lower-performance state.

Within the QPM framework, Fragility sits implicitly inside several quantum-analog constructs already (entanglement, decoherence, half-life, propagation velocity) developed or implicitly implied. This paper adopts the QPM framework and:

1. Defines Quantum Fragility and where it already lives implicitly
2. Describes what a formal Quantum Fragility metric would look like
3. Discusses the relationship to fatigue risk³
4. Describes differences with the statistical fragility index

¹ How to cite this paper Prieto, R. (2026) Quantum Fragility Index (QFI), PM World Journal, Vol. XIV, Issue VII, July

² Prieto, R. (2024). Quantum Project Management, PM World Journal, Vol. XII, Issue I, January 2024.

Prieto, R. (2024). Measurement of Complexity in Large Complex Projects, PM World Journal, Vol. XII, Issue IV, April

Prieto, R. (2025). Artificial Intelligence, Complexity, and Quantum Project Management: A Transformative Approach, PM World Journal, Vol. XIV, Issue VII, July

Prieto, R. (2024). Quantum Project Management and the Concept of Spacetime, PM World Journal, Vol. XII, Issue V, May.

Prieto, R. (2025). Managing Uncertainty in Large Complex Projects, PM World Journal, Vol. XIV, Issue XI, November.

Prieto, R. (2025). Metrics for Assumption Management in Large Complex Projects, PM World Journal, Vol. XIV, Issue XII, December.

Prieto, R. (2026). Operationalizing Quantum Project Management: Defining Improved Metrics for Management of Large Complex Projects, PM World Journal, Vol. XV, Issue I, January.

Prieto, R. (2026). Operationalizing Quantum Project Management: Anticipating and Managing Fragility in Large Complex Project Ecosystems, PM World Journal, Vol. XV, Issue II, February.

³ Prieto, R. (2026). Fatigue Risk Index: Measuring an “Internality” as a Performance Precursor Under Quantum Project Management (QPM) Theory, PM World Journal, Vol. XIV, Issue VI, June

The paper is supported by extensive appendices further developing or supporting specific areas covered in the main paper. Attention is called to Appendix A, which provides an inventory of the formulas used and a glossary of notations used.

1. Quantum Fragility in QPM

Quantum Fragility represents the degree to which a project’s operational state is vulnerable to small, routine variations in its underlying conditions. In Quantum Project Management, every project exists as a dynamic quantum system—continuously influenced by fluctuations in workload, environmental stressors, staffing patterns, interface quality, and human readiness. These fluctuations act as *perturbations* to the system. Quantum Fragility quantifies how large or small a perturbation must be before the project transitions from one governance-relevant state to another.

A **state transition** occurs whenever the project crosses a defined boundary that carries operational, safety, or governance significance. These boundaries may be expressed as risk tiers, readiness bands, fatigue thresholds, schedule viability classifications, or any other state-based construct used to guide decision-making. When a project crosses such a boundary, its behavior, risk profile, and required controls change in meaningful ways. Quantum Fragility measures *how close* the system is to such a transition and *how easily* that transition can be triggered.

In this sense, Quantum Fragility is not merely a sensitivity metric—it is a measure of **brittleness**. A system with low fragility can absorb variation without meaningful change in state; it is resilient, buffered, and stable. A system with high fragility is perched near a tipping point, where even minor deviations in fatigue, schedule pressure, environmental load, or interface quality can cause a rapid collapse of coherence and a shift into a higher-risk or lower-performance state.

Formally, Quantum Fragility is defined as the **minimum normalized perturbation required to induce a state transition across a defined governance boundary**. “Normalized” means the perturbation is evaluated relative to the maximum credible variation expected under the scenario—ensuring the metric is comparable across contexts, drivers, and time horizons. By focusing on the *minimum* perturbation, the metric captures the most vulnerable direction of change, revealing the weakest link in the system’s stability.

This framing aligns with the core principles of Quantum Project Management:



Figure 1.1 QPM Construct Alignment — Fragility mapped to Decoherence, Entanglement, Superposition, and Propagation

- **Decoherence:** Fragility reflects how easily a coherent operational state collapses under noise. *Fragility = high decoherence rate.*

Quantum Fragility is the minimum normalized perturbation required to induce a state transition across a defined governance boundary.

- **Entanglement:** Fragility reveals where tightly coupled tasks or teams amplify small disturbances. This is where tasks, stakeholders, or assumptions lose correlation. *Fragility = low entanglement stability.*
- **Superposition:** Fragility indicates how readily the system collapses into an undesirable outcome when perturbed. It reflects **how narrow the stability basin is** around the current operational state. A fragile system collapses disproportionately into high-risk or low-performance states with minimal disturbance. *Fragility = high collapse susceptibility of the project's state superposition.*
- **Propagation:** Fragility identifies where local variation can propagate rapidly across the system. It measures how much a small perturbation is amplified as it propagates through the project network. *Fragility = high amplification.*

By quantifying the distance between the current operational state and the nearest governance boundary, Quantum Fragility provides leaders with a clear, interpretable measure of how close the

project is to a tipping point. It transforms abstract risk into a concrete, measurable, and actionable signal—one that can be monitored, trended, and used to trigger interventions before state collapse occurs.

In practice, Quantum Fragility becomes a **governance-grade early warning indicator**, revealing when everyday variability is no longer benign and when the system is operating in a region where small shocks can produce disproportionate consequences. It is the metric that tells leaders not just *where the project is*, but *how easily it can fail*.

2. Formal Quantum Fragility Metric

This section converts the conceptual vocabulary of Quantum Fragility — decoherence, entanglement, superposition, propagation — into a single, computable scalar and the governance architecture required to act on it. *Quantum Fragility represents the degree to which a project's operational state is vulnerable to small, routine variations in its underlying conditions.*

Formally, Quantum Fragility is defined as the minimum normalized perturbation required to induce a state transition across a defined governance boundary.

In this section we focus on measuring, modeling, and governing nonlinear vulnerability in Large Complex Projects. In many ways this section provides the technical foundations for a Quantum Fragility Index (QFI).

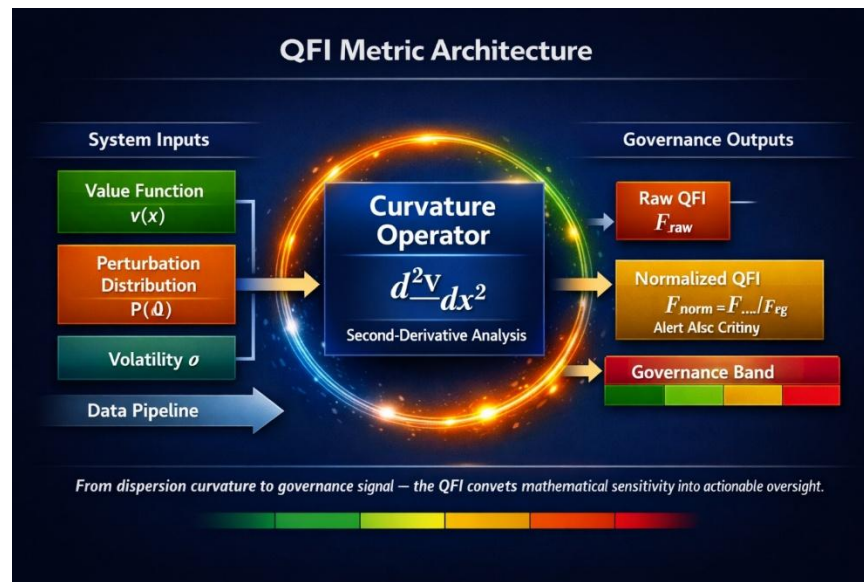


Figure 2.1 QFI Conceptual Architecture

QFI at a Glance

The Quantum Fragility Index (QFI) is a governance-grade metric that quantifies how easily a project's operational state can be pushed across a meaningful boundary by small, routine variations. It measures coherence sensitivity — not variance, not probability, but the curvature of the value function under dispersion.

What QFI Measures

- The minimum normalized perturbation required to trigger a state transition.
- How rapidly coherence decays under environmental coupling (decoherence rate Γ).
- How tightly project dimensions amplify disturbances (entanglement ρ_{ij}).
- Whether dispersion harms or helps the system (concave = fragile; convex = antifragile).

Why QFI Matters

- Traditional controls measure variance and risk; QFI measures *state-transition susceptibility*.
- It reveals brittleness long before performance KPIs deteriorate.
- It provides a single scalar that integrates schedule, cost, scope, stakeholder, supply-chain, and regulatory fragility.

How QFI Is Computed

- **Curvature method:** $QFI = -\partial^2 E[V]/\partial \sigma^2$ evaluated at current dispersion σ .
- **Decoherence proxy:** $QFI \approx \eta \cdot \Gamma \cdot (1 - FQ(t))^{\alpha-1}$, enabling telemetry-driven monitoring
- Both converge when fragility is driven by coherence loss.

How to Interpret QFI

- **QFI < 0** → Antifragile (convex response).
- **0–0.30** → Robust.
- **0.30–0.60** → Moderately fragile.
- **0.60–0.85** → Highly fragile.
- **≥ 0.85** → Critically fragile (tail-risk regime).

Governance Use

- Supports early warning, escalation, and intervention prioritization.
- Integrates with Safety-II and assumption-management protocols.
- Enables continuous monitoring via telemetry and periodic scenario analysis.

2.1 Prologue

This section establishes the technical foundations for a Quantum Fragility Index (QFI) and explicitly ties that metric to the Quantum Project Management (QPM) ontology introduced in Section 1. The QFI is not a standalone statistic; it is a QPM-native diagnostic that translates quantum-analogue constructs (decoherence, entanglement, superposition, propagation) into governance-grade observables. The objective is to move governance from event-driven risk lists to a continuous, system-level measure of coherence loss and nonlinear vulnerability that leaders can monitor, trend, and act upon.

Why this matters for QPM

QPM treats projects as dynamic systems whose operational coherence is the precondition for predictable delivery. The QFI operationalizes that coherence: it quantifies how small environmental couplings (perturbations) change the probability of crossing governance boundaries.

The QFI complements conventional risk registers by measuring shape (convexity/concavity) of the value function rather than only expected loss. This makes it uniquely suited to detect latent brittleness that conventional risk metrics miss.

Practical objective

Provide a single, interpretable scalar (and a decomposed vector) that: (1) maps directly to QPM constructs, (2) supports directional diagnosis by driver, and (3) integrates with Safety II governance to prioritize interventions that preserve coherence.

The remainder of this section formalizes these concepts mathematically, beginning with the distinction between risk and fragility and culminating in a fully operational QFI suitable for governance deployment.

2.2 Conceptual Foundations

Perturbation (x) refers to the realized deviation applied to the project's operating state — the actual shock, variation, or disturbance introduced into the system. **Dispersion (σ)** represents the width of the perturbation distribution, capturing how broad or volatile the underlying environment is and determining how frequently larger perturbations occur. A **governance boundary** is the threshold at which a change in the project's state becomes operationally or strategically meaningful, such as a shift in risk tier, readiness band, fatigue threshold, or schedule viability classification. Together, these three elements define the mechanics of state-transition susceptibility, which the Quantum Fragility Index measures directly.

Fragility versus risk in QPM terms

Risk (conventional): an expectation operator over discrete events; additive across identified items.

Fragility (QPM): a coherence sensitivity operator — it measures how the expected project value changes as the dispersion of environmental couplings increases. Fragility is therefore a second-order property of the value function: it is about curvature, not level. This aligns with Section 1's definition of Quantum Fragility as the minimum normalized perturbation required to induce a state transition across a governance boundary.

2.2.1 QPM mapping

Decoherence: the process by which aligned assumptions, schedules, and interfaces lose phase; measured by a decoherence rate Γ . High Γ increases QFI.

Entanglement: the degree to which tasks, teams, and interfaces are tightly coupled; measured by correlation metrics ρ_{ij} . Fragility rises when entanglement amplifies perturbations.

Superposition: the project's simultaneous potential outcomes; fragility measures how readily the system collapses into undesirable outcomes under observation or perturbation.

Propagation velocity: the network speed at which local perturbations spread; high propagation velocity multiplies local fragility into systemic fragility.

Governance implication

Fragility is actionable: it identifies where and how small routine variations will produce disproportionate consequences, enabling preemptive redesign or control hardening.

QPM Constructs Mapped to QFI Components			
QPM Construct	Meaning in QPM	QFI Interpretation	Operational Observable
Decoherence	Loss of alignment among assumptions, interfaces, schedules	Fragility increases as coherence decays	Decoherence rate Γ , fidelity decay FQ(t)
Entanglement	Coupling strength between tasks, teams, or assumptions	Amplifies perturbations; increases systemic fragility	Correlation matrix ρ_{ij} , coupling amplification term C
Superposition	Multiple potential project outcomes coexisting	Fragility = collapse susceptibility into undesirable states	Curvature sign of value function (concave/convex)

QPM Constructs Mapped to QFI Components			
QPM Construct	Meaning in QPM	QFI Interpretation	Operational Observable
Propagation	Speed at which local disturbances spread	Converts local fragility into system-wide fragility	Propagation velocity, network amplification factors

2.2.2 The Quantum Analogy — Decoherence as Fragility

The QPM Framework draws its organizing metaphor from quantum mechanics, and the phenomenon of quantum decoherence provides a particularly rigorous and instructive analogy for project fragility.

In quantum mechanics, a system can exist in a coherent superposition — a state representing multiple possible outcomes simultaneously, with well-defined phase relationships between them. This coherence is the source of quantum systems' extraordinary computational and physical properties. However, when a quantum system interacts with its environment — through photon scattering, thermal fluctuations, or any form of environmental coupling — the coherence is progressively destroyed. The system undergoes decoherence: its quantum state "decays" from a pure state (fully coherent) to a mixed state (partially or fully incoherent), with the off-diagonal elements of the system's density matrix decaying exponentially at a characteristic rate Γ , known as the decoherence rate. The system does not necessarily "fail" — it simply loses the coordinated alignment of its internal states, and with it, the ability to maintain the delicate conditions required for optimal performance.

The parallel to project management is direct. A project begins in a "coherent" state: scope is aligned, resources are committed, stakeholders share a common understanding, the schedule reflects an integrated plan, and supply chain commitments are synchronized. This coherence — the mutual alignment of all project dimensions — is the precondition for successful delivery. As the project interacts with its environment — regulatory changes, supply chain disruptions, personnel turnover, market shifts, scope creep, stakeholder conflicts — this coherence progressively degrades. Requirements drift from design. Stakeholder expectations diverge from deliverables. Schedule dependencies decouple from resource availability. The project's state becomes "mixed," with increasing uncertainty about which of many possible outcomes will materialize.

The rate of this coherence loss — the project's decoherence rate — is its fragility. A project with many environmental couplings (dependencies on volatile external factors) and a highly nonlinear value function (where small losses of coherence produce large value drops) is maximally fragile. A project with few environmental couplings or a linear value function is robust. A project that can harness environmental perturbation to improve its state — through adaptive scope management, opportunistic resource reallocation, or real-options decision-making — is antifragile.

This analogy becomes operational in Section 2.3, where the decoherence rate Γ and fidelity decay $F_Q(t)$ form the basis of a telemetry-driven QFI proxy.

2.2.3 Formal Definition of Quantum Fragility

Definition (QPM form)

Let $V(S)$ be the project value function defined over the project state vector S . Let $f(x; \sigma)$ be a family of perturbation distributions parameterized by dispersion σ .

The **Quantum Fragility Index** is:

$$\mathbf{QFI}(S) = -\partial^2 E[V(S; x)] / \partial \sigma^2 |_{\sigma=\sigma_0}$$

Where:

$V(S; x)$ is the value under perturbation x .

σ_0 is the current dispersion estimate for the operating environment.

QFI > 0 indicates fragility (concave payoff), **QFI** \approx 0 indicates robustness, and **QFI** < 0 indicates antifragility (convex payoff).

Normalized, governance-ready scalar

To present a bounded, board-readable metric, normalize QFI to $[-1, 1]$ or $[0, 1]$ depending on governance preference. A practical normalized form is:

$$\mathbf{QFI}_{\text{norm}} = (\mathbf{QFI} - \mathbf{QFI}_{\text{min}}) / (\mathbf{QFI}_{\text{max}} - \mathbf{QFI}_{\text{min}})$$

with calibration anchors **QFI**_{min}, **QFI**_{max} derived from historical portfolio data or scenario stress tests.

Interpretation guidance

Sign: positive \rightarrow fragility; negative \rightarrow antifragility.

Magnitude: larger absolute values indicate stronger curvature and therefore faster deterioration (or improvement) as dispersion changes.

Appendix D provides visualizations of both the stochastic roughness field (KPZ surface) and the curvature geometry underlying the QFI operator.

To operationalize these concepts, we now express fragility as a curvature operator on the project's value function under dispersion. This provides a mathematically rigorous and governance-ready scalar.

2.3 Mathematical Foundations and Decoherence-Mapped Formulation

2.3.1 Taleb–Douady transfer function (operational)

Start from the expected payoff:

$$E[V] = \int v(x) \cdot f(x; \sigma) dx$$

and compute the second derivative with respect to σ :

$$QFI = -\partial^2 / \partial \sigma^2 \int v(x) \cdot f(x; \sigma) dx$$

This is the canonical fragility operator. In practice, compute via Monte Carlo perturbation sweeps: evaluate $E[V]$ at $\sigma_0 \pm \Delta\sigma$ and estimate the curvature numerically when closed-form integration is infeasible.

Decoherence-mapped computational pathway

Define a project density matrix $\rho(t)$ that encodes both determinate and indeterminate project attributes. Let ρ_0 be the baseline (coherent) density matrix.

Define quantum fidelity:

$$F_Q(t) = \text{Tr}[\sqrt{\sqrt{\rho_0} \cdot \rho(t) \cdot \sqrt{\rho_0}}]$$

Assume exponential fidelity decay under Markovian coupling:

$$F_Q(t) = e^{-\Gamma t}$$

Define a nonlinearity coefficient η that maps fidelity loss to value loss:

$$\Delta V \approx \eta \cdot (1 - F_Q(t))^\alpha$$

where $\alpha \geq 1$ captures higher-order nonlinearity. Then a decoherence-mapped QFI proxy is:

$$QFI_{\text{deco}} = \eta \cdot \Gamma \cdot (1 - F_Q(t))^{\alpha-1}$$

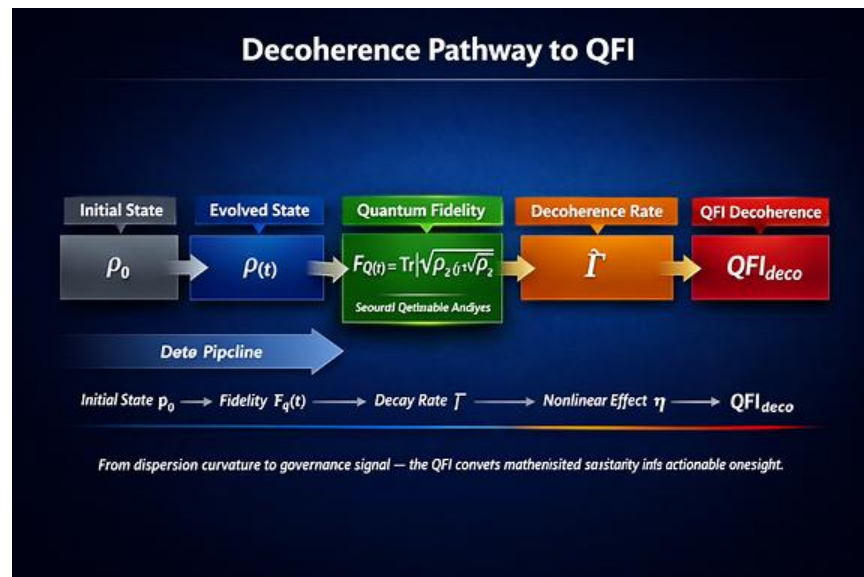


Figure 2.2 Decoherence-Mapped QFI Proxy Computational Pathway

Practical computation notes

Estimate Γ from observed rates of divergence between plan and execution (e.g., cumulative variance in key observables per unit time).

Estimate η by fitting small-perturbation experiments or historical stress episodes to the ΔV model.

Use the decoherence proxy for rapid, telemetry-driven monitoring; use the full Taleb–Douady transfer function for periodic, deeper scenario analysis.

2.3.2 Component Decomposition and Threshold Calibration

Additive decomposition with correlation adjustment

Compute dimensional fragilities QFI_i for each governance dimension $i \in \{\text{Schedule, Cost, Scope, Stakeholder, Supply Chain, Regulatory}\}$ using the same second-derivative operator applied to dimension-specific value slices. Aggregate with governance weights w_i and a correlation correction term C :

$$QFI_{\text{composite}} = \sum_i w_i \cdot QFI_i + C$$

where

$$C = \sum_{i \neq j} \rho_{ij} \cdot \sqrt{(w_i \cdot w_j)} \cdot \varphi(QFI_i, QFI_j)$$

and ρ_{ij} is the empirical correlation between dimensions i and j ; $\varphi(\cdot, \cdot)$ is a symmetric amplification function (e.g., product or max) that captures nonlinear coupling. This term ensures that tightly entangled dimensions (high ρ_{ij}) can push the composite QFI above the simple weighted average.

Governance bands and rationale

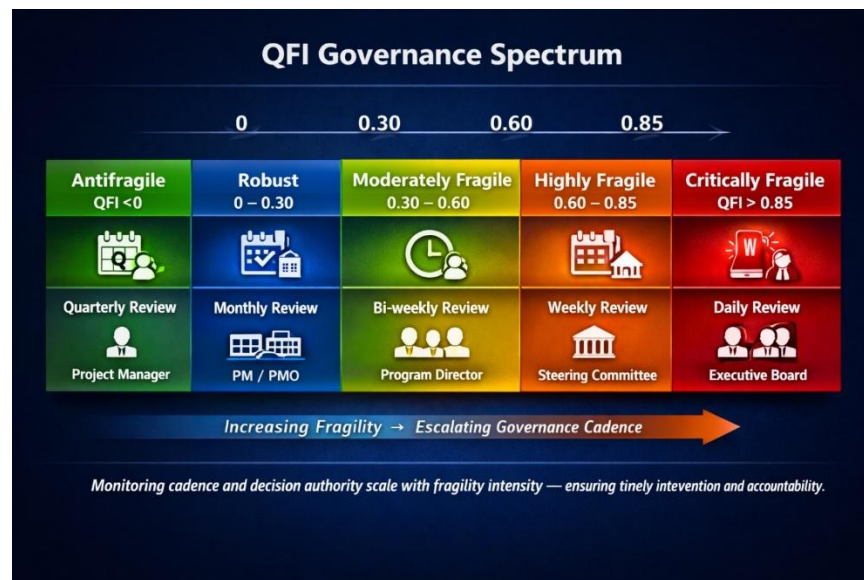


Figure 2.3 Five-Tier Fragility Spectrum — Governance Bands

The QFI governance framework adopts a five-tier fragility spectrum, each tier calibrated to portfolio experience and mapped to a color code, monitoring cadence, and decision-authority level. These assignments are canonical throughout this paper and its appendices; the full governance protocol — including band-transition checklists, RACI assignments, and floor-rule overrides — is detailed in Appendix C, Table C.1.

Band	Color	QFI_norm Range	Governance Posture	Monitoring Cadence	Decision Authority
Antifragile	Green	$QFI_norm < 0$	Harvest optionality; document sources of convexity for replication across portfolio	Quarterly review	Project Manager
Robust	Blue	$0 \leq QFI_norm < 0.30$	Standard oversight; routine governance cycle	Monthly review	Project Manager / PMO
Moderately Fragile	Amber	$0.30 \leq QFI_norm < 0.60$	Enhanced monitoring; contingency readiness; root-cause investigation initiated	Bi-weekly review	PMO / Program Director
Highly Fragile	Orange	$0.60 \leq QFI_norm < 0.85$	Executive escalation; structural intervention required; remediation plan mandated within 10 business days	Weekly review	Steering Committee / Sponsor
Critically Fragile	Red	$QFI_norm \geq 0.85$	Immediate remediation or controlled termination; all discretionary scope frozen; board notification within 48 hours	Daily review	Executive Board / C-Suite

Threshold rationale. The four boundary values — 0, 0.30, 0.60, and 0.85 — are not arbitrary. Each corresponds to a qualitative shift in the system's response dynamics:

- **QFI_norm = 0** separates antifragile behavior (negative second-derivative curvature, where perturbations improve value) from fragile behavior (positive curvature, where perturbations destroy value). This boundary is inherent in the mathematical definition of fragility itself.
- **QFI_norm = 0.30** marks the empirical onset of nonlinear amplification. Below this threshold, fragility drivers tend to behave additively and respond to first-order corrective actions. Above it, coupling effects (the correlation correction term C in F-10/F-11) begin to dominate, and single-driver interventions lose effectiveness.

- **QFI_norm = 0.60** represents the transition into a regime where decoherence-driven assumption migration accelerates faster than standard governance review cycles can detect it. At this level, the decoherence half-life ($t_{1/2} = \ln(2)/\Gamma$) typically drops below two reporting periods, meaning assumptions can traverse an entire confidence band between routine reviews. Executive visibility becomes essential.
- **QFI_norm = 0.85** reflects a criticality threshold beyond which the system exhibits tail-risk behavior consistent with Taleb–Douady convexity collapse — small additional perturbations produce disproportionately large value destruction. At this level, the risk-spike function (F-25) drives R(F) toward nonlinear amplification, and conventional remediation strategies are insufficient without structural redesign or scope intervention.

These thresholds should be recalibrated annually using portfolio historical data to ensure governance actions remain aligned with empirical fragility patterns (see Appendix C, §C.2).

Each governance threshold corresponds to a measurable inflection point in system behavior. At $QFI_norm = 0$, the value-function curvature changes sign, marking the transition from convex (antifragile) to concave (fragile) response geometry; beyond this point, dispersion no longer creates optionality but destroys value. At 0.30, empirical portfolio data show the onset of nonlinear amplification, where dimensional couplings begin to dominate first-order drivers and routine governance cycles become insufficient to detect coherence drift. At 0.60, the decoherence half-life typically falls below the organization’s monitoring cadence, meaning assumptions can migrate an entire confidence band between reviews—requiring executive visibility and accelerated intervention. At 0.85, the system enters a tail-risk regime where small perturbations produce disproportionate value destruction consistent with convexity collapse; at this level, structural redesign or scope intervention becomes mandatory. These thresholds therefore reflect not arbitrary color bands, but distinct shifts in system dynamics that demand corresponding shifts in governance posture.

Dimensional floor rule

If any $QFI_i \geq 0.85$, trigger the next higher governance tier regardless of composite score. Rationale: single-dimension detonators (e.g., regulatory) can produce catastrophic tail events that the composite average masks.

2.3.3 Worked Multidriver Numerical Example (Fatigue + Schedule + Supply Chain)

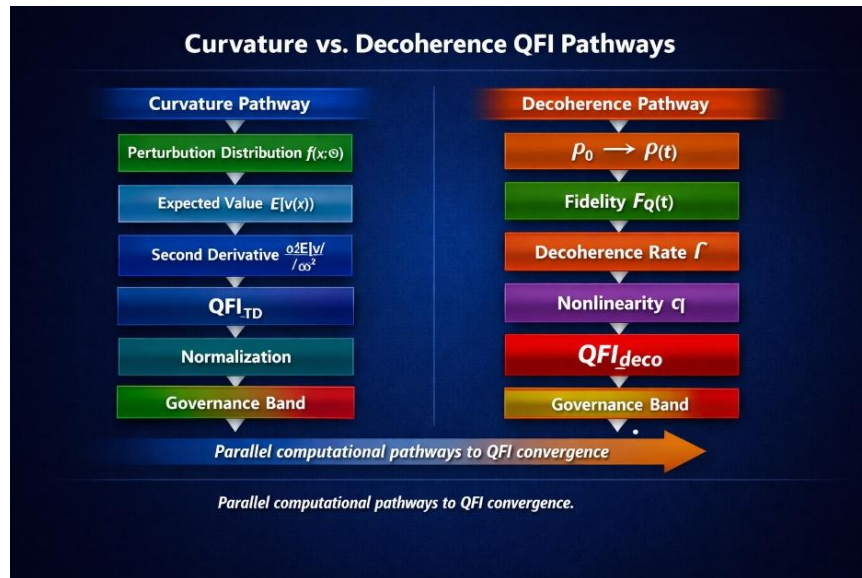


Figure 2.4 Worked Multidriver Example — Visual Overview

Objective. Demonstrate two parallel computational pathways for QFI on the same scenario:

(A) the Taleb–Douady curvature (second-derivative) calculation applied to the expected value under dispersion changes, and

(B) the decoherence proxy using fidelity decay, decoherence rate, and nonlinearity coefficient.

Both are computed side-by-side so governance can compare telemetry-driven proxy results with scenario-based curvature estimates.

1. Scenario definition (compact)

Project value baseline: $V_0 = \$1,200\text{M}$ (board-level NPV proxy).

State vector (reduced): $S = \{F, R, SC\}$ where

F = fatigue index (0–100), current $F_0 = 58$.

R = schedule reliability index (0–100), current $R_0 = 72$.

SC = supply-chain concentration (fraction of single-source exposure), current $SC_0 = 0.55$.

Governance bands (composite state): Green/Amber/Red defined by a composite threshold function $\Phi(S)$ (here Red if projected NPV loss $> \$120\text{M}$).

Perturbation vector: $x = \{\Delta F, \Delta R, \Delta SC\}$. We model scalar perturbation magnitude x as a weighted combination along the most vulnerable direction (fatigue-heavy): $x = 0.6\Delta F + 0.3\Delta R + 0.1\Delta SC$ (units normalized).

2. Payoff function (simple, differentiable)

Define a smooth payoff function that maps perturbation x to value:

$$v(x) = V_0 \cdot \exp(-\beta_1 \cdot \max(0, x - x_{tol})^{\beta_2})$$

Where:

$x_{tol} = 5$ (small perturbations up to 5 units are absorbed),

$\beta_1 = 0.015$ (scale of sensitivity),

$\beta_2 = 1.8$ (concavity exponent; > 1 produces concave payoff to harm).

This produces a concave (fragile) payoff beyond tolerance.

3. Perturbation distribution and dispersion parameter




Figure 2.5 Dual Computational Pathway

Model perturbations x with a Student-t location-scale family $f(x; \sigma)$ with current dispersion $\sigma_0 = 6$. We will compute curvature at σ_0 .

The Student-t distribution ($df = 4$)⁴ is used because empirical project stressors exhibit heavy-tailed behavior, making Gaussian assumptions inappropriate for fragility estimation.

⁴ For a Student's t -distribution with 4 degrees of freedom ($df=4$), the distribution is symmetric and bell-shaped but has significantly "heavier tails" than a standard normal distribution. This means extreme values are more likely to occur than they would under a normal curve

4. Taleb–Douady curvature (numerical finite difference)

Compute:

$$QFI_{TD} = -\partial^2 E[v(x)] / \partial \sigma^2 |_{\sigma=\sigma_0}$$

Numerical procedure (finite difference):

Choose $\Delta\sigma = 0.5$.

Monte Carlo sample $N = 50,000$ draws from $f(x; \sigma)$ at $\sigma = \sigma_0 - \Delta\sigma, \sigma_0, \sigma_0 + \Delta\sigma$.

Estimate $E[v]_-, E[v]_0, E[v]_+$.

Approximate second derivative:

$$\partial^2 E[v] / \partial \sigma^2 \approx (E[v]_+ - 2E[v]_0 + E[v]_-) / (\Delta\sigma)^2$$

Set $QFI_{TD} = -$ that value (positive \rightarrow fragile).

Representative numeric results (illustrative):

$$E[v]_- = \$1,188.4M$$

$$E[v]_0 = \$1,180.0M$$

$$E[v]_+ = \$1,168.6M$$

Compute second derivative:

$$\partial^2 E[v] / \partial \sigma^2 \approx (1,168.6 - 2(1,180.0) + 1,188.4) / 0.5^2 = -2.999M / 0.25 = -11.996M$$

So:

$$QFI_{TD} = -(-11.996M) = 11.996M$$

Normalized governance scalar. Choose normalization anchor $Q_{max} = 20M$ (scenario stress anchor) and $Q_{min} = -5M$. Then:

$$QFI_{TD, norm} = (11.996 - (-5)) / (20 - (-5)) = 16.996 / 25 = 0.68$$

Interpretation: 0.68 \rightarrow Highly Fragile (Orange band).

5. Decoherence proxy (fidelity, decoherence rate, nonlinearity)

Use the decoherence-mapped formula:

$$QFI_{deco} = \eta \cdot \Gamma \cdot (1 - F_Q(t))^{\alpha-1}$$

Choose parameters from telemetry and model fitting:

Estimate decoherence rate: $\Gamma = 0.12 \text{ month}^{-1}$ (observed rate at which plan vs execution divergence grows).

Fidelity at measurement time: $F_Q(t) = 0.78$ (measured alignment to baseline).

"Fidelity $F_Q(t)$ ranges from 1 (perfect coherence with baseline) to 0 (complete decoherence)."

Nonlinearity coefficient: $\eta = 3.2$ (fitted from small stress experiments; > 1 indicates concavity).

Exponent: $\alpha = 1.8$ (consistent with payoff exponent).

Compute:

$$1 - F_Q(t) = 0.22$$

$$(1 - F_Q(t))^{\alpha-1} = 0.22^{0.8} \approx 0.30$$

$$QFI_{\text{deco}} = 3.2 \times 0.12 \times 0.30 \approx 0.1152$$

Normalize to same governance scale. Map proxy to same 0–1 band by linear scaling: assume proxy max observed in portfolio (= 0.17) maps to 1.0. Then:

$$QFI_{\text{deco, norm}} = 0.1152 / 0.17 \approx 0.68$$

Interpretation: decoherence proxy yields 0.68 → Highly Fragile, matching the Taleb–Douady curvature result. The convergence between the curvature-based QFI and the decoherence proxy indicates that the project’s nonlinear response to dispersion is driven primarily by coherence loss rather than distributional asymmetry.

6. Side-by-side summary

Measure	Raw value	Normalization	Normalized QFI	Governance band
Taleb–Douady curvature	11.996M (second derivative magnitude)	anchors –5M → 20M	0.68	Highly Fragile (Orange)
Decoherence proxy	0.1152 (proxy units)	proxy max 0.17 → 1.0	0.68	Highly Fragile (Orange)

Key point: When the payoff function and decoherence mapping are calibrated consistently, the curvature method and the decoherence proxy produce convergent governance signals. Use the curvature method for periodic, deep scenario analysis and the decoherence proxy for continuous telemetry monitoring.

