Metrics for Assumption Management in Large Complex Projects[1]

Bob Prieto

Abstract

Assumptions underpin planning and execution in engineering and construction projects yet migrate over time, compounding uncertainty and driving cost, schedule, and performance deviation. This paper develops a rigorous **Assumption Governance Index (AGI)**[2] and category-level indices that quantify assumption migration, incorporate consequence weighting and confidence decay[3], and support statistical monitoring of mean and dispersion dynamics. The methodology embeds principles from Quantum Project Management[4] (QPM) to address entanglement, emergence, and measurement feedback common to large complex projects. Validation approaches, implementation architecture, and governance decision rules are described. The AGI supports timely rebaselining, targeted mitigations, and adaptive governance for projects where traditional decomposition and deterministic controls are insufficient[5].

The paper also develops a companion **Assumption Diffusion Index (ADI)**[6] that measures how widely and how quickly changes to an assumption's state (migration, rebaseline, or reprioritization) propagate across the assumption register[7] and through entangled clusters. Where AGI summarizes the magnitude and weighted impact of assumptions at a point in time, ADI captures the spread, velocity, and reach of those changes through the system.

---

[1] How to cite this paper: Prieto, R. (2025). Metrics for Assumption Management in Large Complex Projects, *PM World Journal,* Vol. XIV, Issue XII, December.

[2] The AGI is a normalized, consequence-weighted index summarizing per-assumption migration, time-aware confidence, and amplification adjustments to produce a governance-grade KPI; it is intended as a lead indicator for rebaseline and contingency decisions. See Section 1.4 and Appendix B for the formal formula and worked example.

[3] Confidence decay, Ci(t) is the time- and event-modulated confidence for assumption i (default exponential decay) used to increase sensitivity to older, less-tested assumptions. It changes how much a given migration contributes to AGI; see Section 1.3.3 for the decay formula and event amplification notes.

[4] Prieto, R. "Quantum Project Management: A monograph on a new theory for management of large complex projects"; ISBN: 978-1-304-08165-0, December 2024

[5] Prieto, R. "Quantum Project Management." PM World Journal feature paper, January 2024

[6] The ADI quantifies how an assumption change propagates through the entanglement network by reporting footprint, velocity, and reach over a chosen horizon; ADI complements AGI by showing propagation potential rather than instantaneous magnitude. See Section 2.0 and Appendix D (dashboard footprints) for computation and use cases

[7] Even the existence of a comprehensive written Assumption Register is a rarity in too many LCPs

Introduction

Large complex projects (LCPs) in engineering and construction depend on a network of explicit and tacit assumptions spanning technical, environmental, stakeholder, economic, client, and productivity dimensions. These assumptions represent the team's best assessment at a given moment but are inherently uncertain and subject to migration as external and internal conditions evolve. Systematic tracking of assumptions in general, and assumption migration in particular, is essential for governance, rebaselining, contingency allocation, and strategic decision-making. This paper introduces a singular Assumption Governance Index (AGI) plus category-specific indices and statistical measures to monitor migration dynamics and support governance actions in the context of QPM's entanglement-aware view of LCPs[8].

Similarly, it introduces an Assumption Diffusion Index (ADI) that provides significant and important insights into:

- **Propagation dynamics** — shows whether an assumption change is isolated or rapidly influences many others (contagion vs local event).
- **Velocity** — quantifies how fast migrations travel through clusters
- **Network reach** — identifies central assumptions that act as diffusion hubs even if their instantaneous migration is small.
- **Fragility vs concentration** — discriminates between portfolios that have similar AGI levels but very different systemic fragility (one dominated by a few high impact change items versus one where small changes cascade).
- **Early-warning for systemic escalation** — high ADI for a small migration signals risk of future AGI jumps; AGI itself can lag because it is a snapshot of current weighted contributions.

Assumption Migration

Assumption migration is a persistent driver of program instability: small, untracked changes in many assumptions can compound into large, unexpected impacts on cost, schedule, quality, and strategic value. Practitioners (owners, project managers, designers, estimators, etc.) frequently document assumptions in registers, yet few instruments convert these discrete records into a single situational awareness metric usable by executives and governance bodies, including Boards of Directors. Traditional risk and cost uncertainty frameworks provide strong foundations for estimating discrete risk events and cost distributions, but they do not directly quantify the migration of the foundational assumptions that form the baseline model of the project. There is a need for

[8] Prieto, B. "Management of Assumption Infatuation in Large Complex Projects." PM World Journal archive, April 2016. https://pmworldlibrary.net/wp-content/uploads/2016/04/pmwj45-Apr2016-Prieto-Management-of-Assumption-Infatuation-in-Large-Complex-Projects.pdf

a robust governance metric that is normalized, traceable, entanglement-aware, and actionable.

Related Considerations & Metrics

Assumption migration and management has been considered in several related and relevant contexts, including:

- **Assumption management and tracking literature** - Assumption registers and assumption management cycles have been advocated in practitioner literature and guidance as core elements of robust project controls and risk management[9]. Frameworks describe identification, documentation, validation, and periodic re-evaluation, but rarely produce a composite governance index that scales to portfolio oversight and executive decision-making[10].

- **Risk, uncertainty, and cost-estimating traditions** - Quantitative risk analysis, cost uncertainty, and S-curve-based confidence methods provide rigorous approaches for understanding cost and schedule uncertainty and for setting reserves and UFE (Unallocated Future Expense)[11]. These techniques inform AGI design by providing robust methods for propagating uncertainty and calibrating statistical thresholds for governance.

- **Project governance metrics and KPIs** - Project governance research stresses clarity of measurement intent, the risk of perverse incentives from poor metrics, and the importance of combining internal operational metrics with external governance KPIs for trend interpretation[12]. AGI is designed as a governance KPI intended to align with these principles - transparent, traceable, and behaviorally sensible. ADI provides added insights into AGI and key areas of concern and future vulnerability.

- **Quantum Project Management and entanglement considerations** - Quantum Project Management (QPM)[13] recognizes LCPs as dynamic open systems exhibiting entanglement, measurement feedback, and emergent behavior that invalidate assumptions of linear decomposition and independence[14]. QPM's theoretical principles—embracing uncertainty, probabilistic outcomes, and

[9] Making Assumptions Work for Your Project: A Structured Approach for Project Controls; www.logikheim.com.au
[10] Project Assumptions: 30 Examples and How to Manage Them; www.projectpractical.com
[11] Washington State Department of Transportation Project Risk Analysis Model; wsdot.wa.gov
[12] Project Governance: The Art of Measuring Effectively; blueskiesconsulting.com
[13] Quantum Project Management A monograph on a new theory for management of large complex projects (2024); ISBN 978-1-304-08165-0
[14] Quantum Project Management – Lulu; www.lulu.com

systemic entanglement—shape AGI's treatment of correlated assumption migration, cluster amplification, and measurement effects and are an element of supporting ADI metrics..

This paper synthesizes assumption management best practices, quantitative risk and uncertainty methodology, governance measurement design, and QPM theory into an operational AGI with category subindices, statistical trend measures, entanglement diagnostics, and governance actions mapped to thresholds. Validation pathways include Monte Carlo and retrospective case analysis. *(See Appendix C – Monte Carlo Validation and Calibration Protocol)*

References and source materials include contemporary governance guidance, NASA and Washington State Department of Transportation (WSDOT) cost-risk handbooks, practitioner monographs, and seminal practitioner articles on assumption management and QPM.

The developed AGI is then complemented by an ADI that further enhances and interprets the AGI metric. Each will be considered in turn in this paper. A table of contents follows to guide you through the balance of the paper.

The inclusion of extensive mathematical formulation into both the main body of the paper as well as the various appendices was done to ensure rigorous clarity of concepts and relationships.

Table of Contents

Metrics for Assumption Management in Large Complex Projects

- Appendix F: Assumption Mapping and Scores (n = 8); AGI Signal Ranking and Escalation Triggers
- Appendix G: Reproducibility Checklist

1.0 Assumption Governance Index (AGI)

1.1 Conceptual model and principles

The developed AGI considers design principles and high level notation and register requirements. Specifically:

- **Design principles**, including:
  - **Traceability:** each AGI component must link to a transparent evidence trail in the assumption register.
  - **Materiality-normalized measurement:** per-assumption migration is normalized by a project-defined critical tolerance (Δcrit)[15] to reflect material governance impact rather than raw absolute change.
  - **Consequence weighting:** migration of high-consequence assumptions must drive AGI proportionally and visibility.
  - **Time-aware confidence:** baseline confidence decays with time and event exposure; sensitivity to migration increases as confidence decays.

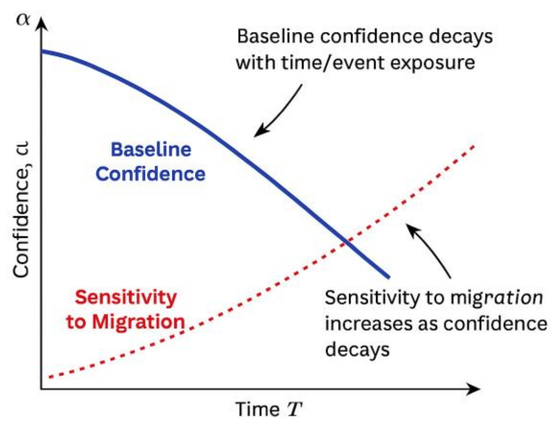

  - **Entanglement-aware aggregation:** correlation among assumption migrations must be measured and treated to capture emergent system-level risk consistent with QPM.
  - **Governance-action mapping:** AGI bands and per-assumption thresholds map to explicit governance actions and workflows.

1.2 High-level notation and register requirements

- Let assumptions be indexed by i = 1..N.

[15] Δcrit is the project-defined migration magnitude beyond which a change is considered materially significant for governance. It is used to normalize per-assumption migration scores so that AGI reflects governance materiality rather than raw unit movement; see Appendix A (Assumption Register schema) for calibration guidance.

- Each register entry includes:
  - unique ID
  - category[16] (k)
  - baseline value V0, i (numeric or categorical)
  - baseline timestamp t0,i
  - current assessed value Vt,i
  - critical tolerance Δcrit,i
  - baseline confidence α0,i
  - volatility parameter λi
  - consequence component scores[17]:
    - ω_cost,i
    - ω_safety,i
    - ω_schedule,i
    - ω_reputation,i
  - owner of the assumption
  - last_update
  - evidence link

Qualitative assumptions are mapped to ordinal scales with a defined distance metric.

1.3 AGI methodology

1.3.1 Per-assumption migration metric (Mi) – this metric would use one of the two following methodologies depending whether the metric is numeric or categorical.

- **Numeric assumptions**[18]**:**

$$M_i(t) = \min\left(1,\ \frac{|V_{t,i} - V_{0,i}|}{\Delta_{crit,i}}\right)$$

- **Categorical/ordinal assumptions:** map categories to positions on a domain-specific ordinal scale and define a normalized distance d(cat0, catt) in [0,1]; then set Mi(t) = d(cat0, catt). The design ensures Mi ∈ [0,1], where 0 indicates no migration and 1 indicates migration beyond the critical tolerance[19].

---

[16] These categories help define the governance domain and are often mapped to escalation protocols, training modules, and entanglement clusters. Category (k) helps identify which domains are tightly coupled. Categories might include economic cost, performance, technical, environmental, client and stakeholder, for example.
[17] Other consequence components may be considered as appropriate.
[18] Metric movements less than the critical tolerance for that metric are scored on that basis while those above the critical threshold are scored at 1, the maximum value for the migration metric.
[19] ∈ means belongs to the set

1.3.2 Consequence weight (Wi):

Compute raw consequence score:

$$W_i = \omega_{cost,i} + \omega_{safety,i} + \omega_{schedule,i} + \omega_{reputation,i}$$

with each component[20] scored against documented criteria (e.g., 0–10).

Normalize weights so sum(Wi) = 1:

$$W_i = \frac{\omega_{cost,i} + \omega_{safety,i} + \omega_{schedule,i} + \omega_{reputation,i}}{\sum_{j=1}^{N} (\omega_{cost,j} + \omega_{safety,j} + \omega_{schedule,j} + \omega_{reputation,j})}$$

or

$$W_i = \omega_i / \Sigma_j \, \omega_j.$$

Periodic peer review of ω components is mandatory to limit subjectivity and gaming.

1.3.3 Confidence and time-decay factor (Ci(t))

Baseline confidence $\alpha_{0,i} \in (0,1]$. Confidence decays exponentially by default:

$$C_i(t) = \alpha_{0,i} \cdot \exp(-\lambda_i \cdot (t - t_{0,i})) \cdot \varphi_i(t)$$

where λ_i is an assumption-specific decay rate reflecting domain volatility, and $\varphi_i(t)$ is an event-driven amplification multiplier ≥ 1 that increases measurement sensitivity after disruptive events (e.g., permit denial, market shock). Cap $C_i(t)$ at 1 for normalization.

1.4 Aggregation and indices

1.4.1 Assumption Governance Index (AGI)

$$AGI(t) = 100 \cdot \Sigma_i \, W_i \cdot C_i(t) \cdot M_i(t)$$

Lower values indicate limited migration among low-consequence assumptions or high retained confidence; higher values indicate substantive migration among consequential assumptions and/or low confidence.

[20] Additional or alternate consequence components may be used as appropriate.

1.4.2 Category Governance Indices (CGI_k)

Group assumptions by category k (Technical, Environmental, Stakeholder, Economic/Cost, Client, Performance/Productivity).

Compute:

$$\mathrm{CGI}_k(t) = 100 \cdot \frac{\sum_{i \in k} W_i \cdot C_i(t) \cdot M_i(t)}{\sum_{i \in k} W_i}$$

or

CGI_k(t) = 100 · Σ*{i∈k} Wi · Ci(t) · Mi(t) /* Σ{i∈k} Wi

CGI_k isolates domain exposure (all assumptions *i* that belong to category *k*) and supports targeted actions.