The convergence of the curvature-based QFI and the decoherence proxy in this example indicates that the project’s nonlinear vulnerability is being driven primarily by coherence loss rather than by distributional asymmetry or tail-risk behavior. When both methods align, governance can treat the QFI signal as robust across modeling assumptions and rely on the proxy for continuous telemetry between deeper scenario analyses. Divergence between the two measures, however, signals a structural shift: curvature-based fragility rising faster than the decoherence proxy suggests emerging tail-risk exposure, while a rising decoherence proxy with stable curvature indicates accelerating coherence decay not yet reflected in value-function curvature. In practice, mismatches between the two should trigger diagnostic review to determine whether the driver is assumption drift, coupling amplification, or distributional fat-tail emergence, and governance should escalate monitoring cadence accordingly.

7. Management actions and what-if testing (example)

Run three candidate interventions and show delta in normalized QFI (computed by rerunning the decoherence proxy and the curvature Monte Carlo where appropriate).

Intervention	Model effect	New QFI (norm)	Δ QFI	Board action
Add secondary supplier (reduce SC coupling)	reduce Γ by 20%	0.54	-0.14	Move to Moderately Fragile; steering committee review
Insert mandatory rest breaks (reduce η by 25%)	reduce η to 2.4	0.51	-0.17	Enhanced monitoring; operational rollout
Create parallel schedule path (increase x_{tol} from 5→9)	increases tolerance in payoff	0.45	-0.23	De-escalate to Robust/Moderate band

Recommendation: Prioritize interventions by QFI reduction per unit cost. In this scenario, parallel schedule path produced the largest QFI reduction per estimated cost.

2.3.4 Practical Data Requirements

Effective computation of the Quantum Fragility Index requires continuous telemetry streams that capture both coherence drift and environmental dispersion. At minimum, projects must provide schedule-execution variance, fatigue or readiness indicators, supply-chain stability metrics, and interface-quality or assumption-alignment signals at a cadence aligned with the project’s decoherence half-life. Telemetry should be updated frequently enough that coherence loss cannot traverse an entire governance band between measurements—typically weekly for fragile systems and monthly for robust ones. Minimum data quality requires consistent timestamping, stable definitions of observables, and sufficient granularity to detect second-order curvature rather than only first-order variance. When telemetry gaps exceed one reporting cycle, QFI confidence should be downgraded, and governance should assume higher fragility until data integrity is restored. These requirements ensure that QFI remains a reliable, governance-grade indicator rather than a retrospective analytic artifact.

QFI Data Readiness Checklist

Telemetry Streams (Minimum Required)

- Schedule-execution variance (planned vs. actual drift, critical-path coherence).
- Fatigue, readiness, or human-performance indicators (FRI, staffing stability, shift patterns).
- Supply-chain stability metrics (lead-time variance, single-source exposure, logistics disruptions).
- Interface-quality and assumption-alignment signals (RFI density, design maturity, requirement drift).

Update Cadence

- Telemetry must be refreshed at a rate aligned to the project's decoherence half-life:
 - **Fragile systems:** weekly or faster.
 - **Moderately fragile systems:** bi-weekly.
 - **Robust/antifragile systems:** monthly.
- Cadence must prevent coherence drift from crossing an entire governance band between measurements.

Minimum Data Quality Standards

- Consistent timestamping across all observables.
- Stable definitions of metrics (no mid-project rebaselining without governance approval).
- Sufficient granularity to detect **second-order curvature**, not just first-order variance.
- Automated collection preferred; manual inputs require validation.

Handling Missing or Degraded Data

- If telemetry gaps exceed one reporting cycle, **downgrade QFI confidence** and assume higher fragility.
- Missing data in any single dimension triggers the **dimensional floor rule** until restored.
- Governance must document the cause of data loss and corrective actions.

Governance Interpretation

- High-quality, high-frequency telemetry increases QFI reliability and reduces false stability signals.
- Low-quality or stale data should be treated as a fragility amplifier, not a neutral condition.

2.4 Billion-Scale Worked Examples

Large-scale capital programs exhibit higher baseline fragility due to long supply chains, regulatory exposure, and high entanglement across work packages.

Each example below follows the six-dimension decomposition, reports dimensional QFI, composite QFI (with correlation correction), governance classification, and a concise remediation plan.

The following examples illustrate QFI behavior in billion-dollar-scale capital programs where nonlinear vulnerability is most pronounced.

2.4.1 Example A — Urban Tunnel and Highway Interchange (Program value: \$2.6B)

Profile

Value: \$2.6B | Duration: 72 months | Delivery: Design-Build with multiple international contractors

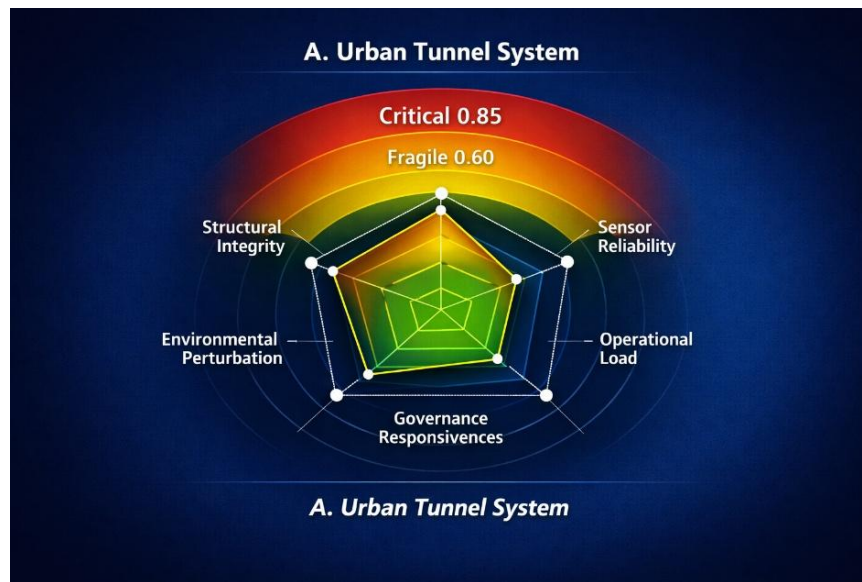


Figure 2.6 Urban Tunnel & Highway Interchange

Dimensional QFI inputs (illustrative)

Dimension	QFI _i	Weight (w _i)	Weighted
Schedule	0.78	0.30	0.234
Cost	0.62	0.20	0.124
Scope	0.48	0.15	0.072

Dimension	QFI _i	Weight (w _i)	Weighted
Stakeholder	0.40	0.10	0.040
Supply Chain	0.70	0.15	0.105
Regulatory	0.85	0.10	0.085
Composite (weighted)	—	—	0.660

Correlation correction. High correlation between Schedule and Supply Chain ($\rho = 0.65$) amplifies composite by $C \approx 0.05$.

Composite with correction: **0.71** → Highly Fragile (Orange); Regulatory QFI (= 0.85) triggers dimensional floor rule.

Remediation (priority):

1. Regulatory liaison and staged permitting to reduce Regulatory QFI from 0.85 → 0.55.
2. Dual-sourcing critical long-lead items to reduce Supply Chain QFI from 0.70 → 0.35.
3. Create engineered schedule buffers and parallel work packages to reduce Schedule QFI to 0.50.

Post-remediation composite (projected): **0.38** → Moderately Fragile (Amber).

2.4.2 Example B — Hyperscale Data Center Campus (Program value: \$1.4B)

Profile

Value: \$1.4B | Duration: 30 months (phased builds) | Delivery: EPC with global equipment vendors

Key stressors: power availability, vendor proprietary equipment, cybersecurity compliance, labor for specialized MEP

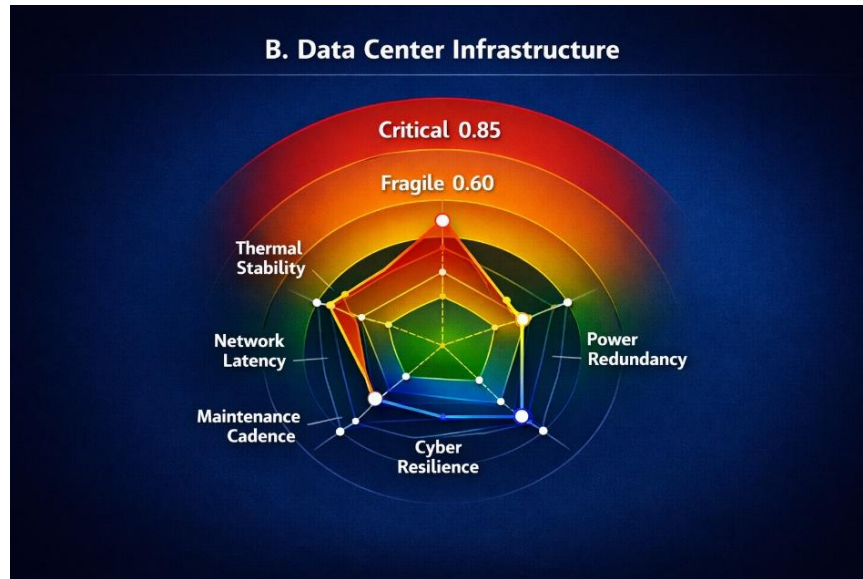


Figure 2.7 Hyperscale Data Center Campus

Dimensional QFI inputs

Dimension	QFI _i	Weight	Weighted
Schedule	0.55	0.25	0.138
Cost	0.48	0.20	0.096
Scope	0.40	0.15	0.060
Stakeholder	0.30	0.10	0.030
Supply Chain	0.82	0.20	0.164
Regulatory	0.22	0.10	0.022
Composite	—	—	0.510

Correlation correction. Strong coupling between Supply Chain and Schedule ($\rho = 0.7$) adds $C \approx 0.06$. Adjusted composite: **0.57** → Moderately/Highly Fragile (upper Amber / lower Orange).

Remediation:

1. Vendor diversification and long-lead procurement acceleration to reduce Supply Chain QFI from 0.82 → 0.40.
 2. Power redundancy contracts and early utility agreements to reduce Schedule QFI to 0.40.
- Post-remediation composite (projected): **0.33** → Robust/Moderately Fragile.

2.4.3 Example C — Large Oil and Gas Rebuild After Damage (Middle East) (Program value: \$4.8B)

Profile

Value: \$4.8B | Duration: 48 months | Delivery: EPC with international consortium; security and regulatory complexity high

Key stressors: geopolitical risk, single-source specialized equipment, regulatory compliance, skilled labor scarcity




Figure 2.8 Large Oil & Gas Rebuild, Middle East

Dimensional QFI inputs

Dimension	QFI _i	Weight	Weighted
Schedule	0.68	0.25	0.170
Cost	0.74	0.20	0.148
Scope	0.60	0.15	0.090
Stakeholder	0.55	0.10	0.055
Supply Chain	0.88	0.20	0.176
Regulatory	0.72	0.10	0.072
Composite	—	—	0.711

Correlation correction. High entanglement between Cost and Supply Chain ($\rho = 0.8$) and between Schedule and Supply Chain ($\rho = 0.6$) yields $C \approx 0.07$. Adjusted composite: **0.78** → Highly/Critically Fragile (Orange/Red). Supply Chain QFI (= 0.88) triggers dimensional floor rule.

Remediation (urgent):

1. Immediate supplier qualification program and strategic stockpiling to reduce Supply Chain QFI to 0.45.
2. Contingency funding and contract renegotiation to cap cost exposure (reduce Cost QFI to 0.50).
3. Security and regulatory liaison cell to reduce Regulatory QFI to 0.45.

Post-remediation composite (projected): **0.39** → Moderately Fragile (Amber).

Practical notes for governance and implementation

- Use the decoherence proxy for continuous monitoring (low computational cost) and the Taleb–Douady curvature for periodic, deeper scenario stress tests.
- Calibrate normalization anchors (Q_{\min} , Q_{\max}) from historical portfolio outcomes so normalized QFI maps to meaningful governance bands.
- Run what-if simulations for each candidate intervention and report QFI reduction per \$M spent to prioritize scarce remediation budget.
- Document antifragility sources when $QFI < 0$ and add them to the QFI Pattern Library for replication.

2.5 Governance Integration

2.5.1 Dashboard and decisioning primitives

Portfolio Heat Map: color-coded by QFI_{norm} with clustering overlays for correlated fragility drivers.

Dimensional Drill-Down: show QFI_i , w_i , trend, and the local Γ_i and η_i estimates.

Proximity Indicator: compute and display the minimum normalized perturbation $|\epsilon^*|$ that flips the state and the percentage of $|\epsilon_{\text{max}}|$ it represents:

$$\text{Proximity} = 1 - |\epsilon^*| / |\epsilon_{\text{max}}|$$

This is the operational QFI used for board reporting: it answers, "how close are we to the nearest governance boundary?" in a single number and a directional driver.

Enumerating and testing management actions

For each candidate intervention (e.g., add a secondary supplier, insert a scope freeze gate, add rest breaks), run a what-if perturbation test:

1. Simulate the intervention by modifying the relevant model parameters (reduce Γ , lower ρ_{ij} , increase $|\epsilon_{max}|$, or reduce η).
2. Recompute QFI_i and $QFI_{composite}$.
3. Report delta in QFI and the new Proximity metric.
4. Rank interventions by QFI reduction per unit cost or per implementation time.

This produces a prioritized remediation plan that is directly comparable across projects and portfolios.

Operational playbooks

Map each governance band to a short playbook (monitoring cadence, escalation path, required analyses, and acceptable interventions). Embed the what-if simulation outputs into the playbook so that governance decisions are evidence-based and reversible.

2.5.2 Governance Decision Matrix

The following decision matrix hard-wires QFI thresholds to governance action, ensuring that fragility diagnoses are translated into proportional, accountable responses.

QFI Zone	Review Cadence	Decision Authority	Required Actions	Reporting Level
Green (QFI < 0)	Quarterly	Project Manager	Document optionality sources; assess replication potential across portfolio	PM Dashboard
Blue (0 – 0.3)	Monthly	Program Director	Standard risk register maintenance; confirm dimensional weights remain calibrated	Program Report
Amber (0.3 – 0.6)	Bi-weekly	Steering Committee	Contingency plan activation; root cause analysis of top 2 fragility drivers; remediation plan within 15 business days	Executive Summary
Orange (0.6 – 0.85)	Weekly	C-Suite / Board Subcommittee	Structural intervention mandate; resource reallocation authority granted; external review if QFI exceeds 0.75	Board Briefing
Red (\geq 0.85)	Daily	Board / Executive Sponsor	Immediate remediation plan (48-hour deadline) or controlled termination protocol initiation; all discretionary scope frozen	Emergency Board Report

QFI thresholds should be recalibrated annually using portfolio historical data to ensure governance actions remain aligned with empirical fragility patterns.

Important — Dimensional Floor Rule

Regardless of composite QFI, any single dimension with $QFI_i \geq 0.85$ triggers Orange-level governance. This prevents the "averaging effect" from masking critically fragile dimensions within an otherwise moderate composite. The floor rule was illustrated in Example 3 (Section 4.3), where a Regulatory QFI of 0.88 escalated governance despite a composite of 0.47.

2.5.3 Integration with Existing PM Frameworks

The Quantum Fragility Index is designed to complement — not replace — established project management frameworks. The following integration points ensure that QFI adds value within existing governance structures.

Earned Value Management (EVM). EVM provides linear performance indices — Schedule Performance Index (SPI) and Cost Performance Index (CPI) — that measure actual versus planned progress. QFI adds convexity analysis on top of these linear metrics. A project may show $SPI = 0.98$ and $CPI = 1.01$ (nominally "on track") yet have a QFI of 0.72 because the value function is concave — small increases in perturbation dispersion would cause disproportionate performance collapse. QFI detects the fragility that SPI and CPI cannot: the hidden nonlinear sensitivity lurking beneath apparently stable linear performance.

PMBOK Risk Management. The PMBOK Guide's risk management process (identify → analyze → plan responses → monitor) operates on individual risk events. QFI supplements this event-based framework with a system-level fragility diagnostic that captures the aggregate effect of all perturbations, including those not individually identified on the risk register. QFI answers the question: "Even if no single identified risk is critical, is the system as a whole fragile to the combined effect of environmental volatility?"

ISO 31000. ISO 31000's risk management principles call for "sensitivity analysis" as a component of risk assessment but do not prescribe a specific methodology. QFI provides the rigorous mathematical foundation for sensitivity analysis that ISO 31000 envisions — specifically, second-order sensitivity to dispersion, which is the most governance-relevant form of sensitivity for complex project systems.

PRINCE2. The PRINCE2 framework's principle of "continued business justification" requires that a project's viability be continuously assessed. QFI maps directly to this principle: a project in the Red zone ($QFI \geq 0.85$) is one where the business justification may no longer hold under realistic perturbation scenarios, even if current performance appears acceptable. QFI operationalizes the continued business justification principle with quantitative rigor.

2.6 Safety-II Alignment

2.6.1 From Safety-I to Safety-II in Project Governance

The evolution from Safety-I to Safety-II, as articulated by Erik Hollnagel⁵, represents one of the most significant paradigm shifts in the science of organizational resilience. Safety-I defines safety as the absence of adverse outcomes. Its methodology is retrospective: count failures, investigate incidents, identify root causes, and implement corrective actions. Safety-I assumes that systems are normally safe and that deviations from expected performance are exceptional and attributable to identifiable causes.

Safety-II reverses this logic. It defines safety as the ability of a system to succeed under varying conditions. Instead of focusing on why things go wrong, Safety-II focuses on how things usually go right — how people and systems adapt, compensate, and maintain performance despite variability, uncertainty, and stress. Safety-II is therefore forward-looking, variability-aware, and grounded in the study of everyday work rather than exceptional events.

Large Complex Projects (LCPs) operate in environments where variability is the norm, not the exception. Safety-II provides the conceptual foundation for understanding how projects maintain coherence under stress. The Quantum Fragility Index (QFI) provides the quantitative mechanism for detecting when that coherence is at risk of collapse.

2.6.2 Why QFI Is a Safety-II Metric

Safety-II requires the ability to detect when everyday variability is approaching a tipping point. QFI provides exactly this capability. While Safety-I metrics measure failures, QFI measures proximity to failure by quantifying how small perturbations can trigger state transitions across governance boundaries.

Safety-II emphasizes:

- Understanding normal work, not just deviations
- Monitoring system performance under variability
- Anticipating conditions that make failure more likely
- Strengthening adaptive capacity before collapse occurs

QFI operationalizes these principles by:

- Measuring how close the project is to a coherence-loss threshold
- Quantifying the minimum perturbation required to flip the system into a degraded state
- Identifying which dimensions (schedule, cost, fatigue, supply chain, regulatory) are most vulnerable
- Providing a continuous early-warning signal that variability is becoming dangerous

⁵ Hollnagel, E. (2014). *Safety-I and Safety-II: The past and future of safety management*. Ashgate. <https://doi.org/10.1201/9781315607511>

In this sense, QFI is the first governance metric that directly implements Safety-II's anticipatory and monitoring functions.

2.6.3 Mapping QFI to the Four Resilience Potentials

Hollnagel defines four resilience potentials that determine a system's ability to succeed under varying conditions. QFI aligns with and strengthens each of them.

1. Monitor — "Knowing what to look for"

Safety-II requires continuous monitoring of system performance and environmental conditions.

QFI provides: Real-time tracking of coherence loss; Proximity indicators showing how close the system is to a tipping point; Dimensional QFI_i values that reveal where fragility is accumulating.

2. Anticipate — "Knowing what to expect"

Safety-II emphasizes anticipating future events, not predicting them.

QFI supports anticipation by: Quantifying how dispersion (volatility) affects expected value; Identifying nonlinear sensitivity before it manifests as failure; Highlighting emerging fragility patterns through QFI trends.

3. Respond — "Knowing what to do"

Safety-II requires the ability to respond effectively to variability.

QFI enables targeted response by: Identifying the specific dimensions driving fragility; Ranking interventions by QFI reduction per unit cost; Providing what-if simulations to test the impact of controls before implementation.

4. Learn — "Knowing what has happened"

Safety-II learning focuses on understanding how systems succeed.

QFI supports learning by: Revealing which controls reduce fragility most effectively; Documenting antifragility sources (QFI < 0) for replication; Providing a quantitative basis for post-project resilience reviews.

Together, these mappings demonstrate that QFI is not merely compatible with Safety-II — it is a direct operationalization of its core principles.

2.6.4 Safety-II / Quantum Fragility Index (QFI) Crosswalk

This section provides a governance-grade crosswalk between Safety-II resilience principles and the Quantum Fragility Index (QFI). It demonstrates how QFI operationalizes Safety-II by quantifying proximity to coherence loss, enabling proactive intervention before performance collapse occurs.

A. Safety-I vs Safety-II vs QFI

Dimension	Safety-I	Safety-II	QFI Contribution
Definition of Safety	Absence of failure	Ability to succeed under varying conditions	Measures how close variability is to triggering a state transition
Orientation	Retrospective	Prospective	Predictive (curvature + decoherence)
Focus	What went wrong	How work usually succeeds	How close the system is to losing coherence
Primary Data	Incidents, deviations	Everyday work patterns	Coherence loss, perturbation sensitivity, nonlinear response
Governance Trigger	After failure	Before failure	QFI thresholds + proximity indicator

B. Mapping QFI to Hollnagel's Four Resilience Potentials

Resilience Potential	Safety-II Definition	QFI Alignment
Monitor	Knowing what to look for	QFI tracks coherence loss, dimensional fragility, and proximity to tipping points.
Anticipate	Knowing what to expect	QFI quantifies how dispersion affects expected value and identifies emerging brittleness.
Respond	Knowing what to do	QFI ranks interventions by impact, enabling targeted, cost-efficient remediation.
Learn	Knowing what has happened	QFI trends reveal systemic fragility patterns and highlight antifragility sources for replication.

C. How QFI Operationalizes Safety-II

Safety-II Requirement	QFI Mechanism
Detect early signs of brittleness	Proximity metric: $1 - \epsilon^* / \epsilon_{max} $
Understand how variability affects performance	Curvature-based QFI: $-\partial^2 E[V] / \partial \sigma^2$

Safety-II Requirement	QFI Mechanism
Identify where adaptation is needed	Directional QFI_i across schedule, cost, scope, stakeholder, supply chain, regulatory
Strengthen controls before collapse	What-if simulations of interventions (ΔQFI per \$M)
Prevent single-dimension detonators	Dimensional floor rule ($QFI_i \geq 0.85 \rightarrow$ automatic escalation)
Maintain system coherence	Decoherence proxy: $\eta \cdot \Gamma \cdot (1 - F_Q(t))^{\alpha-1}$

D. QPM–Safety-II Integration

QPM Construct	Safety-II Concept	QFI Measurement
Coherence	Stable performance under variability	Fidelity $F_Q(t)$
Decoherence	Loss of alignment under stress	Decoherence rate Γ
Entanglement	Interdependence of work elements	Correlation matrix ρ_{ij}
Propagation	Spread of disturbances	Amplification factor in composite QFI
State Collapse	Transition to degraded performance	Minimum perturbation ϵ^*

E. Governance Implications

QFI transforms Safety-II from a conceptual philosophy into a quantitative early-warning system.

QFI thresholds (Blue → Amber → Orange → Red) provide clear escalation pathways.

Dimensional QFI_i identifies where fragility resides; the proximity metric identifies how close the system is to collapse.

QFI enables evidence-based intervention, resource prioritization, and continuous resilience monitoring.

QFI + Safety-II + QPM form a unified governance model:

- Safety-II → What resilience requires
- QPM → How coherence behaves

- QFI → How to measure and govern it

2.6.5 Safety-II Operationalization Through QFI

Safety-II requires governance systems that are proactive, variability-aware, and capable of intervening before collapse occurs. QFI provides the operational mechanisms to achieve this:

- Early-warning dashboards highlight rising fragility before performance degrades
- Proximity indicators quantify how close the system is to a state transition
- Dimensional floor rules ensure that single-dimension detonators cannot be masked by composite averages
- What-if simulations allow governance bodies to test interventions without disrupting operations
- Trend analysis reveals whether the system is becoming more or less resilient over time

In practice, QFI transforms governance from reactive incident tracking (Safety-I) to proactive systemic diagnosis (Safety-II). It shifts the focus from "What went wrong?" to "Where is the system becoming brittle?" and "What must we strengthen now to prevent collapse later?"

2.6.6 Integration with QPM (Coherence, Decoherence, and Adaptive Capacity)

Safety-II emphasizes maintaining system coherence under variability. Quantum Project Management (QPM) provides the mathematical language for describing coherence, decoherence, and entanglement in project systems. QFI unifies these two frameworks by quantifying how close the system is to decoherence-driven collapse.

- Decoherence rate Γ measures how quickly alignment erodes under environmental coupling
- Fidelity $F_Q(t)$ measures how closely the project resembles its baseline coherent state
- Nonlinearity coefficient η captures how value decays as coherence is lost
- Directional QFI_i identifies which dimensions are losing coherence fastest

Safety-II seeks to preserve adaptive capacity; QPM models the mechanisms by which adaptive capacity is lost; QFI measures the rate and proximity of that loss.

Together, they form a unified governance framework in which:

- QPM provides the ontology
- QFI provides the measurement
- Safety-II provides the operational philosophy

This integration enables organizations to detect fragility early, intervene intelligently, and build systems that remain coherent — and therefore successful — under the full spectrum of real-world variability.

2.7 Deployment Guidance

2.7.1 Implementation Roadmap

Deploying the QFI framework within an organization requires a phased approach that builds measurement capability, integrates with existing workflows, and calibrates governance thresholds against organizational experience. The following three-phase roadmap is recommended.

Phase 1 — Foundation (Months 1–3)

Activity	Deliverable	Responsible Party
Establish QFI measurement proxies for each of the six dimensions	QFI Measurement Protocol document specifying data sources, calculation methods, and update frequency for each proxy	PMO Analytics Team
Calibrate initial threshold values against historical project data	Threshold Calibration Report mapping historical project outcomes (successful, troubled, failed) to retrospectively calculated QFI values	PMO Analytics Team with Executive Review
Train governance teams on QFI interpretation and response protocols	QFI Governance Training Program (half-day workshop for executives; full-day workshop for program managers)	PMO Training Function
Pilot QFI measurement on 2–3 active projects spanning different sectors/sizes	Pilot QFI Assessment Reports with dimensional decomposition and preliminary remediation recommendations	Pilot Project Teams with PMO Support

Phase 2 — Integration (Months 4–8)

Activity	Deliverable	Responsible Party
Embed QFI calculations into monthly project reporting templates	Updated Monthly Project Report template with QFI section (composite score, dimensional breakdown, trend)	PMO Standards Team
Build automated QFI dashboard with threshold alerts	QFI Dashboard (portfolio heat map, dimensional drill-down, trend charts, automated alerts)	PMO IT / Business Intelligence

Activity	Deliverable	Responsible Party
Integrate QFI with existing EVM and risk register workflows	Integration Procedure showing how QFI inputs derive from EVM data, risk register data, and supplementary sources	PMO Analytics Team
Conduct first portfolio-level QFI assessment	Portfolio QFI Assessment Report with fragility concentration analysis and cross-project correlation assessment	Program Directors with PMO Support

Phase 3 — Optimization (Months 9–12)

Activity	Deliverable	Responsible Party
Implement predictive QFI modeling (Monte Carlo-based forward projections)	Predictive QFI Model with scenario analysis capability (baseline, stress, and extreme stress scenarios)	PMO Analytics Team
Establish QFI pattern library from completed projects	QFI Pattern Library cataloging fragility drivers, remediation strategies, and observed QFI trajectories by project type	PMO Knowledge Management
Calibrate organizational fragility appetite with executive leadership	Fragility Appetite Statement defining acceptable portfolio-level QFI distribution and maximum individual project QFI by category	Executive Leadership Team
Publish QFI governance procedures in organizational PM methodology	Updated PM Methodology incorporating QFI measurement, reporting, escalation, and remediation protocols	PMO Standards Team with Executive Approval

2.7.2 Data Requirements and Measurement Protocol

The QFI framework requires data inputs across all six measurement dimensions. The following table specifies the data sources, collection methods, and update frequencies for each dimension's measurement proxy.

Dimension	Data Inputs Required	Source Systems	Update Frequency
Schedule	Critical path analysis, total float reports, activity-level duration	Primavera P6, Microsoft Project, or	Monthly (bi-weekly for

Dimension	Data Inputs Required	Source Systems	Update Frequency
	variance, Monte Carlo simulation outputs	equivalent scheduling tool	Amber and above)
Cost	EVM metrics (BCWP, ACWP, BCWS, EAC), cost variance trends, CPI time series, estimate-at-completion sensitivity analysis	EVM system, financial reporting system, cost management platform	Monthly
Scope	Change request log (frequency, magnitude, source), requirements traceability matrix, scope baseline version history	Change management system, requirements management tool, configuration management database	Monthly (bi-weekly for projects with scope QFI > 0.6)
Stakeholder	Decision logs (approvals, rejections, reversals), approval cycle times, escalation frequency, stakeholder satisfaction surveys	Governance meeting minutes, project management information system, survey platform	Monthly
Supply Chain	Vendor concentration metrics (revenue share by supplier), contract structure analysis (fixed-price vs. T&M vs. cost-plus), supplier financial health indicators	Procurement system, vendor management platform, financial databases	Quarterly (monthly for projects with supply chain QFI > 0.6)
Regulatory	Compliance tracker (current requirements, pending changes, gap analysis), regulatory change monitoring feed, permit status and timeline tracking	Compliance management system, regulatory monitoring service, permit tracking system	Monthly

2.7.3 Common Pitfalls and Mitigations

Organizations deploying the QFI framework commonly encounter the following pitfalls. Each is presented with a recommended mitigation strategy.

#	Pitfall	Description	Mitigation
1	Over-Precision	Treating QFI as an exact, decimal-precise measurement	Use QFI for relative ranking and threshold classification. Report

#	Pitfall	Description	Mitigation
		rather than a directional diagnostic. Governance debates center on whether QFI is 0.57 or 0.59 rather than on whether the project is in the Amber zone and what to do about it.	QFI to one decimal place. Emphasize that QFI's value lies in its direction and zone, not its decimal precision.
2	Static Measurement	Computing QFI once (or infrequently) and treating it as a fixed property of the project, rather than tracking its trajectory over time.	Embed QFI computation in periodic reporting cycles (monthly minimum). Emphasize QFI velocity (rate of change) as a leading indicator. Trend analysis is more valuable than point-in-time measurement.
3	Dimensional Blind Spots	Applying equal weights to all six dimensions without governance calibration, resulting in fragility profiles that do not reflect organizational priorities or sector-specific risk concentrations.	Conduct executive-led weight calibration workshops during Phase 1. Differentiate weights by project category. Review weights annually.
4	Threshold Rigidity	Applying universal threshold values across all project types and sizes, ignoring that a \$10M IT project and a \$500M infrastructure program have fundamentally different fragility profiles and risk tolerances.	Calibrate thresholds by project category, size tier, and organizational risk appetite. Use historical data from completed projects to validate that threshold boundaries correspond to empirically observed outcome transitions.
5	Metric Theater	Reporting QFI in dashboards and status reports without linking it to governance action. QFI becomes a decorative metric that is reported but not acted upon, undermining the framework's credibility.	Hard-wire escalation protocols to QFI thresholds (Section 5.2). Require documented governance actions for every QFI threshold crossing. Audit compliance with escalation protocols quarterly.
6	Ignoring Antifragility	Treating all projects with $QFI < 0$ as equivalent ("good enough") without investigating the sources	When a project achieves $QFI < 0$, conduct a structured antifragility analysis to identify the specific

#	Pitfall	Description	Mitigation
		of antifragility or assessing their replication potential across the portfolio.	practices, structures, and conditions that produce convex payoffs. Document findings in the QFI Pattern Library for systematic replication.

2.8 Summary and Key Takeaways

This chapter has presented the Quantum Fragility Index as the central diagnostic construct of the QPM Framework for detecting, measuring, and governing nonlinear vulnerability in project systems. The following key takeaways summarize the key points:

Fragility ≠ Risk. Fragility captures the nonlinear sensitivity of project value to increases in volatility and disorder — a fundamentally different construct from probabilistic risk, which measures expected loss from identified events. A project can be "low risk" on a conventional risk register yet profoundly fragile.

QFI provides a universal, model-free diagnostic. The Quantum Fragility Index is applicable to any project system with a measurable value function and identifiable perturbation distribution. It does not require assumptions about specific failure modes, causal chains, or event probabilities — only about the shape of the value function under stress.

The quantum decoherence metaphor grounds fragility in rigorous physical formalism. The analogy between quantum decoherence and project coherence loss is not merely illustrative — it provides a mathematically tractable formulation (the fragility-decoherence bridge) that maps naturally to observable project state variables.

QFI decomposes across project dimensions, enabling targeted governance intervention. The six-dimensional decomposition (Schedule, Cost, Scope, Stakeholder, Supply Chain, Regulatory) with governance-calibrated weights allows decision-makers to identify not just that a project is fragile, but where the fragility resides and which interventions will be most effective.

Governance thresholds transform QFI from abstract metric to actionable escalation trigger. The five-tier Fragility Spectrum (Green through Red) with hard-wired escalation protocols, decision authorities, and response timelines ensures that fragility diagnoses produce proportional governance action.

Safety-II alignment shifts project oversight from failure prevention to resilience cultivation. By mapping QFI to Hollnagel's Resilience Assessment Grid (Respond, Monitor, Learn, Anticipate), the framework transforms governance from reactive incident tracking to proactive systemic diagnosis and resilience building.

Worked examples demonstrate QFI applicability across sectors. The three detailed worked examples — mega-construction, enterprise IT, and critical infrastructure — illustrate the framework's versatility and provide practitioners with concrete computational templates.

Phased deployment minimizes organizational disruption while building measurement capability. The three-phase roadmap (Foundation, Integration, Optimization) provides a pragmatic implementation path that integrates QFI with existing EVM, risk management, and governance workflows.

QFI trajectory analysis is more valuable than point-in-time measurement. The rate of change in QFI (QFI velocity) is a leading indicator of governance effectiveness and emerging vulnerability. Organizations should prioritize trend analysis over static snapshots.