1.5 Statistical monitoring metrics

The first four metrics are essential for tracking assumption migration and signal volatility in governance monitoring.

- Unweighted mean migration - The simple average of all migration values across assumptions, without applying any importance or confidence weights. Gives a raw sense of how much assumptions are shifting overall, regardless of their governance impact. This is the most basic measure of assumption migration.

$$\mu_M(t) = \frac{1}{N} \sum_{i=1}^{N} M_i(t)$$

- Weighted mean migration - The average migration value where each assumption is weighted by a factor (e.g., confidence level α_i, governance weight ω_i, or footprint ADI_i). Prioritizes shifts in high-impact or high-confidence assumptions, making it more relevant for governance escalation.

$$\mu_{wM}(t) = \sum_{i=1}^{N} W_i \cdot M_i(t)$$

- Unweighted standard deviation - Measures the spread or variability of migration values across assumptions, without weighting. Indicates how volatile or dispersed the assumption shifts are, useful for detecting systemic instability

$$\sigma_M(t) = \sqrt{\frac{1}{N}\sum_{i=1}^{N} \left(M_i(t) - \mu_M(t)\right)^2}$$

- Weighted standard deviation - Measures variability while accounting for the importance or confidence of each assumption. Highlights volatility in high-impact assumptions, helping governance teams focus on meaningful signal noise.

$$\sigma_{wM}(t) = \sqrt{\sum_{i=1}^{N} W_i \cdot \left(M_i(t) - \mu_{wM}(t)\right)^2}$$

- Higher moments: skewness and kurtosis[21] of {Mi(t)} to detect tail risk and concentration (e.g., single assumptions driving destabilization).
- Change metrics over Δt:
  - Δμ = μ_M(t1) − μ_M(t0);
  - Δσ = σ_M(t1) − σ_M(t0).
  - Use bootstrapped[22] confidence intervals for μ and σ to assess statistical significance of observed changes and reduce false alarms.

Governance implications:

- High weighted mean migration → Escalation trigger
- High weighted standard deviation → Signal instability or entanglement volatility

[21] See Appendix D for description of skewness and kurtosis
[22] Bootstrapping is a non-parametric method that builds these intervals by repeatedly resampling the data (with replacement) and recalculating the statistic (μ or σ) each time. This creates a distribution of possible values for μ or σ, from which we can extract the confidence interval — typically the 2.5th and 97.5th percentiles for a 95% interval.

- Low unweighted mean but high weighted mean → Quiet shifts in critical assumptions

1.6 Entanglement and Cluster Diagnostics (Quantum Project Management (QPM) Integration)

To surface latent interdependencies and nonlinear escalation triggers, we apply a structured diagnostic framework grounded in Quantum Project Management (QPM). This approach treats assumptions not as isolated inputs but as entangled elements within a dynamic governance system. By analyzing co-migration patterns *(see box below)* and signal amplification, we identify clusters that drive variance in the Assumption Governance Index (AGI) and require adaptive response.

Co-migration patterns refer to the coordinated or correlated movement of multiple groups that migrate together or in similar directions over time. Recognizing co-migration helps in forecasting future movements.
When tracking **assumption migration**, especially in large-scale programs or digital transformations, you're not just moving one assumption at a time. Instead, **sets of assumptions often shift together**, either because:

- They're **logically linked** (e.g., cost assumptions tied to schedule assumptions)
- They're **embedded in the same model or system**
- They're **triggered by the same external driver** (e.g., regulatory change, technology upgrade)

Co-migration patterns help analysts understand how these assumptions evolve in tandem, which is critical for:

- **Risk tracking**: If one assumption fails, its co-migrating counterparts may also be invalidated.
- **Governance**: Ensures that updates to assumptions are traceable and auditable across domains.
- **Scenario planning**: Helps simulate cascading effects when multiple assumptions shift together.

- **Correlation-Based Entanglement Mapping** We begin by constructing a correlation matrix R, where each entry ρ_{ij}represents the Pearson correlation between the time series of assumptions $M_i$and M_j. High absolute values of ρ_{ij}(i.e., $|\rho_{ij}|$) indicate strong entanglement—suggesting that the assumptions co-migrate due to shared drivers, feedback loops, or systemic exposure. These entangled pairs form the basis for cluster identification.

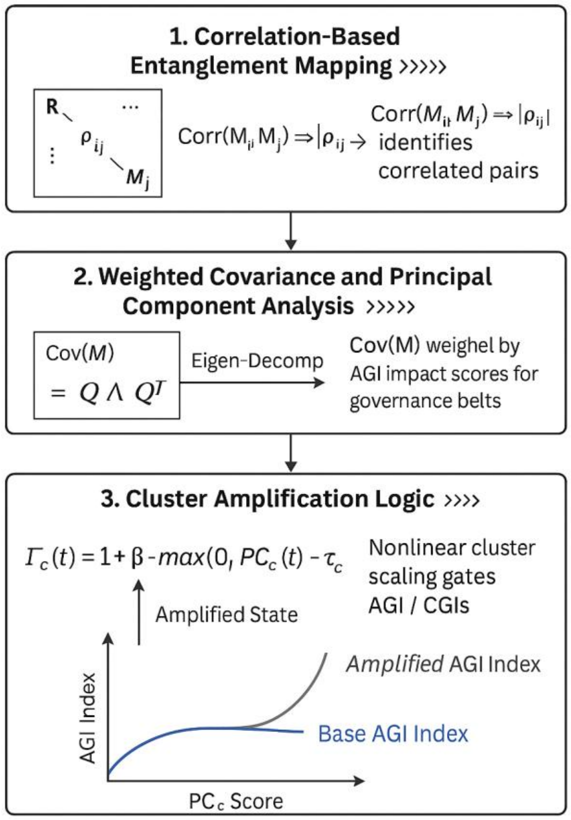

Entanglement and Cluster Diagnostics in QPM

- **Weighted Covariance and Principal Component Extraction** To prioritize governance-relevant signals, we compute a weighted covariance matrix Σ_w:

$$\Sigma_w = \text{diag}(W) \cdot \text{Cov}(M) \cdot \text{diag}(W)$$

Here, $W$is a vector of governance weights (e.g., derived from AGI scores or domain-specific impact factors), and Cov$(M)$is the standard covariance matrix of assumption time series. This weighting ensures that assumptions with greater governance significance exert proportionally more influence on cluster formation.

We then perform eigen-decomposition:

$$\Sigma_w = Q\Lambda Q^T$$

This yields a set of principal components (PCs)[23]—orthogonal vectors that capture dominant modes of co-migration. Each PC represents a latent cluster of assumptions whose synchronized behavior contributes significantly to AGI variance.

- **Cluster Amplification and Nonlinear Signal Adjustment** To detect emergent nonlinear behavior, we define a cluster amplification function $\Gamma_c(t)$for each principal component c:

$$\Gamma_c(t) = 1 + \beta \cdot \max{(0, PC_c(t) - \tau_c)}$$

Where:

- $PC_c(t)$is the score of principal component c at time t
- $\tau_c$is a sensitivity threshold calibrated to detect meaningful co-migration
- $\beta$is an amplification factor tuned via validation

When $PC_c(t)$exceeds its threshold τ_c, the cluster enters an amplified state, indicating emergent risk or systemic volatility. In response, the AGI signal—or specific component governance indices (CGIs)—are multiplied by $\Gamma_c(t)$to reflect nonlinear escalation consistent with QPM logic.

This mechanism ensures that governance responses are dynamically scaled in proportion to entangled cluster behavior, enabling proactive intervention and adaptive signal calibration.

1.7 Measurement feedback and second-order migration

QPM highlights that measurement can alter system state. This is a fundamental property of all quantum systems.

Capture measurement-feedback by tracking secondary migration events following updates. This is aided by a system of systems perspective. Implement an observation-amplification factor in φi(t)[24] when rapid sequential updates correlate with stakeholder behavior changes and log these dynamics for governance review.

[23] Principal components (PCs) are new, uncorrelated variables formed as linear combinations of original variables in a dataset, designed to capture the maximum variance possible. Principal components arise from Principal Component Analysis (PCA), a technique used to reduce the dimensionality of data while preserving as much variability as possible.
[24] As part of the event amplification factor.

1.8 Governance thresholds and operational rules

AGI interpretation bands are defined below. These may be modified by organizational risk tolerance and project criticality if appropriate[25].

- Green: 0–20 — Monitoring; no immediate action required.
- Yellow: 20–50 — Management review and targeted mitigation plans.
- Amber: 50–75 — Program-level review, re-evaluate contingencies, and consider rebaseline.
- Red: 75–100 — Executive steering committee required; consider contract, scope, or schedule action.

Per-assumption action thresholds

- $M_i \geq 0.25$: Owner notification; update mitigation register and evidence.
- $M_i \geq 0.50$: Mandatory impact assessment; quantify changes in cost and schedule and propose mitigation with budget implications.
- $M_i \geq 0.80$: Sponsor-level escalation; prepare procurement/contract remedies and potential re-scope actions.

Systemic triggers

- $\Delta\mu$ over reporting period > $\Delta\mu\_crit$ triggers governance review.
- $\Delta\sigma$ rising above $\Delta\sigma\_crit$ suggests growing heterogeneity; triggers cluster analysis.
- Top-3 contributors to AGI > 50% triggers focused mitigation on those assumptions.

Reweighting and recalibration cadence

- Monthly weight (W_i) revalidation for active projects or immediate reweighting after major events. All reweighting must be documented and peer-reviewed.

Auditable decision rules - All governance actions initiated by AGI or CGI triggers must be logged with timestamped rationales, evidence, corrective measures, owner assignment, and acceptance criteria for closure.

1.9 Implementation architecture and data integrity

Assumption register schema (recommended fields are shown in Table 1).

[25] A later example in this paper used a score of 40 for Amber.

Table 1 Assumption Register Schema
ID
Category
Baseline Value
Baseline Timestamp
Current Value
Units
$\Delta crit$
$\alpha 0$
λ
ω_cost
ω_safety
ω_schedule
ω_reputation
Owner
Last Update
Evidence Link
Update Notes

Appendix B provides a worked example of AGI computation in practice,

Computation pipeline approach:

- ETL[26] processes standardize units, normalize scales, and validate consistency.
- Deterministic computation scripts (R[27] or Python) run AGI and CGI calculations and archive inputs and outputs for auditability.
- Dashboard presents AGI, CGIs, μ_wM, σ_wM, entanglement heatmap (correlation matrix), and top contributors with drill-through to evidence.
- Role-based access controls limit data editing; edits require justification and approver signature logged in the register.

Data quality controls and governance of the metric

- Mandatory evidence attachments for any $M_i > 0.25$.
- Outlier detection and peer adjudication for $\Delta crit$ and ω component changes.

[26] ETL stands for Extract, Transform, Load. It involves extracting data from various sources, transforming it into a suitable format, and loading it into a target data warehouse or data repository. ETL processes offer several benefits, such as improved data quality, better data integration, increased data security, and enhanced scalability.
[27] **R** is an open-source programming language widely used by statisticians, data scientists, and researchers.

- Periodic data audits to ensure registration completeness and unit consistency.

1.10 Validation strategy

Three validation approaches are considered. These include:

- Synthetic validation via Monte Carlo experiments
- Retrospective empirical validation
- Case study protocol (recommended)

Each is discussed in turn.

1.10.1 Synthetic validation via Monte Carlo experiments

Construct ensembles where assumption baseline values and migration processes are simulated under configurable volatility, correlation structures, and event distributions.

Steps:

- Generate synthetic assumption sets with N assumptions, assigned Wi and Δcrit distributions.
- Simulate Mi(t) time series under stochastic processes (e.g., Ornstein-Uhlenbeck[28] for mean-reverting variables, geometric Brownian motion[29] for prices) and embed correlation structure via copulas.
- Compute AGI trajectories and evaluate sensitivity to λ, β, τ_c, and φ parameters[30].

[28] The Ornstein-Uhlenbeck process is a stochastic process that models the velocity of a particle undergoing Brownian motion. It is characterized by its tendency to revert to a long-term mean, making it useful in various applications, particularly in finance for modeling mean-reverting behaviors such as interest rates and commodity prices.

[29] A geometric Brownian motion (GBM), also known as an exponential Brownian motion, is a continuous-time stochastic process in which the logarithm of the randomly varying quantity follows a Brownian motion with drift. The GBM process is characterized by its ability to produce continuous sample paths, meaning values do not jump discontinuously.

[30] λ – Confidence decay rate; volatility parameter that defines how fragile an assumption is or how likely to change over time; β – Cluster amplification parameter used to scale the impact of entangled assumption clusters when their joint migration exceeds a threshold. It reflects the nonlinear risk escalation that occurs when multiple correlated assumptions migrate together — a key concept from Quantum Project Management (QPM); τ_c – Cluster sensitivity threshold defines how much joint migration (via principal component score) is needed before amplification kicks in. It ensures that only material entanglement triggers nonlinear escalation.; and ϕ – Event-Driven Amplification Multiplier adjusts the confidence decay rate in response to disruptive events. After a shock (e.g., permit denial, market crash), assumptions become more fragile — even if time hasn't passed. This reflects measurement feedback and system perturbation from QPM.

- Use Receiver Operating Characteristic (ROC) analysis[31] to calibrate AGI thresholds for timely detection of true governance events while controlling false positives.

1.10.2 Retrospective empirical validation

Apply AGI methodology to historical LCP datasets where assumption registers or reconstructed assumption proxies exist. Analyze whether elevated AGI or CGI predicted rebaseline, cost overrun, or schedule slippage and quantify lead time improvement over existing governance indicators.

Validation metrics:

- true positive rate,
- false positive rate,
- time-to-detection delta, and
- reduction in unanticipated contingency draws.