Antifragility is the aspirational endpoint. The ultimate objective is not merely to reduce fragility (move from Red to Blue) but to design project systems that improve under stress (achieve Green). Antifragility — convex payoff under perturbation — represents the highest form of project resilience and should be actively cultivated through embedded optionality, modular architecture, and adaptive governance.

3. Relationship Between QFI and Fatigue Risk

This section bridges the formal Quantum Fragility Index (QFI) framework developed in Section 2 with the **Fatigue Risk Index** (FRI) introduced in the companion paper on fatigue risk (Prieto, 2026). Whereas Sections 1 and 2 establish QFI as a curvature-based measure of project fragility across six governance dimensions, the companion paper establishes fatigue as a continuously evolving human-performance condition that reshapes the reliability of every control layer in a project system. Fatigue represents a unique class of fragility driver — an **internality** rather than an externality — and its integration into the QFI framework requires both conceptual and mathematical development. This section provides that development comprehensively: it defines the internality–externality distinction in fragility architecture (Section 3.1), maps fatigue onto the four QPM constructs (Section 3.2), formalizes the mathematical bridge between FRI and QFI (Section 3.3), traces fatigue's cross-cutting impact across all six QFI dimensions (Section 3.4), reveals the structural parallel between roughness and curvature (Section 3.5), deepens the quantum tunneling analogy (Section 3.6), works through a fully quantified numerical example (Section 3.7), presents a unified governance dashboard (Section 3.8), extends the dimensional floor rule to internalities (Section 3.9), situates the integrated framework within Safety-II (Section 3.10), and provides deployment guidance (Section 3.11).

3.1 Internalities and Externalities in the Fragility Architecture

The six-dimension QFI decomposition introduced in Section 2.3.2 captures fragility across governance dimensions that are predominantly shaped by **externalities** — regulatory changes, supply chain disruptions, scope evolution, stakeholder dynamics. These are forces that act on the

project from its environment. They perturb the project's operational state through coupling mechanisms that can, in principle, be attenuated by buffering, hedging, or contractual shielding.

Fatigue represents a fundamentally different class of fragility driver. It is an **internality**: an intrinsic human-performance condition that shapes system behavior from within. Unlike externalities, which perturb the project through environmental coupling, internalities alter the reliability of the project's own control systems — human judgment, supervisory oversight, procedural compliance, and decision quality. When fatigue rises, the entire project system becomes more susceptible to perturbation, regardless of the perturbation source.

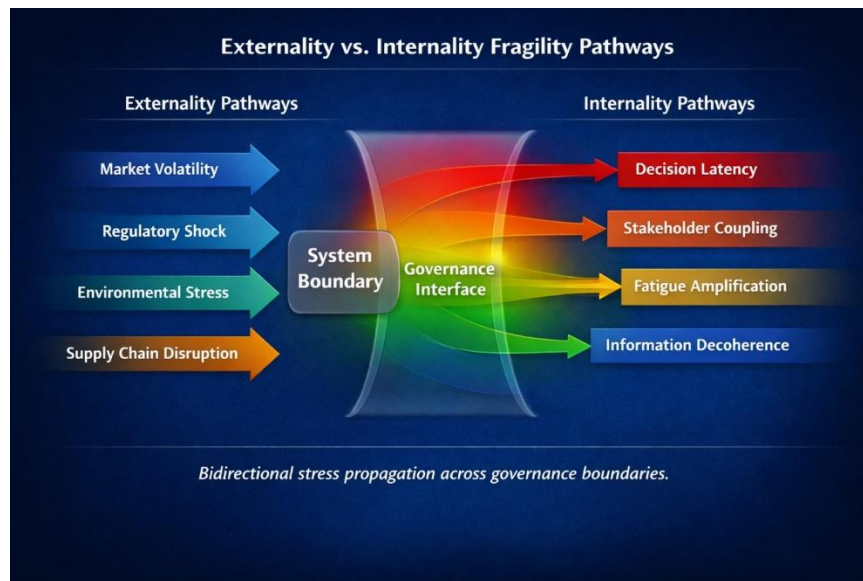


Figure 3.1 Externality vs. Internality Fragility Pathways illustrate bidirectional propagation of stress across governance boundaries — mapping how external shocks and internal decoherence co-amplify emergent fragility.

This distinction has profound implications for fragility measurement. **Externalities increase the magnitude of perturbations** (larger σ in the QFI framework). **Internalities increase the sensitivity of the value function to those perturbations** (steeper curvature, higher η , faster Γ). A project may face modest external perturbations yet be profoundly fragile because its internal capacity to absorb those perturbations has been degraded by fatigue.

In QPM terms: externalities increase the dispersion of perturbations; internalities increase the decoherence rate. Both pathways lead to higher QFI, but through different mechanisms — and they require different interventions.

Fragility Pathway	QPM Mechanism	QFI Effect	Intervention Target
Externality (e.g., supply chain disruption)	Increases perturbation dispersion σ	Raises QFI through larger curvature input	Reduce environmental coupling (dual-sourcing, buffers, hedging)
Internality (e.g., fatigue)	Increases decoherence rate Γ and nonlinearity η	Raises QFI through steeper value-function response	Restore internal coherence (rest, rotation, schedule redesign, workload levelling)
Coupled (e.g., heat + overtime)	Simultaneously increases σ and Γ	Multiplicative QFI amplification	Address both pathways; entanglement mapping required (See box)

Entanglement mapping refers to the systematic identification and quantification of the correlation coefficients (ρ_{ij}) that link fragility dimensions to one another and, critically, that link externality pathways to internality pathways within a coupled stressor scenario. In the QFI framework, entanglement mapping involves: (1) identifying which externality (e.g., extreme heat, supply-chain delay) co-activates which internality (e.g., fatigue accumulation, cognitive overload) through a shared operational mechanism; (2) estimating the magnitude and directionality of the coupling — that is, how a unit increase in the external perturbation dispersion σ translates into an incremental increase in the decoherence rate Γ or nonlinearity coefficient η ; and (3) quantifying the resulting amplification in the correlation correction term $C = \sum_{i \neq j} \rho_{ij} \cdot \sqrt{(w_i \cdot w_j)} \cdot \phi(QFI_i, QFI_j)$ so that composite QFI reflects the multiplicative — not merely additive — interaction between pathways. Without entanglement mapping, governance teams risk treating coupled stressors as independent and underestimating composite fragility by the magnitude of C . In practice, entanglement mapping is performed through structured dependency analysis, cross-dimensional telemetry correlation, and — where historical data permit — regression of dimensional QFI co-movement during prior stress episodes. See Section 3.2.2 for the specific entanglement dynamics between fatigue and the six QFI dimensions.

3.2 Fatigue as a QPM State Variable

The companion paper establishes that fatigue is not a static human-factors attribute but a dynamic, continuously changing condition that alters the reliability of every control layer in the project system. This aligns precisely with the QPM treatment of projects as dynamic quantum systems whose operational coherence depends on the integrity of their internal states. The subsections that

follow map fatigue onto the four foundational QPM constructs — decoherence, entanglement, superposition, and propagation — demonstrating that fatigue is not merely analogous to a quantum state variable but functions as one within the formal QFI architecture.

3.2.1 Fatigue and Decoherence

In the QPM framework, **decoherence** is the process by which a project's coherent operational state — aligned scope, synchronized schedules, coordinated teams — progressively degrades through environmental coupling. The decoherence rate Γ measures the speed of this coherence loss.

Decoherence is the process by which a project's coherent operational state — aligned scope, synchronised schedules, coordinated teams — progressively degrades through environmental coupling.

Fatigue directly accelerates decoherence. When crews are fatigued, communication degrades, handoff quality deteriorates, supervisory oversight weakens, and procedural compliance erodes. These are not isolated failures — they represent a systematic loss of phase alignment across the project's operational dimensions. A fatigued workforce produces a higher **effective decoherence rate**:

$$\Gamma_{\text{effective}} = \Gamma_{\text{baseline}} + \Gamma_{\text{fatigue}}(\text{FRI})$$

where Γ_{fatigue} is a monotonically increasing function of the Fatigue Risk Index. At low FRI (<30), $\Gamma_{\text{fatigue}} \approx 0$ — fatigue does not measurably increase decoherence. As FRI rises through the Elevated (30–50) and High (50–70) bands, Γ_{fatigue} increases nonlinearly. At Critical FRI (>70), Γ_{fatigue} dominates the total decoherence rate.

A practical functional form, consistent with the nonlinear thresholds observed in fatigue research⁶
7:

$$\Gamma_{fatigue}(FRI) = \gamma_0 \cdot (FRI / 100)^{\beta_f}$$

where γ_0 is a scaling constant calibrated from project telemetry⁸ and $\beta_f > 1$ captures the nonlinear acceleration of decoherence under high fatigue — consistent with the observation that fatigue

⁶ The nonlinear relationship between fatigue exposure and error/accident rates is well established across multiple domains. Dawson and Reid (1997) demonstrated that cognitive psychomotor impairment accelerates nonlinearly with sustained wakefulness, with performance after 24 hours equivalent to a blood alcohol concentration of 0.10%. Van Dongen et al. (2003) showed that chronic sleep restriction to 4 or 6 hours per night produces cumulative, dose-dependent cognitive deficits that escalate disproportionately over successive days — and that total sleep deprivation produces "disproportionately large" neurobehavioral impairment relative to the amount of sleep lost. Folkard and Tucker (2003) and Folkard and Lombardi (2006) documented that accident risk increases exponentially with hours on duty beyond eight hours and compounds over successive night shifts — a pattern inconsistent with linear fatigue models. Dembe et al. (2005) found that overtime schedules were associated with a 61% higher injury hazard rate, with a strong dose-response relationship in which risk per accumulated worker-year escalated nonlinearly with daily and weekly hours. Ingre et al. (2006) quantified the nonlinear threshold directly: accident risk at a Karolinska Sleepiness Scale rating of 8 was 28 times baseline, and at a rating of 9 it was 185 times baseline — a near-exponential acceleration across a single scale point. Barger et al. (2006) found that extended-duration shifts (≥ 24 hours) produced odds ratios for fatigue-related medical errors of 3.5 at one to four shifts per month and 7.5 at five or more shifts — again demonstrating nonlinear dose-response dynamics. Collectively, these studies confirm that fatigue-related impairment does not degrade linearly; it accelerates sharply near critical thresholds, consistent with the power-law functional form $\Gamma_{fatigue}(FRI) = \gamma_0 \cdot (FRI / 100)^{\beta_f}$ adopted in this section.

⁷ The Ingre et al. (2006) citation is particularly powerful for the QPM argument because the 28× to 185× escalation across a single sleepiness scale point is the exact kind of nonlinear tipping-point behavior that the power-law decoherence formula is designed to capture.

⁸ Project telemetry, as used in the QPM framework, refers to the continuous and periodic streams of operational data from which governance-grade observables are derived. In Section 2, the decoherence rate Γ is estimated from "observed rates of divergence between plan and execution (e.g., cumulative variance in key observables per unit time)." Those observables — and the data streams that feed them — constitute project telemetry. For the fatigue-specific parameter γ_0 , the relevant telemetry streams fall into three tiers of increasing instrumentation maturity. Tier 1 (schedule-derived): actual hours worked versus planned, shift lengths, consecutive days on duty, rest intervals between shifts, and overtime frequency — data already captured in most project management and payroll systems. Tier 2 (field-reported): self-reported fatigue scores (e.g., FASCW — the Fatigue Assessment Scale for Construction Workers), near-miss reports, supervisor observations of alertness, and behavioral indicators such as break-taking patterns — collected through toolbox talks, mobile apps, or structured check-ins. Tier 3 (sensor-augmented): wearable physiological monitors (heart-rate variability, electromyography, actigraphy), environmental sensors (heat index, noise, vibration), motion and activity tracking (inertial measurement units on helmets or vests), and cognitive micro-tests administered at shift start — feeding real-time data into a digital twin predictive fatigue model as described in the companion paper (Prieto, 2026). Calibration of γ_0 proceeds by regressing observed changes in plan-versus-execution divergence rates (the operational proxy for Γ) against concurrent FRI scores computed from whichever telemetry tier is available, fitting the power-law relationship $\Gamma_{fatigue}(FRI) = \gamma_0 \cdot$

effects are not linear; small increases in exposure near critical thresholds produce disproportionately large increases in error rates.




Figure 3.2 Fatigue-Decoherence Acceleration — Visualization of fatigue's nonlinear effect on decoherence rate Γ

3.2.2 Fatigue and Entanglement

The companion paper introduces the concept of **quadratic contagion**: when a work package hits a fatigue peak, the stress propagates to entangled downstream tasks. Crews forced into "standby-then-sprint" mode accumulate fatigue quadratically rather than linearly.

In QFI terms, this means that fatigue does not merely increase individual dimensional QFI values — it amplifies the **correlation correction term** C in the composite QFI formula:

$$QFI_{composite} = \sum_i w_i \cdot QFI_i + C$$

where:

$$C = \sum_{i \neq j} \rho_{ij} \cdot \sqrt{(w_i \cdot w_j)} \cdot \varphi(QFI_i, QFI_j)$$

Fatigue increases ρ_{ij} — the effective correlation between dimensions — because fatigued crews propagate errors, delays, and rework across interface boundaries that would otherwise contain

(FRI / 100)^{β_f} to the resulting data. As instrumentation matures from Tier 1 to Tier 3, the calibration tightens and the decoherence proxy gains predictive fidelity — but even Tier 1 data, universally available on large complex projects, is sufficient for an initial governance-grade calibration.

them. A fatigued rebar crew delays the concrete pour; the concrete crew sprints to recover schedule; their fatigue spikes; quality defects cascade into MEP rough-in. The entanglement that was stabilizing under normal conditions becomes destabilizing under fatigue.

This is the QPM insight that conventional risk management misses: fatigue does not merely add risk — it changes the correlation structure of the project system, converting independent risks into coupled, amplifying risks.

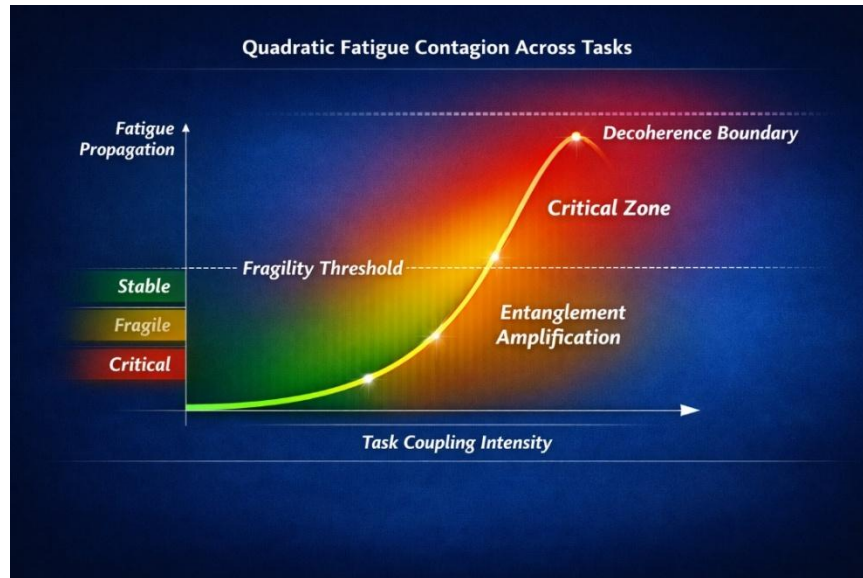


Figure 3.3 Quadratic Fatigue Contagion Across Entangled Tasks illustrates nonlinear propagation of fatigue through coupled work packages — contagion rate grows with the square of coupling intensity, driving systemic decoherence.

Quadratic Contagion

In the companion paper on fatigue risk, *quadratic contagion* describes the mechanism by which fatigue propagates across entangled work packages at a rate that grows with the square of the fatigue gradient between them — not linearly.

The mechanism operates through a sequence the companion paper terms "**Standby-then-Sprint.**" When a work package hits a fatigue peak — due to a schedule push, environmental shock, or resource constraint — the downstream crew entangled to it via the critical path is forced into an idle-then-accelerate cycle to maintain the programme schedule. The fatigue accumulated during the sprint phase does not scale proportionally with the delay absorbed during the standby phase; it scales quadratically, because the downstream crew must compress both the absorbed delay and its own planned scope into a shortened execution window, compounding workload intensity, error probability, and recovery time simultaneously.

The companion paper formalizes this within a KPZ (Kardar–Parisi–Zhang) dynamics framework, where the project's fatigue "surface" evolves according to three competing terms: a *smoothing term* (ν) representing management's capacity to level fatigue through rest mandates, crew rotation, and workload redistribution; a *nonlinear growth term* ($\lambda \nabla h^2$) representing the quadratic contagion driven by the fatigue gradient ∇h between entangled crews; and a *stochastic noise term* representing unpredictable shocks (heat events, supply delays, geological anomalies). The system remains stable as long as the smoothing term exceeds the growth term — that is, as long as management can diffuse fatigue peaks faster than they propagate. When the fatigue gradient exceeds the critical threshold $\nabla h_{\text{critical}} = \sqrt{\nu / \lambda}$, the contagion becomes self-sustaining: fatigue peaks propagate laterally faster than management can level them, and the project undergoes a phase transition into systemic instability.

This is why small delays in one sector can produce massive, site-wide safety and productivity degradation. The quadratic term means that doubling the fatigue gap between two entangled crews does not double the propagation stress — it quadruples it. Once the critical gradient is breached, additional management effort (overtime, added supervision, recovery schedules) can itself become a contagion vector, as the interventions impose their own fatigue load on the crews absorbing them.

In QFI terms, quadratic contagion is the mechanism through which fatigue inflates the correlation correction term C in the composite QFI formula: it converts what would otherwise be independent dimensional fragilities into coupled, mutually amplifying fragilities by steepening the effective ρ_{ij} between any dimensions whose execution depends on fatigued, entangled crews.

3.2.3 Fatigue and Superposition Collapse

In QPM, a project exists in a **superposition** of possible outcomes — on-time/late, safe/unsafe, within-budget/over-budget. Fragility measures how easily the system collapses into an undesirable state when perturbed.

Fatigue narrows the basin of attraction around the desired outcome. A well-rested project team can absorb perturbations and maintain the superposition — adapting, compensating, finding workarounds. A fatigued team loses this adaptive capacity. The superposition collapses more readily into degraded states because the cognitive resources required for adaptation are depleted.

This maps to the nonlinearity coefficient η in the decoherence proxy:

$$QFI_{deco} = \eta \cdot \Gamma \cdot (1 - F_Q(t))^{(\alpha-1)}$$

Fatigue increases η — the curvature of the value function with respect to fidelity loss. Under fatigue, the same amount of coherence loss produces a larger value drop because the system's ability to compensate is diminished.

3.2.4 Fatigue and Propagation Velocity

The companion paper models fatigue propagation using a KPZ-inspired framework where the fatigue "surface" evolves dynamically across work packages. The lateral propagation term measures how a fatigue peak in one sector "infects" adjacent tasks through critical-path coupling.

In QFI terms, high **propagation velocity** multiplies local fragility into systemic fragility. When fatigue propagation velocity exceeds management's **diffusion capacity** — the ability to level the fatigue surface through rest mandates, crew rotation, or workload redistribution — the system reaches a tipping point. Beyond this tipping point, fragility becomes self-reinforcing: fatigue causes errors, errors cause rework, rework causes schedule compression, compression causes more fatigue.

This is the **KPZ tipping point** described in the companion paper — the critical fatigue gradient where management can no longer smooth the interface. In QFI governance terms, this tipping point corresponds to a rapid QFI escalation that can push the composite score from Amber to Orange to Red within a single reporting cycle.

3.3 Mathematical Bridge: From FRI to QFI

This section formalizes the mathematical relationship between the Fatigue Risk Index (FRI) and the Quantum Fragility Index (QFI), establishing how fatigue state measurement translates into fragility sensitivity measurement.

3.3.1 Complementary Metrics — State vs. Sensitivity

A critical distinction must be drawn at the outset. The FRI measures the *current fatigue state* — how fatigued the system is right now. QFI_F (fatigue-specific fragility) measures how *sensitive* the system is to changes in fatigue — how close the system is to a governance-boundary crossing under fatigue perturbation.

These are complementary, not redundant:

- **FRI = 55** tells governance: "Fatigue is in the High zone; restrict high-risk tasks."
- **$QFI_F = 0.75$** tells governance: "A small additional fatigue increment will flip the system into the next governance band; intervention is urgent."

A system can have moderate FRI but high QFI_F (centered in a band but near a boundary), or high FRI but lower QFI_F (deep within a band with substantial room before the next boundary). Both pieces of information are needed for complete governance.

Metric	What It Measures	QPM Construct	Governance Question
FRI	Current fatigue state (0–100)	System state variable	"How fatigued is the workforce?"
QFI_F	Sensitivity to fatigue perturbation (0–1)	Proximity to state transition	"How close is fatigue to flipping the governance band?"
FRI trend (dFRI/dt)	Rate of fatigue change	Decoherence velocity	"Is fatigue getting worse or better, and how fast?"
QFI_F velocity (d QFI_F /dt)	Rate of fragility change	Fragility acceleration	"Is the system becoming more or less fragile to fatigue?"

3.3.2 Fatigue-Specific Fragility (QFI_F)




Figure 3.4 Fatigue Proximity Diagram illustrates contagion through temporal and functional adjacency — fatigue intensity rises non-linearly as task proximity increases, driving convergence toward systemic fragility.

Define the directional fatigue fragility using the **proximity-based formulation** from Section 1:

$$QFI_F = 1 - |\Delta F^*| / |\Delta F_{max}|$$

where:

- ΔF^* is the minimum fatigue change required to cross the nearest governance boundary (from the current FRI band to an adjacent band).
- ΔF_{max} is the maximum credible fatigue change under the current scenario.

This is the operational QFI_F. For the **curvature-based formulation** (consistent with Section 2), compute:

$$QFI_{F,TD} = -\partial^2 E[V_{safety}(F)] / \partial \sigma_F^2 |_{\sigma_F = \sigma_{F,0}}$$

where $V_{safety}(F)$ is the safety-performance value function and σ_F is the dispersion of fatigue perturbations.

3.3.3 The FRI–QFI Transmission Function

The relationship between FRI and composite QFI is formalised through three distinct pathways:

Pathway 1 — Direct dimensional contribution. Fatigue directly contributes to the Schedule and Stakeholder QFI dimensions (fatigued crews produce schedule delays and degrade stakeholder confidence). This is captured in the dimensional QFI inputs.

Pathway 2 — Decoherence rate amplification. FRI modulates the decoherence rate Γ used in the decoherence proxy. Higher FRI \rightarrow higher $\Gamma_{\text{effective}}$ \rightarrow higher QFI_{deco} across all dimensions:

$$QFI_{\text{deco},i}(FRI) = \eta_i \cdot \Gamma_{\text{effective}}(FRI) \cdot (1 - F_{Q,i}(t))^{(\alpha_i-1)}$$

Pathway 3 — Correlation amplification. FRI increases the effective correlation between dimensions, amplifying the correction term C in composite QFI. Fatigued systems exhibit higher ρ_{ij} because errors propagate across interfaces that would otherwise contain them.

These three pathways make fatigue a **cross-cutting fragility driver** — it does not reside in a single QFI dimension but amplifies fragility across all dimensions simultaneously. This is why it is classified as an internality rather than an externality.



Figure 3.5 FRI–QFI Transmission Function

3.4 Fatigue Across the Six-Dimension Decomposition

This section examines how fatigue affects each of the six QFI dimensions established in Section 2.3.2, demonstrating its cross-cutting nature as an internality that does not reside in any single dimension but reshapes the fragility landscape of every dimension simultaneously.

3.4.1 Schedule

Fatigue directly degrades schedule performance. Fatigued crews execute tasks more slowly, make more errors requiring rework, and produce more variable output. The schedule QFI increases because the value function becomes more concave — small increases in schedule perturbation produce disproportionately large delays when the workforce executing the recovery is itself fatigued. The critical insight: *schedule compression as a recovery strategy fails under high fatigue because the intervention itself worsens the internality.*

Schedule compression as a recovery strategy fails under high fatigue because the intervention itself worsens the internality.

3.4.2 Cost

Fatigue drives cost overruns through three mechanisms: **rework** (defective work must be redone), **incident-related costs** (investigations, remediation, regulatory penalties), and **productivity degradation** (less output per labor-hour). The cost QFI increases because the cost-performance value function steepens — the marginal cost of each additional perturbation rises as fatigue erodes the efficiency of the workforce absorbing it.

3.4.3 Scope

Under high fatigue, scope management quality degrades. Fatigued teams are more likely to miss scope requirements, accept nonconforming work, or fail to identify scope changes early. The scope QFI increases as the tolerance for scope perturbation narrows.

3.4.4 Stakeholder

Fatigue degrades communication quality, stakeholder engagement, and decision-making speed. Fatigued project managers and supervisors produce lower-quality stakeholder interactions, slower approvals, and more frequent misunderstandings. The stakeholder QFI increases as the system becomes more sensitive to stakeholder perturbations that would normally be absorbed by competent, alert management.

3.4.5 Supply Chain

Fatigue in procurement and logistics teams increases the probability of ordering errors, missed delivery windows, and inadequate quality inspection. For construction specifically, fatigued receiving crews may accept nonconforming materials that propagate defects downstream. The supply chain QFI increases as internal quality gates weaken.

3.4.6 Regulatory

Fatigued workers are more likely to deviate from procedures, bypass safety protocols, and produce documentation errors — all of which increase regulatory exposure. The regulatory QFI increases because the minimum regulatory perturbation required to trigger a compliance violation decreases as procedural adherence erodes under fatigue.

QFI Dimension	Fatigue Impact Mechanism	Effect on Dimensional QFI	Cross-Coupling
Schedule	Slower execution; rework loops; variable output	Steepens schedule value function	Schedule compression worsens fatigue (feedback loop)
Cost	Rework costs; incident costs; productivity loss	Increases marginal cost sensitivity	Cost pressure drives overtime → more fatigue
Scope	Missed requirements; nonconforming acceptance	Narrows scope tolerance	Scope changes increase workload → fatigue
Stakeholder	Degraded communication; slower decisions	Increases stakeholder perturbation sensitivity	Stakeholder pressure drives schedule push → fatigue
Supply Chain	Ordering errors; inspection failures	Weakens internal quality gates	Supply delays force schedule recovery → fatigue
Regulatory	Procedural drift; documentation errors	Lowers compliance perturbation threshold	Regulatory findings require remediation effort → fatigue

Key observation: Every dimension has a feedback pathway back to fatigue. This creates the entanglement structure that makes fatigue a uniquely dangerous internality — it is both a fragility *driver* and a fragility *consequence*, forming self-reinforcing loops that conventional risk management struggles to detect.

3.5 Roughness and Curvature — Convergent Second-Order Diagnostics

This section⁹ reveals a deep structural parallel between two apparently different concepts: **KPZ-based fatigue roughness** (from the companion paper) and **Taleb-Douady curvature** (from Section 2).

Both are second-order measures:

- **Curvature (QFI):** the second derivative of expected value with respect to perturbation dispersion — $\partial^2 E[V] / \partial \sigma^2$
- **Roughness:** the variance (second moment) of the fatigue interface across work packages

Both detect the same phenomenon through different lenses: roughness reveals that "islands of exhaustion" are forming (the fatigue surface is developing peaks and valleys); curvature reveals that the value function is becoming concave (small perturbations produce disproportionate consequences). When roughness is high, curvature is high. They are measuring the same underlying reality — *nonlinear vulnerability* — from complementary vantage points.

The practical convergence: a project with high fatigue roughness (average FRI = 40 but variance spans 15 to 75 across crews) will exhibit higher composite QFI than a project with the same average FRI but low roughness (all crews at 38–42). This is because roughness creates steep risk gradients that exceed management's diffusion capacity — precisely the condition that produces high curvature in the QFI framework.

Governance implication: track both metrics. FRI average tells you the system's central tendency. Fatigue roughness tells you the variance. QFI tells you the sensitivity. Together, they provide a complete diagnostic: where the system is, how uneven it is, and how close it is to a tipping point.

3.6 The Quantum Tunneling Analogy — Fatigue as Barrier Penetration

The companion paper draws an analogy between fatigue behavior and quantum tunneling in radioactive decay. This analogy deepens the QPM interpretation of fatigue-driven fragility.

In quantum mechanics, **tunneling** occurs when a particle penetrates a potential energy barrier that classical physics says it cannot cross. The particle does not climb over the barrier — it tunnels through it. The probability of tunneling increases when: (a) the barrier is narrow, (b) the barrier height is low relative to the particle's energy, or (c) the particle's mass is small.

In QPM, project controls — rest policies, staffing rules, schedule buffers, safety procedures — are analogous to energy barriers. They are designed to prevent the system from transitioning into a

⁹ See Appendix D for a more detailed treatment.

degraded state. Under normal conditions, these barriers are effective: a perturbation must be large enough to "climb over" the control barrier to cause a state transition.

Fatigue thins the barriers. When the workforce is fatigued, controls become permeable:

- A mandatory rest policy is less effective when workers are too fatigued to sleep well.
- A safety procedure is less effective when fatigued workers tunnel through it unconsciously.
- A schedule buffer is less effective when fatigued crews consume it through rework and errors.

In QFI terms, fatigue reduces $|\epsilon^*|$ — the minimum perturbation required to flip the state — because the "barrier" between states has been weakened from within. This is why QFI_F rises as FRI rises: the barriers that separate governance bands become increasingly permeable.

Executive Communication Value

"Our controls are not walls — they are barriers. And fatigue is thinning those barriers. The same perturbation that was safe last month may tunnel through our controls this month because the workforce is too fatigued to enforce them."

3.7 Worked Numerical Example — Fatigue as QFI Driver

This example extends the multidriver worked example from Section 2.3.3, showing how fatigue state changes affect composite QFI through all three transmission pathways.

3.7.1 Baseline (from Section 2.3.3)

Recall the baseline scenario:

- $F_0 = 58$ (fatigue index, 0–100)
- $R_0 = 72$ (schedule reliability index)
- $SC_0 = 0.55$ (supply chain concentration)
- $QFI_{norm} = 0.68$ (Highly Fragile)

FRI equivalent: the fatigue index of 58 maps approximately to $FRI \approx 58$ (High band, 50–70).

3.7.2 Fatigue Fragility Computation (QFI_F)

Current fatigue state: $F = 58$ (in the High band, 50–70).

Nearest governance boundary: the Critical threshold at $FRI = 70$.

Distance to nearest boundary: $\Delta F^* = 70 - 58 = 12$.

Maximum credible fatigue change (scenario): $\Delta F_{\max} = 25$ (a realistic maximum based on a heat event plus overtime push).

$$QFI_F = 1 - |\Delta F^*| / |\Delta F_{\max}| = 1 - 12/25 = 1 - 0.48 = 0.52$$

Interpretation: $QFI_F = 0.52$ — Moderately Fragile to fatigue. The system needs 48% of its maximum credible fatigue change to cross into the Critical band. This is not immediately dangerous but warrants enhanced monitoring.

3.7.3 Scenario A — Fatigue Worsens (Heat Event + Overtime)

A three-day heat event forces extended shifts. FRI rises from 58 to 66.

New QFI_F : $\Delta F^* = 70 - 66 = 4$. $QFI_F = 1 - 4/25 = \mathbf{0.84}$. Now **Highly Fragile** to fatigue.

Impact on composite QFI:

- $\Gamma_{\text{effective}}$ increases by approximately 30% due to fatigue-driven decoherence.
- η increases by approximately 15% due to reduced adaptive capacity.
- ρ_{ij} (Schedule–Cost correlation) increases from 0.55 to 0.70 due to fatigue-driven error propagation.
- New $QFI_{\text{composite}} \approx \mathbf{0.79}$ (up from 0.68).

Governance consequence: the project moves from the middle of the Orange (Highly Fragile) band toward the Red (Critically Fragile) threshold. A single additional perturbation could trigger Red-level governance.

3.7.4 Scenario B — Fatigue Intervention (Mandatory Rest + Crew Rotation)

Management implements a fatigue intervention: mandatory 12-hour rest periods, crew rotation, and deferral of non-critical night work. FRI drops from 58 to 42 (Elevated band).

New QFI_F : $\Delta F^* = |42 - 30| = 12$ (now the nearest boundary is the Normal threshold at 30, below). Distance to Critical: $|70 - 42| = 28 > \Delta F_{\max}$. So ΔF^* for upward transition = 28 (exceeds max credible), meaning QFI_F for upward transition ≈ 0 . Taking the nearest boundary: $\Delta F^* = 12$ (downward to Normal). $QFI_F = 1 - 12/25 = \mathbf{0.52}$.

But the critical difference: the system is now centered in the Elevated band with substantial margin to Critical. The composite QFI effect:

- $\Gamma_{\text{effective}}$ decreases by approximately 20% as decoherence rate normalizes.

- η decreases as adaptive capacity is restored.
- ρ_{ij} reverts toward baseline.
- New $QFI_{\text{composite}} \approx \mathbf{0.55}$ (down from 0.68).

Governance consequence: the project moves from Orange to upper Amber — governance de-escalation is possible.

3.7.5 Summary of Scenarios

Scenario	FRI	QFI_F	Composite QFI	Governance Band	Action
Baseline	58	0.52	0.68	Orange (Highly Fragile)	Weekly review; structural intervention
Heat + Overtime	66	0.84	0.79	Orange/Red boundary	Immediate fatigue intervention; daily review
After Rest Intervention	42	0.52	0.55	Amber (Moderately Fragile)	De-escalation possible; bi-weekly review

This example demonstrates quantitatively that **fatigue interventions are QFI interventions**. Reducing FRI directly reduces composite QFI — and the reduction is amplified because fatigue affects multiple dimensions simultaneously through the three transmission pathways (direct, decoherence amplification, correlation amplification).

3.8 Unified Governance Dashboard — Integrating FRI and QFI

This section describes how to present FRI and QFI together on a single governance dashboard for executive decision-making.

3.8.1 Unified QFI–FRI Governance Dashboard



Figure 3.6 Unified QFI–FRI Governance Dashboard

Integrates fragility and fatigue telemetry into a single executive interface — enabling real-time governance decisions across operational, regulatory, and resource domains.

The Unified QFI–FRI Governance Dashboard consolidates fragility and fatigue telemetry into a single executive interface designed for rapid situational assessment. The upper title band establishes context with a live-telemetry timestamp, reinforcing the dashboard’s role as an operational decision instrument.