1.10.3 Case study protocol (recommended)

- Select representative LCPs from different sectors (energy, data center hyperscale, infrastructure).
- Reconstruct assumption migration series via archived registers, procurement records, and stakeholder minutes.
- Compute AGI at historical reporting cadences and correlate AGI spikes to documented governance events (e.g., rebaselines, change orders).
- Use survival analysis[32] to estimate hazard ratio[33] of rebaseline conditional on AGI band.

1.11 Worked Example: AGI Trajectory and Cluster Amplification (n = 8)

This example models eight assumptions A_1 through A_8 across five governance categories: EconomicCost, Performance, Technical, Environmental, and Client. Each assumption is

[31] Receiver Operating Characteristic (ROC) analysis is a graphical representation used to assess the performance of a binary classification model. It plots the True Positive Rate (TPR), also known as sensitivity, against the False Positive Rate (FPR) at different threshold levels. The ROC curve helps in visualizing the trade-offs between sensitivity and specificity, allowing for the selection of optimal models and the comparison of different classifiers. (Wikipedia)

[32] Survival analysis is used to understand how likely a project is to undergo a rebaseline (a major governance event) depending on which AGI band it falls into (Green, Yellow, Amber, Red).

[33] The hazard ratio (HR) compares the risk of an event (here, rebaseline) between two groups. For example, if projects in the Red AGI band have a hazard ratio of 3.0 compared to those in the Green band, it means they are three times more likely to rebaseline at any given time.

tracked over a 12-month horizon, with monthly migration values $M_i(t)$, confidence decay $C_i(t)$, and governance weights ωapplied to compute AGI and CGI trajectories.

Step 1. Per-Assumption Migration and Weighting

Each assumption $A_i$is assigned:

- A baseline and current value to compute migration $M_i(t)$
- A decay rate $\lambda_i$and initial confidence $\alpha_{0,i}$
- Governance weights $\omega_{\text{cost}}, \omega_{\text{safety}}, \omega_{\text{schedule}}, \omega_{\text{reputation}}$

Migration is normalized against delta thresholds and scaled by confidence decay:

$$C_i(t) = \alpha_{0,i} \cdot e^{-\lambda_i t} \text{and} \text{ADI}_i(t) = C_i(t) \cdot \frac{|\, M_i(t) \,|}{\Delta_{\text{crit},i}}$$

AGI is computed as:

$$\text{AGI}_i(t) = \sum_{j \in \{\text{cost, safety, schedule, reputation}\}} \omega_j \cdot \text{ADI}_i(t)$$

CGIs are derived by aggregating AGI_i values within each governance category.

Step 2. Confidence Decay and AGI Trajectory

Over 12 months, assumptions with higher decay rates (e.g., $\lambda = 0.04$) show sharper drops in confidence, reducing their ADI footprint unless offset by migration events. For example:

- A3 (Technical) with $\lambda = 0.01$retains ~86% confidence at month 12
- A5 (Client Funding) with $\lambda = 0.04$drops to ~61% confidence by month 12

AGI trajectories reflect this decay unless amplified by entanglement triggers.

Step 3. Commodity Price Shock and Cluster Amplification

At month 9, a commodity price shock affects assumptions A1 and A6 (both in EconomicCost). Correlation matrix R shows:

$$\rho_{1,6} = 0.91$$

Principal component analysis identifies a dominant cluster PC_{ECON} exceeding its sensitivity threshold $\tau_{\text{ECON}} = 0.75$. Cluster amplification is triggered:

$$\Gamma_{\text{ECON}}(t) = 1 + \beta \cdot \max\left(0, PC_{\text{ECON}}(t) - \tau_{\text{ECON}}\right) \text{with} \beta = 0.6$$

At month 9, $PC_{\text{ECON}} = 0.90$, yielding $\Gamma = 1.09$. AGI values for A1 and A6 are scaled accordingly.

Step 4. Governance Signal Escalation

The aggregate AGI index crosses the Amber threshold at month 10, driven by amplified EconomicCost assumptions and decaying confidence in Client and Performance categories. This signal precedes a sponsor-level rebaseline decision at month 13—three reporting cycles later—demonstrating AGI's role as a lead indicator, consistent with assumption management literature.

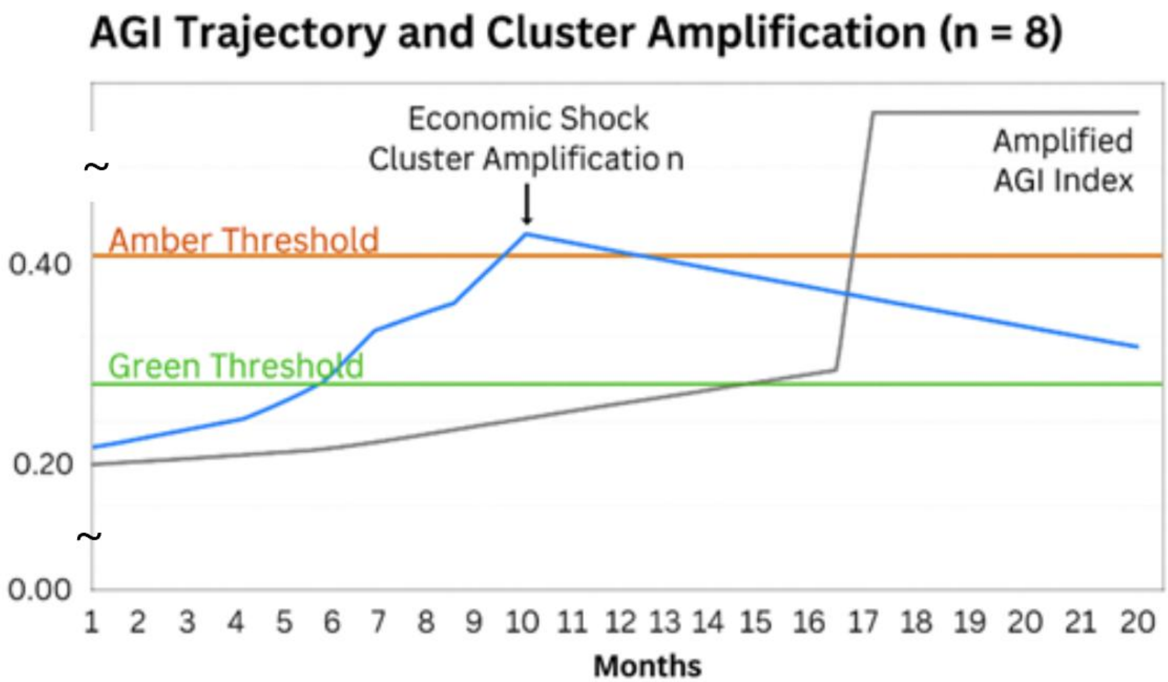

The base AGI index (blue) reflects normalized migration and confidence decay across eight assumptions. At Month 9, a commodity price shock triggers cluster amplification in the EconomicCost group (A1, A6), elevating the AGI index (gray) above the Amber threshold at Month 10. A sponsor-level rebaseline follows at Month 13, three reporting cycles later, demonstrating AGI's role as a lead indicator consistent with assumption management literature.

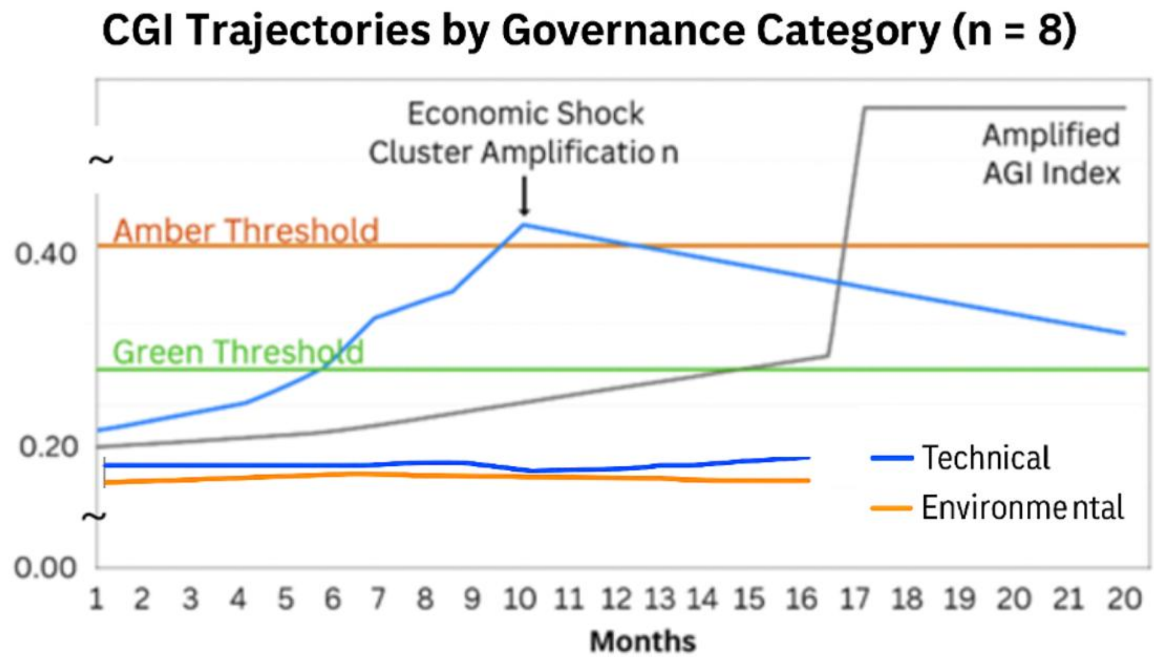

This figure plots Component Governance Indices (CGIs) for four governance categories over 12 months: EconomicCost, Performance, Technical, and Environmental. EconomicCost shows a noticeable spike at Month 9 consistent with the commodity-price shock; other CGIs show modest changes or gradual trends. The Amber threshold and sponsor rebaseline (Month 13) are included for alignment with the AGI timeline.

Assumption Mapping and Scores (n = 8) and AGI Signal Ranking and Escalation Triggers may be found in Appendix F

1.12 Understanding the AGI Framework

The AGI (Assumption Governance Index) framework is a structured analytic tool designed to transform granular project assumptions into traceable, decision-ready insights. It plays a pivotal role in construction safety, project delivery, and executive governance by translating raw assumption records into a single, auditable KPI (Key Performance Indicator). This KPI is not just a summary—it's a dynamic reflection of materiality, confidence, and consequence across time

The AGI framework provides numerous strengths in project governance as a leading indicator. Having said that, its limitations, implementation risks and good practice must be recognized and understood. These are discussed below.

Strengths of the AGI Framework

- **Governance-Ready KPI from Assumption Records** - AGI synthesizes detailed assumption logs—often buried in technical documentation—into a single, governance-ready metric. This enables executive teams to monitor risk exposure and assumption volatility without parsing raw data.
- **Materiality and Consequence Encoding** - Using Δcrit (criticality delta) and Wi (weighted impact), AGI embeds both the importance and consequence of each assumption. This dual encoding ensures that high-impact assumptions are prioritized in governance reviews.
- **Time-Aware Confidence Decay and Amplification** - AGI models how confidence in assumptions degrades over time (decay) and how emergent events can amplify their relevance. This temporal sensitivity reflects real-world dynamics, where yesterday's assumptions may no longer hold.
- **Nonlinearity and Entanglement Diagnostics** - By leveraging correlation matrices and principal component analysis (PCA), AGI identifies entangled assumptions and nonlinear risk propagation. These diagnostics are aligned with Quantitative Project Management (QPM) principles.
- **Traceable, Auditable Outputs** - Every AGI output is designed for auditability. This means governance bodies can justify decisions, trace back to source assumptions, and defend actions under scrutiny.

Limitations and Implementation Risks

- **Subjectivity in Weighting and Δcrit** - The assignment of weights and criticality thresholds can introduce bias. Without governance oversight, these parameters may be gamed or misused.
- **Data Completeness Challenges** - Early-stage projects often lack robust assumption registers. As a result, initial AGI outputs may be noisy or misleading until data maturity improves.
- **Parameter Calibration Complexity** - AGI relies on parameters like λ (confidence decay rate), β (amplification factor), and τ (time horizon). These must be calibrated using Monte Carlo simulations and historical fitting. Poor calibration can distort signal quality.
- **Cultural Risk of Metric Overreliance** - AGI is a powerful tool, but it must not replace expert judgment. Governance bodies must resist the temptation to treat AGI as a singular truth source.

Mitigations and Good Practices

- **Peer Review and Evidence Requirements** - High-impact changes to AGI inputs should undergo peer review and require documented evidence. This builds trust and reduces manipulation risk.
- **Bootstrapped and Smoothing**[34] - Statistical techniques like bootstrapping and smoothing help distinguish real signals from noise, especially in sparse datasets.
- **Governance Charter Integration** - AGI should be explicitly referenced in governance charters, with defined actions, thresholds, and audit trails to ensure accountability.
- **Periodic Recalibration** - AGI parameters should be recalibrated using empirical outcomes from past projects. This keeps the model aligned with evolving realities.

Alignment with QPM and contemporary governance thought

AGI operationalizes QPM's call to manage entangled, probabilistic projects by measuring assumption migration, identifying entangled clusters, and reacting adaptively to emergent system dynamics. It complements existing risk and cost uncertainty approaches (e.g., S-curve and PRAM[35]) and addresses the unique governance need to monitor the foundations (assumptions) upon which those quantitative assessments rest.

---

[34] AI can play a transformative role in both determining bootstrapped confidence intervals and applying statistical smoothing to distinguish noise from signal — especially in complex, high-dimensional governance environments like assumption migration analysis. AI can aid in bootstrapped confidence intervals:

- **Automate the resampling process** across thousands of iterations, dynamically adjusting sample size and stratification based on domain (e.g., Technical vs Stakeholder assumptions).
- **Detect and adapt to distributional anomalies** (e.g., heavy tails, multimodal clusters) using unsupervised learning.
- **Generate interpretable confidence bands** for metrics like μ_M, σ_M, or weighted migration μ_{wM}, and flag statistically significant shifts.
- **Integrate temporal decay** (via λ) and event-driven amplification (via ϕ) into bootstrapped models, making them governance-aware.

It can aid in smoothing:

- **Select optimal smoothing techniques** based on data characteristics (e.g., volatility, autocorrelation, entanglement).
- **Apply adaptive smoothing** — increasing sensitivity during shocks (high ϕ) and dampening noise during stable periods.
- **Use machine learning models** (e.g., Gaussian processes, recurrent neural networks) to forecast migration trajectories and detect inflection points.
- **Integrate smoothing with AGI computation**, so governance dashboards show both raw and smoothed metrics with confidence overlays. *(See Appendix D – Governance Dashboard and Audit Architecture)*

[35] **Project Risk Analysis and Management (PRAM)** is a methodology developed to help project teams systematically address uncertainty. It was formalized by the Association for Project Management (APM) and is

1.13 Conclusion

Assumption migration is a quiet but powerful driver of project instability in LCPs. The Assumption Governance Index proposed offers a reproducible, auditable, entanglement-aware instrument that converts granular assumption registers into executive-grade governance signals while providing domain-specific diagnostics for targeted mitigation. By aligning with QPM principles and established cost-risk practices, AGI supports adaptive governance, earlier detection of foundation erosion, and more accountable rebaselining decisions.

2.0 Assumption Diffusion Index

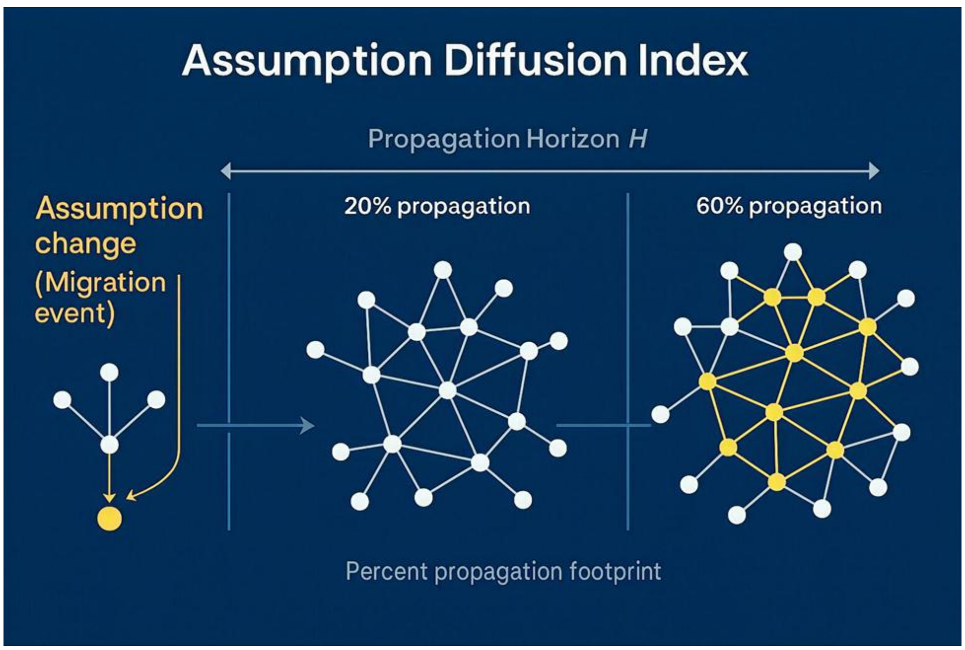

2.1 Introduction

The Assumption Diffusion Index (ADI) is a dynamic metric developed to complement the Assumption Governance Index (AGI). While AGI measures the current weighted contribution of assumptions to portfolio risk, ADI quantifies how changes to one or more assumptions propagate through the entanglement network over a defined horizon. ADI captures three governance-relevant properties that AGI does not:

- spread (how many and which nodes are affected)
- velocity (how quickly those effects materialize)

widely used across industries to improve decision-making, increase transparency, and enhance the likelihood of project success.

- network reach (which nodes act as diffusion hubs).

This section fully specifies ADI — conceptual framing, mathematical formulation, estimation and calibration procedures, implementation pseudocode, dashboard outputs, and governance interpretation. It has been structured to the extent possible to be free standing but has been consciously integrated with the Assumption Governance Index (AGI) recognizing its value added effect on that index. It aids in operationalizing AGI and transforming assumption migration into a valuable, predictive metric for large complex projects (LCPs).

2.2 Conceptual framing

ADI answers the operational question: if assumption X changes now, what fraction of weighted consequence across the register will be exposed, how fast will exposure grow, and which other assumptions are likely to become affected?

In complex project environments, assumptions are not isolated declarations—they form a dynamic network of interdependent beliefs that evolve over time. The **Assumption Diffusion Index (ADI)** is a governance-ready metric designed to quantify how assumptions propagate, influence one another, and expose projects to systemic risk. It complements the Assumption Governance Index (AGI) by offering a distinct lens focused on **spread, velocity, and reach**, rather than magnitude alone.

2.3 Structural Foundations: Nodes, Edges, and Temporal Influence

At the heart of ADI is a **graph-based model** where:

- **Nodes represent individual assumptions** tracked in the canonical assumption register[36]; each node carries a consequence weight on the same scale used in AGI. Each node carries a **consequence weight** aligned with the same scale used in AGI, ensuring consistency across governance metrics (Hollmann, 2016; Project Management Institute, 2017).
- **Edges encode entanglement coefficients**, which measure the **strength and sign of influence** between assumptions. These coefficients reflect how a

[36] A canonical register is the single authoritative inventory of assumptions used for governance, analysis, and decision-making. It is the trusted source of record that centralizes each assumption's identity, metadata, quantitative measures, provenance, and lifecycle events so that downstream indices (e.g., AGI, ADI), reports, and controls read from and reconcile to one consistent dataset. Its purpose is to:

- Ensure a single truth for assumption tracking across teams and tools.
- Provide traceability from observations and decisions back to the originating assumption and its evidence.
- Support reproducible computation of governance metrics, dashboards, and audit artifacts.
- Enable consistent escalation, remediation, and review workflows tied to named assumptions.

change in one assumption may affect others—positively, negatively, or ambiguously (Jolliffe & Cadima, 2016).

- **State changes**—such as migrations, rebaselines, or shocks—at a source node may **propagate influence to neighboring nodes**, attenuated by a **temporal kernel**[37]. that models decay over time. This mirrors confidence decay mechanisms found in AGI and supports time-aware governance (Efron & Tibshirani, 1994; ISO, 2018).

Aggregation and Normalization

ADI aggregates the **consequence-weighted exposure** reached by propagation from all potential source events (or a curated subset of seeded events). This cumulative exposure is then **normalized to the portfolio consequence scale**, allowing governance bodies to compare diffusion risk across projects, phases, or portfolios (COSO, 2017; NASA, n.d.).

Governance Applications

The governance value of ADI lies in its ability to surface latent risks and guide proactive intervention:

- **Early Warning**: ADI can flag **systemic exposures** before the AGI snapshot shifts materially. This enables governance teams to act on diffusion signals even when magnitude metrics remain stable.
- **Triage**: By identifying **diffusion hubs**—nodes with high outbound influence—ADI supports **edge-level interventions** *(see box below).* These may include assumption revalidation, constraint imposition, or targeted controls to reduce propagation.

[37] A temporal kernel is a function that governs how influence or effect decays (or persists) over time in a dynamic propagation model. In diffusion or contagion models, the kernel K(Δt) maps a time lag Δt ≥ 0 to a nonnegative scalar multiplier that weights how strongly an earlier event contributes to influence at a later time. It may take any one of a number of functional forms. Key properties to consider include:

- Decay rate - Controls how quickly past events lose influence; higher decay parameters compress propagation into earlier steps.
- Tail behavior - Determines whether rare, long-lag effects remain material (power law) or effectively vanish (exponential).
- Normalization - Whether kernel integrates/sums to 1 (probability interpretation) or <1 (attenuating transfer).
- Nonnegativity and sign - Whether kernel is strictly nonnegative (typical) or signed when oscillatory dynamics are modelled.

Edge-level interventions are targeted actions taken at the periphery or boundaries of a system—where risk, change, or failure is most likely to emerge—in order to prevent escalation, contain disruption, or reinforce system resilience.
In systems thinking, the "edge" refers to **points of interaction, transition, or vulnerability**—such as:

- **Frontline operations** (e.g., field crews, customer service, first responders)
- **Interfaces between systems** (e.g., handoffs between contractors and clients)
- **Thresholds of risk** (e.g., near-miss zones, early warning indicators)
- **Emerging or unstable zones** (e.g., new technologies, regulatory gray areas)

These are the places where small failures can propagate, or where early signals of systemic stress often appear.
Edge-level interventions look are typically granular, localized, and proactive, such as:

- **Field-level checklists** triggered by threshold crossings (e.g., heat index, noise exposure)
- **Micro-training modules** deployed after a near miss or deviation
- **Real-time alerts** when leading indicators (e.g., AGI Index, cluster amplification) cross predefined thresholds
- **Localized governance nudges** (e.g., supervisor prompts, signage, or digital prompts at decision points)

These interventions are often automated, modular, and traceable, designed to interrupt drift or amplify protective behaviors before systemic escalation.
Edge-level interventions facilitate:

- **Containment**: They act as circuit breakers, preventing small issues from cascading into major failures.
- **Responsiveness**: They enable rapid, context-specific action without waiting for top-down directives.
- **Learning**: They generate data for feedback loops, enabling continuous improvement and adaptive governance.
- **Empowerment**: They equip frontline teams with tools to act decisively within their span of control.

Construction safety example:
A project's AGI Index is rising due to clustered deviations in confined space entries. An edge-level intervention might include:

- A localized pause for rebriefing at the affected workface
- A QR-code-triggered microlearning on atmospheric monitoring
- A temporary escalation protocol requiring dual sign-off for reentry

These actions are targeted, time-bound, and traceable, reinforcing system integrity at the edge.

- **Prioritization**: ADI introduces a **second axis of prioritization** based on spread, velocity, and reach. This complements AGI's magnitude axis and enables more nuanced monitoring and mitigation strategies.

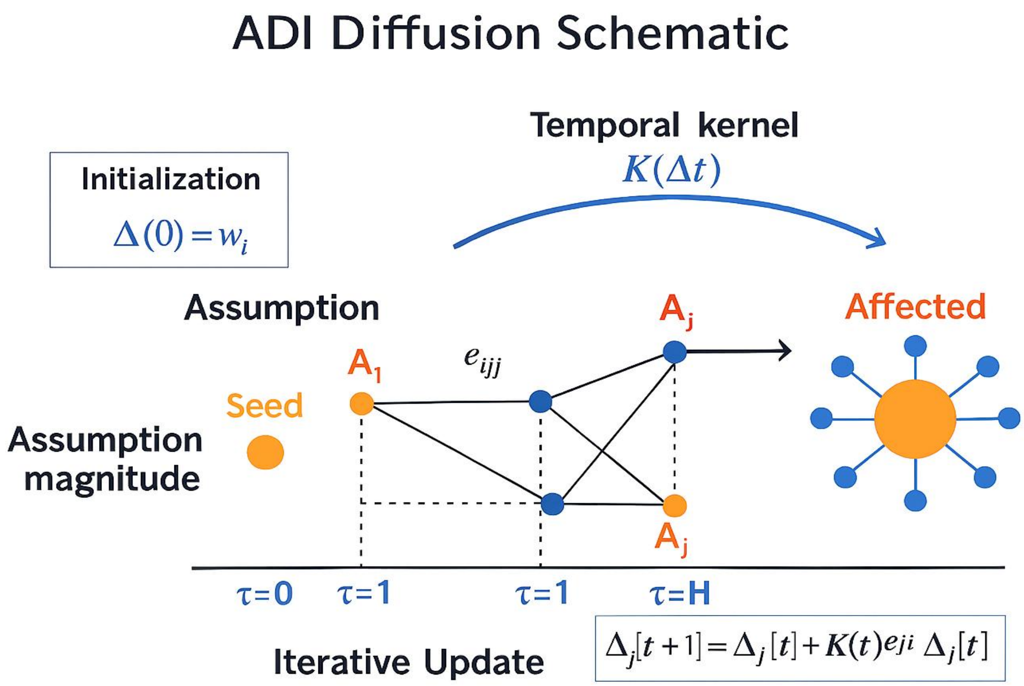

Integration with Existing Frameworks

ADI aligns with established risk modeling tools such as WSDOT's **Project Risk Analysis Model (PRAM)** and NASA's **Cost Risk and Uncertainty Methodologies**, both of which emphasize assumption tracking and scenario propagation. It also supports ISO 31000's call for integrated risk governance and COSO ERM's emphasis on performance-linked metrics.

2.4 Notation and core definitions

This section formalizes the computational structure of the Assumption Diffusion Index (ADI), offering a rigorous foundation for modeling how individual assumptions propagate influence across a project's assumption network. By treating assumptions as nodes within a directed graph, ADI captures both the magnitude of state changes and the directional entanglement between assumptions. Each node carries a consequence weight consistent with AGI metrics, while edges encode influence coefficients that may amplify or attenuate downstream effects.

The model incorporates a temporal decay kernel to reflect the diminishing strength of influence over time, and a propagation horizon to bound the evaluation window. A transfer function governs how influence scores are transmitted, and node-specific thresholds determine when an assumption is considered "affected." The resulting scalar—

instantaneous diffusion potential—quantifies the immediate systemic exposure triggered by a source event, normalized to the portfolio's total consequence weight. This formulation enables governance teams to detect latent risks, prioritize interventions, and justify decisions with traceable, time-aware metrics. It serves as a critical bridge between assumption-level volatility and portfolio-level consequence management.