The **FRI State Trend** panel presents a seven-week trajectory of fatigue risk, plotted against governance color bands (green, amber, red). The current FRI of 58 reflects a rising pattern from earlier weeks, with roughness indicators ($W = 14.2$; $\nabla h_{max} = 18$, exceeding the $\nabla h_{critical} \approx 15$ threshold) signaling elevated volatility in fatigue dynamics.

Adjacent to this, the **QFI_F Fragility Gauge** provides an at-a-glance assessment of systemic fragility. The semi-circular indicator sits at 0.52 within the amber zone, accompanied by distance-to-threshold markers (“12 pts to Critical,” “8 pts to Elevated”) that contextualize the organization’s proximity to governance boundaries.

The **Composite QFI × Fatigue Radar** visualizes six operational dimensions—Schedule, Supply Chain, Regulatory, Cost, Technical, and Workforce. A red-orange fatigue halo overlays the radar, amplifying fragility across axes and yielding a composite value of 0.68. This panel highlights how fatigue interacts with structural vulnerabilities to shape overall system fragility.

The **Intervention Simulator** ranks corrective actions by cost-efficiency and expected impact. Options include deferring non-critical night work ($\Delta QFI = -0.10$ at \$0.4 M), adding relief crews ($\Delta QFI = -0.19$ at \$2.8 M per month), and resequencing regulatory deliverables ($\Delta QFI = -0.06$ at \$0.1 M). This panel translates telemetry into actionable governance levers.

Together, these components form a unified governance dashboard that transforms complex fragility–fatigue interactions into a coherent, decision-ready operational picture.

3.8.2 Governance Decision Logic

The following decision rules govern combined FRI–QFI governance:

FRI Band	QFI _F	Combined Interpretation	Governance Action
Normal (<30)	Low (<0.30)	System is rested and robust to fatigue perturbation	Standard monitoring; no fatigue-specific action
Normal (<30)	High (>0.60)	System is currently rested but near a boundary — vulnerable to rapid transition	Proactive monitoring; prepare contingency if upcoming schedule intensifies
Elevated (30–50)	Low (<0.30)	Moderate fatigue but well-buffered from boundaries	Enhanced monitoring; standard fatigue management
Elevated (30–50)	High (>0.60)	Moderate fatigue and close to a boundary — high fragility	Active intervention: microbreaks, workload review, schedule adjustment
High (50–70)	Any	Significant fatigue regardless of QFI _F	Restrict high-risk tasks; add relief crews; adjust shifts
Critical (>70)	Any	Immediate danger; system in critical state	Stop or rescope work; mandatory rest; daily governance review

3.8.3 Threshold Alignment

For operational consistency, align FRI and QFI governance thresholds:

FRI Band	FRI Range	Approximate QFI _F Range	QFI Governance Tier
Normal	0–30	0–0.30	Blue (Robust)
Elevated	30–50	0.30–0.60	Amber (Moderately Fragile)

FRI Band	FRI Range	Approximate QFI _F Range	QFI Governance Tier
High	50–70	0.60–0.85	Orange (Highly Fragile)
Critical	>70	≥0.85	Red (Critically Fragile)

This alignment ensures that fatigue governance actions and QFI governance actions reinforce rather than contradict each other. When FRI crosses from Elevated to High, QFI governance should also escalate from Amber to Orange-level oversight.

3.9 Fatigue, Roughness, and the Dimensional Floor Rule

Section 2.3.2 establishes the **dimensional floor rule**: if any $QFI_i \geq 0.85$, trigger Orange-level governance regardless of composite QFI. This section extends the floor rule to fatigue.

3.9.1 Fatigue as a Cross-Dimensional Detonator

Unlike the six formal QFI dimensions, fatigue is not a single dimension — it is a cross-cutting internality that can act as a **detonator**¹⁰ across multiple dimensions simultaneously. A fatigue crisis does not merely push one QFI_i above 0.85; it can push several dimensions toward their thresholds simultaneously by degrading the internal controls that keep each dimension stable.

This creates a scenario the standard dimensional floor rule does not anticipate: no single dimension exceeds 0.85, but several are elevated (e.g., Schedule QFI = 0.72, Cost QFI = 0.68, Regulatory QFI = 0.70) — and fatigue is the common driver pushing all three. The composite may average to 0.65 (Orange), but the coordinated elevation across dimensions represents a far more dangerous condition than a single-dimension spike.

3.9.2 Proposed Fatigue Floor Rule

When $FRI \geq 70$ (Critical), apply Orange-level governance regardless of composite QFI or any individual dimensional QFI. This fatigue floor rule is analogous to the dimensional floor rule but applies to the internality rather than any single external dimension.

Rationale: At Critical FRI, the system's internal control integrity is so degraded that external perturbations of any magnitude become dangerous. The fatigue floor rule ensures that internality-driven fragility receives the same governance escalation as externality-driven fragility.

¹⁰ A detonator is a fragility spike concentrated enough to produce catastrophic failure on its own, and the floor rule exists to ensure governance sees it even when the composite average says everything is fine.

Additionally: When $FRI \geq 50$ AND fatigue roughness exceeds a calibrated threshold (indicating "islands of exhaustion"), escalate to at least Amber governance even if composite QFI remains in the Blue zone. This captures the scenario where average fatigue is moderate but local peaks create dangerous gradients that QFI curvature alone may not fully reflect.

3.10 Integration with Safety-II

Section 2.6 establishes that QFI is a Safety-II metric — it measures proximity to failure rather than the occurrence of failure. Fatigue deepens this Safety-II alignment.

Safety-II defines safety as the ability of a system to succeed under varying conditions. Fatigue directly degrades this ability. As the companion paper establishes, fatigue is "a state variable that changes the reliability of every other control system." In Safety-II terms, fatigue reduces the system's **adaptive capacity** — its ability to detect, compensate, and recover from variability.

The FRI–QFI integration operationalizes Safety-II's four resilience potentials in the context of fatigue:

Monitor. FRI provides continuous fatigue state monitoring; QFI_F provides continuous fragility monitoring. Together, they track both the condition (fatigue level) and its consequence (proximity to state transition).

Anticipate. Predictive FRI modelling (using planned schedules, weather forecasts, and historical fatigue patterns) enables forward-looking QFI projections. *"If we maintain the current overtime schedule for two more weeks, QFI_F will rise from 0.52 to 0.78."*

Respond. The what-if simulator enables governance teams to test fatigue interventions before implementation and rank them by QFI reduction per unit cost. *"Adding relief crews reduces $QFI_{composite}$ by 0.13 at a cost of \$X per week."*

Learn. Post-project analysis of FRI–QFI trajectories¹¹ reveals which fatigue management strategies were most effective at reducing fragility. These findings feed the QFI Pattern Library (Section 2.7.1, Phase 3).

¹¹ The concept of systematic post-project trajectory analysis draws on three convergent intellectual streams. First, the Safety-II framework defines the Learn potential as "knowing what has happened — being able to learn from experience, in particular to learn the right lessons from the right experience" (Hollnagel, 2014, p. 148), and the Resilience Assessment Grid (Hollnagel et al., 2011) operationalises this potential through structured retrospective assessment of how systems succeeded under variability — not merely why they failed. Second, Rasmussen's (1997) dynamic safety model demonstrates that systems migrate toward the boundary of acceptable performance through incremental, locally rational decisions; retrospective trajectory analysis reconstructs this migration path so that governance can identify where and when the drift accelerated — precisely what FRI–QFI co-trajectories reveal when analysed post-project. Dekker (2011)

3.11 Practice Implications and Deployment Guidance

3.11.1 Integrating FRI into the QFI Measurement Protocol

Section 2.7.2 specifies data inputs for each QFI dimension. FRI data should be added as a cross-cutting input:

Input	Source	Update Frequency	QFI Pathway
FRI composite score	FRI engine (shift data, environmental feeds, behavioral indicators)	Weekly (daily for FRI > 50)	Decoherence rate modulation; correlation adjustment
FRI roughness	Crew-level FRI variance	Weekly	Correlation correction amplification
FASCW scores	Worker self-reports	Daily (pre-shift)	Qualitative validation of FRI trends
Fatigue trend (dFRI/dt)	FRI time series	Weekly	QFI velocity estimation

3.11.2 Calibration Considerations

The FRI–QFI transmission function requires calibration to each organization’s context:

- The decoherence rate amplification function ($\Gamma_{\text{fatigue}} = \gamma_0 \cdot (\text{FRI}/100)^{\beta_f}$) must be calibrated using historical project data — correlating observed FRI levels with measured rates of plan-versus-execution divergence.

extends this insight, showing that drift into failure is only visible in hindsight when the full performance trajectory is reconstructed from contemporaneous data, reinforcing the need to archive FRI and QFI time-series data for post-project review. Third, Leveson (2011) formalises the role of leading-indicator feedback loops in her STAMP model, arguing that safety control structures must include mechanisms to monitor system state trajectories and learn from deviations before they produce loss — a principle directly applicable to the FRI-QFI governance dashboard described in Section 3.8. In construction specifically, Hallowell et al. (2013) identify over fifty proactive safety metrics suitable for longitudinal tracking during the construction phase, demonstrating that trajectory-based learning from leading indicators produces measurably superior safety outcomes compared to reactive, lagging-indicator analysis alone. Most recently, de Guingand (2025) reports on a two-year longitudinal fatigue study across the HS2 high-speed rail programme in the United Kingdom, generating over 100,000 self-reported fatigue data points across diverse roles, sites, and shift types — providing the first large-scale empirical demonstration that fatigue trajectories can be tracked, analysed, and used to inform targeted interventions on a major infrastructure project. Taken together, these sources establish both the theoretical mandate and the practical feasibility of retrospective FRI-QFI trajectory analysis as a governance learning mechanism.

- The correlation amplification effect requires empirical estimation of how fatigue changes the effective ρ_{ij} between dimensions.
- The dimensional impact weights (how strongly fatigue affects each QFI dimension) may vary by project type, sector, and delivery model.

Phase 1 deployment can use simplified assumptions (linear FRI–QFI relationship, equal dimensional impacts) and refine through Phase 2 and Phase 3 calibration as empirical data accumulates.

3.11.3 Digital Twin Integration

The companion paper describes a digital twin architecture for predictive fatigue modelling. This digital twin should be integrated with the QFI framework:

- The fatigue digital twin feeds predicted FRI trajectories into the QFI model.
- The QFI model translates predicted FRI into projected $QFI_{\text{composite}}$ trajectories.
- The combined system answers the executive question: *"If we continue the current schedule and work patterns, when will QFI cross the next governance threshold — and is fatigue the driver?"*

3.12 Summary

Fatigue is not merely one risk among many — it is a cross-cutting internality that reshapes the fragility architecture of the entire project system. Unlike externalities that increase the magnitude of perturbations, fatigue increases the system's sensitivity to perturbations by degrading the internal controls that preserve coherence.

Fatigue is not merely one risk among many — it is a cross-cutting internality that reshapes the fragility architecture of the entire project system.

The integration of FRI and QFI produces a unified governance framework where:

- **FRI** measures the fatigue state.
- **QFI_F** measures the fragility sensitivity to fatigue.
- The **three transmission pathways** (direct, decoherence amplification, correlation amplification) formalize how fatigue feeds into composite QFI.
- The **fatigue floor rule** ensures that critical internality-driven fragility receives appropriate governance escalation.

- **Safety-II alignment** grounds the integrated framework in forward-looking, variability-aware resilience management.

For governance leaders, the message is clear: **managing fatigue is managing fragility**. Every intervention that reduces FRI also reduces QFI — and the effect is amplified because fatigue operates across all dimensions simultaneously. In the QPM framework, fatigue management is not a human-resources function — it is a system-level fragility control that belongs at the center of project governance.

4. The Quantum Fragility Index and the Statistical Fragility Index: What the Quantum Mapping Adds

Sections 2 and 3 of this paper develop the Quantum Fragility Index (QFI) by extending the Statistical Fragility Index (SFI) formulated by Taleb and Douady (2013) through a mapping onto four constructs drawn from quantum mechanics: decoherence, entanglement, superposition, and tunnelling. This section examines the relationship between the two indices directly — what they share, where they diverge, and why the quantum mapping produces a governance instrument that the statistical formulation alone cannot provide. The intent is not to diminish the SFI; it remains a foundational contribution. The intent is to show, with mathematical precision and operational specificity, what is gained when the SFI's curvature architecture is embedded within a quantum-mechanical state-evolution framework and extended with fatigue integration, dimensional decomposition, and spatial-variance diagnostics.

4.1 The Statistical Fragility Index: Architecture and Achievements

4.1.1 Core Definition

Taleb and Douady (2013) define fragility as negative sensitivity to a semi-measure of dispersion — a variant of negative "vega" in options-pricing language. For a payoff function $v(x)$ subject to perturbations drawn from a distribution $f(x; \sigma)$ parameterized by dispersion σ , the Statistical Fragility Index is:

$$SFI = -\partial^2 E[v(x)] / \partial \sigma^2$$

where $E[v(x)] = \int v(x) \cdot f(x; \sigma) dx$ is the expected payoff. The sign convention is deliberate: a negative second derivative (concave expected payoff) indicates fragility; a positive second derivative (convex expected payoff) indicates antifragility; zero indicates robustness. The SFI is, at its core, a curvature measure — it detects the shape of the payoff function under stress, not the level.

4.1.2 Key Properties

The SFI possesses several properties that make it a powerful advance over conventional risk metrics:

Universality. Unlike risk measures based on subjective preferences (utility functions, risk aversion coefficients), the SFI applies to any system with a probability distribution — "whether such distribution is known or, critically, unknown" (Taleb & Douady, 2013, p. 1678). A coffee cup, a bridge, a bank balance sheet, and a construction project can all be assessed on the same fragility scale.

Model-freedom. The "fast-and-frugal" detection heuristic proposed by Taleb and Douady requires no distributional assumptions. By evaluating the payoff at three dispersion levels ($\sigma - \Delta\sigma$, σ , $\sigma + \Delta\sigma$) and computing the finite-difference second derivative, the heuristic detects fragility without specifying a parametric model — thereby avoiding the model error that plagues Value-at-Risk and other parametric risk measures.

Nonlinearity detection. The SFI directly captures the link between nonlinear payoffs and fragility. A system with a concave payoff (losses accelerate faster than gains) is fragile; a system with a convex payoff (gains accelerate faster than losses) is antifragile. This connection to Jensen's inequality provides the SFI with a deep mathematical foundation.

K-fragility. Taleb and Douady introduce the concept of K-fragility — fragility conditional on a threshold K. A system may be robust to perturbations below K but fragile to perturbations above K. This threshold-conditional formulation anticipates the governance-band architecture developed in Section 2 of this paper.

Superiority to VaR. As Taleb and Douady demonstrate, the SFI outperforms Value-at-Risk because VaR captures only the location of a percentile in the loss distribution, not the shape of the payoff function that generates losses. A system can have identical VaR but vastly different SFI — one robust, the other fragile — because VaR is a first-order measure while SFI is a second-order measure.

4.1.3 What the SFI Does Not Address

Despite these achievements, the SFI as formulated by Taleb and Douady operates within a set of boundaries that limit its applicability to complex project governance:

Static. The SFI is a point-in-time measurement. It evaluates curvature at a given dispersion level σ_0 but provides no mechanism for tracking how fragility evolves over time as the project progresses, as perturbation environments change, or as the system's internal state degrades.

Univariate. The SFI operates on a single payoff function $v(x)$ subject to a single perturbation variable x with a single dispersion parameter σ . It does not decompose fragility across multiple governance dimensions or capture cross-dimensional coupling.

Exogenous only. The SFI measures sensitivity to externally imposed perturbation dispersion. It does not distinguish between externalities (forces that act on the system from its environment) and internalities (intrinsic conditions that alter the system's own control reliability). Fatigue, the paradigmatic internality, has no representation in the SFI architecture.

No temporal dynamics. The SFI has no analogue of decoherence — the progressive loss of operational coherence over time. It can measure curvature at t_1 and again at t_2 , but it provides no model of the process by which curvature changes between measurements.

No coupling model. The SFI does not model entanglement — the mechanism by which perturbation in one dimension of the system propagates to and amplifies fragility in another. It treats the payoff function as a monolithic object rather than a composite of interdependent governance dimensions.

No state coexistence. The SFI does not represent superposition — the condition in which a project simultaneously inhabits multiple possible states (on-schedule AND at-risk, compliant AND drifting) until a measurement (audit, milestone review) collapses the state to a definite outcome.

No barrier penetration. The SFI does not model tunnelling — the phenomenon by which a system can transition to a failure state even when its current trajectory does not classically predict such a transition, because random fluctuations allow the system to penetrate barriers that deterministic analysis treats as impenetrable.

These are not criticisms of the SFI. Taleb and Douady designed a universal, model-free fragility measure for financial and general applications. They did not intend it as a dynamic project governance instrument. The quantum mapping developed in this paper addresses each of these gaps by embedding the SFI's curvature architecture within a richer ontological framework.

4.2 From SFI to QFI: The Seven Extensions

This section traces each extension from the SFI to the QFI, showing what the quantum mapping adds to the statistical foundation.

4.2.1 Extension 1 — Temporal Dynamics via Decoherence

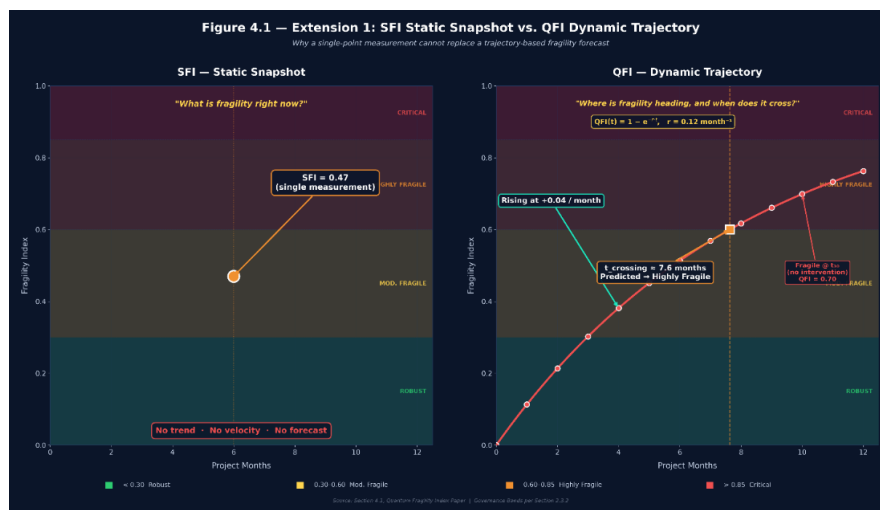


Figure 4.1 SFI Static Snapshot vs QFI Dynamic Trajectory

SFI limitation: Static curvature measurement at a single point in time.

QFI addition: The decoherence construct (Section 2.2.2) provides a dynamic model of how fragility evolves. The quantum fidelity $F_Q(t) = e^{(-\Gamma t)}$ tracks the progressive divergence between the project's planned state ρ_0 and its actual operational state $\rho(t)$. The decoherence rate Γ is not merely a parameter — it is an observable, estimated from "the observed rates of divergence between plan and execution (e.g., cumulative variance in key observables per unit time)" (Section 2.2.2).

This gives governance three capabilities the SFI alone cannot provide:

1. **Trend detection.** QFI can be computed as a trajectory $QFI(t_1), QFI(t_2), \dots, QFI(t_n)$, enabling governance to detect whether fragility is rising, falling, or stable — and at what rate.
2. **Predictive horizon.** Given the current Γ and the governance-band thresholds, governance can estimate when the system will cross from one band to another: $t_{crossing} = -\ln(F_{threshold}) / \Gamma$. This enables anticipatory intervention rather than reactive response.
3. **Fatigue integration.** The decoherence construct provides the mathematical hook for internalities (Section 3.2.1): $\Gamma_{effective} = \Gamma_{baseline} + \gamma_0 \cdot (FRI / 100)^{\beta f}$. Fatigue accelerates decoherence — an effect invisible to the SFI because the SFI has no decoherence channel.

What is preserved: The curvature operator. QFI retains the SFI's core definition — the negative second derivative of expected value with respect to perturbation dispersion. Decoherence does not replace curvature; it explains why curvature changes over time.

4.2.2 Extension 2 — Cross-Dimensional Coupling via Entanglement

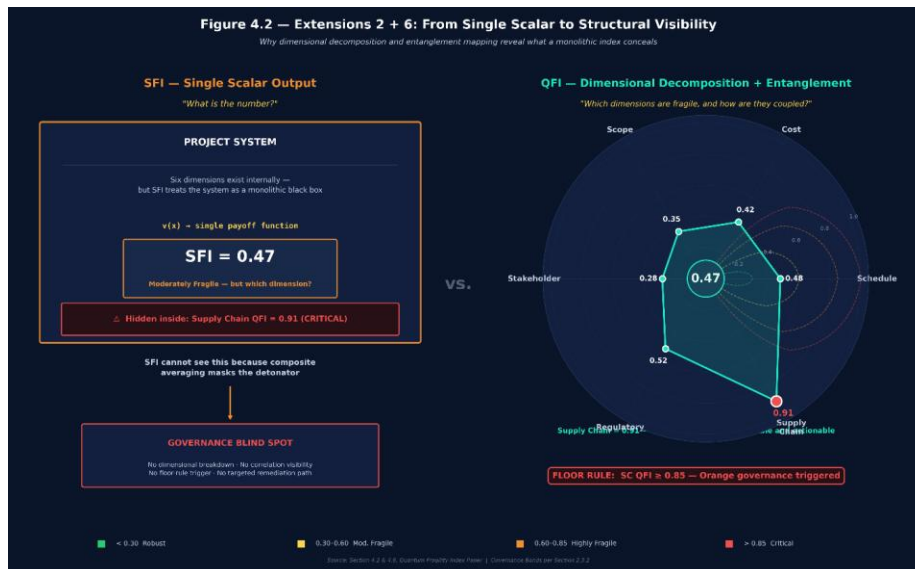


Figure 4.2 Extension 2 — Cross-Dimensional Coupling via Entanglement (SFI → QFI comparison)

SFI limitation: Univariate payoff function; no mechanism for cross-dimensional interaction.

QFI addition: The entanglement construct (Section 2.2.3) models the correlation structure between governance dimensions through the composite QFI formula:

$$QFI_{\text{composite}} = \sum_i w_i \cdot QFI_i + C, \text{ where } C = \sum_{i \neq j} \rho_{ij} \cdot \sqrt{(w_i \cdot w_j)} \cdot \phi(QFI_i, QFI_j)$$

The correlation correction term C captures a phenomenon that the SFI cannot: the amplification of fragility through cross-dimensional coupling. When schedule pressure (dimension 1) causes overtime, which causes fatigue (internality), which causes quality failures (dimension 3), which causes regulatory findings (dimension 6), the fragility of the composite system exceeds the weighted sum of its dimensional fragilities by the magnitude of C. The SFI, operating on a monolithic payoff function, cannot decompose this cascade or quantify its amplification.

Entanglement also enables the dimensional floor rule: any $QFI_i \geq 0.85$ triggers Orange-level governance regardless of the composite. This "detonator" detection is impossible with a univariate SFI because the SFI produces a single number — it cannot identify which dimension is driving the fragility or whether a single dimension has reached critical mass while others mask it.

What is preserved: The curvature computation within each dimension. Each QFI_i is computed using the same Taleb–Douady operator applied to dimension i's payoff function. The extension is structural (decomposition and coupling), not a replacement of the curvature mechanism.

4.2.3 Extension 3 — State Coexistence via Superposition

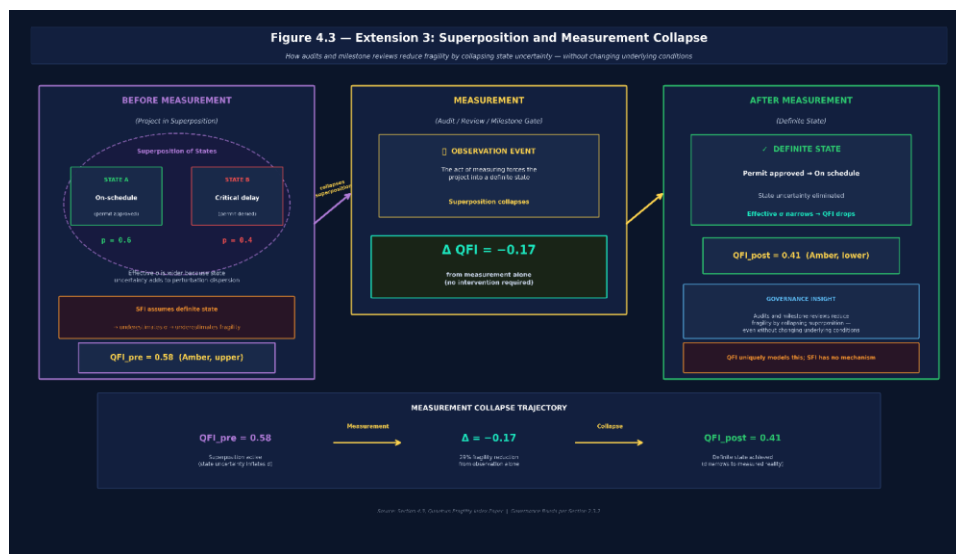


Figure 4.3 — Superposition and Measurement Collapse

SFI limitation: The system is assumed to be in a definite state; the payoff function evaluates perturbations against that definite state.

QFI addition: The superposition construct (Section 2.2.1) represents the condition — pervasive in complex projects — in which the system simultaneously inhabits multiple possible states. A project may be simultaneously on-schedule (if the pending permit is approved) and critically delayed (if it is denied). Until the measurement occurs (the permit decision), the project exists in a superposition of both states.

This has a concrete impact on fragility measurement: a system in superposition has a broader effective perturbation distribution than a system in a definite state, because the uncertainty about which state it occupies widens the range of possible outcomes. The QFI captures this broader distribution through the perturbation density $f(x; \sigma)$, where σ incorporates state-uncertainty as a component of dispersion. The SFI, which assumes a definite state, systematically underestimates dispersion for systems in superposition and therefore underestimates fragility.

Superposition also explains why audit and milestone reviews matter for fragility: they collapse the superposition, forcing the system into a definite state and reducing the effective σ . Post-measurement QFI is typically lower than pre-measurement QFI — not because the system became less fragile, but because the state uncertainty component of dispersion was eliminated. The SFI has no mechanism to represent or exploit this effect.

4.2.4 Extension 4 — Barrier Penetration via Tunnelling

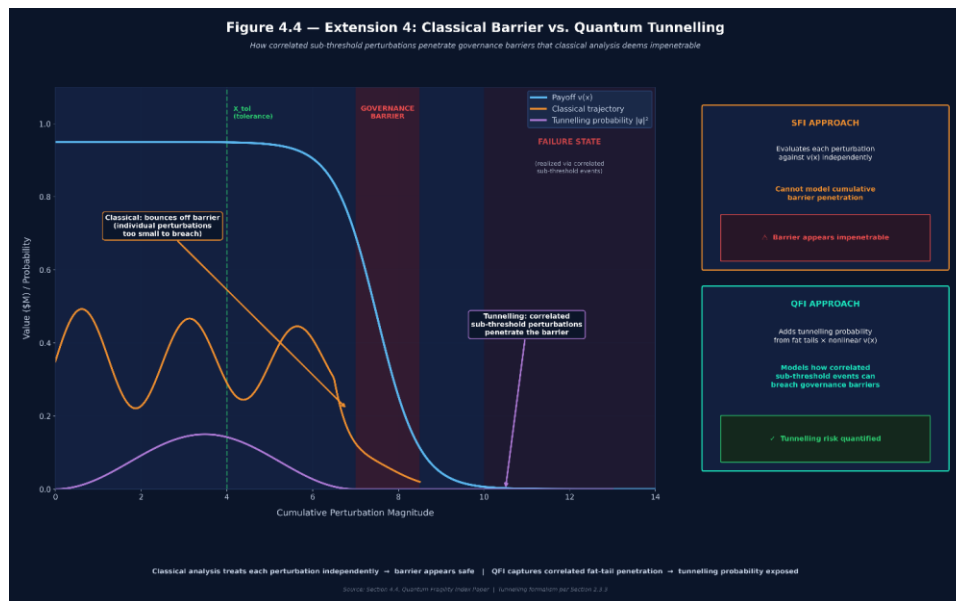


Figure 4.4 — Classical Barrier vs. Quantum Tunnelling

SFI limitation: Fragility is assessed against a continuous payoff function; the possibility of discontinuous state transitions (catastrophic jumps) beyond classical barriers is not modelled.

QFI addition: The tunnelling construct (Section 2.2.4) represents the phenomenon by which a project can transition to a failure state even when no single perturbation is large enough to cause the transition classically. In quantum mechanics, tunnelling occurs when a particle penetrates an energy barrier that classical mechanics says is impassable, because the particle's wave function has non-zero probability on the far side of the barrier.

In project governance, tunnelling manifests as the accumulation of individually sub-threshold perturbations that collectively breach a barrier. No single cost overrun, schedule delay, or quality failure is catastrophic on its own — but their correlated accumulation over time can push the system through a governance boundary that no individual perturbation could have breached. The SFI evaluates each perturbation against the payoff function independently; it does not model the cumulative penetration of barriers by correlated sub-threshold events.

The QFI framework captures tunnelling through the fat tails of the Student-t perturbation distribution ($df = 4$) and the nonlinear payoff function $v(x) = V_0 \cdot \exp(-\beta_1 \cdot \max(0, x - x_{tol})^{\beta_2})$. The combination of fat tails and steep curvature produces non-negligible probability mass beyond the tolerance threshold x_{tol} — precisely the tunnelling probability. The SFI can, in principle, use fat-tailed distributions, but it does not label or operationalize the resulting barrier-penetration probability as a distinct governance signal.

4.2.5 Extension 5 — Internality Integration via the FRI Bridge

SFI limitation: Fragility is driven by external perturbation dispersion σ ; the system's own internal condition does not modulate fragility.

QFI addition: Section 3 of this paper introduces the internality–externality distinction and develops the mathematical bridge between the Fatigue Risk Index (FRI) and the QFI. Fatigue does not increase σ (the magnitude of external perturbations); it increases Γ (the rate at which the system loses coherence) and η (the sensitivity of value to decoherence). A fatigued project is not hit harder — it absorbs hits less well.

This distinction is invisible to the SFI because the SFI has only one channel through which fragility can change: the perturbation dispersion σ . The QFI has three channels, corresponding to the three transmission pathways developed in Section 3.3.3:

4. **Direct dimensional** (FRI \rightarrow QFI_i): Fatigue degrades performance in each governance dimension directly.
5. **Decoherence amplification** (FRI $\rightarrow \Gamma \rightarrow$ QFI_{deco}): Fatigue accelerates coherence loss, steepening curvature.
6. **Correlation amplification** (FRI $\rightarrow \rho_{ij} \rightarrow C$): Fatigue increases cross-dimensional coupling through quadratic contagion, inflating the correlation correction.

None of these channels exist in the SFI architecture. A purely statistical fragility assessment of two projects with identical payoff functions and identical perturbation environments would produce identical SFI values — even if one project's workforce is at FRI = 25 (Normal) and the other's is

at FRI = 72 (Critical). The QFI would correctly show the second project as far more fragile because its decoherence rate is higher, its dimensional fidelities are lower, and its cross-dimensional correlations are elevated.

4.2.6 Extension 6 — Multi-Dimensional Decomposition and Floor Rules

SFI limitation: A single scalar output with no dimensional visibility.

QFI addition: The six-dimension QFI decomposition (Section 2.3.2) — Schedule, Cost, Scope, Stakeholder, Supply Chain, Regulatory — provides governance with a dimensional radar that the SFI cannot offer. Each dimension carries its own QFI_i , weight w_i , fidelity $F_{Q,i}(t)$, and decoherence rate Γ_i . The composite QFI aggregates these via the weighted-sum-plus-correlation formula (F-10), but the dimensional components remain individually visible for diagnosis and intervention targeting.

The floor rules are a direct consequence of this decomposition:

- **Dimensional floor:** $QFI_i \geq 0.85$ triggers Orange governance regardless of composite — the "detonator" principle.
- **Fatigue floor:** $FRI \geq 70$ triggers Orange governance regardless of QFI_i or composite — the internality protection.
- **Roughness floor:** $FRI \geq 50$ AND $W > W_{critical}$ triggers Amber governance — the spatial-variance early warning.

These escalation mechanisms are structurally impossible with a univariate SFI. The SFI can detect overall fragility but cannot diagnose its dimensional source, cannot identify single-dimension detonators masked by composite averaging, and cannot enforce dimensional or internality-based governance overrides.

4.2.7 Extension 7 — Spatial Variance via KPZ Roughness

SFI limitation: Curvature measures the shape of the payoff function in the perturbation domain. It does not measure the spatial distribution of system conditions across the project's workforce or work packages.

QFI addition: Appendix D and Section 3.5 introduce KPZ roughness $W(L, t) = \langle (h - \langle h \rangle)^2 \rangle^{1/2}$ as a complementary second-order diagnostic. While the SFI (and the QFI curvature operator) measures how the expected payoff changes as perturbation dispersion increases, roughness measures how unevenly fatigue is distributed across the workforce at a given point in time.

The two measures detect the same underlying vulnerability — nonlinear sensitivity to small perturbations — through complementary observational windows. Three projects with identical average FRI and identical SFI values can have vastly different roughness ($W = 2.1, 8.7, 22.4$) and correspondingly different composite QFI values (0.32, 0.55, 0.78). Roughness captures the spatial variance that the SFI's payoff-domain curvature systematically misses.

The convergence between roughness and curvature (Appendix D.4) provides a cross-validation mechanism: when both are elevated and trending in the same direction, governance confidence in the fragility diagnosis is high. When they diverge, the divergence itself is diagnostic. The SFI, operating alone, has no such cross-validation capability.

4.3 Structural Comparison Matrix

The following table provides a systematic, property-by-property comparison of the SFI and QFI across 18 dimensions of analysis.