Computational details have been included for completeness, but the essential thought is that ADI can be measured and tracked and together with AGI provides directional insights with predictive value in LCP. **LCP, like all quantum systems, are influenced by the systems within which they are embedded.** Other complementary metrics look at the surrounding project ecosystem considering changes in those systems, entanglements driven by complexity, inherent quantum uncertainty and other QPM related aspects. The combination of AGI and ADI looks at the continuing validity of foundational assumptions and their current trajectories. Together they aid in the predictive management of LCP.

Let's turn now to the assumption register which contains n assumptions indexed by i = 1,…,n. Each assumption is associated with a compilation of consequences, signals, entanglement (reflecting the quantum nature of LCP), decay function describing the time value of signal strength, and other influence and threshold scores. These are summarized here and reflect the need for a deeper understanding of the nature of the characteristics that describe a given assumption.

- **w_i** : consequence weight for assumption i (same units as AGI components).
- **$s_i(t)$** : state-change signal for i at time t (continuous migration magnitude or binary event indicator).
- **$E = [e_{ij}]$** : entanglement matrix with $e_{ij} \in [-1,1]$, representing influence of node i on node j; matrix may be directed.
- **$K(\Delta t)$**[38] : temporal decay kernel controlling transfer strength at lag Δt.
- **H** : discrete propagation horizon (number of time steps to evaluate post-event spreading).
- **$g(\cdot)$** : transfer function applied to source or intermediate influence scores (linear or saturating).
- **θ_j** : per-node threshold used to define when node j becomes "affected."
- **$W_{tot} = \Sigma_i w_i$** : portfolio total consequence weight used for normalization.

To quantify the immediate systemic exposure triggered by a change in any given assumption, the Assumption Diffusion Index (ADI) introduces the concept of **instantaneous diffusion potential**. This scalar value reflects how a state change at

[38] K(Δt) is the temporal weighting function that determines how influence attenuates (or persists) with lag; kernel choice (exponential, power-law, boxcar, delayed) materially affects ADI velocity and reach. See Section 2.7.1 for guidance on kernel selection and calibration

node i at time t may influence other assumptions in the network, based on the magnitude of the change and the strength of outbound entanglements. Formally, it is defined as:

$$DP_i(t) = | s_i(t) | \cdot \sum_{\substack{j=1 \\ j \neq i}}^{n} | e_{ij} |$$

Here, $s_i(t)$represents the state change signal for assumption $i$at time t, which may be a binary event indicator or a continuous migration magnitude. The summation term aggregates the absolute values of the entanglement coefficients e_{ij}, capturing the total influence that node $i$exerts on all other nodes $j \neq i$. The product of these two components yields a scalar that can be used to rank assumptions by their potential to propagate risk, enabling governance teams to identify high-impact nodes and prioritize monitoring or intervention accordingly.

2.5 Propagation Model and ADI Formulas

To operationalize the Assumption Diffusion Index (ADI), we use a **discrete-time iterative propagation model**[39] that simulates how influence spreads across an assumption network following a source event. This model enables governance teams to quantify not only the magnitude of individual assumption changes but also their systemic reach and velocity over time. The following subsections define the core mechanics and metrics of ADI; each grounded in traceable mathematical logic and aligned with consequence-based governance[40].

[39] A discrete-time iterative propagation model is a stepwise computational framework used to simulate how a quantity—such as risk, influence, energy, or information—spreads or evolves over time across a system, using fixed time intervals and repeated update rules.

[40] **Notes on Implementation**

- **Stochastic vs. Deterministic Modeling**: Probabilities arise when propagation is modeled stochastically—e.g., randomized thresholds, noisy edges, or Monte Carlo seed variations. Deterministic implementations use binary affected states directly in footprint calculations.
- **Dashboard Integration**: ADI values can be rescaled to a 0–100 range for dashboard visualization, enabling intuitive tracking of systemic assumption risk.

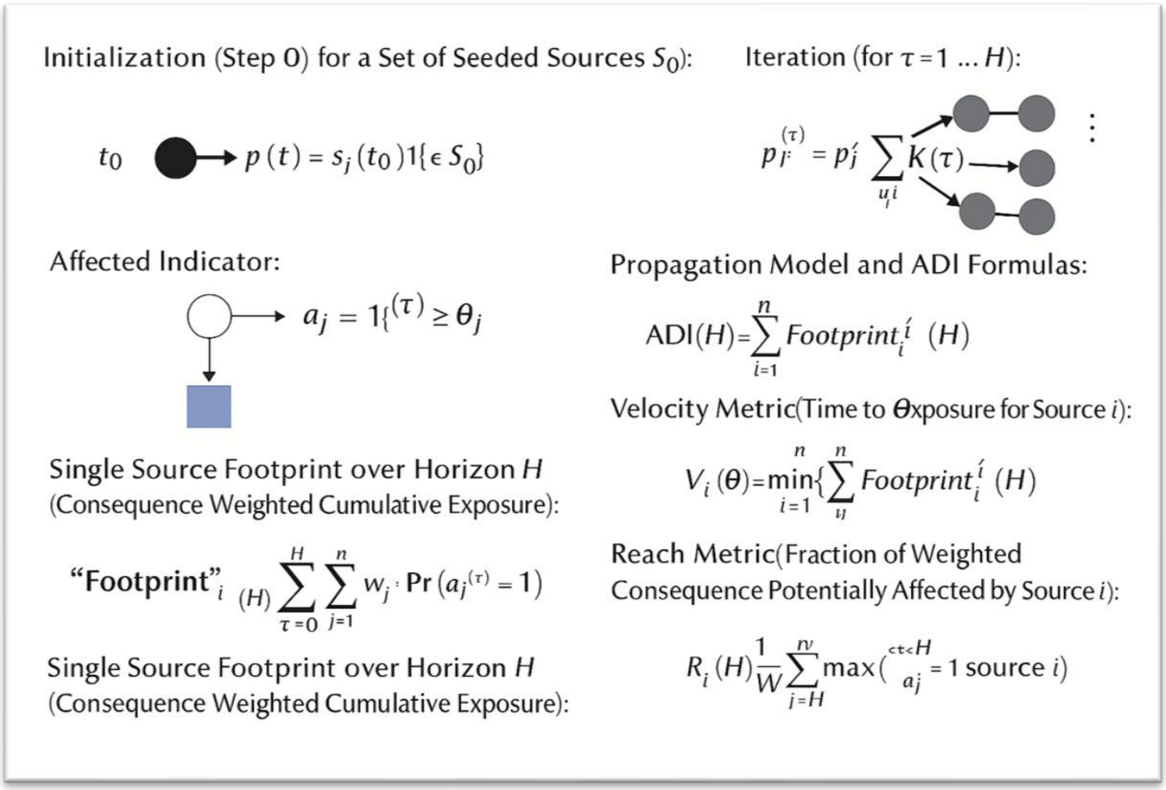

2.5.1 Initialization of Influence Scores

At time step $\tau = 0$, the model initializes influence scores for a set of seeded source assumptions S_0. Each source node begins with a non-zero influence score, while all other nodes are set to zero. This establishes the starting condition for propagation.

$$p_j^{(0)} = \begin{cases} 1, & \text{if } j \in S_0 \\ 0, & \text{otherwise} \end{cases}$$

This initialization reflects the immediate activation of selected assumptions due to a triggering event (e.g., migration, rebaseline, or shock). It sets the stage for evaluating how influence spreads through entangled assumptions over the propagation horizon H.

2.5.2 Iterative Propagation Update Rule

For each subsequent time step $\tau = 1,2, \dots, H$, influence scores evolve based on entanglement coefficients, a temporal decay kernel, and a transfer function. The update rule is:

$$p_j^{(\tau)} = p_j^{(\tau-1)} + \sum_i K(\tau) \cdot e_{ij} \cdot g(p_i^{(\tau-1)})$$

Here:

- $K(\tau)$is the **temporal decay kernel** that attenuates influence over time.
- $e_{ij} \in [-1,1]$is the **entanglement coefficient** representing the directional influence from node i to node j.
- $g(\cdot)$is a **transfer function**, which may be linear or saturating, used to modulate influence transmission.

This formulation allows the model to simulate realistic propagation dynamics, where influence weakens over time and may be amplified or dampened depending on network structure and transfer behavior.

2.5.3 Affected Indicator Function

To determine whether a node becomes "affected" by propagated influence, the model uses a threshold-based indicator:

$$a_j^{(\tau)} = \mathbf{1}\{p_j^{(\tau)} \geq \theta_j\}$$

Each node j has a threshold θ_j, and the indicator function returns 1 if the cumulative influence exceeds this threshold. This binary signal is used to track which assumptions are activated at each time step, forming the basis for exposure metrics.

2.5.4 Single Source Footprint

The **footprint** of a single source assumption over the propagation horizon $H$is defined as the cumulative consequence-weighted exposure it generates:

$$\text{Footprint}_i(H) = \sum_{\tau=0}^{H} \sum_{j=1}^{n} w_j \cdot \Pr(a_j^{(\tau)} = 1 \mid \text{source } i)$$

This metric aggregates the probability-weighted impact across all affected nodes, scaled by their consequence weights w_j. It captures the total systemic exposure attributable to a single assumption, enabling governance teams to rank sources by their diffusion potential.

2.5.5 Portfolio-Level ADI

To evaluate systemic exposure across the entire assumption portfolio, the model computes a normalized index:

$$ADI(H) = \frac{1}{W_{\text{tot}}} \sum_{i=1}^{n} w_i \cdot \text{Footprint}_i(H)$$

Here, $W_{\text{tot}} = \sum_i w_i$ is the total portfolio consequence weight. This normalization ensures that ADI is expressed on the same scale as AGI, allowing for integrated dashboarding and comparative analysis. ADI can be rescaled to a 0–100 range for readability in executive reporting.

2.6 ADI Metrics For Governance

Velocity Metric

The **velocity** of a source assumption quantifies how quickly its influence reaches a specified fraction θ of its total footprint:

$$V_i(\theta) = \min \{\Delta t: \sum_{\tau=0}^{\Delta t} \sum_{j} w_j \cdot \Pr(a_j^{(\tau)} = 1) \geq \theta \cdot \text{Footprint}_i(H)\}$$

This metric supports governance triage by identifying fast-spreading assumptions that may require immediate intervention or monitoring.

Reach Metric

The **reach** of a source assumption measures the fraction of total portfolio consequence weight that it potentially affects over the horizon H:

$$R_i(H) = \frac{1}{W_{\text{tot}}} \sum_{j} w_j \cdot \Pr(\max_{\tau \leq H} a_j^{(\tau)} = 1 \mid \text{source } i)$$

This metric highlights assumptions with broad systemic influence, even if their velocity is low. It is particularly useful for identifying latent risks that may not trigger immediate concern but have long-term governance implications.

2.7 Choice of kernels, transfer functions, and thresholds

2.7.1 Kernel in the ADI context

A **kernel** in the ADI (Assumption Diffusion Index) context is the time-dependent weighting function that controls how influence from a source node is transferred to other nodes as a function of the time lag since the source event. The kernel determines the strength and temporal profile of propagation: it scales and shapes the contribution from a past influence score p_i at lag Δt so that near-term effects can be emphasized, long-term effects attenuated, or specific temporal patterns (delay, burst, persistence) can be modeled.

Why the kernel matters

- **Controls decay and persistence** — it encodes whether influence drops off quickly, slowly, or remains steady over the propagation horizon H.
- **Shapes velocity and reach** — faster decaying kernels compress diffusion into early steps (higher velocity, lower long-term reach); long-tailed kernels spread influence slowly but broadly (lower velocity, greater reach).
- **Enables realistic dynamics** — kernels let you represent domain behaviors such as information latency, review cycles, approval delays, or learning effects in assumption validation.
- **Governance signal tuning** — selecting a kernel directly affects early-warning sensitivity and triage prioritization because it changes when and how much downstream nodes are affected.

Common kernel choices and their behaviors

- **Exponential decay** $K(\Delta t) = \alpha \cdot \exp(-\beta \Delta t)$
  - Behavior: strong immediate impact that decays smoothly.
  - Use when influence fades continuously (e.g., initial estimates losing relevance).
- **Power law (long-tailed)** $K(\Delta t) = \alpha \cdot (\Delta t + 1)^{-\gamma}$
  - Behavior: slower decay with long persistence.
  - Use when effects have memory or long-running dependencies.
- **Finite window / boxcar** $K(\Delta t) = \alpha$ for $\Delta t \leq \tau_0$, else 0
  - Behavior: uniform transfer within a fixed delay window, then zero.
  - Use for batch review cycles or contractual milestones.
- **Delayed (shifted) kernel** $K(\Delta t) = \alpha \cdot f(\Delta t - \delta)$ for $\Delta t \geq \delta$
  - Behavior: little or no immediate transfer, then onset after delay δ.
  - Use where downstream influence occurs only after a known lag (e.g., procurement lead time).
- **Custom multimodal kernel** combination of the above
  - Behavior: models multiple phases (immediate spike + long tail + delayed secondary pulse).

- Use for complex processes with several distinct propagation modes.

In ADI modeling the first three kernel types dominate.

Practical parameter guidance

- **Scale α** should be chosen so that aggregated transfer magnitudes are meaningful relative to node thresholds θ_j and consequence weights w_j.
- **Decay rate (β or γ)** calibrates velocity: higher values → faster attenuation. Calibrate using historical event propagation or expert elicitation.
- **Horizon alignment** ensure kernel support is consistent with propagation horizon H; truncate kernels beyond H to bound computation.
- **Normalization** optionally normalize K(·) so that total transferred influence per source over H is interpretable (e.g., integrates to 1) and comparable across scenarios.

Implementation and governance notes

- **Stochastic vs deterministic** implementations can apply kernels deterministically or sample kernel parameters in Monte Carlo runs to capture temporal uncertainty.
- **Explainability**: record chosen kernel form and parameters in governance artifacts so ADI outputs remain traceable and auditable.
- **Sensitivity testing**: run ADI under alternate kernel families to surface how detection, triage, and prioritization change with temporal assumptions.
- **Operational use**: tune kernels to align ADI alerting thresholds with organizational review cadences to reduce false alarms and focus controls where propagation is most active.

2.7.2 Transfer function g(·)

The **transfer function g(·)**[41] in the Assumption Diffusion Index (ADI) framework is the node-level transformation applied to a source or intermediate influence score before it is transmitted across an edge. It maps a node's raw influence state into the effective signal that flows to neighbors, controlling saturation, nonlinearity, sign handling, and sensitivity to small or extreme signals.

[41] g(·) maps a node's raw influence into the effective signal transmitted across edges (linear, saturating, dead-zone, or super-linear forms are possible). Choice of g(·) changes ADI sensitivity and must be recorded for audit; see Section 2.7.2 and calibration notes in Section 2.10

Purpose and intuitive role

- **Modulate amplitude** — scale or clip raw influence so downstream propagation respects physical, contractual, or cognitive limits.
- **Introduce nonlinearity** — create thresholds, saturation, or amplification regimes so small perturbations do not necessarily propagate while larger shocks do.
- **Preserve or transform sign** — ensure negative/positive influences behave appropriately when edges carry signed entanglement (e.g., mitigation vs. amplification).
- **Increase interpretability** — convert internal scores to a bounded, comparable range used by governance thresholds and kernels.

Common functional forms and behaviors

- **Linear**: $g(x) = \alpha x$
  - Behavior: proportional transfer; preserves sign and scale.
  - Use when influence transmits without thresholding or saturation.
- **Saturating (logistic or tanh)**: $g(x) = \gamma \cdot \tanh(\alpha x)$ or $g(x) = \gamma / (1 + \exp(-\alpha(x - x0)))$
  - Behavior: small inputs attenuated; outputs limited to a finite range; smooth transition around inflection.
  - Use when downstream systems have capacity limits (saturation) or when very large upstream shocks shouldn't produce unbounded downstream effects.
- **Rectified or half-wave**: $g(x) = \max(0, \alpha x - \beta)$
  - Behavior: floor/threshold behavior; negative or sub-threshold values produce no transfer.
  - Use when only increases (or only decreases) matter for propagation.
- **Piecewise or dead zone**: $g(x) = 0$ for $|x| \le \delta$; linear/saturating outside
  - Behavior: ignores noise and small migrations; transmits only meaningful deviations.
  - Use to avoid false positives from measurement noise or routine fluctuation.
- **Power law / super linear**: $g(x) = \alpha \cdot |x|^\kappa \cdot \mathrm{sign}(x)$ $(\kappa > 1)$
  - Behavior: amplifies large shocks relative to small ones.
  - Use when systemic effects accelerate nonlinearly with source magnitude.
- **Binary gating**: $g(x) = 1\{x \ge \gamma\}$
  - Behavior: transmits only when the source exceeds a hard gate.
  - Use in scenario analysis or event-driven propagation.

The first two functional forms (linear; logistics or saturating) are most applicable in the ADI context.

Key properties to design for

- **Monotonicity**: ensures larger source states do not produce smaller transfers unless intentional (useful for interpretability).
- **Boundedness**: keeps propagated influence within interpretable limits for threshold comparisons.
- **Differentiability** (optional): helpful for sensitivity analysis and gradient-based calibration.
- **Symmetry or sign-awareness**: decide whether positive and negative changes should be treated symmetrically or differently.

Governance and operational implications

- **Alert sensitivity**: choice of g(·) directly affects early-warning behavior—more aggressive g increases ADI sensitivity and false-alarm risk; conservative g reduces noise but can delay detection.
- **Control design**: where g imposes dead zones or saturation, governance can map these mechanics to policy actions (e.g., review thresholds, automatic holds).
- **Traceability**: record the chosen g(·) and parameters in governance artifacts so ADI outputs remain auditable.
- **Calibration**: calibrate using historical propagation events, expert elicitation, or Monte Carlo tuning to align ADI outputs with observed systemic responses.

Practical implementation notes

- **Normalization**: align g(·) output range with edge weights and kernel scale so combined contributions are meaningful relative to thresholds θ_j and consequence weights w_j.
- **Stochastic variants**: sample parameters of g(·) in Monte Carlo runs to capture model uncertainty or heterogeneity across organizations.
- **Computational efficiency**: prefer simple closed-form g(·) for large networks; reserve complex forms for focused scenario analyses.
- **Sensitivity testing**: run ADI across alternate g(·) families to quantify impact on early-warning, velocity, and reach metrics.

2.7.3 Thresholds

Thresholds θ_j are the node-level cutoffs that determine when a propagated influence score $p_j^{(\tau)}$ is considered to have meaningfully affected assumption j. They convert continuous influence trajectories into binary governance signals used for escalation, logging in the canonical register, and feeding footprint and ADI calculations. Properly designed thresholds balance sensitivity to true systemic propagation against resilience to noise and routine fluctuation.

Principles for selecting thresholds

- **Governance alignment** — set thresholds so that an "affected" flag corresponds to an operationally meaningful event (e.g., triggers a review, revalidation, or control). Thresholds should map to governance actions, not just statistical quirks.
- **Risk proportionality** — lower thresholds for high-consequence nodes and higher thresholds for low-consequence nodes so that signals reflect potential impact as well as spread.
- **Comparability with AGI scale** — choose threshold units consistent with the consequence and state scales used across AGI/ADI so signals and normalized indices remain interpretable.
- **Robustness to noise** — ensure thresholds ignore ordinary measurement noise and minor migrations while capturing substantive deviations that matter to delivery or cost.

Practical selection methods

- **Empirical percentile calibration**
  - Derive θ_j from historical distributions of p_j (for example, the 90th percentile) so that only unusually large, propagated influence triggers an affected state.
  - Useful when ample historical propagation data exist and when the organization prefers probabilistic control of false alarm rates.
- **Consequence sensitivity calibration**
  - Start with a baseline percentile then scale by consequence: $\theta_j = \theta_{\text{base}} \cdot f(w_j)$ where f reduces the cutoff for critical w_j values.
  - Ensures critical assumptions are monitored more aggressively.
- **Cost-benefit / decision-theoretic calibration**
  - Choose thresholds by balancing the expected cost of missed detections (undetected diffusion leading to loss) against the cost of false positives (unnecessary reviews).
  - Useful where action costs and failure costs can be approximated.
- **Receiver operating characteristic style tuning**
  - If labeled propagation outcomes exist, sweep θ_j and plot detection rate versus false alarm rate to select operating points aligned with governance tolerance.
- **Bayesian / empirical Bayes priors**
  - For nodes with limited data, combine weak empirical priors with domain priors to obtain a posterior-based threshold that shrinks extreme values toward plausible defaults.
- **Expert elicitation and consensus rules**

  - Convene SME panels to set thresholds where data are sparse or where contextual factors (contracts, regulations, safety margins) dominate. Capture rationale in the register for auditability.
- **Hybrid hierarchical approach**
  - Use a tiered scheme: project-level default thresholds, category adjustments (e.g., safety, cost, schedule), and node-specific overrides for known special cases.

Empirical and expert elicitation are common choices with respect to ADI modeling.

Dynamic and adaptive thresholding

- **Time-varying thresholds**
  - Allow $\theta_j(t)$to change across project phases (design, procurement, construction) to reflect shifting tolerances and inspection cadences.
- **Adaptive thresholds via feedback**
  - Use online performance metrics (e.g., precision of prior alerts) to tighten or relax thresholds automatically, with governance approval and logged parameter changes.
- **Contextual gating**
  - Implement conditional thresholds that depend on contextual factors (e.g., vendor risk level, milestone proximity) to reduce spurious alerts.

Implementation and validation best practices

- **Document rationale**[42] — each threshold entry must include the selection method, data sources, date, owner, and intended governance action so the assumption register remains auditable.
- **Sensitivity testing** — run ADI under multiple threshold scenarios to quantify how detection, velocity, and reach metrics shift; surface any brittle nodes that flip state with small threshold changes.
- **Monte Carlo calibration** — simulate stochastic propagation with varied kernels, edge noise, and transfer functions to estimate expected false alarm and missed detection rates for candidate θ_j.
- **Operational tuning window** — establish a governance review cadence (e.g., monthly or per milestone) to reassess thresholds based on new data, incidents, or changes in consequence weights.
- **Fail-safe defaults** — when in doubt, apply conservative defaults for critical domains (err on the side of earlier review for safety or high financial exposure).

[42] See Appendix A - Structuring the Assumption Register for Governance Integrity

- **Traceability and version control** — store threshold versions and the outputs they produced so post-event forensics can link decisions to the threshold regime in force at the time.

Examples and numeric guidance

- **Example default**: set project default θ_{base}at the 95th percentile of historical $p_j$values; for nodes in the top 10 percent by w_j, reduce to the 85th percentile.
- **Example decision threshold**: choose $\theta_j$so that expected number of governance reviews per quarter remains operationally feasible (control for alert volume).
- **Example adaptive rule**: if false alarm rate > 30% over three review cycles, increase $\theta_j$by 10%; if missed event cost exceeds expected review cost, decrease $\theta_j$by 10%.

Governance implications and why thresholds matter

Thresholds are the fulcrum between sensitivity and stability in ADI. Poorly chosen thresholds either flood governance with false positives or hide slowly diffusing but consequential exposures until they become crises. Well-constructed thresholds make ADI actionable: they turn propagation signals into timely reviews, targeted controls, and auditable decision points that integrate seamlessly with AGI magnitude insights. Recording threshold logic and performing routine recalibration converts ADI from a purely analytic metric into a living governance control that adapts with the project lifecycle.

2.8 Estimating the entanglement matrix E[43] (practitioner guidance)

The entanglement matrix E = [e_{ij}] is the backbone of ADI: it encodes how a change (migration, shock, rebaseline) in assumption i influences assumption j. Reliable ADI depends on well-estimated e_{ij} values that are interpretable, auditable, and fit for governance use. Below is a robust, practitioner-oriented narrative that defines terms, explains methods, shows worked examples, and gives concrete recommendations for estimation, post-processing, validation, and governance use.

[43] E is the directed, signed matrix of pairwise transfer coefficients (e_ij ∈ [−1,1]) encoding how a change in assumption i influences assumption j. Because ADI results are sensitive to E, each edge must include provenance and a confidence tier (see Section 2.8 and the matrix-hygiene checklist in Appendix E).

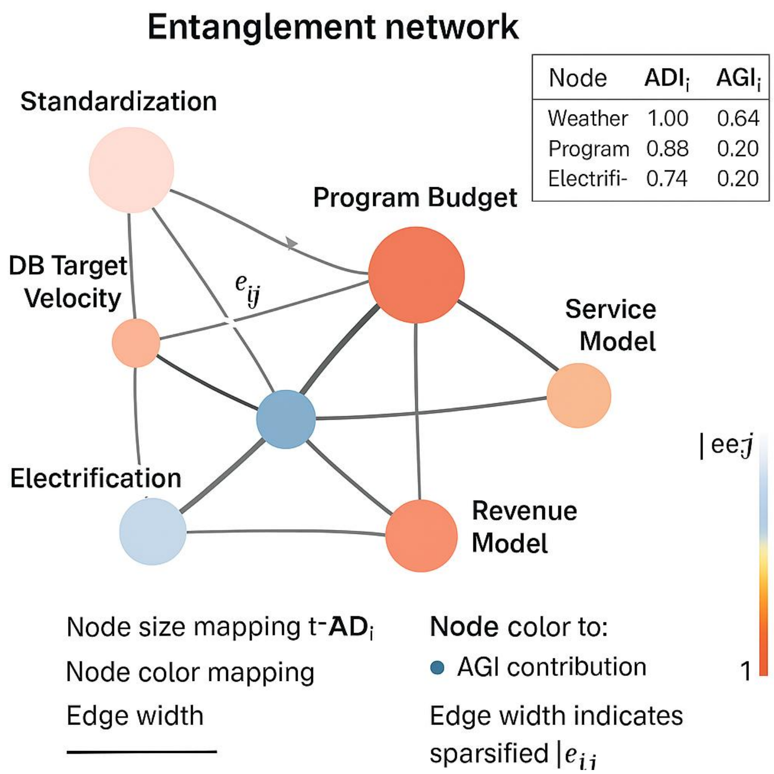

What e_{ij} represents

- e_{ij} ∈ [−1, 1] is a signed, bounded coefficient that measures the *direction*, *strength*, and (where appropriate) *directionality* of influence from node i (source) to node j (target).
  - Sign: positive means an increase/migration in i tends to increase the state or risk exposure of j; negative means it tends to reduce j or push it opposite.
  - Magnitude: absolute value near 1 indicates a strong influence; value near 0 indicates little or no influence.
  - Directionality: e_{ij} need not equal e_{ji}; directed edges capture asymmetric causal influence.
- Interpretation: think of e {ij} as a normalized sensitivity or transfer coefficient used in propagation updates p_j ← p_j + ∑ K e_{ij} g(p_i).

2.8.1 Data-driven methods

1. Historical co-migration frequency (discrete event approach)
   - Goal: estimate P(j migrates | i migrates) from past event logs and map that probability to a bounded signed scale.
   - Steps:

1. Define a consistent event window following an i-event (e.g., 30 days after i migrates) within which j is considered to have co-migrated.
2. Compute conditional probability p̂ = count(j co-migrates within window | i migrates) / count(i migrates).
3. Map p̂ ∈ [0,1] to e_{ij} ∈ [−1,1]. Simplest mapping: e_{ij} = 2·p̂ − 1 (maps 0→−1, 0.5→0, 1→+1). Alternative: preserve sign if you have directional labels (e.g., if i's positive migration is known to reduce j, flip sign).

- Interpretation: p̂ close to 1 → strong positive coupling; p̂ close to 0 → strong negative in the linear mapping (check sign semantics first). If you require only nonnegative coupling, map p̂ to [0,1] and keep sign separate.
- When to use: event-rich datasets with many migrations and consistent event definitions.