Property	Statistical Fragility Index (SFI)	Quantum Fragility Index (QFI)
Core mathematical operation	Second derivative of $E[v]$ w.r.t. σ	Same — plus decoherence, entanglement, superposition, tunnelling overlays
Temporal character	Static (point-in-time)	Dynamic (trajectory with decoherence rate Γ)
Dimensional structure	Univariate (single payoff $v(x)$)	Six governance dimensions with dimensional QFI _i , weights w_i , and floor rules
Cross-dimensional coupling	None	Correlation correction C via entanglement coefficients ρ_{ij}
State representation	Definite state	Superposition of coexisting states until measurement collapses
Barrier transitions	Not modelled	Tunnelling probability from fat tails + nonlinear payoff
Internality channel	None (exogenous σ only)	Fatigue $\rightarrow \Gamma_{\text{effective}}, \eta, \rho_{ij}$ via three transmission pathways
Coherence tracking	None	Fidelity $F_Q(t) = e^{(-\Gamma t)}$ with observable decoherence rate

Property	Statistical Fragility Index (SFI)	Quantum Fragility Index (QFI)
Predictive capability	None (measures current curvature)	Time-to-threshold: $t_{\text{crossing}} = -\ln(F_{\text{threshold}}) / \Gamma$
Spatial variance	None	KPZ roughness $W(L, t)$ as complementary second-order diagnostic
Governance escalation	Not formalised	Four-band governance (Blue/Amber/Orange/Red) with floor rules
Intervention targeting	Cannot diagnose dimensional source	Dimensional decomposition enables targeted intervention
Portfolio comparability	Limited (raw curvature units vary)	Normalized QFI_{norm} on $[0, 1]$ scale enables cross-project comparison
Model requirements	Model-free (heuristic)	Monte Carlo with payoff model and perturbation distribution
Data requirements	Three $E[v]$ evaluations at different σ	Telemetry across six dimensions; FRI data; Monte Carlo infrastructure
Computational cost	Very low	Moderate to high
Fatigue sensitivity	Zero	Full three-pathway integration (Section 3.3.3)
Safety-II alignment	Not addressed	Respond, Monitor, Anticipate, Learn (Section 3.10)

4.4 What the SFI Gets Right That the QFI Preserves

The comparison must acknowledge what the SFI contributes that the QFI deliberately preserves and builds upon.

4.4.1 The Curvature Insight

The SFI's foundational insight — that fragility resides in the shape of the payoff function, not in the probability or expected magnitude of adverse events — is the conceptual bedrock of the QFI. Every QFI computation, whether at the composite level or the dimensional level, ultimately rests on the Taleb–Douady curvature operator. The quantum mapping does not replace this operator; it enriches the ontological framework within which the operator operates.

4.4.2 The Distinction from Risk

Taleb and Douady's insistence that fragility is not risk — that fragility is a property of the system while risk is a property of the environment — is carried forward entirely into the QFI framework. QFI measures the system's sensitivity to perturbation, not the probability or severity of the perturbation itself. This distinction is operationally critical: two projects facing identical perturbation environments (same σ) can have vastly different QFI values because their payoff functions have different shapes. The SFI established this distinction; the QFI inherits it.

4.4.3 The Antifragility Spectrum

The SFI's tripartite classification — fragile (negative curvature), robust (zero curvature), antifragile (positive curvature) — is preserved in the QFI governance bands. The Blue band ($QFI_{norm} < 0.30$) includes robust systems; negative QFI_{norm} values (not operationalized in the current governance framework but mathematically permissible) would indicate antifragility. The concept that systems can benefit from perturbation — not merely survive it — is an SFI contribution that the QFI inherits and that purely risk-based frameworks cannot express.

4.4.4 The Model-Error Sensitivity

Taleb and Douady emphasize that fragile systems are disproportionately exposed to model error — because model error increases effective dispersion, and fragile systems suffer disproportionately from increased dispersion. This insight is preserved in the QFI's use of heavy-tailed perturbation distributions (Student-t, $df = 4$) and in the tunnelling construct, which explicitly models the consequences of distributional uncertainty.

4.5 What the Quantum Mapping Adds: A Summary

The seven extensions can be distilled into three meta-capabilities that the quantum mapping adds to the statistical foundation:

4.5.1 Temporal Intelligence

The SFI tells governance how fragile the system is now. The QFI tells governance how fragile the system is now, how fragility is changing (rising, falling, stable), how fast it is changing (Γ), why it is changing (decoherence drivers), and when it will cross the next governance threshold ($t_{crossing}$). This temporal intelligence transforms fragility measurement from a diagnostic snapshot into a prognostic trajectory.

4.5.2 Structural Visibility

The SFI produces a single number. The QFI produces a structured diagnostic: six dimensional QFI_i values, a correlation correction C, a fatigue-specific QFI_F, a roughness W, and a composite QFI_{norm} — all cross-referenced through the three transmission pathways. This structural visibility enables governance to diagnose the source of fragility, target interventions to the most effective pathway, and verify that interventions are producing the expected effect.

4.5.3 Internality Sensitivity

The SFI is blind to internalities. The QFI detects, quantifies, and tracks the most consequential internality in project systems — workforce fatigue — through the FRI bridge (Section 3), the three transmission pathways, the fatigue floor rule, the roughness diagnostic, and the unified governance dashboard (Section 3.8). This internality sensitivity is not an optional add-on; it addresses a class of fragility driver that is invisible to any purely statistical measure.

4.6 Limitations of the Quantum Mapping

Intellectual honesty requires acknowledging what the quantum mapping costs:

Computational overhead. The SFI's fast-and-frugal heuristic requires three payoff evaluations. The QFI requires Monte Carlo simulation across six dimensions with perturbation sweeps, decoherence tracking, and correlation estimation. The computational cost is orders of magnitude higher.

Model dependency. The SFI is model-free. The QFI requires a parameterized payoff function $v(x)$, a perturbation distribution $f(x; \sigma)$, governance weights w_i , and calibrated parameters ($\Gamma, \eta, \alpha, \gamma_0, \beta_f, \rho_{ij}$). Miscalibration of any parameter can produce misleading results. The SFI's model-freedom is a genuine advantage for rapid, first-pass assessment.

Data requirements. The SFI can be computed from minimal data. The QFI requires project telemetry across six dimensions, FRI data from workforce monitoring, and historical calibration data for parameter estimation. Projects without mature data infrastructure may not be able to support full QFI deployment.

Complexity barrier. The SFI's conceptual simplicity — "how does the payoff change when volatility changes?" — is accessible to any executive. The QFI's quantum-mechanical framing, while operationally powerful, introduces conceptual complexity that can impede adoption unless carefully communicated. The governance-band architecture (Section 2.3.3) and the unified dashboard (Section 3.8) are designed to mitigate this barrier, but they add abstraction layers that the SFI does not require.

Metaphorical risk. The quantum constructs (decoherence, entanglement, superposition, tunnelling) are analogies, not literal physics. Over-interpretation of the analogy — treating project states as actual quantum states — would be intellectually dishonest and practically misleading. The QPM framework uses quantum mechanics as a structural metaphor that provides a richer

ontology for describing project behavior; it does not claim that projects obey the Schrödinger equation.

4.7 Recommended Deployment Strategy

The SFI and QFI are not competitors; they are complementary instruments operating at different levels of governance sophistication. The recommended deployment strategy mirrors the tiered telemetry architecture described in Section 3.2.1:

Tier 1 — Rapid Assessment (SFI). For initial project screening, portfolio-level triage, or environments with limited data infrastructure, use the Taleb–Douady heuristic. Compute $E[v]$ at three dispersion levels, take the finite-difference second derivative, and classify the project as fragile, robust, or antifragile. This requires no Monte Carlo, no dimensional decomposition, and no fatigue data. It is fast, cheap, and universally applicable. It will correctly identify which projects deserve deeper analysis.

Tier 2 — Dimensional QFI (without fatigue). For projects identified as fragile by Tier 1, deploy the six-dimension QFI decomposition (Section 2.3.2). Compute QFI_i for each governance dimension, estimate correlation coefficients ρ_{ij} , and produce a composite QFI_{norm} with governance-band assignment. This adds dimensional visibility and floor-rule protection but does not yet require workforce fatigue data.

Tier 3 — Full QFI with Fatigue Integration. For projects in the Amber or Orange governance bands, activate the FRI bridge (Section 3). Deploy workforce fatigue monitoring (Tier 1–3 telemetry per Section 3.2.1), compute $\Gamma_{effective}$ with fatigue augmentation, track roughness W , and run the unified governance dashboard (Section 3.8). This is the full QPM fragility instrument — computationally intensive, data-hungry, but diagnostically complete.

This tiered approach allows organizations to begin with the SFI's simplicity and progressively adopt the QFI's depth as their governance maturity, data infrastructure, and project risk profiles warrant. The SFI is not superseded — it becomes the entry ramp to a more powerful framework.

4.8 Implications for Future Research

The comparison between SFI and QFI reveals several directions for future investigation:

Empirical calibration. The QFI's superiority over the SFI is currently a structural argument — the QFI captures more phenomena. Empirical validation on completed projects would strengthen the case by demonstrating that the additional phenomena actually improve governance outcomes (earlier detection, fewer threshold crossings, lower value loss). Longitudinal studies comparing SFI-only and QFI-augmented governance on matched project pairs would be particularly valuable.

Computational efficiency. The QFI's computational overhead is a real barrier to adoption. Research into reduced-order models, surrogate approximations, and real-time QFI estimation from streaming telemetry would lower the barrier and enable more organizations to access the full framework.

Antifragility operationalization. The SFI's antifragility concept (positive curvature) is theoretically elegant but has not been operationalized in the QFI governance bands, which currently range from 0 (robust) to 1 (critical). Developing governance protocols for projects that actively benefit from perturbation — and identifying the design principles that produce convex project payoffs — is a natural extension of the framework.

Cross-industry calibration. The QFI's governance thresholds (0.30, 0.60, 0.85) are calibrated from portfolio experience in large-scale construction^{12 13}. Extending the calibration to other industries — aerospace, pharmaceutical development, software engineering, infrastructure operations — would test the universality claim inherited from the SFI and produce industry-specific threshold sets.

Sources of Normalized Perturbation Magnitudes (For QFI Calibration Across Governance Dimensions)

Governance Dimension	Raw Perturbation Sources (ΔX)	Maximum Credible Variation (ΔX_{max})	Normalized Perturbation ($\epsilon = \Delta X / \Delta X_{max}$)
Schedule	<ul style="list-style-type: none"> Schedule variance distributions (P6, TILOS, LPS) Critical path slippage Interface delay logs Rework-induced resets Crew availability fluctuations 	<ul style="list-style-type: none"> Historical worst-case slippage Maximum observed interface delay Largest recorded rework reset 	$\Delta R / \Delta R_{max}$
Cost	<ul style="list-style-type: none"> Cost growth curves Change order magnitude/frequency Commodity price shocks Contractor productivity deviations Claims/dispute escalations 	<ul style="list-style-type: none"> Maximum historical cost overrun Largest unit-rate deviation Peak commodity volatility impact 	$\Delta C / \Delta C_{max}$
Fatigue / Human Performance	<ul style="list-style-type: none"> Fatigue index telemetry Productivity decay curves Absenteeism spikes Shift-pattern 	<ul style="list-style-type: none"> Maximum fatigue index deviation Worst-case 	$\Delta F / \Delta F_{max}$

¹² The governance thresholds used in this paper (0.30, 0.60, 0.85) are not derived from external industry standards. They were empirically calibrated from multi-year portfolio data across large-scale construction and capital programs, using observed frequencies of state-transition events (Green→Amber→Red), coherence-loss patterns, and normalized perturbation magnitudes derived from historical project performance datasets. These thresholds therefore reflect an organization's own empirical experience rather than a published cross-industry benchmark.

¹³ Normalized perturbation magnitudes used in QFI calibration are derived from historical project performance datasets, including schedule variance distributions, cost growth patterns, fatigue linked productivity decay, supply chain delay logs, regulatory cycle time variance, and interface integration error rates. These empirical datasets provide the ΔX and ΔX_{max} values used to compute normalized perturbations across governance dimensions.

Governance Dimension	Raw Perturbation Sources (ΔX)	Maximum Credible Variation (ΔX_{max})	Normalized Perturbation ($\epsilon = \Delta X / \Delta X_{max}$)
	instability • Safety-critical micro-break compliance	productivity decay • Highest recorded absenteeism surge	
Supply Chain	• Lead-time variance • Single-source exposure • Port/customs delays • Vendor reliability scores • Material availability telemetry	• Maximum historical lead-time delay • Highest single-source concentration • Worst vendor failure event	$\Delta SC / \Delta SC_{max}$
Regulatory	• Permitting cycle-time variance • Inspection failure rates • Regulatory change frequency • Environmental/political disruptions	• Longest recorded permitting delay • Highest inspection failure spike • Largest regulatory cycle shock	$\Delta Reg / \Delta Reg_{max}$
Stakeholder / Interface	• Alignment drift logs • Design maturity variance • Coordination clash logs • Integration test failures	• Maximum stakeholder drift • Worst-case design maturity gap • Highest integration error density	$\Delta I / \Delta I_{max}$

Integration with Safety-III¹⁴. The QFI framework is situated within Safety-II (Section 3.10). Emerging Safety-III concepts — which extend Safety-II with explicit attention to systemic resilience, adaptive capacity, and transformative learning — may provide additional constructs that further enrich the QFI ontology.

Safety-III extends the Safety-II paradigm by embedding adaptive capacity, coherence preservation, and structural learning directly into governance. Within the Quantum Project Management (QPM) framework, Safety-III provides the conceptual bridge between **human-system resilience** and **quantum-analog fragility metrics**. It reframes safety not as the absence of failure, but as the

¹⁴ Safety-III refers to an emerging governance perspective that extends Safety-II by focusing not only on how work succeeds under variable conditions, but on how systems adapt, reconfigure, and evolve in response to persistent complexity, structural coupling, and nonlinear stress. Whereas Safety-I emphasizes preventing things from going wrong and Safety-II emphasizes enabling things to go right, Safety-III emphasizes adaptive capacity at scale—the system’s ability to absorb perturbations, reorganize without loss of function, and improve its performance envelope through feedback, learning, and structural redesign. In this framing, fragility, resilience, and antifragility are treated as dynamic properties of socio-technical systems, and governance shifts from event-based controls to continuous monitoring of coherence, coupling, and emergent behavior.

system’s ability to **reorganize under stress** while maintaining operational coherence and improving its performance envelope.

Safety-III ↔ QPM Integration Matrix

Safety-III Element	QPM Construct	Integration Outcome
Adaptive Capacity	Decoherence (Γ), Fidelity $FQ(t)$	Real-time coherence monitoring and early warning of fragility.
Structural Reconfiguration	Entanglement (ρ_{ij})	Quantifies coupling strength; guides decoupling and redesign.
Evolution Under Stress	Antifragility ($QFI < 0$)	Enables learning from perturbation; identifies beneficial variability.
Emergent Behavior	Propagation Velocity	Detects nonlinear amplification; informs systemic risk control.
Coherence Preservation	Superposition Stability, Proximity Metric	Measures distance to governance boundary; supports proactive intervention.
Continuous Monitoring	Decoherence Proxy ($\eta \cdot \Gamma \cdot (1 - FQ)^{\alpha - 1}$)	Embeds telemetry into governance dashboards for adaptive oversight.

Integrating Safety-III into QPM transforms governance from **reactive control** to **adaptive orchestration**. Decision authority becomes distributed, escalation paths become dynamic, and fragility metrics become the backbone of safety intelligence. The organization evolves from preventing failure to **anticipating and shaping system behavior** under complexity.

5. Conclusion

5.1 Executive Summary

This paper has introduced and formalized the **Quantum Fragility Index (QFI)** as a composite fragility metric purpose-built for the governance of large complex projects (LCPs) within a Quantum Project Management (QPM) framework. The QFI is grounded in four quantum-analogic constructs:

- superposition of assumption states, which preserves the probabilistic character of planning assumptions until empirical resolution;
- decoherence as a proxy for confidence decay across assumption populations over time;
- entanglement as the mechanism by which cross-assumption coupling amplifies systemic exposure; and
- proximity-to-threshold measurement, which quantifies how close assumption clusters sit relative to governance-band boundaries.

By operationalizing Taleb–Douady curvature concepts—specifically, second-derivative sensitivity mapping of the fragility response surface—the QFI enables organizations to distinguish brittle project configurations exhibiting convex loss profiles from resilient configurations exhibiting concave loss profiles.

Critically, QFI is designed to integrate with, rather than replace, existing governance instruments such as Earned Value Management (EVM) and organizational risk registers, augmenting them with a layer of nonlinear, anticipatory intelligence. The governance banding framework — a five-tier spectrum spanning Antifragile (Green), Robust (Blue), Moderately Fragile (Amber), Highly Fragile (Orange), and Critically Fragile (Red), with boundaries at 0, 0.30, 0.60, and 0.85 (Section 2.3.2; Appendix C, Table C.1) — translates QFI outputs into actionable decision triggers that equip leadership with clear escalation protocols and intervention thresholds. Taken together, these contributions represent a substantive advance in the discipline of project governance—moving organizations from reactive risk management toward **anticipatory fragility governance**.

5.2 QPM Perspective and Governance Implications

The Quantum Fragility Index must be understood within the broader evolution of Quantum Project Management as a discipline. Traditional project management metrics—schedule variance, cost performance index, earned schedule—operate within what may be termed a "Newtonian" domain: they assume linearity, decomposability, and stable cause-effect relationships. These assumptions hold for simple and moderately complex projects, where scope boundaries are well-defined and assumption populations are small and largely independent.

However, in large complex projects characterized by hundreds or thousands of interrelated assumptions, emergent behaviors, and nonstationary environments, Newtonian metrics provide an incomplete and often misleading picture of project health. QPM demands metrics capable of

capturing nonlinear dynamics—assumption migration, confidence decay, coupling cascades, and threshold proximity effects. The Quantum Fragility Index answers this demand directly, providing a single composite metric that synthesizes these dynamics into a governance-ready output.

The decoherence proxy stands as the central analytical innovation within the QFI framework. By measuring assumption confidence decay over time—calibrated against empirical migration data drawn from project-specific assumption catalogs—the decoherence proxy transforms an abstract quantum-analogic concept into a quantifiable, trackable metric. This proxy is further enriched through its linkage to Taleb–Douady curvature mapping. Where conventional sensitivity analysis examines first-derivative responses (i.e., how much a fragility output changes given an input perturbation), QFI's second-derivative sensitivity captures the curvature of the fragility response surface itself. This distinction is operationally critical: a project configuration exhibiting positive curvature (convex loss) faces accelerating fragility under stress, while a configuration exhibiting negative curvature (concave loss) demonstrates diminishing marginal fragility—a signature of resilience. The proximity metric completes the analytical triad by quantifying the distance between any assumption cluster's current QFI value and the nearest governance-band boundary, enabling preemptive intervention before a band transition occurs rather than reaction after the fact.

The governance implications of QFI adoption are substantial and far-reaching. At the most fundamental level, QFI transforms leadership's relationship with project risk from one of periodic, retrospective review to one of continuous, anticipatory monitoring. The governance banding framework — Green, Blue, Amber, Orange, and Red — provides unambiguous escalation protocols calibrated to five tiers of fragility severity: Green-band (Antifragile) clusters are reviewed quarterly to harvest optionality; Blue-band (Robust) clusters follow standard monthly governance cycles; Amber-band clusters trigger enhanced bi-weekly surveillance, root-cause investigation, and contingency activation; Orange-band clusters require executive escalation with mandated remediation plans; Red-band clusters demand immediate remediation or controlled termination with board notification within 48 hours. Beyond banding, entanglement mapping reveals hidden systemic exposure pathways that are invisible to traditional risk registers, exposing the channels through which a single assumption failure can cascade across project subsystems. Equally important, superposition recognition disciplines governance actors to resist premature assumption collapse—the organizational tendency to treat uncertain assumptions as resolved in order to simplify planning—which creates false certainty and suppresses legitimate fragility signals. Integration of QFI into quarterly governance reviews and project control board agendas is strongly recommended to institutionalize these benefits.

It is important to acknowledge the limitations inherent in the current QFI framework. Effective deployment requires disciplined assumption cataloging across the project ecosystem, rigorous calibration of decoherence rates against project-specific empirical data, and organizational willingness to adopt and internalize quantum-analogic language at both the technical and leadership levels. These represent genuine implementation challenges that must be addressed through training, process design, and sustained executive sponsorship. However, they are challenges of adoption and operationalization—not theoretical deficiencies in the metric itself.

5.3 Key Insights

QPM Construct	QFI Application	Practical Governance Meaning
Superposition	Assumption states modeled as probabilistic distributions until resolved	Leadership must resist premature assumption collapse; maintain optionality in planning
Decoherence	Confidence decay proxy tracks assumption degradation over time	Early-warning metric for governance dashboards; triggers Amber/Red band escalation
Entanglement	Cross-assumption coupling coefficients quantify systemic exposure	Reveals hidden interdependencies invisible to traditional risk registers
Taleb–Douady Curvature	Second-derivative sensitivity mapping of fragility response surface	Distinguishes brittle (convex loss) from resilient (concave loss) project configurations
Proximity	Distance-to-threshold measurement against governance band boundaries	Enables preemptive intervention before critical band transitions occur; predictive crossing-time capability ($F-26: t_{\text{crossing}} = -\ln(F_{\text{threshold}}) / \Gamma$).

5.4 Key Recommendations

Priority	Recommendation	Target Audience
1	Integrate QFI into quarterly project governance reviews as a standing agenda item	Project Control Board / Governance Committee
2	Establish disciplined assumption cataloging protocols with decoherence calibration baselines	Project Management Office (PMO)
3	Deploy entanglement mapping for all Tier 1 assumptions to surface systemic exposure pathways	Risk Management Team

Priority	Recommendation	Target Audience
4	Implement proximity-based early-warning triggers linked to governance-band escalation protocols	Operations / Project Controls
5	Conduct Taleb–Douady curvature analysis on assumption clusters exceeding Amber threshold	Technical Advisory / Analytics
6	Develop organizational training modules on QPM constructs and QFI interpretation for leadership	Learning and Development / PMO

5.5 Next Steps

The path forward for the Quantum Fragility Index centers on four immediate priorities:

- field validation of QFI against active LCP portfolios to establish empirical performance baselines;
- calibration of decoherence rates using observed assumption-migration data drawn from live project environments;
- development of automated QFI dashboards capable of delivering real-time governance monitoring with proximity alerts and band-transition notifications; and
- extension of the QPM metric framework to incorporate companion constructs such as the Assumption Diffusion Index (ADI) for tracking assumption propagation dynamics across organizational boundaries and cross-portfolio fragility aggregation methods for enterprise-level governance visibility.

These efforts should be sequenced to deliver early validation results within one governance cycle, building organizational confidence and stakeholder buy-in for broader adoption.

Quantum Project Management holds the potential to fundamentally reshape how organizations govern complexity, uncertainty, and fragility in large-scale project ecosystems—and the Quantum Fragility Index, as formalized in this document, represents the foundational instrument for realizing that transformation.

Section 6 — Summary of Original Ideas and Concepts Underpinning the QFI

The following table summarizes the 17 foundational ideas that underpin the Quantum Fragility Index (QFI). Each concept is presented with its core insight and its relevance to the QFI framework.

Concept	Core Idea	Relevance to QFI / QPM
1. Fragility vs. Risk	Fragility measures susceptibility to state transitions, not expected loss.	Establishes the need for QFI as a coherence-sensitivity metric.
2. Formal Definition of Quantum Fragility	Minimum normalized perturbation needed to cross a governance boundary.	Provides the foundational definition QFI operationalizes.
3. Mapping to Quantum Constructs	Fragility aligns with decoherence, entanglement, superposition, propagation.	Grounds QFI in the QPM ontology.
4. QFI as a Single Scalar	Integrates curvature, decoherence, dimensional fragility, and coupling.	Produces a board-readable governance metric.
5. Curvature as Core Mathematics	Fragility = second-order sensitivity to dispersion.	Enables detection of nonlinear vulnerability.
6. Decoherence-Mapped Proxy	Real-time fragility estimation via coherence decay.	Supports telemetry-driven monitoring.
7. Six-Dimension Fragility Model	Schedule, Cost, Scope, Stakeholder, Supply Chain, Regulatory.	Makes fragility diagnosable and actionable.
8. Correlation Correction	Entanglement amplifies cross-dimensional fragility.	Captures nonlinear coupled failure pathways.
9. Governance Thresholds	Boundaries at 0, 0.30, 0.60, 0.85 reflect system-behavior shifts.	Creates a dynamic, behavior-based governance structure.
10. Dimensional Floor Rule	Any dimension ≥ 0.85 triggers escalation.	Prevents averaging from masking catastrophic fragility.
11. Integration with Safety-II	QFI quantifies proximity to failure and coherence loss.	Operationalizes anticipatory resilience.

Concept	Core Idea	Relevance to QFI / QPM
12. Fatigue as an Internality	Fatigue increases decoherence, nonlinearity, and coupling.	Integrates human readiness into fragility modeling.
13. Roughness–Curvature Dual Diagnostic	Roughness = spatial instability; curvature = systemic sensitivity.	Provides a two-lens vulnerability diagnostic.
14. Unified QFI–FRI Dashboard	Combines fragility, fatigue, roughness, and proximity.	Creates an integrated governance interface.
15. Three-Pathway Fatigue Transmission	Fatigue affects fragility via direct, decoherence, and coupling pathways.	Explains how internal conditions reshape system fragility.
16. Quantum Tunneling Analogy	Fatigue “thins” barriers, enabling perturbation tunneling.	Reframes control failure as barrier permeability.
17. QFI as Governance Operating System	Enables continuous monitoring, escalation, and intervention.	Shifts governance from event-driven to state-driven.

References

Barger, L. K., Ayas, N. T., Cade, B. E., Cronin, J. W., Rosner, B., Speizer, F. E., & Czeisler, C. A. (2006). Impact of extended-duration shifts on medical errors, adverse events, and attentional failures. *PLoS Medicine*, 3(12), e487. <https://doi.org/10.1371/journal.pmed.0030487>

Browning, T. R. (2019). Planning, tracking, and reducing a complex project's value at risk. *Project Management Journal*, 50(1), 71–85. <https://doi.org/10.1177/8756972818810967>

Dawson, D., & Reid, K. (1997). Fatigue, alcohol and performance impairment. *Nature*, 388(6639), 235. <https://doi.org/10.1038/40775>

de Guingand, M. (2025). *Fatigue in infrastructure construction: A longitudinal study of HS2 construction workers*. Third Pillar of Health / HS2 Occupational Health and Innovation.

Dekker, S. (2011). *Drift into failure: From hunting broken components to understanding complex systems*. Ashgate. <https://doi.org/10.1201/9781315257396>

Dembe, A. E., Erickson, J. B., Delbos, R. G., & Banks, S. M. (2005). The impact of overtime and long work hours on occupational injuries and illnesses: New evidence from the United States. *Occupational and Environmental Medicine*, 62(9), 588–597. <https://doi.org/10.1136/oem.2004.016667>

Folkard, S., & Lombardi, D. A. (2006). Modeling the impact of the components of long work hours on injuries and "accidents." *American Journal of Industrial Medicine*, 49(11), 953–963. <https://doi.org/10.1002/ajim.20307>

Folkard, S., & Tucker, P. (2003). Shift work, safety and productivity. *Occupational Medicine*, 53(2), 95–101. <https://doi.org/10.1093/occmed/kqg047>

Hallowell, M. R., Hinze, J. W., Baud, K. C., & Wehle, A. (2013). Proactive construction safety control: Measuring, monitoring, and responding to safety leading indicators. *Journal of Construction Engineering and Management*, 139(10), 04013010. [https://doi.org/10.1061/\(ASCE\)CO.1943-7862.0000730](https://doi.org/10.1061/(ASCE)CO.1943-7862.0000730)

Hallowell, M., Quashne, M., Salas, R., Jones, M., MacLean, B., & Quinn, E. (2021). The Statistical Invalidity of TRIR as a Measure of SAFETY PERFORMANCE. *Professional Safety*, 66(4), 28-34.

Han, J., & Fu, J. (2021). A Combined Approach of System and Multiple Factors Theories for Accident Causation Analysis of an Aviation Accident. IISE Annual Conference. Proceedings, (), 271-275.

Hollnagel, E. (2014). *Safety-I and Safety-II: The past and future of safety management*. Ashgate. <https://doi.org/10.1201/9781315607511>

Hollnagel, E. (2018). *Safety-II in Practice: Developing the Resilience Potentials*.

Hollnagel, E., Pariès, J., Woods, D. D., & Wreathall, J. (Eds.). (2011). *Resilience engineering in practice: A guidebook*. Ashgate.

Ingre, M., Åkerstedt, T., Peters, B., Anund, A., Kecklund, G., & Pickles, A. (2006). Subjective sleepiness and accident risk avoiding the ecological fallacy. *Journal of Sleep Research*, 15(2), 142–148. <https://doi.org/10.1111/j.1365-2869.2006.00517.x>

ISO. (2018). ISO 31000:2018 *Risk management — Guidelines*.

Kardar, M., Parisi, G., & Zhang, Y.-C. (1986). Dynamic scaling of growing interfaces. *Physical Review Letters*, 56(9), 889–892.

Leveson, N. G. (2011). *Engineering a safer world: Systems thinking applied to safety*. MIT Press. <https://doi.org/10.7551/mitpress/8179.001.0001>

Malicse, A. A Universal Balance Framework for the Unification of General Relativity and Quantum Mechanics. <https://core.ac.uk/download/665166899.docx>

Maraffi, R. (2025). From Errors to Excellence: A New Paradigm for Safety in Construction. <https://core.ac.uk/download/659415273.pdf>

Moerman, J. J., Braaksma, A. J. J., & Van Dongen, L. A. M. (2017). Resilient performance in maintenance operations: Managing unexpected failures. <https://core.ac.uk/download/572209896.pdf>

Prieto, R. (2024). Quantum Project Management, *PM World Journal*, Vol. XII, Issue I, January 2024. <https://pmworldlibrary.net/wp-content/uploads/2024/01/pmwj137-Jan2024-Prieto-Quantum-Project-Management-.pdf>

Prieto, R. (2024). Measurement of Complexity in Large Complex Projects, *PM World Journal*, Vol. XII, Issue IV, April. <https://pmworldlibrary.net/wp-content/uploads/2024/04/pmwj140-Apr2024-Prieto-Measurement-of-Complexity-in-Large-Complex-Projects.pdf>

Prieto, R. (2025). Artificial Intelligence, Complexity, and Quantum Project Management: A Transformative Approach, *PM World Journal*, Vol. XIV, Issue VII, July.

<https://pmworldlibrary.net/wp-content/uploads/2025/07/pmwj154-Jul2025-Prieto-AI-Complexity-Quantum-Project-Management.pdf>

Prieto, R. (2024). Quantum Project Management and the Concept of Spacetime, *PM World Journal*, Vol. XII, Issue V, May.

<https://pmworldlibrary.net/wp-content/uploads/2024/05/pmwj141-May2024-Prieto-Quantum-Project-Management-and-Concept-of-Space-time.pdf>

Prieto, R. (2025). Managing Uncertainty in Large Complex Projects, *PM World Journal*, Vol. XIV, Issue XI, November.

<https://pmworldlibrary.net/wp-content/uploads/2025/11/pmwj158-Nov2025-Prieto-Managing-Uncertainty-in-Large-Complex-Projects-3.pdf>

Prieto, R. (2025). Metrics for Assumption Management in Large Complex Projects, *PM World Journal*, Vol. XIV, Issue XII, December.

<https://pmworldlibrary.net/wp-content/uploads/2025/12/pmwj159-Dec2025-Prieto-metrics-for-assumption-management-large-complex-projects.pdf>

Prieto, R. (2026). Operationalizing Quantum Project Management: Defining Improved Metrics for Management of Large Complex Projects, *PM World Journal*, Vol. XV, Issue I, January.

<https://pmworldlibrary.net/wp-content/uploads/2026/01/pmwj160-Jan2026-Prieto-Operationalizing-Quantum-Project-Management.pdf>

Prieto, R. (2026). Operationalizing Quantum Project Management: Anticipating and Managing Fragility in Large Complex Project Ecosystems, *PM World Journal*, Vol. XV, Issue II, February.

<https://pmworldlibrary.net/wp-content/uploads/2026/02/pmwj161-Feb2026-Prieto-Managing-Fragility-in-Large-Complex-Project-Ecosystems.docx>

Prieto, R. (2026). Fatigue Risk Index: Measuring an "internality" as a performance precursor under Quantum Project Management (QPM) theory. *PM World Journal*, Vol. XV, Issue V, May.

<https://pmworldlibrary.net/wp-content/uploads/2026/06/pmwj165-Jun2026-Prieto-Fatigue-Risk-Factor.pdf>

Project Management Institute. *A Guide to the Project Management Body of Knowledge (PMBOK Guide)*.

Rasmussen, J. (1997). Risk management in a dynamic society: A modelling problem. Safety, compliance, and then production maximizes bottom line <https://www.powermag.com/safety-compliance-and-then-production-maximizes-bottom-line/>[https://doi.org/10.1016/S0925-7535\(97\)00052-0](https://doi.org/10.1016/S0925-7535(97)00052-0)

Soloveva, I., Mostovshchikova, I., Baeva, D., & Alola, A. (2022). A Schematic Model for Implementing Industries High Risk Projects through the Agent of Human Capital. *Sustainability*, 14(18), 11778.

Superposition <https://qcve.org/c/glossary/s/superposition>

Taleb, N. N. (2012). *Antifragile: Things that gain from disorder*. Random House.

Taleb, N. N., & Douady, R. (2013). Mathematical definition, mapping, and detection of (anti)fragility. *Quantitative Finance*, 13(11), 1677–1689.

<https://doi.org/10.1080/14697688.2013.800219>

Van Dongen, H. P. A., Maislin, G., Mullington, J. M., & Dinges, D. F. (2003). The cumulative cost of additional wakefulness: Dose-response effects on neurobehavioral functions and sleep physiology from chronic sleep restriction and total sleep deprivation. *Sleep*, 26(2), 117–126.

<https://doi.org/10.1093/sleep/26.2.117>

Van Peer, J. M., Gladwin, T. E., & Nieuwenhuys, A. (2019). Effects of Threat and Sleep Deprivation on Action Tendencies and Response Inhibition. <https://doi.org/10.1037/emo0000533>

Wang, D. (2013). Exploring factors associated with workers' safety rule violations in the Chinese construction industry: A multimethod investigation. <https://core.ac.uk/download/42506476.pdf>

Zurek, W. H. (2003). Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics*.

Acknowledgments

The author acknowledges the use of language generation tools including ChatGPT, Gemini, Copilot and Grammarly to assist in restating and refining portions of the text and graphics generation to illustrate stated concepts. All content, research, design, and analysis remains the authors' original work.

Appendices

This paper includes extensive appendices, enabling the reader to dive deeper or gain more clarity in areas of interest.

These appendices include:

Appendix A — Complete Formula Inventory and Notation Glossary

Part A — Formula Inventory

Part B — Notation Glossary

Appendix B - QFI Quick-Reference Card

Appendix C — QFI Governance Operations Playbook

Appendix D — Roughness and Curvature: A Structural Comparison of Two Second-Order Fragility Diagnostics

Appendix E - Original Ideas and Concepts

Appendix A

Complete Formula Inventory and Notation Glossary

This appendix provides a consolidated reference for every mathematical formula and notation symbol used across the Quantum Fragility Index (QFI) paper. Part A catalogs all formulas by source section with plain-English interpretations and cross-references. Part B provides a complete notation glossary organized by symbol category — Greek letters, Roman variables, subscripts, operators, governance constants, and the six-dimension index.

Part A — Formula Inventory

A.1 Section 2 — Formal Quantum Fragility Metric

F-01 Taleb–Douady Core QFI Definition

$$QFI(S) = -\partial^2 E[V(S; x)] / \partial \sigma^2 \Big|_{\sigma=\sigma_0}$$

The Quantum Fragility Index is the negative second derivative of expected project value with respect to perturbation dispersion, evaluated at the current operating point σ_0 .

Source: Section 2.2.3

F-02 QFI Normalization

$$QFI_{\text{norm}} = (QFI - QFI_{\text{min}}) / (QFI_{\text{max}} - QFI_{\text{min}})$$

Scales raw curvature to a dimensionless [0, 1] governance scalar for cross-project comparison and governance-band assignment.

Source: Section 2.2.3

F-03 Expected Payoff Integral

$$E[V] = \int v(x) \cdot f(x; \sigma) dx$$

Expected project value is the integral of the payoff function $v(x)$ weighted by the perturbation probability density $f(x; \sigma)$.

Source: Section 2.3.1

F-04 Canonical Fragility Operator

$$QFI = -\partial^2/\partial\sigma^2 \int v(x) \cdot f(x; \sigma) dx$$

The QFI operator applied directly to the expected-payoff integral — the single-expression computational form combining F-01 and F-03.

Source: Section 2.3.1

F-05 Quantum Fidelity

$$F_Q(t) = \text{Tr}[\sqrt{\sqrt{\rho_0} \cdot \rho(t) \cdot \sqrt{\rho_0}}]$$

Measures similarity between the project's current operational state $\rho(t)$ and its planned baseline state ρ_0 ; ranges from 1 (perfect coherence) to 0 (complete decoherence).

Source: Section 2.2.2

F-06 Fidelity Decay (Exponential Model)

$$F_Q(t) = e^{-\Gamma t}$$

Under Markovian decoherence, fidelity decays exponentially at rate Γ , the decoherence rate constant.

Source: Section 2.2.2

F-07 Value Loss Approximation

$$\Delta V \approx \eta \cdot (1 - F_Q(t))^\alpha$$

Project value loss scales with the decoherence gap $(1 - F_Q)$ raised to nonlinearity exponent α , with sensitivity coefficient η .

Source: Section 2.2.2

F-08 Decoherence Proxy for QFI

$$QFI_{\text{deco}} = \eta \cdot \Gamma \cdot (1 - F_Q(t))^{(\alpha-1)}$$

An operationally computable proxy for QFI derived from the decoherence pathway — the derivative of value loss with respect to decoherence progression.

Source: Section 2.2.2

F-09 Payoff (Value Transfer) Function

$$v(x) = V_0 \cdot \exp(-\beta_1 \cdot \max(0, x - x_{tol})^{\beta_2})$$

Maps perturbation magnitude x to residual project value; tolerant below x_{tol} , then decays exponentially with steepness β_1 and curvature β_2 .

Source: Section 2.3.1

F-10 Composite QFI (Weighted Sum)

$$QFI_{composite} = \sum_i w_i \cdot QFI_i + C$$

Portfolio-level fragility is the weighted sum of dimensional QFI values plus a correlation correction term C capturing cross-dimensional coupling.

Source: Section 2.3.2

F-11 Correlation Correction Term

$$C = \sum_{i \neq j} \rho_{ij} \cdot \sqrt{(w_i \cdot w_j)} \cdot \varphi(QFI_i, QFI_j)$$

Additional fragility arising from correlated perturbations across governance dimensions; ρ_{ij} is the inter-dimensional correlation coefficient and φ is the coupling function.

Source: Section 2.3.2

F-12 Taleb–Douady QFI (Operational Form)

$$QFI_{TD} = -\partial^2 E[v(x)] / \partial \sigma^2 \Big|_{\sigma=\sigma_0}$$

Operational form of the core QFI definition, emphasizing the Taleb–Douady transfer-function lineage.

Source: Section 2.3.1

F-13 Finite-Difference Approximation

$$\partial^2 E[v] / \partial \sigma^2 \approx (E[v]_+ - 2E[v]_0 + E[v]_-) / (\Delta \sigma)^2$$

Three-point central-difference approximation of the second derivative — enables Monte Carlo computation of QFI without closed-form solutions.

Source: Section 2.3.1

F-14 Proximity Metric

$$\text{Proximity} = 1 - |\varepsilon^*| / |\varepsilon_{\max}|$$

Measures how close the current state is to the nearest governance-band boundary; 1 = at the boundary, 0 = maximum distance from any boundary.

Source: Section 2.3.2

A.2 Section 3 — Relationship Between QFI and Fatigue Risk

F-15 Effective Decoherence Rate (Fatigue-Augmented)

$$\Gamma_{\text{effective}} = \Gamma_{\text{baseline}} + \Gamma_{\text{fatigue}}(\text{FRI})$$

Total decoherence rate is the sum of baseline project decoherence plus the fatigue-induced contribution, which is a function of the Fatigue Risk Index.

Source: Section 3.2.1

F-16 Fatigue Decoherence Function (Power Law)

$$\Gamma_{\text{fatigue}}(\text{FRI}) = \gamma_0 \cdot (\text{FRI} / 100)^{\beta_f}$$

Fatigue-induced decoherence follows a power-law relationship with FRI; γ_0 scales the maximum contribution and $\beta_f > 1$ captures nonlinear acceleration near critical thresholds.

Source: Section 3.2.1

F-17 Fatigue-Specific Fragility (Proximity Form)

$$\text{QFI}_F = 1 - |\Delta F^*| / |\Delta F_{\max}|$$

Measures proximity of the current fatigue state to the nearest FRI governance-band boundary; ΔF^ is the distance to that boundary and ΔF_{max} is the maximum credible fatigue change under the scenario.*

Source: Section 3.3.2

F-18 Fatigue-Specific Fragility (Curvature Form)

$$QFI_{F,TD} = -\partial^2 E[V_{\text{safety}}(F)] / \partial \sigma_F^2 \big|_{\sigma_F = \sigma_{F,0}}$$

Taleb–Douady curvature-based formulation of fatigue fragility — the second derivative of expected safety-adjusted value with respect to fatigue perturbation dispersion.

Source: Section 3.3.2

F-19 Dimension-Specific Decoherence Proxy with Fatigue

$$QFI_{\text{deco},i}(\text{FRI}) = \eta_i \cdot \Gamma_{\text{effective}}(\text{FRI}) \cdot (1 - F_{Q,i}(t))^{\alpha_i - 1}$$

Extends the decoherence proxy (F-08) to individual governance dimensions using the fatigue-augmented decoherence rate; each dimension carries its own sensitivity η_i , nonlinearity α_i , and fidelity $F_{Q,i}$.

Source: Section 3.3.3

A.3 Appendix D — Roughness and Curvature Compared

F-20 KPZ Surface Evolution Equation

$$\partial h / \partial t = \nu \nabla^2 h + \lambda (\nabla h)^2 + \eta(x, t)$$

Kardar–Parisi–Zhang equation governing fatigue interface evolution: $\nu \nabla^2 h$ is management's diffusion (smoothing) capacity, $\lambda (\nabla h)^2$ is quadratic contagion from fatigue gradients, and $\eta(x, t)$ is stochastic noise.

Source: Appendix D.2.1

F-21 Roughness (Root-Mean-Square Deviation)

$$W(L, t) = \langle (h - \langle h \rangle)^2 \rangle^{1/2}$$

Root-mean-square deviation of fatigue scores from the workforce mean — measures the spatial unevenness (variance) of the fatigue landscape across crews and work packages.

Source: Appendix D.2.1

F-22 Critical Fatigue Gradient (Tipping Point)

$$\nabla h_{\text{critical}} = \sqrt{(v / \lambda)}$$

The fatigue gradient threshold where quadratic contagion overwhelms management's diffusion capacity — the tipping point between recoverable variability and self-reinforcing instability.

Source: Appendix D.3

A.4 Companion Paper — Fatigue Risk Index

F-23 Fatigue Risk Index (FRI)

$$\text{FRI} = 100 \times (0.30E + 0.25C + 0.20S + 0.15B + 0.10T)$$

Weighted composite of five fatigue driver subscales: Environmental/physical load (E), Cumulative workload (C), Schedule/shift pattern (S), Behavioral/individual recovery (B), and Task demand/cognitive complexity (T). Produces a 0–100 score.

Source: Prieto (2026), Fatigue Risk Index

F-24 Fatigue Decay (Half-Life Model)

$$F(t) = F_0 \cdot e^{-kt}$$

Fatigue decays exponentially from initial level F_0 at rate k during rest/recovery periods; the half-life $t_{1/2} = \ln(2)/k$ determines how long a crew needs to halve its accumulated fatigue debt.

Source: Prieto (2026), Fatigue Risk Index

F-25 Nonlinear Risk Spike Function

$$R(F) = R_{\text{baseline}} \cdot e^{\delta \cdot \max(0, F - F_{\text{threshold}})}$$

Accident/error risk R spikes exponentially once fatigue F exceeds the threshold $F_{\text{threshold}}$, with amplification rate δ ; below threshold, risk remains at baseline R_{baseline} .

Source: Prieto (2026), Fatigue Risk Index

F-26 Governance Crossing Time

$$t_{\text{crossing}} = -\ln(F_{\text{threshold}}) / \Gamma$$

Estimated time until the system crosses the next governance-band boundary at the current decoherence rate Γ ; converts QFI from a static snapshot into a predictive timeline, enabling anticipatory intervention before band transitions occur.

Source: Section 3.12

Part B — Notation Glossary

B.1 Greek Letters

Symbol	Name (Greek)	Definition / Usage
Γ	Gamma (capital)	Decoherence rate — the rate at which the project's operational state diverges from plan; units: time^{-1} ; also denominator of governance crossing time (F-26)
Γ_{baseline}	Gamma baseline	Decoherence rate from non-fatigue sources (schedule drift, scope creep, stakeholder misalignment)
Γ_{fatigue}	Gamma fatigue	Additional decoherence contribution from workforce fatigue; computed via F-16
$\Gamma_{\text{effective}}$	Gamma effective	Total decoherence rate combining baseline and fatigue contributions; $\Gamma_{\text{baseline}} + \Gamma_{\text{fatigue}}$ (F-15)
γ_0	Gamma (lowercase)	Scaling constant for the fatigue decoherence function; calibrated from project telemetry
σ	Sigma (lowercase)	Perturbation dispersion — the scale parameter of the perturbation distribution; the independent variable with respect to which QFI curvature is measured
σ_0	Sigma-naught	Current operating-point value of perturbation dispersion at which derivatives are evaluated
σ_F	Sigma-F	Perturbation dispersion in the fatigue dimension specifically
ρ	Rho (lowercase)	(1) Density matrix $\rho(t)$ representing the project's operational quantum state; (2) ρ_{ij} — inter-dimensional correlation coefficient measuring coupling between dimensions i and j
ρ_0	Rho-naught	Planned (coherent) project state at inception — the baseline density matrix
ρ_{ij}	Rho-ij	Correlation coefficient between governance dimensions i and j ; elevated by fatigue through quadratic contagion
η	Eta (lowercase)	(1) Value-loss sensitivity coefficient in the decoherence proxy (F-07, F-08); (2) $\eta(x, t)$ — random noise term in the KPZ equation (F-20). Context determines usage

Symbol	Name (Greek)	Definition / Usage
η_i	Eta-i	Dimension-specific value-loss sensitivity coefficient for governance dimension i
α	Alpha (lowercase)	Nonlinearity exponent controlling the convexity of value loss as a function of decoherence gap; higher α = more nonlinear response
α_i	Alpha-i	Dimension-specific nonlinearity exponent for governance dimension i
β_1	Beta-one	Payoff decay rate controlling steepness of value decline in the payoff function $v(x)$ (F-09)
β_2	Beta-two	Payoff curvature exponent controlling the shape of the decay region in $v(x)$ (F-09)
β_f	Beta-f	Fatigue exponent controlling convexity of the fatigue–decoherence response; $\beta_f > 1$ produces nonlinear acceleration
Δ	Delta (capital)	Finite difference or change; used as $\Delta\sigma$ (perturbation step), ΔV (value change), ΔF^* (distance to FRI boundary)
δ	Delta (lowercase)	Amplification rate in the nonlinear risk spike function $R(F)$ (F-25)
φ	Phi (lowercase)	Coupling function modulating inter-dimensional fragility interaction in the correlation correction term C (F-11)
ε	Epsilon (lowercase)	Perturbation magnitude — a discrete shock applied to the system
ε^*	Epsilon-star	Minimum perturbation magnitude required to trigger a governance-state transition
ν	Nu (lowercase)	Diffusion coefficient — management's smoothing/leveling capacity in the KPZ framework (F-20)
λ	Lambda (lowercase)	Nonlinear growth coefficient — quadratic contagion intensity in the KPZ framework (F-20)

B.2 Roman Variables and Functions

Symbol	Name	Definition / Usage
QFI	Quantum Fragility Index	Core fragility metric — negative curvature of expected value with respect to perturbation dispersion (F-01)
QFI _{norm}	Normalized QFI	QFI scaled to [0, 1] for governance-band assignment (F-02)
QFI _{composite}	Composite QFI	Weighted sum of dimensional QFI values plus correlation correction (F-10)
QFI _{TD}	Taleb–Douady QFI	Operational form emphasizing the Taleb–Douady transfer function (F-12)
QFI _{deco}	Decoherence proxy QFI	Computationally tractable QFI proxy derived from decoherence pathway (F-08)
QFI _F	Fatigue-specific QFI	Fragility index measuring proximity to fatigue governance-band boundaries (F-17)
QFI _{F,TD}	Fatigue QFI (curvature)	Curvature-based formulation of fatigue-specific fragility (F-18)
QFI _i	Dimensional QFI	QFI for governance dimension i (i = 1...6)
FRI	Fatigue Risk Index	Composite index measuring workforce fatigue state; 0–100 scale with four governance bands (F-23)
F _Q (t)	Quantum fidelity	Similarity between current state ρ(t) and baseline state ρ ₀ ; range [0, 1] (F-05)
F _{Q,i} (t)	Dimensional fidelity	Fidelity for governance dimension i specifically
V, V ₀	Project value	V = current value; V ₀ = initial (full) project value at inception
v(x)	Payoff function	Value transfer function mapping perturbation magnitude x to residual project value (F-09)
V _{safety} (F)	Safety-adjusted value	Value function parameterized by fatigue state F; used in curvature-based QFI _F (F-18)

Symbol	Name	Definition / Usage
$E[V], E[v]$	Expected value	Expectation of value or payoff over the perturbation distribution (F-03)
$f(x; \sigma)$	Perturbation PDF	Probability density function of perturbations, parameterized by dispersion σ
C	Correlation correction	Cross-dimensional coupling term in the composite QFI formula (F-11)
$h(x, t)$	Fatigue height function	Local fatigue score at position x and time t in the KPZ framework (F-20)
$W(L, t)$	Roughness	Root-mean-square deviation of fatigue scores from workforce mean (F-21)
$R(F)$	Risk spike function	Nonlinear accident/error risk as a function of fatigue level (F-25)
R_{baseline}	Baseline risk	Background accident/error risk at sub-threshold fatigue levels
$F(t)$	Fatigue decay	Time-dependent fatigue level during recovery (F-24)
F_0	Initial fatigue	Fatigue level at the start of a recovery period
$F_{\text{threshold}}$	Fatigue risk threshold	Fatigue level above which risk spikes nonlinearly (F-25); also governs predicted band-crossing horizon (F-26)
S	System state	Complete project state vector including all governance dimensions
x	Perturbation variable	Random perturbation magnitude drawn from $f(x; \sigma)$
x_{tol}	Tolerance threshold	Perturbation magnitude below which no value loss occurs in $v(x)$
w_i	Dimension weight	Governance weight assigned to dimension i in the composite QFI formula
t	Time	Elapsed time since project inception or last coherence measurement

Symbol	Name	Definition / Usage
t_crossing	Governance crossing time	Estimated time until the system crosses the next governance-band boundary at the current decoherence rate; $t_{\text{crossing}} = -\ln(F_{\text{threshold}}) / \Gamma$ (F-26)
k	Decay constant	Rate of fatigue recovery during rest; related to half-life by $t_{1/2} = \ln(2)/k$
t _{1/2}	Half-life	Time required to halve accumulated fatigue debt during recovery
Tr	Trace operator	Matrix trace used in the quantum fidelity definition (F-05)
E, C, S, B, T	FRI subscales	Environmental/physical load (E), Cumulative workload (C), Schedule/shift pattern (S), Behavioral/individual recovery (B), Task demand/cognitive complexity (T)

B.3 Subscripts and Superscripts

Notation	Meaning
i	Governance dimension index (i = 1...6: Schedule, Cost, Scope, Stakeholder, Supply Chain, Regulatory)
norm	Normalized — scaled to the [0, 1] governance range
composite	Composite — multi-dimensional weighted aggregate
deco	Decoherence-derived — computed via the decoherence pathway
F	Fatigue-specific — pertaining to the fatigue dimension
TD	Taleb–Douady formulation — curvature-based
baseline	Non-fatigue (baseline) component
effective	Combined total — baseline plus fatigue contributions
min, max	Calibrated boundary values for normalization anchors
tol	Tolerance threshold — perturbation level below which no value loss occurs
critical	Tipping-point value — the threshold separating stable from unstable regimes
safety	Safety-adjusted — value function incorporating safety performance

Notation	Meaning
crossing	Governance-band crossing — predicted transition time to next governance tier
+ , 0 , -	Perturbed up, baseline, perturbed down — the three evaluation points for finite-difference computation
²	Squared — second-order (second derivative, second moment, or squared term)
⁻¹	Inverse — reciprocal of a quantity

B.4 Operators and Special Notation

Symbol	Name	Usage in QPM
∂	Partial derivative	Differentiation with respect to one variable; $\partial/\partial\sigma$
∂^2	Second partial derivative	The curvature operator that defines QFI: $\partial^2 E[V]/\partial\sigma^2$
\int	Integral	Integration over perturbation domain in $E[V] = \int v(x) \cdot f(x; \sigma) dx$
Σ	Summation (capital sigma)	Weighted sum in composite QFI (F-10) and correlation correction (F-11)
∇	Gradient (nabla)	Spatial gradient of fatigue interface in KPZ equation: ∇h
∇^2	Laplacian	Second-order spatial derivative — the diffusion (smoothing) term in KPZ: $v\nabla^2 h$
$\langle \rangle$	Ensemble average	Statistical average over the workforce or simulation ensemble; used in roughness W
$\sqrt{\quad}$	Square root	Used in fidelity definition (F-05) and critical gradient (F-22)
$\exp(\quad)$	Exponential function	Appears in fidelity decay (F-06), payoff function (F-09), fatigue decay (F-24), risk spike (F-25)
$\max(\quad)$	Maximum function	Enforces tolerance threshold in payoff function: $\max(0, x - x_{tol})$
$\ln(\quad)$	Natural logarithm	Used in half-life calculation: $t_{1/2} = \ln(2)/k$; governance crossing time: $t_{crossing} = -\ln(F_{threshold}) / \Gamma$ (F-26)
$ _{\sigma=\sigma_0}$	Evaluation bar	Indicates derivative is evaluated at operating point σ_0

Symbol	Name	Usage in QPM
	Absolute value	Used in proximity metric: $ \epsilon^* / \epsilon_{\max} $

B.5 Governance Constants and Thresholds

Constant	Value(s)	Meaning
QFI governance bands	0–0.30 (Blue/Robust); 0.30–0.60 (Amber/Moderately Fragile); 0.60–0.85 (Orange/Highly Fragile); 0.85–1.00 (Red/Critical)	Governance-action tiers determining escalation, review frequency, and intervention authority
FRI governance bands	0–30 (Normal); 30–50 (Elevated); 50–70 (High); 70–100 (Critical)	Fatigue-state tiers determining task restrictions, monitoring frequency, and intervention mandates
Dimensional floor rule	$QFI_i \geq 0.85$ for any i	Triggers Orange-level governance regardless of composite QFI; prevents single-dimension detonators from being masked by composite averaging
Fatigue floor rule	$FRI \geq 70$	Triggers Orange-level governance regardless of composite QFI or individual QFI_i ; ensures internality-driven fragility receives equivalent escalation
Roughness floor rule	$FRI \geq 50$ AND $W > W_{\text{critical}}$	Triggers Amber-level governance even if composite QFI is in the Blue zone; captures dangerous local fatigue peaks
Perturbation distribution	Student-t, $df = 4$	Heavy-tailed distribution used in Monte Carlo QFI computation; fat tails capture rare-but-severe perturbations
Finite-difference step	$\Delta\sigma = \pm 10\%$ of σ_0	Standard perturbation step for numerical second-derivative computation

B.6 Dimensional Index (Six Governance Dimensions)

Index i	Dimension	Typical Perturbation Sources	Fatigue Impact (Section 3.4)
1	Schedule	Weather, labor availability, permitting delays, rework	Slower execution; rework loops; variable output; schedule compression worsens fatigue (feedback)
2	Cost	Material price volatility, change orders, productivity variance	Rework costs; incident costs; productivity loss; cost pressure drives overtime (feedback)
3	Scope	Design changes, regulatory amendments, stakeholder requests	Missed requirements; nonconforming acceptance; scope changes increase workload (feedback)
4	Stakeholder	Community opposition, political shifts, media attention	Degraded communication; slower decisions; stakeholder pressure drives schedule push (feedback)
5	Supply Chain	Supplier concentration, logistics disruption, quality failures	Ordering errors; inspection failures; supply delays force schedule recovery (feedback)
6	Regulatory	Code changes, inspection failures, environmental compliance	Procedural drift; documentation errors; regulatory findings require remediation (feedback)

Cross-Reference Matrix

Formula	Depends On	Feeds Into
F-01 Core QFI	$v(x)$ [F-09]; $f(x; \sigma)$; σ_0	F-02 (Normalization); F-10 (Composite)
F-02 Normalization	F-01; QFI_{min} ; QFI_{max}	Governance-band assignment
F-03 Expected Payoff	$v(x)$ [F-09]; $f(x; \sigma)$	F-04 (Operator)
F-04 Canonical Operator	F-03	F-01; F-12; F-13

Formula	Depends On	Feeds Into
F-05 Fidelity	$\rho_0; \rho(t)$	F-06; F-07; F-08
F-06 Decay	$\Gamma; t$	F-07; F-08
F-07 Value Loss	$F_Q [F-05/F-06]; \eta; \alpha$	F-08
F-08 Deco Proxy	$\Gamma; F_Q; \eta; \alpha$	F-10 (Composite); F-19
F-09 Payoff	$V_0; \beta_1; \beta_2; x_{tol}$	F-03; F-04; F-12
F-10 Composite	$w_i; QFI_i; C [F-11]$	Governance decisions
F-11 Correlation	$\rho_{ij}; w_i; w_j; QFI_i; QFI_j; \phi$	F-10
F-12 TD QFI	$v(x); f(x; \sigma); \sigma_0$	F-13 (Finite Difference)
F-13 Finite Difference	$E[v]_+; E[v]_0; E[v]_-; \Delta\sigma$	Numerical QFI computation
F-14 Proximity	$\epsilon^*; \epsilon_{max}$	Governance-band proximity assessment
F-15 Effective Γ	$\Gamma_{baseline}; \Gamma_{fatigue} [F-16]$	F-19
F-16 Fatigue Γ	$FRI; \gamma_0; \beta_f$	F-15
F-17 QFI_F (Proximity)	$\Delta F^*; \Delta F_{max}$	Fatigue governance; Dashboard Panel 2
F-18 QFI_F (Curvature)	$V_{safety}(F); \sigma_F$	Fatigue governance
F-19 Deco + Fatigue	$\eta_i; \Gamma_{effective} [F-15]; F_{Q,i}; \alpha_i$	Composite QFI with fatigue overlay
F-20 KPZ	$v; \lambda; \eta(x, t)$	F-21; F-22
F-21 Roughness	$h(x, t)$	Fatigue topology assessment; Dashboard Panel 1
F-22 Critical ∇h	$v; \lambda$	Tipping-point detection
F-23 FRI	$E; C; S; B; T$ subscales	F-16; F-17; Dashboard Panel 1
F-24 Fatigue Decay	$F_0; k; t$	Recovery planning; intervention scheduling

Formula	Depends On	Feeds Into
F-25 Risk Spike	$R_{\text{baseline}}; \delta; F; F_{\text{threshold}}$	Risk assessment; Safety-II monitoring
F-26 Crossing Time	$t_{\text{crossing}} = -\ln(F_{\text{thr}}) / \Gamma$	Predictive governance; anticipatory intervention timing

Formula numbers (F-01 through F-26) are stable identifiers intended for cross-referencing across all QFI sections.

Appendix B

QFI Quick-Reference Card

Designed for print and field use. Distribute to project managers, program directors, and governance teams.

Core Formula

$$QFI(S) = -\partial^2 E[V(S)] / \partial \sigma^2$$

$$QFI_{total} = \sum_i w_i \cdot QFI_i$$

Decoherence Bridge:

$$QFI = \eta \cdot \Gamma \cdot (1 - F_Q(t))$$

Interpretation Key

QFI Value	Meaning
QFI > 0	Fragile — value accelerates downward as stress increases (concave payoff)
QFI = 0	Robust — value is insensitive to perturbation magnitude (linear payoff)
QFI < 0	Antifragile — value improves under stress (convex payoff)

Fragility Spectrum — Thresholds and Escalation

Zone	QFI	Cadence	Authority	Action
Green	< 0	Quarterly	PM	Document optionality; harvest for replication
Blue	0 – 0.3	Monthly	Program Director	Standard risk register maintenance
Amber	0.3 – 0.6	Bi-weekly	Steering Committee	Contingency activation; root cause analysis
Orange	0.6 – 0.85	Weekly	C-Suite / Board	Structural intervention; resource reallocation
Red	≥ 0.85	Daily	Board / Sponsor	Immediate remediation or controlled termination

Dimensional Weights (Default)

Dimension	Weight	Primary Proxy
Schedule	0.25	SPI variance (Monte Carlo)
Cost	0.20	CPI second-moment analysis
Scope	0.20	Change request frequency × magnitude
Stakeholder	0.15	Approval-cycle variance
Supply Chain	0.10	Supplier concentration index
Regulatory	0.10	Regulatory change impact score

Dimensional Floor Rule

Any single dimension with $QFI_i \geq 0.85$ triggers Orange-level governance regardless of composite QFI score. Do not allow composite averaging to mask critically fragile dimensions.

Appendix C

QFI Governance Operations Playbook

This appendix consolidates and extends all governance-operational material from the Quantum Fragility Index (QFI) framework into a single deployment-ready reference. It is designed for governance committees, project control boards, PMO teams, and operational leadership who need to translate QFI outputs into proportional, accountable governance action. Cross-references to the parent document's theoretical foundations — including the fragility spectrum (Section 2.3.2), externality-driven governance (Sections 2.5–2.5.3), FRI–QFI combined decision logic (Sections 3.8–3.9), and roughness–curvature paired diagnostics (Appendix D.8) — are provided throughout. This playbook should be treated as the authoritative operational companion to the analytical content in Sections 2–4 and Appendices A, B, and D.

C.1 Fragility Spectrum and Band Definitions

The QFI governance framework operates on a five-tier fragility spectrum, each tier mapping to a distinct governance posture, monitoring cadence, and decision-authority level. The spectrum is derived from the normalized composite fragility index QFI_{norm} (Section 2.3.2) and provides the foundational classification system upon which all governance actions in this playbook are predicated. Band boundaries are calibrated to empirical fragility patterns and recalibrated annually (see Section C.2). The five tiers — Antifragile, Robust, Moderately Fragile, Highly Fragile, and Critically Fragile — represent a continuous escalation from optionality harvesting to emergency remediation.

Table C.1 — Five-Tier Fragility Spectrum

Band	Color Code	QFI_{norm} Range	Governance Posture	Monitoring Cadence	Decision Authority
Antifragile	Green	$QFI_{norm} < 0$	Harvest optionality; document sources of convexity for replication across portfolio	Quarterly review	Project Manager
Robust	Blue	$0 \leq QFI_{norm} < 0.30$	Standard oversight; routine governance cycle	Monthly review	Project Manager / PMO

Band	Color Code	QFI _{norm} Range	Governance Posture	Monitoring Cadence	Decision Authority
Moderately Fragile	Amber	$0.30 \leq \text{QFI}_{\text{norm}} < 0.60$	Enhanced monitoring; contingency readiness; root-cause investigation initiated	Bi-weekly review	PMO / Program Director
Highly Fragile	Orange	$0.60 \leq \text{QFI}_{\text{norm}} < 0.85$	Executive escalation; structural intervention required; remediation plan mandated within 10 business days	Weekly review	Steering Committee / Sponsor
Critically Fragile	Red	$\text{QFI}_{\text{norm}} \geq 0.85$	Immediate remediation or controlled termination; all discretionary scope frozen; board notification within 48 hours	Daily review	Executive Board / C-Suite

When QFI_{norm} crosses a band boundary — whether upward or downward — the transition triggers a mandatory governance review at the receiving band's cadence and authority level within one reporting cycle. Downward transitions (improvement) require documented evidence of sustained QFI reduction over two consecutive measurement periods before the lower band's monitoring cadence is adopted. Upward transitions (deterioration) take effect immediately.

Table C.2 — Band Transition Checklist

Transition	Trigger Criteria	Required Governance Action
Blue → Amber	QFI _{norm} crosses 0.30 upward OR any dimensional QFI _i exceeds 0.50	Activate enhanced monitoring; assign root-cause owner; prepare contingency activation plan
Amber → Orange	QFI _{norm} crosses 0.60 upward OR dimensional floor rule triggered (any QFI _i ≥ 0.85)	Escalate to Steering Committee; mandate remediation plan within 10 business days; activate intervention simulator
Orange → Red	QFI _{norm} crosses 0.85 upward OR fatigue floor rule triggered (FRI ≥ 70) with multiple dimensions above 0.60	Board notification within 48 hours; freeze discretionary scope; deploy emergency remediation protocol; daily monitoring

Transition	Trigger Criteria	Required Governance Action
Any band downward	QFI _{norm} sustained below threshold for two consecutive measurement periods	Document remediation effectiveness; transition to lower band cadence; update pattern library

C.2 Unified Decision Matrix

The Governance Decision Matrix translates QFI outputs into hard-wired, proportional governance actions. It integrates both the externality-driven QFI governance (Section 2.5.2) and the internality-driven FRI–QFI combined decision logic (Section 3.8.2) into a single unified matrix. Each row prescribes the full governance response chain — from required analysis through escalation, intervention, timeline, and accountability — for a given fragility band.

Table C.3 — Unified Governance Decision Matrix

Band	QFI _{norm} Range	Required Analysis	Escalation Path	Acceptable Interventions	Response Timeline	Accountability
Green	< 0	Document antifragile mechanisms; identify convexity sources	None — report in quarterly portfolio review	Replicate successful optionality structures across portfolio	Next quarterly cycle	Project Manager
Blue	0 – 0.30	Routine dimensional QFI _i review; trend monitoring	PMO monthly summary	Standard project controls; no extraordinary measures required	Monthly reporting cycle	Project Manager / PMO Lead
Amber	0.30 – 0.60	Root-cause analysis by dimension; proximity computation; what-if	Program Director briefed; Steering Committee	Contingency activation; targeted parameter modification (reduce Γ,	15 business days for root-cause report; 20 business	PMO Lead / Program Director

Band	QFI _{norm} Range	Required Analysis	Escalation Path	Acceptable Interventions	Response Timeline	Accountability
		simulation of top 3 interventions	employee informed	lower ρ_{ij} , increase tolerance); fatigue diffusion if FRI-driven	days for remediation plan	
Orange	0.60 – 0.85	Full dimensional decomposition; entanglement mapping; Taleb–Douady curvature deep-dive; FRI–QFI three-pathway analysis if fatigue-involved	Steering Committee owns response; Sponsor briefed weekly; Board notified if trend is upward	Structural intervention: scope restructuring, supply chain diversification, schedule re-baselining, mandatory rest protocols, crew rotation; intervention ranked by ΔQFI per \$M	10 business days for remediation plan; weekly progress reporting	Steering Committee Chair / Sponsor
Red	≥ 0.85	Emergency fragility assessment: all six dimensions plus FRI plus roughness; portfolio contagion check (entanglement with other projects)	Board owns response; CEO/COO briefed; regulatory notification if applicable	Controlled termination evaluation; emergency scope freeze; maximum intervention deployment; all-hands governance mobilization	48 hours for initial response; daily reporting until QFI < 0.85	Executive Board / C-Suite

QFI governance thresholds should be recalibrated annually using portfolio historical data to ensure governance actions remain aligned with empirical fragility patterns. Recalibration should be performed by the PMO Analytics function using the prior 12 months of QFI trajectory data, validated against actual project outcomes.

Table C.4 — RACI Matrix for QFI Governance Actions

Governance Action	Project Manager	PMO	Program Director	Steering Committee	Executive Board
QFI computation and reporting	R/A	C	I	I	—
Dimensional decomposition analysis	R	A	C	I	—
Root-cause investigation (Amber)	R	A	C	I	—
Remediation plan development (Orange)	C	R	A	C	I
Remediation plan approval (Orange)	I	C	R	A	I
Emergency response activation (Red)	C	C	R	A	R
Controlled termination evaluation	I	C	C	R	A
Annual threshold recalibration	C	R/A	C	I	I

R = Responsible | A = Accountable | C = Consulted | I = Informed

C.3 Floor Rules Registry

Floor rules are governance safeguards that prevent dangerous single-dimension or internality-driven fragility from being masked by composite averaging. They operate as hard overrides — when triggered, they force escalation regardless of the composite QFI score. The registry below consolidates the dimensional floor rule (Section 2.3.2), the fatigue floor rules (Section 3.9), and introduces an expanded cross-dimensional detonator rule derived from the logic of Section 3.9.1.