2. Copula or rank-correlation of migration magnitudes (continuous co-movement)
   - Goal: measure continuous co-movement of migration magnitudes while robust to nonlinearity and outliers.
   - Steps:

     1. Collect paired time series of migration magnitudes Δp_i(t), Δp_j(t) on the same cadence.
     2. Compute a rank-based association (Spearman ρ) or use a copula fit to capture tail dependence.
     3. Convert correlation to e_{ij} by rescaling (e_{ij} = ρ if ρ ∈ [−1,1]). Optionally dampen via shrinkage.
   - Interpretation: captures whether large movements in i co-occur with large moves in j (and whether in same or opposite direction). Copulas allow separate treatment of tail co-movement (important for systemic shocks).

3. Granger-style causality[44] (time-series directional inference)
   - Goal: infer directed predictive influence from i → j using time series testing.
   - Steps:

[44] Granger-style causality is a statistical method used to determine whether one time series can predict another, based on the idea that causes precede effects. In Granger's framework, a variable X is said to *Granger-cause* another variable Y if past values of X contain information that helps predict future values of Y—beyond what is possible using past values of Y alone. It's a form of predictive causality. (See Wikipedia Granger causality - Wikipedia)

1. Ensure sufficient sample length and appropriate stationarity or difference the series.
2. Fit vector autoregression (VAR) or Granger causality tests to assess whether past values of i improve prediction of j beyond j's own history.
3. Where Granger significance exists, extract effect sizes (impulse response functions or VAR coefficients) and rescale to [−1,1] for e_{ij}.

- Interpretation: gives evidence of temporal precedence and predictive influence; best used where sampling cadence and history permit robust inference.

4. Frequency-domain / spectral coherence (for periodic or cyclical propagation)
   - Goal: quantify coupling that is concentrated at particular periodicities (review cycles, weekly cadence).
   - Steps: compute coherence between series at relevant frequency bands and map coherence to e_{ij}.
   - Use when propagation is expected to follow cadence patterns.

2.8.2 Expert elicitation and structured judgment

- When data are sparse, biased, or nonrepresentative, use structured elicitation:
   - Define clear prompts (what constitutes an i-event; time window; sign conventions).
   - Use bounded numeric scales (e.g., map experts' P(j|i) estimates to [−1,1] via the same mapping used for historical $\hat{p}$).
   - Apply formal aggregation (Cooke's method or weighted averages with calibrated experts) and capture ranges (min/median/max) to represent uncertainty.
   - Always record provenance: who provided the estimate, date, and rationale.

Worked example (expert): SME estimates P(j migrates | i migrates) = 0.7 and judge the direction as positive → $e_{ij} = 2 \cdot 0.7 - 1 = 0.4$.

Hybrid estimation (recommended default)

- Combine data and expert priors with shrinkage toward zero for noisy links. Typical workflow:
   1. Compute data-driven estimate $\hat{e}_{data}$ (from one or more methods above).
   2. Obtain expert prior e_{prior} and a confidence weight $\alpha \in [0,1]$ (higher α → more trust in data).

3. Shrink/merge: $e_{ij} = \lambda \cdot \hat{e}_{data} + (1 - \lambda) \cdot e_{prior}$ where λ = function(sample size, signal-to-noise). For small sample sizes, set λ small (lean on prior); for large sample sizes, $\lambda \to 1$.
4. Variance estimate: produce a confidence interval for e_{ij} to use in stochastic Monte Carlo runs.

- Recommendation: encode λ rules in governance (e.g., $\lambda = \min(1, n_{obs}/50)$) and document choices. *(See the following box)*

Encode λ rules in governance

Definition To "encode λ rules in governance" means formally prescribing how the weight λ is calculated and applied when combining data-driven estimates with expert priors for model parameters (here, the entanglement coefficient e_{ij}). The rule specifies a deterministic mapping from observable inputs (for example, the number of independent observations n_obs) to a numeric weight $\lambda \in [0,1]$ that controls how much trust the system places in empirical data versus prior judgment. Embedding the rule in governance makes the choice auditable, repeatable, and subject to change control.

Why this matters

- Reproducibility: A written rule prevents ad-hoc or opaque adjustments and ensures identical inputs produce identical λ values.
- Bias control: A transparent rule reduces subjective over-reliance on sparse data or on authoritative experts.
- Auditability: When ADI drives governance actions, reviewers must be able to trace why a particular e_{ij} was chosen and how much came from data versus priors.
- Adaptivity: Rules can be tuned centrally so the organization adjusts how quickly it leans on data as instrumentation improves.

The example rule explained: $\lambda = \min(1, n_{obs} / 50)$

- Purpose: convert the raw count of relevant observations n_obs into a mixing weight λ that increases with data availability and caps at 1.
- Mechanics:
  - If $n_{obs} = 0 \rightarrow \lambda = 0$ (no trust in data; rely entirely on the prior).
  - If $0 < n_{obs} < 50 \rightarrow \lambda = n_{obs} / 50$ (gradual, linear increase in reliance on data).
  - If $n_{obs} \geq 50 \rightarrow \lambda = 1$ (full reliance on data; prior is ignored).
- Intuition: the rule says, "use data proportionally to the amount of evidence, but require roughly 50 independent observations to treat the data estimate as fully reliable."

Worked examples

- $n_{obs} = 0 \rightarrow \lambda = \min(1, 0/50) = 0$. Result: $e_{ij} = 1 \cdot e_{prior}$ (pure prior).
- $n_{obs} = 10 \rightarrow \lambda = 10/50 = 0.2$. Result: $e_{ij} = 0.2 \cdot e_{data} + 0.8 \cdot e_{prior}$.
- $n_{obs} = 50 \rightarrow \lambda = 50/50 = 1$. Result: $e_{ij} = 1 \cdot e_{data} + 0 \cdot e_{prior}$ (data dominates).
- $n_{obs} = 120 \rightarrow \lambda = \min(1, 120/50) = 1$. Result: same as 50 or above; rule caps at full data trust.

Design considerations and extensions

- Choice of denominator (50) should reflect domain specifics: event rarity, signal-to-noise, and acceptable sample size for stable estimates. Use empirical calibration or power analysis to set it.
- Nonlinear mappings: consider log or square-root transforms to reflect diminishing returns of extra observations, e.g., $\lambda = \min(1, \log(1 + n_{obs})/\log(1 + N_{ref}))$.
- Confidence weighting: incorporate data quality by multiplying n_obs by a quality factor $q \in [0,1]$, giving $\lambda = \min(1, q \cdot n_{obs} / N_{ref})$.
- Hierarchical rules: allow different N_ref per edge class (critical edges require more data) or allow prior overrides for safety-critical links.
- Uncertainty propagation: rather than a point λ, store a distribution or CI for λ when n_obs is noisy and use it in Monte Carlo analyses

Encode λ rules in governance (continued)

Governance recommendations (operational)

- Document the exact formula, rationale for parameter choices, and where it is applied in the pipeline.
- Record n_obs, q, λ, e_data, e_prior, and the final fused e_{ij} in the assumption register for each update.
- Periodically validate the rule by back testing: compare fused estimates against out-of-sample events and adjust the denominator or function form if miscalibrated.
- Treat the λ rule as versioned policy: any change requires approval, impact analysis, and re-run of high-leverage ADI scenarios.

Quick checklist for adoption

- Choose a baseline mapping function and reference sample size N_ref.
- Define what counts as an independent observation (consistent event windowing and de-duplication).
- Decide whether to include a data quality factor.
- Encode the formula in the parameter registry and automate its computation.
- Log provenance and review performance on a regular cadence.

2.8.3 Post-processing and matrix hygiene

This section describes the practical, governance-facing steps you should take after estimating the entanglement matrix E = [e_{ij}]. The goal is to turn raw estimates into a reliable, auditable, and computationally tractable matrix that produces interpretable ADI outputs. For each substep I define the objective, explain the rationale, show worked examples or numeric guidance where helpful, and give governance recommendations you can action immediately.

Standardize and clip

- Purpose
    - Ensure every coefficient lies on the same, bounded scale so interpretation is consistent and propagation math behaves predictably.
- What to do
    - Enforce the constraint **e_{ij} ∈ [−1, 1]** by clipping any values outside that interval.
    - If different estimation methods produce scores on different ranges (for example, raw probabilities, correlation coefficients, or VAR impulse

responses), apply a documented, reversible mapping to a common scale before combining or comparing values.

- Example mappings
    - Probability $\hat{p} \rightarrow e$: $e = 2 \cdot \hat{p} - 1$ maps $\hat{p} \in [0,1]$ to $e \in [-1,1]$.
    - Pearson or Spearman ρ already lies in $[-1,1]$ and can be used directly.
    - VAR coefficients or impulse responses: normalize by a reference scale or by the maximum absolute response observed across candidate edges, then clip.
- Recommendations and governance notes
    - Record the mapping function and the rationale in the assumption register so transforms are auditable.
    - Prefer linear, monotonic transforms that preserve sign and ordering.
    - When clipping occurs frequently for many edges, revisit upstream estimation choices or consider re-scaling inputs rather than blanket clipping.

Sparsification and pruning weak links

- Purpose
    - Remove spurious, low-magnitude edges that create artificial propagation chains, increase computational cost, and reduce interpretability.
- What to do
    - Apply a sparsification rule: set $\mathbf{e_{ij}} \leftarrow \mathbf{0}$ when $|e_{ij}| < \tau$. Choose τ using one of the following approaches:
        - Absolute floor (practical default): $\tau = 0.05$ or 0.03 depending on noise level.
        - Significance-based: τ is the smallest $|e_{ij}|$ that is statistically distinguishable from zero at chosen confidence level.
        - Relative/adaptive: $\tau = c \cdot \text{median}(|E| \text{ nonzero})$ with c in $[0.05, 0.25]$ to preserve signal in sparse matrices.

- Optionally apply group-wise sparsification where edges for high-confidence pairs are pruned with a lower τ than low-confidence pairs.

- Worked example
  - If many edges cluster around ±0.02 and a few are ±0.4, choose $\tau = 0.05$ to remove background noise while retaining meaningful links.

- Rationale and governance tradeoffs
  - Too low τ: matrix remains noisy and ADI footprints inflate via long chains.
  - Too high τ: you may remove weak but real propagation paths, underestimating reach.
  - Governance action: codify τ selection rules, require review for edges removed that exceed a secondary scrutiny threshold (e.g., any removed edge with |e_{ij}| ≥ 0.1 before shrinkage triggers an expert check).

Sign and directionality handling

- Purpose
  - Resolve or explicitly flag uncertainty about whether influence is symmetric, asymmetric, positive, or negative so downstream users understand limitations.

- Options and procedures
  - If directionality is well supported by data or tests, keep directed values (e_{ij} ≠ e_{ji}).
  - If directionality is uncertain but you have estimates both ways, create a symmetric edge by averaging: **e_{sym} = (ê*{ij}* + *ê*{ji}) / 2** and flag the edge as undirected.
  - If sign is ambiguous (opposite signs from different methods or experts), preserve both candidate signs and record the disagreement, or introduce a three-state encoding: positive, negative, ambiguous with a confidence score.
  - Maintain explicit metadata per edge: source method, sample size, confidence level, timestamp, transformation applied, and rationale for directionality decision.

- Worked example
    - Data gives $\hat{e}_{ij} = 0.35$ and $\hat{e}_{ji} = 0.05$ with low confidence on the second estimate. Governance rule: keep directed $e_{ij} = 0.35$; set $e_{ji} = 0.05$ but label as low confidence and include in low-priority stochastic sampling only.
- Governance recommendations
    - Surface asymmetric high-impact edges for expert review.
    - For reporting, visually distinguish directed edges from undirected ones and annotate sign ambiguity on dashboards.

Normalization across projects and portfolios

- Purpose
    - Ensure ADI comparisons across projects, business units, or time are meaningful when data scales or event frequencies differ.
- Approaches
    - Option A: **Global normalization** — scale rows or the whole matrix so a standard reference (for example, the median absolute outgoing strength) matches across projects.
    - Option B: **Annotated transforms** — keep local E unscaled but store explicit normalization metadata so ADI can be computed on either local or normalized basis.
    - Option C: **Consequence-aware normalization** — adjust e_{ij} scaling by project exposure metrics so ADI comparisons reflect both coupling and absolute consequence levels.
- Worked example
    - Project A has dense event logging and many small migrations; Project B logs few large migrations. Without normalization, Project A's raw E may appear more connected. Normalize each E by its median nonzero $|e_{ij}|$ before computing portfolio ADI, or compute ADI on normalized matrices and report both raw and normalized ADI.

- Governance notes
  - Document the chosen normalization approach and show both normalized and raw ADI in executive reports when cross-project comparison is required.
  - If you normalize, include sensitivity checks to ensure normalization does not hide true systemic concentration in a single project.

2.9 Matrix hygiene and verification steps

- Confidence tagging
  - Label every nonzero edge with a confidence tier (High/Medium/Low), provenance, and last calibrated date.
- Version control and changelog
  - Treat E as a versioned governance artifact. Log each update, the driver (new data, expert update), and its impact on ADI snapshots.
- Diagnostic analytics
  - Run diagnostics after post-processing: distribution of |e_{ij}|, sparsity fraction, top outbound influence nodes, and change-impact analysis versus prior versions. Flag large shifts for review.
- Sensitivity and backtesting
  - Perform sensitivity sweeps on τ, clipping thresholds, and mapping functions to quantify how ADI, velocity, and reach respond. Backtest where historical co-migration events exist.
- Computational considerations
  - Sparse E enables efficient propagation computation. Use sparse matrix data structures and document any numerical approximations (truncation of kernels, iterative tolerance) used in ADI runs.