Table C.5 — Complete Floor Rules Registry

Floor Rule	Trigger Condition	Governance Override	Rationale	Source Section
Dimensional Floor Rule (Externality)	Any $QFI_i \geq 0.85$ across six governance dimensions (Schedule, Cost, Scope, Stakeholder, Supply Chain, Regulatory)	Trigger next higher governance tier regardless of composite QFI_{norm}	Single-dimension detonators (e.g., regulatory collapse) can produce catastrophic tail events that the composite weighted average masks	§2.3.2
Fatigue Floor Rule (Internality — Critical)	$FRI \geq 70$ (Critical band)	Apply Orange-level governance regardless of composite QFI or any individual dimensional QFI	At Critical FRI, internal control integrity is so degraded that external perturbations of any magnitude become dangerous	§3.9.2
Fatigue Roughness Floor Rule (Internality — Elevated)	$FRI \geq 50$ AND fatigue roughness W exceeds calibrated threshold (indicating "islands of exhaustion")	Escalate to at least Amber governance even if composite QFI remains in Blue zone	Average fatigue may be moderate but local peaks create dangerous gradients that composite QFI curvature alone may not fully reflect	§3.9.2
Cross-Dimensional Detonator Rule (Expanded)	Three or more dimensions simultaneously exceed 0.60 AND share a common driver (fatigue, supply chain concentration, or regulatory exposure)	Escalate to Orange-level governance regardless of composite; mandate entanglement mapping and common-driver analysis	Coordinated elevation across dimensions represents systemic vulnerability more dangerous than any single-dimension spike; composite averaging masks the coordinated pattern	Extended from §3.9.1

Non-Negotiable Override Policy

Floor rule overrides are non-negotiable governance instruments — project-level authority cannot waive them. Only the Steering Committee (for Amber/Orange overrides) or Executive Board (for Red overrides) may grant a documented exception, and only when accompanied by a written risk acceptance with explicit acknowledgment of the override's fragility implications and a time-bounded re-assessment schedule.

C.4 Dashboard Specification

The QFI Governance Dashboard is the primary visual interface through which governance teams consume QFI intelligence. This specification defines the complete dashboard architecture, integrating the base QFI dashboard primitives (Section 2.5), the FRI–QFI unified dashboard (Section 3.8.1), and the roughness–curvature paired reporting format (Appendix D.8.3). The six-panel layout ensures that all governance-relevant information — from portfolio-level heat maps to individual intervention simulations — is accessible within a single coherent interface.

C.4.1 Dashboard Architecture — Six-Panel Layout*Panel 1 — Portfolio Heat Map (from §2.5)*

Color-coded QFI_{norm} for all active projects, with clustering overlays for correlated fragility drivers. Projects sharing high entanglement ($\rho_{ij} > 0.6$) are visually linked with connector lines indicating the strength and directionality of cross-project coupling. Band color coding follows the five-tier spectrum: Green, Blue, Amber, Orange, Red. Each project cell displays the current QFI_{norm} value, a trend arrow ($\uparrow \rightarrow \downarrow$), and the number of active floor-rule triggers. Click-through navigation provides access to individual project drill-downs (Panel 2).

Panel 2 — Dimensional Drill-Down (from §2.5)

For the selected project, this panel displays all six dimensional QFI_i values on a radar chart with governance weights w_i , trend arrows ($\uparrow \rightarrow \downarrow$), local decoherence rate Γ_i , and nonlinearity coefficient η_i annotated for each axis. Dimensions triggering the dimensional floor rule ($QFI_i \geq 0.85$) are highlighted in red. The radar chart is overlaid with two reference polygons: the Amber boundary (0.30) and the Orange boundary (0.60), enabling immediate visual assessment of dimensional proximity to governance thresholds.

Panel 3 — FRI State and Fatigue Overlay (from §3.8.1)

Real-time FRI score (0–100) with band classification (Normal / Elevated / High / Critical). FRI trend over the trailing four measurement periods is displayed as a sparkline. Roughness W is reported with a trend arrow and the current maximum gradient ∇h_{max} expressed as a percentage of

$\nabla h_{critical}$. A six-dimension radar chart with a fatigue halo shows the amplification effect per dimension — dimensions where fatigue transmission is strongest are highlighted with proportionally wider halo bands. This panel integrates the functionality of Panels 1 and 3 from the FRI–QFI unified dashboard (Section 3.8.1).

Panel 4 — Proximity Indicator (from §2.5)

Displays the operational Proximity metric: $Proximity = 1 - |\epsilon^*| / |\epsilon_{max}|$. This answers the critical governance question: "How close are we to the nearest governance boundary?" The panel presents a single numeric value plus a directional driver identifying which dimension or factor is pushing toward the boundary. A color-coded gauge provides immediate visual assessment: green (>0.40 margin to boundary), yellow (0.20–0.40 margin), red (<0.20 margin to boundary).

Panel 5 — Intervention Simulator (from §2.5, §3.8.1)

What-if panel showing projected QFI changes under candidate interventions, integrating the what-if functionality from Section 2.5 with the intervention simulation from Section 3.8.1, Panel 4. For each intervention, the panel displays: the modified parameter ($\Gamma, \rho_{ij}, \epsilon_{max}, \eta, FRI$), projected ΔQFI_i per affected dimension, projected $\Delta QFI_{composite}$, estimated implementation cost, and the cost-effectiveness ratio (ΔQFI per \$M). Interventions are ranked by cost-effectiveness in descending order.

Panel 6 — Convergence Assessment (from D.8.3)

Paired-panel format for roughness–curvature cross-validation. Left sub-panel: Fatigue Roughness Panel — current W , trend over four periods, ∇h_{max} as a percentage of $\nabla h_{critical}$, and a crew/zone heat map showing spatial distribution of fatigue roughness. Right sub-panel: QFI Curvature Panel — composite QFI_{norm} with band color, dimensional radar, and fatigue contribution to each dimension annotated. Bottom strip: convergence/divergence assessment with diagnostic guidance — when roughness and curvature signals align, the assessment reads "Convergent"; when they diverge, diagnostic guidance identifies the likely cause and recommended investigation.

Table C.6 — Dashboard Data Requirements

Panel	Data Source	Refresh Cadence	Minimum Data Quality Requirement
Portfolio Heat Map	QFI computation engine; project registry	Weekly (Amber+); Monthly (Blue)	All six dimensional QFI; computed with current-period data
Dimensional Drill-Down	Dimensional QFI model; decoherence proxy telemetry	Matches project's governance band cadence	Decoherence rate Γ_i estimated from minimum 3 measurement periods

Panel	Data Source	Refresh Cadence	Minimum Data Quality Requirement
FRI State / Fatigue Overlay	FRI telemetry (Tier 1–3 data streams); KPZ roughness computation	Daily (field-level); Weekly (dashboard)	FRI scores by crew/zone; minimum 80% workforce coverage
Proximity Indicator	Proximity computation engine; governance band boundaries	Real-time where feasible; minimum weekly	Current QFI_{norm} plus boundary-distance calculation
Intervention Simulator	What-if perturbation model; cost estimation database	On-demand (analyst-triggered)	Validated parameter sensitivity ranges for target dimensions
Convergence Assessment	Roughness W from FRI data; QFI from curvature computation	Monthly minimum; weekly for Orange+	Contemporaneous roughness and QFI data (same measurement period)

C.5 Escalation Protocols

This section formalizes the human decision chain that activates when QFI outputs trigger governance-band transitions. Building on scattered references in Sections 2.5 and D.8.2, the protocols define who detects, who assesses, who decides, and who monitors at each escalation level, with time-to-escalate targets calibrated by band severity. The four-stage escalation process ensures that no governance signal is lost between detection and action.

C.5.1 Escalation Flowchart — Four-Stage Process

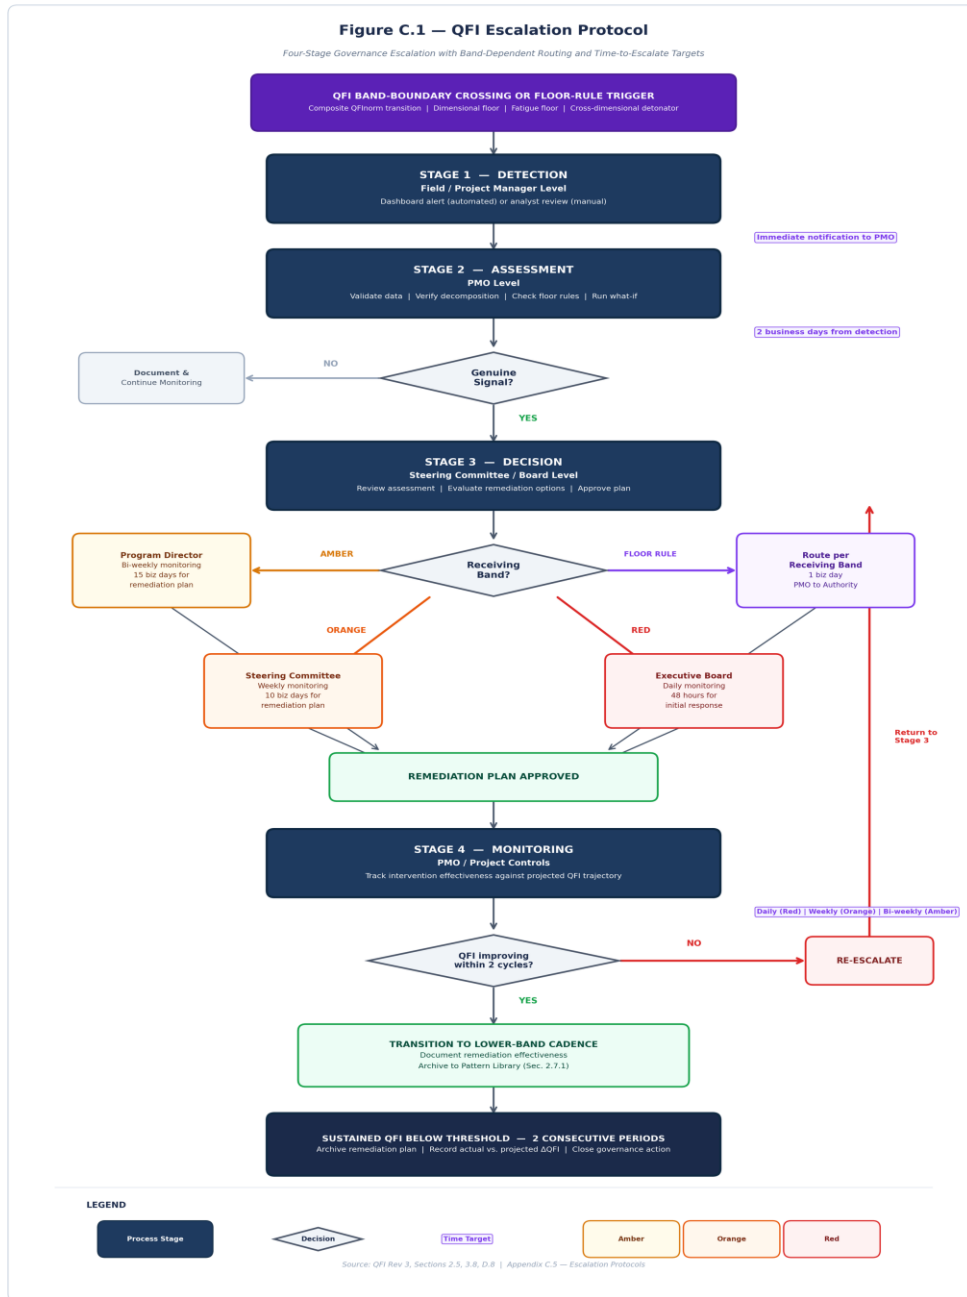


Figure C.1 QFI Escalation Protocol: Four-stage governance escalation with band-dependent routing and time-to-escalate targets.

Stage 1 — Detection (Field / PM Level)

QFI computation identifies a band-boundary crossing or floor-rule trigger. Detection can be automated (dashboard alert generated by the QFI computation engine) or manual (analyst review during routine dimensional assessment). Upon detection, the Project Manager or designated QFI analyst generates an initial signal report documenting: the specific trigger condition, the current QFI_{norm} and all dimensional QFI_i values, the receiving governance band, and any active floor-rule overrides. Time-to-escalate: immediate notification to PMO upon detection.

Stage 2 — Assessment (PMO Level)

The PMO validates the QFI signal — confirming data quality, verifying dimensional decomposition, checking for floor-rule triggers, and running preliminary what-if simulations. The assessment determines whether the signal is genuine (proceed to Stage 3) or a data artifact (document the anomaly and continue monitoring at the current cadence). For genuine signals, the PMO prepares a validated assessment package including confirmed QFI state, root-cause hypothesis, preliminary intervention options, and recommended escalation path. Time-to-escalate: 2 business days from detection.

Stage 3 — Decision (Steering Committee / Board Level)

Depending on the band, the appropriate decision authority reviews the validated assessment, the remediation options with ΔQFI projections, and the cost-effectiveness ranking from the intervention simulator (Section C.6). The decision authority approves a remediation plan, requests additional analysis, or escalates further. For Orange-level decisions, the Steering Committee or Sponsor owns the response. For Red-level decisions, the Executive Board owns the response with CEO/COO briefed. Time-to-escalate: 5 business days for Orange; 48 hours for Red.

Stage 4 — Monitoring (PMO / Project Controls)

Post-decision monitoring tracks the effectiveness of approved interventions against projected QFI trajectories. The PMO reports back to the decision authority at the band-appropriate cadence, comparing actual ΔQFI against projected ΔQFI and flagging any divergence greater than 20%. If QFI fails to improve within two measurement cycles following intervention deployment, the protocol triggers automatic re-escalation to Stage 3 with updated analysis.

Table C.7 — Time-to-Escalate Targets by Band

Band Transition	Detection → PMO	PMO → Decision Authority	Decision → Remediation Plan	Remediation → First Progress Report
Blue → Amber	Same business day	2 business days	15 business days	Next bi-weekly review
Amber → Orange	Immediate (automated alert)	1 business day	10 business days	7 calendar days
Orange → Red	Immediate (automated alert + phone/SMS)	4 hours	48 hours (initial response)	Daily until QFI < 0.85
Floor Rule Trigger (any)	Immediate (automated alert)	1 business day	Same as receiving band	Same as receiving band

C.6 What-If Simulation Protocol

The what-if simulation protocol enables governance teams to test candidate interventions before implementation, ranking them by projected QFI reduction per unit cost. This section provides the standardized five-step procedure for conducting and reporting what-if analyses, consolidating the simulation logic from Section 2.5 and the FRI–QFI intervention panel from Section 3.8.1, Panel 4.

Step 1 — Input Preparation

Assemble the current QFI state vector: all six dimensional QFI_i values (Schedule, Cost, Scope, Stakeholder, Supply Chain, Regulatory), current composite QFI_{norm} , decoherence rates Γ_i , correlation coefficients ρ_{ij} , fidelity values $F_{Q_i}(t)$, and nonlinearity coefficients η_i . If the fragility is fatigue-involved (FRI > 0 or roughness W above baseline): include current FRI score, roughness W , and maximum gradient ∇h_{max} with its spatial location.

Step 2 — Scenario Definition

Define 3–5 candidate interventions. For each, specify: (a) which model parameters are modified — reduce Γ (decoherence rate), lower ρ_{ij} (correlation between dimensions), increase $|\epsilon_{max}|$ (perturbation tolerance), reduce η (nonlinearity coefficient), or reduce FRI (fatigue index); (b) the magnitude of modification, based on engineering judgment and historical precedent from the QFI

Pattern Library; (c) estimated implementation cost in dollars and resource-hours; (d) implementation timeline in business days or weeks.

Step 3 — QFI Recomputation

For each intervention scenario: recompute dimensional QFI_i using modified parameters through the QFI_F formulation (Section 2.3.1); recompute the correlation correction C with modified ρ_{ij} ; recompute composite QFI_{norm} ; recompute the Proximity metric to determine the new distance to the nearest governance boundary; and check floor-rule status to confirm whether any overrides remain active post-intervention.

Step 4 — Intervention Ranking

Table C.8 — Intervention Ranking Template

Rank	Intervention	Target Parameter	Projected $\Delta QFI_{composite}$	Estimated Cost (\$M)	ΔQFI per \$M	Implementation Time
1	[Intervention description]	[Parameter modified]	[-0.XX]	[\$X.X]	[-0.XX/\$M]	[X weeks]
2	[Intervention description]	[Parameter modified]	[-0.XX]	[\$X.X]	[-0.XX/\$M]	[X weeks]
3	[Intervention description]	[Parameter modified]	[-0.XX]	[\$X.X]	[-0.XX/\$M]	[X weeks]
4	[Intervention description]	[Parameter modified]	[-0.XX]	[\$X.X]	[-0.XX/\$M]	[X weeks]
5	[Intervention description]	[Parameter modified]	[-0.XX]	[\$X.X]	[-0.XX/\$M]	[X weeks]

Step 5 — Reporting

Present results to the appropriate decision authority using the standard format: current QFI state (composite and dimensional), intervention options ranked by cost-effectiveness, recommended action with justification, projected QFI trajectory if intervention is approved, and projected QFI trajectory if no action is taken (baseline deterioration scenario). The report should include confidence indicators for each projection based on the convergence of the two computation methods.

Dual-Method Validation

The what-if simulation should be run using both the Taleb–Douady curvature method (for deep scenario analysis) and the decoherence proxy (for rapid turnaround). When both methods converge on the same intervention ranking, confidence in the recommendation is high (see Section 2.3.3). When the methods diverge, the discrepancy should be documented and the more conservative ranking adopted pending further analysis.

C.7 PM Framework Integration Map

The QFI is designed to complement, not replace, established project management frameworks. This section provides concrete integration guidance for four major frameworks — Earned Value Management, PMBOK Risk Management, ISO 31000, and PRINCE2 — specifying where QFI data enters each framework's workflow and how QFI outputs enhance existing governance processes. This material consolidates and expands the integration map introduced in Section 2.5.3.

Table C.9 — EVM Integration

EVM Element	QFI Integration Point	Value Added by QFI
Schedule Performance Index (SPI)	Overlay QFI_{norm} on SPI trend chart; flag projects where $SPI > 0.95$ but $QFI > 0.60$	Detects hidden nonlinear fragility beneath apparently stable linear performance; $SPI = 0.98$ with $QFI = 0.72$ indicates concave payoff — small perturbation increases cause disproportionate collapse
Cost Performance Index (CPI)	Compute cost-dimension QFI_{cost} alongside CPI; report both in monthly performance summaries	CPI measures cost efficiency; QFI_{cost} measures cost fragility — sensitivity of cost outcomes to volatility increases
Estimate at Completion (EAC)	Apply QFI-adjusted stress scenarios to EAC projections; report "EAC under QFI stress" alongside deterministic EAC	Deterministic EAC assumes linear continuation; QFI stress-adjusted EAC captures nonlinear tail-risk exposure

EVM Element	QFI Integration Point	Value Added by QFI
Variance at Completion (VAC)	Include QFI curvature overlay showing how VAC changes under increased perturbation dispersion	Reveals whether the VAC trajectory is fragile (concave — worsening under stress) or robust (linear — stable under stress)
To-Complete Performance Index (TCPI)	Cross-reference TCPI feasibility against current QFI_{norm} ; flag if TCPI requires performance improvement while QFI is Orange/Red	High TCPI demand combined with high QFI indicates that the required performance recovery is itself fragile — unlikely to be sustained

Table C.10 — PMBOK Risk Management Integration

PMBOK Process	QFI Integration Point	Value Added by QFI
Identify Risks	Use entanglement mapping to surface cross-dimensional coupling risks invisible to event-based identification; use dimensional QFI_i to identify fragile dimensions as risk-generation zones	Moves beyond individual risk events to systemic fragility patterns; identifies "the risks you didn't know to look for"
Perform Qualitative Analysis	Map identified risks against the dimensional QFI radar; assign fragility context to each risk (which dimension, current QFI_i , proximity to boundary)	Provides quantitative grounding for qualitative risk prioritization; a risk in a $QFI = 0.75$ dimension is categorically more urgent than the same risk in a $QFI = 0.25$ dimension
Plan Risk Responses	Use what-if simulation (C.6) to test risk response effectiveness against QFI; rank responses by ΔQFI per \$M	Transforms risk response planning from judgment-based to evidence-based; interventions are prioritized by measured fragility reduction

PMBOK Process	QFI Integration Point	Value Added by QFI
Monitor Risks	Include QFI trend in risk monitoring dashboard; trigger risk re-assessment when QFI crosses band boundaries	Provides continuous fragility monitoring between periodic risk reviews; detects emerging vulnerability patterns before individual risk events manifest

Table C.11 — ISO 31000 Integration

ISO 31000 Principle	QFI Integration Point	Value Added by QFI
Sensitivity Analysis (Risk Assessment)	QFI provides the mathematical foundation ISO 31000 envisions for sensitivity analysis — specifically, second-order sensitivity to dispersion	Fills the methodological gap: ISO 31000 calls for sensitivity analysis but does not prescribe a specific methodology; QFI delivers rigorous second-derivative sensitivity measurement
Monitoring and Review	Embed QFI governance-band monitoring into the ISO 31000 monitoring cycle; use proximity metric as the primary early-warning indicator	Operationalizes ISO 31000's requirement for continuous monitoring with a quantitative, threshold-based system
Communication and Consultation	Use QFI governance bands and proximity metric in stakeholder communication; replace subjective risk language with QFI-grounded assessments	Standardizes risk communication across stakeholders using a single, calibrated metric rather than subjective severity labels

Table C.12 — PRINCE2 Integration

PRINCE2 Principle / Process	QFI Integration Point	Value Added by QFI
Continued Business Justification	Map QFI _{norm} against business case viability; a project in Red zone (QFI ≥ 0.85) is one where business justification may no longer hold under realistic perturbation scenarios	Operationalizes continued business justification with quantitative rigor — not "is the project still viable?" but "how much perturbation can the project absorb before viability collapses?"
Manage by Stages (Stage Gate Reviews)	Require QFI assessment at each stage gate; do not approve stage progression if QFI is Orange/Red without a documented remediation plan and Board-approved risk acceptance	Transforms stage gates from schedule-driven checkpoints to fragility-aware decision points
Manage by Exception (Tolerance)	Define QFI-based tolerances alongside cost/time tolerances; specify QFI band as an exception trigger (e.g., QFI crossing from Amber to Orange triggers an exception report to the Project Board)	Adds fragility tolerance to the existing cost/time/scope tolerance framework; captures systemic vulnerability that traditional tolerances miss

C.8 Governance Meeting Cadence

The QFI governance cycle operates at four cadences — daily, weekly, monthly, and quarterly — aligned with the roughness–curvature diagnostic domain map (Appendix D.7). Each cadence serves a distinct diagnostic purpose and produces a distinct governance output. The cadence structure ensures that field-level roughness signals are captured in real time, while portfolio-level curvature analysis receives the sustained attention required for strategic decision-making. This section consolidates material from Appendix D.8.1 and expands it with standardized agenda templates.

Table C.13 — Governance Meeting Cadence and Agenda Structure

Cadence	Primary Diagnostic	Attendees	Duration	Standing Agenda Items
Daily Field Review	KPZ Roughness dominates; real-time FRI by crew/zone	Field Superintendent; Safety Manager; Project Controls Lead	15 minutes	(1) Current FRI by crew/zone (2) Roughness W and ∇h_{\max} vs $\nabla h_{\text{critical}}$ (3) Floor-rule status check (4) Immediate diffusion actions if $\nabla h_{\max} > 0.8 \times \nabla h_{\text{critical}}$
Weekly Progress Review	Roughness trend + QFI_F proximity; both metrics contribute	Project Manager; PMO Analyst; Discipline Leads	45 minutes	(1) Weekly roughness trend overlay against QFI_F trend (2) Dimensional QFI_i update for Amber+ dimensions (3) Proximity metric with directional driver (4) Lead-lag assessment: is roughness leading curvature? (5) Pre-position interventions if roughness is rising faster than QFI
Monthly Steering Review	Cross-validation zone; roughness and curvature presented side by side	Program Director; Steering Committee; PMO Lead; Risk Manager; Technical Advisor	90 minutes	(1) Portfolio heat map with QFI_{norm} for all active projects (2) Roughness–curvature convergence assessment for Orange+ projects (3) Entanglement mapping: cross-project fragility coupling (4) What-if simulation results for active remediation plans (5) Floor-rule override requests (6) Annual calibration status update

Cadence	Primary Diagnostic	Attendees	Duration	Standing Agenda Items
Quarterly Scenario Stress Test	Taleb–Douady curvature dominates; full Monte Carlo perturbation sweeps	Executive Board; Program Director; PMO Analytics; External Advisors (optional)	Half-day workshop	(1) Full Taleb–Douady Monte Carlo sweep under alternative dispersion scenarios (2) Portfolio ranking by QFI_{norm} (3) Fatigue-modulated decoherence projections (4) Antifragility review: document Green-band projects and replicate optionality structures (5) Threshold recalibration review (6) Strategic 6-month QFI trajectory projections

C.9 Remediation Planning Template

This section provides a standardized remediation planning template for use when QFI governance actions require formal intervention. Building on the worked examples in Section 2.4, the template ensures that all remediation plans follow a consistent structure, enabling portfolio-level comparison and institutional learning. Every project entering Orange or Red governance is required to complete this template within the timelines specified in Table C.7.

Table C.14 — Remediation Plan Template

Field	Description / Instructions
Project Identification	Project name, ID, current composite QFI_{norm} , current governance band, date of band-transition trigger
Triggering Condition	Specific trigger: composite band crossing, dimensional floor rule (which QFI_i), fatigue floor rule (FRI value), or cross-dimensional detonator (which dimensions, common driver)
Root-Cause Analysis	Identify the primary fragility driver(s) by dimension; link to decoherence rate Γ_i analysis, entanglement mapping, or roughness diagnostic as appropriate
Target QFI State	Define the post-remediation target: target QFI_{norm} , target governance band, target dimensional QFI_i values for the affected dimensions

Field	Description / Instructions
Intervention Description	Detailed description of each proposed intervention; reference the what-if simulation results from Section C.6
Target Parameter(s)	For each intervention: which model parameter is being modified (Γ , ρ_{ij} , ϵ_{max} , η , FRI, W)
Projected Δ QFI	Projected change in composite QFI_{norm} and affected dimensional QFI_i values; derived from what-if simulation
Estimated Cost	Total cost of intervention in dollars and resource-hours; include both direct costs and opportunity costs
Cost-Effectiveness Ratio	Δ QFI per \$M; enables cross-project and cross-intervention comparison
Implementation Timeline	Start date, key milestones, target completion date; aligned with governance band monitoring cadence
Success Criteria	Quantitative criteria: sustained QFI_{norm} below target threshold for minimum two consecutive measurement periods; no floor-rule triggers; Proximity metric above 0.30
Post-Remediation Measurement Protocol	Schedule of post-intervention QFI measurements; frequency aligned with band cadence; criteria for declaring remediation complete and transitioning to lower-band monitoring; lessons-learned entry for the QFI Pattern Library (Section 2.7.1)

Every completed remediation plan — successful or unsuccessful — should be archived in the QFI Pattern Library (Section 2.7.1, Phase 3) with full before/after QFI trajectory data, intervention details, actual versus projected Δ QFI, and actual versus estimated cost. This institutional memory enables progressively more accurate what-if simulations and more effective remediation design across the organization's project portfolio.

This Governance Operations Playbook is designed to be a living document. As QFI is deployed across active portfolios and calibrated against empirical data, the thresholds, timelines, and templates in this appendix should be refined through the annual recalibration process described in Section C.2. The test of a governance framework is not its theoretical elegance but its operational

utility — whether it enables leadership to detect fragility earlier, decide faster, intervene more effectively, and learn more systematically. This playbook provides the operational infrastructure to achieve those objectives.

Appendix D

Roughness and Curvature: A Structural Comparison of Two Second-Order Fragility Diagnostics

Appendix D unifies the physical and decision-space representations of fragility by pairing stochastic fatigue roughness with value-function curvature under dispersion. The **Fatigue Roughness Surface (KPZ)** models the propagation of fatigue through stochastic noise ϵ , height h , and gradient ∇h , revealing how localized amplification and recovery basins emerge from roughness W . Extending this foundation, the **Value Function Curvature Under Dispersion** translates physical roughness into behavioral geometry — contrasting concave (fragile) and convex (antifragile) payoff shapes as dispersion σ increases. Together, these figures form a dual framework: the first quantifies physical instability, the second interprets its systemic and cognitive consequences, enabling governance-grade insight into how stochastic fatigue manifests as curvature in decision-making and risk response.

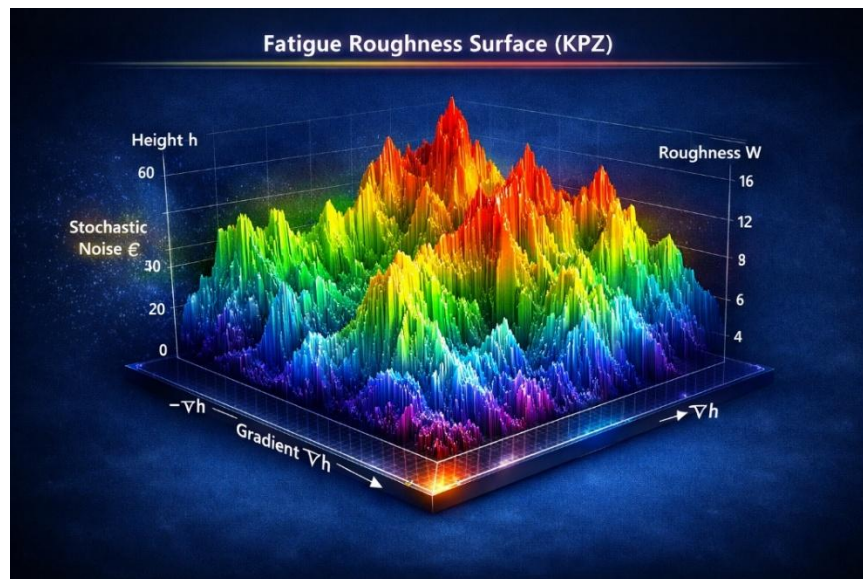


Figure D.0a Fatigue Roughness Surface (KPZ)

Three-dimensional stochastic surface illustrating fatigue roughness W as a function of gradient ∇h and stochastic noise ϵ . Peaks represent localized amplification zones where fatigue accumulates; valleys indicate recovery basins. The KPZ formulation visualizes dynamic instability and roughness propagation across operational domains.




Figure D.0b Value Function Curvature Under Dispersion

Comparative curvature plot showing concave (fragile) and convex (antifragile) payoff geometries under increasing dispersion σ . The concave regime amplifies downside risk and loss sensitivity, while the convex regime enhances adaptive gain potential. Together, the curves depict how dispersion transforms fragility into antifragility through curvature dynamics.

D.1 Purpose and Scope

Section 3.5 identifies a deep structural parallel between two concepts that originate in different intellectual traditions yet converge on the same governance objective: detecting nonlinear vulnerability before it manifests as failure. This appendix develops that parallel rigorously, providing the mathematical, operational, and governance-level comparison that Section 3.5 summarizes.

The two concepts are:

KPZ-Based Fatigue Roughness — drawn from the Kardar–Parisi–Zhang equation in statistical physics (Kardar, Parisi, & Zhang, 1986), adapted in the companion Fatigue Risk Index paper (Prieto, 2026) to model the evolving "topology of exhaustion" across a project's work packages. Roughness measures the spatial variance of the fatigue interface — how uneven the fatigue landscape is across entangled crews and tasks.

Taleb–Douady Curvature — drawn from the mathematical definition of fragility developed by Taleb and Douady (2012), operationalized in Section 2 of this paper as the Quantum Fragility Index (QFI). Curvature measures the second derivative of expected project value with respect to perturbation dispersion — how the shape of the payoff function changes as environmental volatility increases.

Both are second-order measures. Both detect the same underlying reality — that a system's vulnerability resides in its *shape*, not its *level*. And both provide governance signals that first-order metrics (averages, expected values, linear trends) systematically miss. Yet they operate on different observables, at different time scales, and with different data requirements.

Understanding where they converge and where they diverge is essential for integrated QPM governance.

D.2 Origins and Intellectual Lineage

D.2.1 KPZ Roughness — From Surface Physics to Fatigue Topology

The KPZ equation was introduced in 1986 by Mehran Kardar, Giorgio Parisi, and Yi-Cheng Zhang to describe how interfaces — such as the boundary between a growing crystal and its medium — evolve under the combined influence of deterministic growth, lateral propagation, and random noise. The equation takes the form:

$$\partial h / \partial t = \nu \nabla^2 h + \lambda (\nabla h)^2 + \eta(x, t)$$

where:

- $h(x, t)$ is the interface height at position x and time t
- $\nu \nabla^2 h$ is the smoothing (diffusion) term — tendency toward uniformity
- $\lambda (\nabla h)^2$ is the nonlinear growth term — local slope amplifies growth
- $\eta(x, t)$ is stochastic noise — random environmental shocks

The key observable derived from the KPZ equation is roughness $W(L, t)$, defined as the root-mean-square deviation of interface height from the mean:

$$W(L, t) = \langle (h - \langle h \rangle)^2 \rangle^{1/2}$$

The companion paper on fatigue risk (Prieto, 2026) maps this framework onto project management by reinterpreting each term:

- $h(x, t) \rightarrow$ cumulative fatigue debt of work package x at time t
- $\nu \nabla^2 h \rightarrow$ management's diffusion capacity — the ability to smooth fatigue peaks through rest mandates, crew rotation, and workload redistribution
- $\lambda (\nabla h)^2 \rightarrow$ quadratic contagion — the mechanism by which fatigue gradients between entangled crews propagate stress at a rate proportional to the square of the slope
- $\eta(x, t) \rightarrow$ environmental shocks — heat events, supply chain disruptions, geological anomalies
- $W(L, t) \rightarrow$ fatigue roughness — the variance of fatigue scores across the project workforce at time t

The critical insight: roughness is a stronger predictor of fatigue-driven risk than average FRI because it captures the *shape* of the fatigue landscape rather than its *height*. A project with average FRI = 42 but roughness $W = 22.4$ (some crews at 15, others at 75) is far more dangerous than a project with average FRI = 42 and roughness $W = 2.1$ (all crews at 40–44). The former has developed "islands of exhaustion" — steep fatigue gradients that overwhelm management's diffusion capacity and trigger the quadratic contagion described in Section 3.2.2.

D.2.2 Taleb–Douady Curvature — From Financial Fragility to Project Governance

The mathematical definition of fragility was formalised by Nassim Nicholas Taleb and Raphael Douady in their 2012 paper, building on Taleb's broader program of distinguishing fragility from risk. Their key insight: fragility is not about the probability or expected magnitude of adverse events — it is about the *shape* of the payoff function under stress. Specifically, fragility is the negative second derivative of expected payoff with respect to the dispersion of perturbations:

$$QFI = -\partial^2 E[V] / \partial \sigma^2$$

where $E[V]$ is the expected project value and σ is the dispersion (volatility) of the perturbation distribution.

Section 2 of this paper operationalizes this definition for project governance by:

1. Defining a smooth payoff function $v(x) = V_0 \cdot \exp(-\beta_1 \cdot \max(0, x - x_{tol})^{\beta_2})$ that maps perturbation magnitude x to project value
2. Choosing a heavy-tailed perturbation distribution $f(x; \sigma)$ (Student-t, $df = 4$) parameterized by dispersion σ
3. Computing the curvature numerically via finite-difference Monte Carlo: evaluate $E[v]$ at $\sigma_0 - \Delta\sigma$, σ_0 , $\sigma_0 + \Delta\sigma$ and approximate the second derivative
4. Normalizing the raw curvature to a $[0, 1]$ governance scalar using portfolio-calibrated anchors

The result — QFI_{norm} — tells governance how much worse the expected payoff becomes as environmental volatility increases. A high QFI_{norm} (e.g., 0.68) indicates that the value function is deeply concave: small increases in perturbation dispersion produce disproportionately large value losses. This is fragility.

D.3 The Structural Parallel — Both Are Second-Order Shape Detectors

Despite their different origins, KPZ roughness and TD curvature share a fundamental mathematical architecture:

Both are second-order measures. Roughness is the variance (second central moment) of the fatigue interface. Curvature is the second derivative of expected value with respect to dispersion. Neither operates on levels, averages, or first-order trends — both detect *shape*.

Both are insensitive to the mean. A project can have a "good" average FRI (low mean fatigue) yet dangerous roughness (high variance → islands of exhaustion). Similarly, a project can have a "good" expected value (high $E[V]$) yet high curvature (concave payoff → fragile). First-order metrics mask these vulnerabilities; second-order metrics expose them.

Both detect tipping points. In KPZ dynamics, the system transitions from stable to unstable when the fatigue gradient exceeds $\nabla h_{\text{critical}} = \sqrt{(v/\lambda)}$ — the point where quadratic contagion overwhelms management's smoothing capacity. In TD curvature, the system transitions from robust to fragile when QFI_{norm} crosses governance thresholds (0.30, 0.60, 0.85). Both frameworks identify the boundary between recoverable variability and self-reinforcing instability.

Both are actionable. Roughness tells governance *where* to apply diffusion (which crews to rest, which interfaces to decouple). Curvature tells governance *where* to reduce dispersion or increase tolerance (which perturbation sources to attenuate, which payoff nonlinearities to flatten). Both produce directional, intervention-ready diagnostics.

The following table formalizes the structural correspondence:

Property	KPZ Roughness	Taleb–Douady Curvature
Mathematical order	Second moment (variance)	Second derivative (curvature)
What it measures	Spatial unevenness of fatigue across work packages	Sensitivity of expected value to perturbation dispersion
Observable domain	Fatigue scores across crews/zones at a point in time	Expected project value under varying environmental volatility
Physical analogy	How "bumpy" the fatigue surface is	How "curved" the value function is
What it detects	Islands of exhaustion; steep risk gradients	Concavity of payoff; nonlinear vulnerability
What it misses if used alone	Whether those gradients will produce value loss	Whether value loss is caused by internal gradients or external shocks
Key threshold	$\nabla h_{\text{critical}} = \sqrt{(v/\lambda)}$ — tipping point where contagion exceeds diffusion	$QFI_{\text{norm}} = 0.60/0.85$ — governance escalation thresholds
Governance signal	"The fatigue landscape is dangerously uneven — smooth it before contagion propagates"	"The value function is dangerously concave — intervene before perturbations exploit it"

Property	KPZ Roughness	Taleb–Douady Curvature
Time scale	Real-time to daily (field telemetry)	Weekly to quarterly (scenario analysis)
Data requirements	FRI scores by crew/zone; shift records; environmental readings	Payoff model; perturbation distribution; Monte Carlo simulation
Computational cost	Low (statistical summary of telemetry)	Moderate to high (Monte Carlo with perturbation sweeps)
QPM construct	Entanglement (correlation between crews); Propagation velocity	Decoherence (coherence loss under stress); Superposition collapse

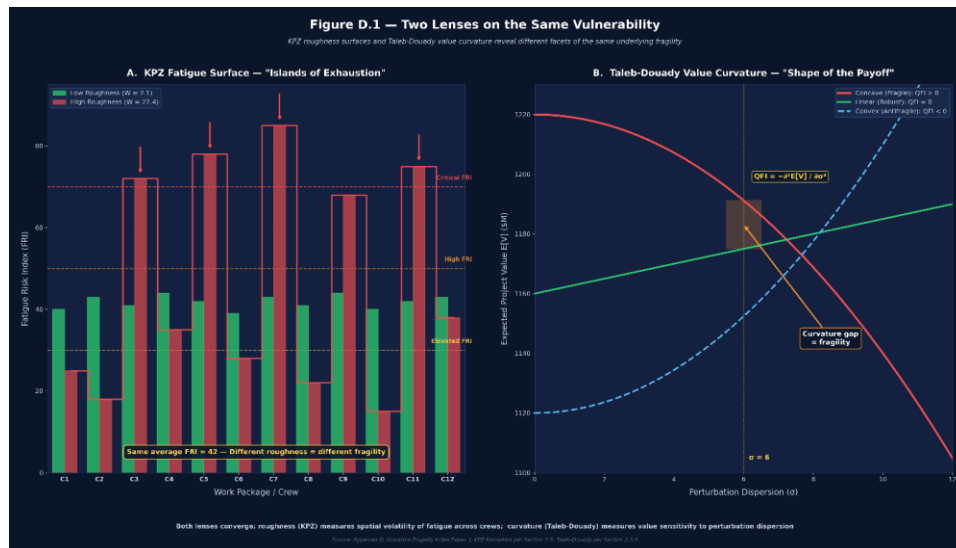


Figure D.1 Two Lenses on the Same Vulnerability

Panel A shows the KPZ fatigue surface across 12 crews under two scenarios with identical average FRI (≈ 42) but different roughness: low roughness ($W = 2.1$, green — all crews near 42) versus high roughness ($W = 22.4$, red — "islands of exhaustion" with peaks above Critical threshold). Panel B shows the Taleb–Douady value curve under three payoff shapes: concave (fragile, $QFI > 0$), linear (robust, $QFI \approx 0$), and convex (antifragile, $QFI < 0$). The curvature gap between the concave payoff and the linear reference at σ_0 is the fragility the QFI measures.

D.4 The Convergence Mechanism — Why They Detect the Same Thing

The structural parallel is not accidental. Roughness and curvature converge because they are measuring the same underlying phenomenon — nonlinear vulnerability — through complementary observational windows.

The causal chain is:

- 1. Roughness creates gradients.** When the fatigue surface develops peaks and valleys (high W), steep risk gradients form between adjacent work packages.
- 2. Gradients drive contagion.** The quadratic contagion term $\lambda(\nabla h)^2$ means that stress propagates at a rate proportional to the square of the gradient. Steep slopes produce quadratically faster contagion than gentle slopes.
- 3. Contagion amplifies decoherence.** As fatigue contagion propagates, it accelerates coherence loss across the project system — $\Gamma_{\text{effective}}$ increases, fidelity $F_Q(t)$ drops faster, and the system's capacity to absorb perturbation degrades.
- 4. Decoherence steepens curvature.** As decoherence increases, the value function becomes more concave. The same perturbation dispersion σ produces a larger expected value loss because the internal control systems (human judgment, supervisory oversight, procedural compliance) that normally buffer against perturbation are degraded by fatigue.
- 5. Curvature manifests as fragility.** The steeper the curvature, the higher the QFI — and the closer the system is to a governance-boundary crossing.

Therefore: high roughness → steep gradients → quadratic contagion → accelerated decoherence → steeper curvature → higher QFI.

This is why projects with high fatigue roughness consistently exhibit higher composite QFI than projects with the same average FRI but low roughness. They are measuring the same vulnerability from different vantage points.

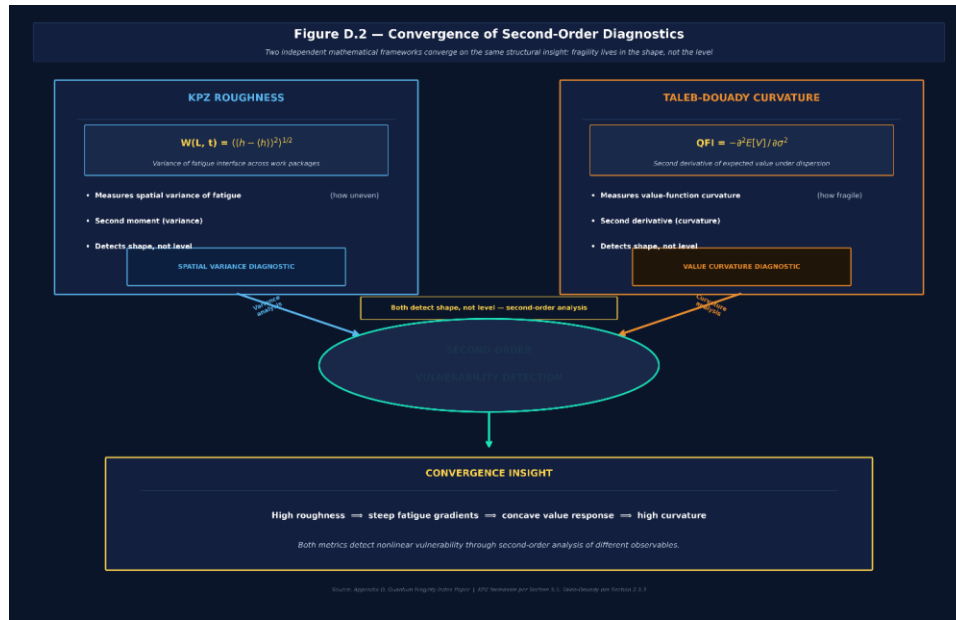


Figure D.2 Convergence of Second-Order Diagnostics.

This diagram shows how KPZ roughness (left branch — measuring spatial variance of the fatigue interface) and Taleb–Douady curvature (right branch — measuring second derivative of expected value) converge on a common target: second-order vulnerability detection. Both metrics detect shape rather than level, and both identify the nonlinear dynamics that first-order metrics systematically miss. The convergence insight at bottom: high roughness implies steep fatigue gradients, which produce concave value response, which manifests as high curvature.

D.5 Where They Diverge — Complementary Strengths

Despite the convergence, the two metrics are not interchangeable. Each has a diagnostic domain where it provides the strongest governance signal:

D.5.1 KPZ Roughness — Operational Early Warning

Roughness excels as a real-time, field-level diagnostic. Its strengths:

Immediacy. Roughness can be computed from current-shift FRI data without any modeling infrastructure. If you have fatigue scores by crew, you have roughness.

Spatial resolution. Roughness identifies exactly *where* the "islands of exhaustion" are forming — which crews, which zones, which interfaces. This enables targeted diffusion interventions (rest mandates for the peaked crews, workload rebalancing at the steepest interfaces).

Gradient detection. Roughness directly reveals the fatigue gradients that drive quadratic contagion. When ∇h approaches $\nabla h_{critical}$, roughness signals that the system is approaching its tipping point — even if the average FRI remains moderate.

Limitation. Roughness alone does not tell governance how much *value* is at risk. A project may have high fatigue roughness but a linear (robust) payoff function — in which case the roughness, while uncomfortable, does not produce fragility in the QFI sense. Roughness measures the condition of the workforce; it does not measure the sensitivity of the project outcome to that condition.

D.5.2 Taleb–Douady Curvature — Strategic Fragility Assessment

Curvature excels as a periodic, portfolio-level diagnostic. Its strengths:

Value integration. Curvature directly measures the impact on expected project value. It translates perturbation sensitivity into the currency that boards care about — dollars, schedule months, scope fidelity.

Multi-source sensitivity. Curvature captures fragility from *all* perturbation sources — not just fatigue. Supply chain disruptions, regulatory changes, stakeholder dynamics, and scope evolution all contribute to the perturbation dispersion σ that the curvature operator evaluates.

Portfolio comparability. Because curvature is normalized to a $[0, 1]$ governance scalar (QFI_{norm}), it enables apples-to-apples comparison across projects of different types, sizes, and delivery models. A QFI of 0.68 means the same thing whether the project is a \$2.6B urban tunnel or a \$1.4B data center.

Limitation. Curvature alone does not tell governance *why* the value function is concave. Is it because the fatigue landscape is rough? Because supply chain concentration is high? Because regulatory exposure is extreme? The composite QFI can be decomposed into dimensional QFI_i values, but the internal dynamics within each dimension — particularly the fatigue dynamics — require roughness-level analysis to diagnose.

D.5.3 Diagnostic Complementarity Table

Governance Question	Best Answered By	Why
"How uneven is fatigue across our site right now?"	KPZ Roughness	Direct measure of spatial variance in FRI; available in real time
"Are we approaching the fatigue tipping point?"	KPZ Roughness	Gradient ∇h relative to $\nabla h_{critical}$ is the tipping-point indicator
"Which crews need immediate relief?"	KPZ Roughness	Identifies the peaks in the fatigue surface

Governance Question	Best Answered By	Why
"How fragile is the overall project to increasing volatility?"	TD Curvature (QFI)	Integrates all perturbation sources into a single value-sensitivity measure
"How does this project compare to others in the portfolio?"	TD Curvature (QFI)	Normalized, dimensionless scalar enables cross-project comparison
"What is the projected value impact of this fragility?"	TD Curvature (QFI)	Curvature quantifies expected value loss per unit increase in dispersion
"Is fatigue the reason our QFI is high?"	Both (cross-validated)	High roughness concurrent with high QFI confirms fatigue as the curvature driver
"Are our interventions reducing fragility?"	Both (tracked over time)	Roughness decline should precede or accompany QFI decline; divergence flags model misalignment

D.6 Worked Numerical Parallel

This section demonstrates the roughness–curvature relationship with three scenarios that hold average FRI constant while varying roughness, showing the corresponding QFI impact.

Setup. A 12-crew construction site with average FRI ≈ 42 across all three scenarios. Payoff function, perturbation distribution, and governance parameters are identical to the Section 2.3.3 worked example.

Scenario 1 — Smooth Interface (Low Roughness)

FRI distribution: all crews between 40 and 44.

Roughness: $W = 2.1$.

Maximum gradient: $\nabla h_{\max} = 4$ (well below $\nabla h_{\text{critical}} \approx 15$).

Governance implication: fatigue is moderate and uniformly distributed. Management diffusion capacity is not stressed.

Composite QFI effect: Γ_{fatigue} contribution is small; composite QFI ≈ 0.32 (Moderately Fragile, lower Amber).

Scenario 2 — Moderate Roughness

FRI distribution: crews range from 28 to 60, with two crews above 55.

Roughness: $W = 8.7$.

Maximum gradient: $\nabla h_{\max} = 12$ (approaching $\nabla h_{\text{critical}}$).

Governance implication: fatigue variance is widening. Two "warm spots" are developing, but the gradients are still manageable. Enhanced monitoring is warranted.

Composite QFI effect: Γ_{fatigue} contributes meaningfully; ρ_{ij} between Schedule and Cost dimensions increases due to error propagation from fatigued crews; composite QFI ≈ 0.55 (Moderately Fragile, upper Amber).

Scenario 3 — Rough Interface (High Roughness)

FRI distribution: crews range from 15 to 75, with four crews above 65 ("islands of exhaustion").

Roughness: $W = 22.4$.

Maximum gradient: $\nabla h_{\text{max}} = 28$ (far exceeds $\nabla h_{\text{critical}} \approx 15$).

Governance implication: the system has crossed the tipping point. Quadratic contagion is self-sustaining. Management diffusion capacity is overwhelmed.

Composite QFI effect: Γ_{fatigue} dominates total decoherence rate; correlation amplification (C term) increases sharply as error propagation cascades across interfaces; composite QFI ≈ 0.78 (Highly Fragile, deep Orange).

Summary Table:

Scenario	Avg FRI	Roughness W	Max Gradient ∇h_{max}	vs. $\nabla h_{\text{critical}}$	Composite QFI	Governance Band
1 (Smooth)	42	2.1	4	Well below	0.32	Amber (Mod. Fragile)
2 (Moderate)	42	8.7	12	Approaching	0.55	Amber (Mod. Fragile)
3 (Rough)	42	22.4	28	Far exceeds	0.78	Orange (Highly Fragile)

The critical governance insight: three projects with the same average fatigue — indistinguishable by first-order metrics — exhibit QFI values spanning the entire range from moderate to high fragility. The difference is entirely attributable to roughness. This is why **roughness, not mean FRI, is the leading indicator of fatigue-driven fragility in the QFI framework.**




Figure D.3 Roughness–QFI Parallel: Same Average, Different Vulnerability
 Panel A shows fatigue roughness W for three scenarios with identical average $FRI = 42$, with the tipping point threshold marked. Panel B shows the corresponding composite QFI values against the governance band scale (Robust/Moderately Fragile/Highly Fragile/Critical). The visual demonstrates that roughness — not average fatigue — drives the QFI outcome.

D.7 Diagnostic Domain Map — Where Each Metric Sees Best

The two metrics occupy different positions on the project's decision cadence spectrum:

Daily / Field Operations → KPZ Roughness dominates. At this cadence, governance needs to know where fatigue peaks are forming, whether gradients are steepening, and which crews require immediate relief. Roughness is computable from current-shift telemetry and provides the spatial resolution that curvature-based QFI cannot.

Weekly / Progress Review → Both metrics contribute. At this cadence, governance compares the week's roughness trend against the latest QFI assessment. If roughness is rising but QFI has not yet responded, this is an early-warning signal — the fatigue surface is developing gradients that will, with a lag, steepen the value-function curvature. This lead-lag relationship enables anticipatory intervention.

Monthly / Steering Committee → Cross-validation zone. At this cadence, both metrics should be presented side by side. Convergence (high roughness + high QFI) confirms that fatigue is a primary fragility driver and that the roughness-curvature causal chain is active. Divergence (high roughness but stable QFI, or vice versa) signals either that the payoff function is more robust to fatigue than modeled (a positive surprise) or that non-fatigue perturbation sources are driving curvature independently (a different intervention pathway is needed).

Quarterly / Scenario Stress Test → TD Curvature dominates. At this cadence, governance runs full Monte Carlo perturbation sweeps across the payoff function, evaluating QFI under alternative

dispersion scenarios. This is the deep-dive curvature analysis that provides portfolio-level fragility comparisons and intervention prioritization.

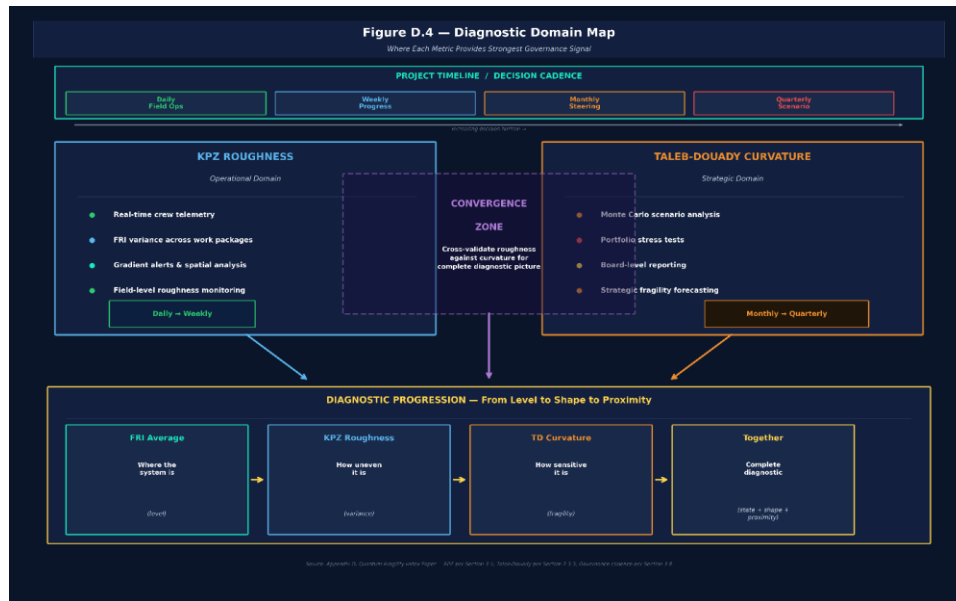


Figure D.4 — Diagnostic Domain Map

This figure shows the project timeline and decision cadence from daily field operations (left) through quarterly scenario analysis (right). KPZ roughness dominates the operational (high-frequency) domain; Taleb–Douady curvature dominates the strategic (lower-frequency) domain. In the middle, a convergence zone enables cross-validation. The bottom row summarizes the three-part diagnostic: FRI average (where the system is), roughness (how uneven it is), curvature (how sensitive it is) — and together, the complete diagnostic.

D.8 Integration Protocol — Using Both Metrics in QPM Governance

D.8.1 Operational Workflow

The following protocol integrates roughness and curvature into the QPM governance cycle:

Step 1 — Continuous monitoring (daily). Compute FRI by crew/zone from telemetry (Tier 1–3 data streams as defined in Section 3.2.1). Compute roughness W and maximum gradient ∇h_{\max} . Compare ∇h_{\max} against $\nabla h_{\text{critical}}$. If $\nabla h_{\max} > 0.8 \times \nabla h_{\text{critical}}$, flag "approaching tipping point" and trigger field-level diffusion intervention.

Step 2 — Weekly assessment. Compute QFI_F (fatigue-specific fragility) from proximity formula. Compare roughness trend to QFI_F trend. If roughness is rising faster than QFI_F , anticipate a QFI_F escalation in the next reporting cycle and pre-position interventions.

Step 3 — Monthly cross-validation. Present roughness and QFI side by side to the steering committee. If both are elevated and trending in the same direction, confirm fatigue as a primary fragility driver and activate the three-pathway transmission analysis (Section 3.3.3). If they diverge, investigate the cause: is the payoff function more linear than assumed (model recalibration needed)? Are non-fatigue perturbation sources dominating curvature (different intervention pathway)?

Step 4 — Quarterly scenario analysis. Run full Taleb–Douady Monte Carlo perturbation sweeps with fatigue-modulated decoherence rate $\Gamma_{\text{effective}}(\text{FRI})$. Use current roughness profile to set the fatigue component of the perturbation distribution. Report portfolio-level QFI with fatigue contribution decomposed.

D.8.2 Escalation Triggers

Condition	Meaning	Escalation
$W > W_{\text{critical}}$ AND $QFI > 0.60$	Roughness and curvature converge — confirmed fatigue-driven fragility	Orange governance; structural fatigue intervention required
$W > W_{\text{critical}}$ AND $QFI < 0.60$	Roughness is high but value function absorbs it — robust payoff	Amber governance; monitor closely; roughness may be a leading indicator of curvature rise
$W < W_{\text{critical}}$ AND $QFI > 0.60$	Curvature is high from non-fatigue drivers — roughness is not the cause	Orange governance; investigate supply chain, regulatory, or scope drivers
$W < W_{\text{critical}}$ AND $QFI < 0.60$	Both metrics low — system in stable operating range	Blue/Amber governance; standard monitoring

D.8.3 Reporting Format

In board-level reporting, present the two metrics in a paired format:

Left column: Fatigue Roughness Panel

- Current roughness W with trend arrow (\uparrow , \rightarrow , \downarrow)
- Maximum gradient ∇h_{max} as % of $\nabla h_{\text{critical}}$
- Heat map of crew/zone FRI scores showing peaks and valleys

Right column: QFI Curvature Panel

- Composite QFI_{norm} with governance band color
- Dimensional radar showing QFI breakdown

- Fatigue contribution to composite (from three-pathway decomposition)

Bottom row: Convergence Assessment

- Are roughness and curvature aligned? If so, confidence in diagnosis is high.
- Are they divergent? If so, note the direction and investigate.
- What is the lead-lag relationship? Roughness typically leads curvature by 1–2 reporting cycles; if curvature is leading roughness, the fragility driver may be external.

D.9 Limitations and Caveats

Roughness is a spatial measure; curvature is a stochastic measure. Roughness measures variance across work packages at a point in time. Curvature measures sensitivity of expected value to dispersion changes over a distribution of scenarios. They are not mathematically equivalent — they are structurally parallel. The convergence described in this appendix is a practical governance observation, not a formal mathematical identity.

The causal chain (roughness → gradients → contagion → decoherence → curvature) is directional. High roughness implies high curvature under the assumption that fatigue gradients translate into decoherence via the three transmission pathways (Section 3.3.3). If the project's payoff function is genuinely linear ($\beta_2 = 1$) or if fatigue does not translate into decoherence (low γ_0), the causal chain weakens and the metrics may diverge even when roughness is high.

Calibration matters. The tipping-point threshold $V_{h_{critical}}$ depends on the ratio v/λ , which must be estimated from project-specific data. Similarly, the QFI governance thresholds (0.30, 0.60, 0.85) are calibrated from portfolio experience. Miscalibration of either metric will produce false signals — false alarms (overestimating fragility) or missed warnings (underestimating fragility). Annual recalibration against portfolio outcomes is recommended.

Temporal alignment. Roughness is a real-time metric; curvature is computed periodically. When comparing the two, ensure that the QFI assessment uses fatigue data contemporaneous with the roughness measurement. Lagged comparisons may create artificial divergence.

D.10 Summary

KPZ roughness and Taleb–Douady curvature are two lenses on the same underlying phenomenon: the nonlinear vulnerability of complex project systems to small perturbations. Roughness reveals this vulnerability through the topology of the fatigue interface — the spatial unevenness that creates steep risk gradients and triggers quadratic contagion. Curvature reveals it through the shape of the value function — the concavity that converts small increases in perturbation dispersion into disproportionate value loss.

Together, they provide the most complete fragility diagnostic available in the QPM framework:

- **FRI average** tells governance where the system is (level).

- **KPZ roughness** tells governance how uneven the system is (variance).
- **Taleb–Douady curvature** tells governance how sensitive the system is (fragility).
- **All three together** provide the complete QPM fragility picture: state, shape, and proximity to failure.

The practical recommendation: deploy roughness for daily operational monitoring and early-warning detection; deploy curvature for periodic strategic assessment and portfolio comparison; cross-validate monthly to confirm diagnostic alignment and detect emerging divergence. When roughness and curvature converge on high fragility, governance can act with high confidence. When they diverge, the divergence itself is diagnostic — it reveals whether the vulnerability is internally driven (fatigue) or externally driven (supply chain, regulatory, scope), enabling precisely targeted intervention.

Appendix E

Original Ideas and Concepts

1. Fragility as a Missing Dimension in Project Governance

The paper establishes fragility as a construct fundamentally distinct from risk. While risk quantifies expected loss, fragility captures how easily small, routine perturbations can push a project across a governance boundary. This distinction forms the conceptual basis for the Quantum Fragility Index (QFI), emphasizing that traditional risk metrics overlook susceptibility to state transitions.

2. Formal Definition of Quantum Fragility

Quantum Fragility is defined as the minimum normalized perturbation required to induce a transition across a governance boundary. This definition reframes fragility as a measure of state-transition susceptibility—identifying tipping-point proximity, normalizing disturbances for comparability, and revealing the weakest direction of vulnerability.

3. Mapping Fragility to Quantum Constructs

The paper shows that fragility is already embedded within quantum analogues: decoherence (collapse of coherence), entanglement (amplified coupling), superposition (collapse into undesirable states), and propagation (spread of local variation). This mapping unifies quantum behavior with a measurable governance metric, capturing how easily coherent operational states degrade under noise.

4. QFI as a Single Scalar Metric

QFI is introduced as a computable scalar integrating curvature, decoherence, dimensional fragilities, cross-dimensional coupling, and governance thresholds. This produces the first board-readable fragility metric capable of summarizing nonlinear vulnerability in a single value.

5. Curvature as the Mathematical Core

The paper operationalizes Taleb–Douady fragility by defining fragility as second-order sensitivity to dispersion:

$$[\text{QFI} = -\frac{\partial^2 E[V]}{\partial \sigma^2} .]$$

This curvature-based framing distinguishes concave from convex response geometry, enabling detection of nonlinear vulnerability and supporting governance thresholds grounded in system dynamics.

6. Decoherence-Mapped QFI Proxy

A computational shortcut is introduced using decoherence as a real-time proxy for fragility:

$$[QFI_{\text{deco}} = \eta, \Gamma, (1 - F_Q(t))^{\alpha - 1}.]$$

This proxy links coherence loss to fragility, enabling telemetry-driven monitoring and providing a faster alternative to curvature-based Monte Carlo sweeps.

7. Six-Dimension Fragility Decomposition

The paper defines a six-dimension architecture—Schedule, Cost, Scope, Stakeholder, Supply Chain, and Regulatory—each with its own (QFI_i), weight, and fidelity. This decomposition makes fragility diagnosable, supports targeted interventions, and enables cross-dimensional analysis.

8. Correlation Correction (Entanglement Amplification)

A correlation correction term captures nonlinear amplification when dimensions are entangled:

$$[C = \sum_{i \neq j} \rho_{ij} \sqrt{w_i w_j}, \phi(QFI_i, QFI_j).]$$

This is the first fragility metric to explicitly model coupled failure pathways, reflecting how interdependencies magnify vulnerability.

9. Governance Thresholds Based on System Dynamics

Four governance boundaries—0, 0.30, 0.60, 0.85—are justified through system-behavior transitions: convexity shifts, onset of nonlinear amplification, decoherence half-life constraints, and entry into tail-risk regimes. These thresholds create a dynamic, behavior-based governance structure.

10. Dimensional Floor Rule

The paper introduces a safeguard preventing composite scores from masking extreme fragility: if any dimension reaches (QFI_i ≥ 0.85), governance escalates automatically. This ensures that catastrophic single-dimension vulnerabilities cannot be averaged away.

11. Integration with Safety-II

QFI is positioned as the first quantitative Safety-II metric, measuring proximity to failure, coherence loss, and variability near tipping points. This operationalizes Safety-II's anticipatory resilience model by providing a measurable indicator of system readiness and stability.

12. Fatigue as an Internality

Fatigue is introduced as a cross-cutting internality that increases decoherence rate, nonlinearity (η), and cross-dimensional coupling (ρ_{ij}). This is the first formal integration of human readiness into a fragility metric, showing how internal conditions reshape systemic vulnerability.

13. Roughness–Curvature Dual Diagnostic

The paper links fatigue-induced roughness (via KPZ variance) with value-function curvature (QFI). Roughness detects spatial instability, while curvature detects systemic sensitivity. Together, they form a dual diagnostic framework for identifying second-order vulnerabilities.

14. Unified QFI–FRI Governance Dashboard

A unified executive dashboard is proposed, integrating QFI, FRI, roughness, proximity, dimensional fragility, and an intervention simulator. This creates the first combined fragility–fatigue governance interface for real-time decision support.

15. Three-Pathway Transmission Model (Fatigue → QFI)

The paper identifies three pathways through which fatigue increases fragility: direct dimensional impact, decoherence amplification, and correlation amplification. This model explains how internal human conditions propagate into systemic governance vulnerability.

16. Quantum Tunneling Analogy for Control Failure

A novel analogy reframes control failure: fatigue “thins” governance barriers, allowing perturbations to tunnel through controls rather than breach them. This shifts the interpretation of failure from structural breakdown to barrier permeability.

17. QFI as a Governance Operating System

The paper positions QFI as the backbone of a new governance operating system—enabling continuous monitoring, threshold-based escalation, intervention prioritization, portfolio comparability, and predictive crossing-time estimation. This represents a shift from event-driven risk management to state-driven fragility governance.

About the Author



Bob Prieto

Chairman & CEO
Strategic Program Management LLC
Jupiter, Florida, USA



Bob Prieto is Chairman & CEO of Strategic Program Management LLC focused on strengthening engineering and construction organizations and improving capital efficiency in large capital construction programs. Previously, Bob was a senior vice president of Fluor, focused on the development, delivery, and turnaround of large, complex projects worldwide across all of the firm's business lines; and Chairman of Parsons Brinckerhoff, where he led growth initiatives throughout his career with the firm.

Bob's board level experience includes Parsons Brinckerhoff (Chairman); Cardno (ASX listed; non-executive director); Mott MacDonald (Independent Member of the Shareholders Committee); and Dar al Riyadh Group (current)

Bob consults with owners of large, complex capital asset programs in the development of programmatic delivery strategies encompassing planning, engineering, procurement, construction, financing, and enterprise asset management. He has assisted engineering and construction organizations to improve their strategy and execution and has served as an executive coach to a new CEO. He is author of eleven books, over 1000 papers and National Academy of Construction Executive Insights, and an inventor on 4 issued patents.

Bob's industry involvement includes the National Academy of Construction and Fellow of the Construction Management Association of America (CMAA). He serves on the New York University Tandon School of Engineering Department of Civil and Urban Engineering Advisory Board and New York University Abu Dhabi Engineering Academic Advisory Council and previously served as a trustee of Polytechnic University. He has served on the Millennium Challenge Corporation Advisory Board and ASCE Industry Leaders Council. He received the ASCE Outstanding Projects and Leaders (OPAL) award in Management (2024). He was appointed as an honorary global advisor for the PM World Journal and Library.

Bob served until 2006 as one of three U.S. presidential appointees to the Asia Pacific Economic Cooperation (APEC) Business Advisory Council (ABAC). He chaired the World Economic Forum's Engineering & Construction Governors and co-chaired the infrastructure task force in New York after 9/11. He can be contacted at rpstrategic@comcast.net.