Practical checklist for matrix hygiene

[ ] Map raw estimates to [−1,1] and log mapping.
[ ] Clip and record any clipped values.
[ ] Apply sparsification with governance-defined τ and log edges removed.
[ ] Resolve or flag directionality and sign ambiguity with metadata.
[ ] Choose and document normalization policy for cross-project use.
[ ] Tag edges with confidence and provenance.
[ ] Version E and produce a delta report of changes.
[ ] Run sensitivity tests and backtests where feasible.

Closing insight

Post processing is where statistical estimation meets governance discipline. Clear, auditable rules for clipping, sparsification, directionality, and normalization preserve ADI's interpretability and operational value. Conservative pruning combined with targeted instrumentation for high-leverage edges keeps the matrix compact and trustworthy while enabling the organization to focus data collection and expert review where it will most improve detection, triage, and mitigation.

Validation, calibration, and uncertainty quantification

- **Backtest**: where historical events exist, simulate propagation using the estimated E and compare predicted affected sets/time to observed co-migrations. Measure precision, recall, and timing errors.
- **Sensitivity analysis**: vary e_{ij} by ±X% or sample from posterior distributions and observe ADI, velocity, reach variance to find brittle edges.
- **Monte Carlo**: sample stochastic elements (thresholds, noisy edges, kernels) to produce distributions of footprints and ADI rather than point estimates.
- **Diagnostics**: identify edges that dominate ADI outputs (high leverage) and prioritize those for data collection or expert review.

Worked example (small backtest):

- Network of 3 nodes A,B,C; estimated E yields strong e_{AB}=0.7; historical logs show 8 of 10 A events followed by B events → good calibration. If predicted B activations are far more frequent than observed, reduce e_{AB} via shrinkage.

Governance, auditability, and operational recommendations

1. **Provenance and version control**
   - Record the method, data windows, expert panel, mapping functions, and parameter choices for each e_{ij}. Version E whenever the matrix is updated.
2. **Tiering edges by confidence**
   - Label edges as High/Medium/Low confidence. Use only High+Medium in deterministic governance triggers; treat Low as exploratory in Monte Carlo sensitivity runs.
3. **Invest in targeted data collection**
   - For high-leverage edges, build instrumentation (event tagging, consistent windows, standard migration magnitude measures) so future estimates move from expert priors to data-driven.
4. **Operational thresholds for sparsification**
   - Define and governance-approve the sparsification threshold τ. Periodically review τ as data richness changes.
5. **Audit and explainability**
   - Produce a short narrative for each high-impact e_{ij} explaining why it's nonzero, how estimated, and what controls exist if it drives ADI alerts.

Practical checklist for estimating E (quick reference)

[ ] Define event windows and migration magnitude definitions consistently.
[ ] Select primary estimation method appropriate to data richness (frequency vs magnitude).
[ ] Map raw statistics to [−1,1] using documented, reversible transforms.
[ ] Apply shrinkage/hybrid fusion rules when combining data and expert priors.
[ ] Clip and sparsify with governance-approved thresholds.
[ ] Label edges with confidence metadata and provenance.
[ ] Backtest and run sensitivity analyses; adjust as needed.
[ ] Version control E and capture rationale for changes.

Estimating E is as much organizational design as statistical modeling. Data can expose robust coupling, but expert judgment and governance rules are necessary to make the matrix actionable and auditable. Use conservative sparsification to keep ADI interpretable, prioritize instrumentation for edges that materially change ADI outputs, and treat E as a living artifact: versioned, reviewed, and routinely recalibrated to reflect new evidence and changing project dynamics.

2.10 Calibration, Monte Carlo, and uncertainty quantification

This section defines how to align the ADI model with real-world behavior, quantify uncertainty in ADI outputs, and produce governance-ready confidence statements. Calibration turns raw model structure (kernels, transfer function, thresholds, and E) into a probabilistic generator whose outputs can be compared to observed propagation events. Monte Carlo and uncertainty quantification then propagate input uncertainty through the model to produce distributions for footprints, ADI, velocity, and reach rather than brittle point estimates. Together these steps convert ADI from a single deterministic metric into a traceable, auditable risk signal with measurable reliability.

Calibration objectives (what good looks like)

- Match observed propagation statistics where historical shocks exist:
  - Mean path length: average number of hops or propagation steps observed following a seed event.
  - Fraction affected: proportion of nodes that become affected in empirical post-event windows.
  - Time to spread: distribution of delays from seed event to downstream activation.
- Set priors for kernels and thresholds that reflect observed lead/lag behavior:
  - Ensure kernel shapes and decay rates produce delays and persistence consistent with logged events.
  - Ensure thresholds θ_j generate affected flags at realistic magnitudes and frequencies.
- Produce posterior parameter distributions that capture both parameter uncertainty and structural ambiguity so governance can make probabilistic decisions (e.g., “there is a 90% credible interval that ADI exceeds X under scenario Y”).

Key terms explained

- **Calibration:** the process of adjusting model parameters so that simulated outputs reproduce key statistics of historical or synthetic events.
- **Prior:** the initial probability distribution (belief) over a parameter before seeing calibration data.
- **Posterior**: the updated probability distribution over a parameter after conditioning on observed data.
- **MCMC (Markov Chain Monte Carlo):** an algorithmic family for sampling from posterior distributions when closed-form solutions are unavailable.
- **Monte Carlo simulation**: repeated random sampling from parameter and stochastic model distributions to produce an empirical distribution of model outputs.
- **Credible interval:** Bayesian analogue of a confidence interval; e.g., a 90% credible interval contains 90% of the posterior mass for a metric.
- **Coverage:** the frequency with which true values fall inside estimated intervals in repeated trials (used to validate calibration).
- **Synthetic injection experiment:** deliberately simulating seeded events with known properties to test whether the calibrated model reproduces expected propagation and interval coverage

Recommended calibration practice — expanded step-by-step guide

This section turns the calibration workflow into a practical, auditable playbook. Each step explains purpose, concrete actions, design choices, worked examples, and governance checks so teams can implement reproducible Bayesian[45] calibration and posterior-driven Monte Carlo for ADI.

Step 1. Define calibration targets and data windows

Purpose

- Identify the observable quantities the model must reproduce so calibration is focused, testable, and aligned with governance needs.

What to do

- Select 3–5 summary statistics that capture core propagation behavior. Typical targets:
    - **Mean path length**: average number of hops from seed to affected nodes.
    - **Fraction affected**: proportion of nodes flagged within a fixed horizon.
    - **Time-to-threshold quantiles**: e.g., median and 90th percentile time-to-first-impact.
    - **Distributional moments**: variance of footprint or ADI across events.

[45] Bayesian prediction is distributional.

- Define event windows clearly (e.g., 30-day, 90-day) and the seeding rule (what constitutes a seed event, de-duplication, and exclusion windows).
- Ensure the same event definitions were used when estimating E, mapping g(·), and thresholds θ_j to avoid mismatched assumptions.

Worked tip

- If historical shocks are heterogeneous, stratify targets by seed type (e.g., vendor failure vs. scope change) and calibrate separate parameter subsets or hierarchical priors.

Governance check

- Log the chosen targets, window lengths, and event inclusion criteria in the calibration record.

Step 2. Specify model priors grounded in domain knowledge

Purpose

- Encode plausible parameter ranges so calibration is stable, interpretable, and conservative where data are weak.

What to do

- For each parameter family choose priors that reflect operational reality:
  - **Kernel parameters** (scale α, decay β, tail γ): set priors tied to temporal rhythms (e.g., review cycles, lead times). Example: β ~ Exponential(λ) with mean matching typical decay per review cycle.
  - **Transfer function parameters** (logistic slope α_g, midpoint x0, saturation γ): use priors that prevent unrealistic super-linear amplification[46] (e.g., α_g ~ Normal(μ, σ) with μ near 1).

[46] **Super-linear amplification** refers to a dynamic where small changes in input lead to disproportionately large changes in output—especially when the amplification factor (α_g) is drawn from a distribution centered near 1, such as α_g ~ Normal(μ, σ) with μ ≈ 1.
In systems modeling, **α_g** often represents a *governance amplification factor*, risk multiplier, or influence coefficient. When it's modeled as a **random variable from a normal distribution**, like:

$$\alpha_g \sim \text{Normal}(\mu, \sigma)$$

and **μ is close to 1**, the system is poised near a **critical threshold**:

- **If α_g < 1:** The system dampens deviations—risk or influence decays over time.
- **If α_g > 1:** The system amplifies deviations—risk or influence grows exponentially.

When μ ≈ 1, even small fluctuations (due to σ) can push α_g above 1, triggering **super-linear amplification**—a regime where **effects grow faster than inputs**.

    - **Thresholds θ_j**: hierarchical priors centered on empirical percentiles (e.g., group mean at 90th percentile) with node-level variance reflecting domain heterogeneity.
- Use informative priors when expert knowledge is reliable; use weakly informative priors when uncertainty is high.

Worked example

- If procurement lead times average 14 days, choose kernel decay such that median effective lag ≈ 14 days (translate to prior on β). Document mapping.

Governance check

- Record prior choices, rationale, and alternative priors considered. Store prior predictive checks showing simulated behavior before seeing data.

Step 3. Estimate posterior distributions using Bayesian calibration (MCMC recommended)

Purpose

- Produce a posterior distribution over parameters that reflects both data evidence and prior beliefs, preserving parameter dependencies.

What to do

- Define a likelihood that links model outputs to calibration targets:
    - Counts (number affected) → **Poisson** or **Binomial** likelihood.
    - Continuous summaries (mean path length) → **Gaussian** likelihood with estimated observation noise.
    - Quantiles → use likelihoods based on empirical cumulative distributions or asymmetric losses.
- Choose a sampler (e.g., Hamiltonian Monte Carlo, NUTS, or well-tuned Metropolis) and run sufficient chains/samples. Typical guidance: 4 chains, 1,000–2,000 warmup + 2,000–5,000 posterior draws, adjusted to effective sample size needs.

Sampler	Core Mechanism	Best For	Key Advantage
HMC	Physics-based trajectories	High-dimensional models	Efficient exploration

Sampler	Core Mechanism	Best For	Key Advantage
NUTS	Adaptive HMC with no U-turns	Bayesian inference	No manual tuning
Metropolis	Random walk proposals	Simple models	Easy to implement

- Monitor diagnostics: trace plots, $\hat{R} \approx 1.00$, effective sample size (ESS) per parameter, and posterior pair plots to reveal strong correlations.
- Preserve joint posterior structure: export joint samples to drive forward simulations rather than fitting marginal distributions separately.

Worked checklist

- Run prior predictive checks, run MCMC, inspect diagnostics, and if mixing is poor adjust parameterization or rescale priors.

Governance check

- Archive sampler settings, random seeds, convergence diagnostics, and final posterior sample set for audit.

Step 4. Run Monte Carlo forward simulations across posterior samples

Purpose

- Propagate parameter uncertainty through the stochastic propagation model to produce distributional ADI outputs that support probabilistic decision making.

What to do

- For each posterior draw:
  - Run the discrete propagation model multiple times if there is intrinsic stochasticity (e.g., randomized thresholds, noisy edges). Suggested: 100–1,000 replicates per draw for stable tail estimates, fewer for exploratory runs.
  - Collect metrics: **Footprint_i(H)**[47], **ADI(H)**, **V_i(θ)**, **R_i(H)**, and their time series distributions.

[47] A source's footprint over horizon H is the consequence-weighted cumulative exposure its change produces in the portfolio (probability-weighted where stochastic). Footprints are the basis for ADI ranking and triage; see Section 2.5.4 and dashboard examples in Appendix D

- Aggregate results to produce medians, means, and credible intervals (e.g., 50% and 90%). Visualize using fan charts, violin plots, and empirical CDF overlays.

Worked example

- 1,000 posterior draws × 200 stochastic replicates → 200k forward runs. Compute median ADI(H) and the 5th/95th percentile band for dashboarding.

Governance check

- Set computational budgets and a reproducible pipeline that logs which posterior draws were used for which dashboard snapshot.

Step 5. Validate with synthetic injection and backtesting

Purpose

- Demonstrate posterior predictive validity, check interval coverage, and identify model misspecification or overconfidence.

What to do

- Synthetic injection experiments:
  - Inject seeded events with known properties into historical network snapshots and simulate forward draws from the posterior to check whether credible intervals cover observed propagation behavior.
  - Measure **coverage** (fraction of injected trials where observed metrics fall inside predicted intervals) and tune model if coverage deviates from nominal levels.
- Backtesting:
  - Hold out a subset of historical seed events. Calibrate on the remainder and predict held-out events. Compute performance metrics: precision, recall, F1, mean absolute timing error, and calibration curves.
- Failure analysis:
  - Where predictions fail, inspect which parameters, edges, or thresholds contributed most to the mismatch and prioritize data collection or model refinement.

Worked example

- Calibrate on 80% of events, predict the remaining 20%. If observed fraction affected falls outside the 90% credible interval in 40% of held-out events, investigate underestimation of kernel tails or underestimated edge strengths.

Governance check

- Define acceptance criteria (e.g., 90% credible intervals should have ≥85% empirical coverage) and record remediation actions when criteria are not met.

Diagnostics, sensitivity, and reporting

- Sensitivity sweeps: vary key scalar parameters (kernel decay, transfer slope, threshold offsets) across plausible ranges and report how ADI and velocity metrics change. Present tornado plots to identify high-leverage parameters.
- Uncertainty decomposition: attribute variance in ADI outputs to parameter groups (kernels vs. transfer vs. E estimation) using ANOVA-style or variance-based methods.
- Reporting format: always present a central estimate plus at least two uncertainty bands (50% and 90%) and supply downloadable posterior samples for audit.

Governance check

- Produce an executive summary with key calibration choices, coverage diagnostics, and prioritized data collection actions derived from sensitivity analysis.

Practical recommendations and tooling notes

- Automate pipelines: implement repeatable workflows that run calibration, posterior simulation, diagnostics, and report generation with versioned inputs.
- Use scalable compute: posterior-driven Monte Carlo can be expensive; schedule runs and provide incremental dashboards (quick exploratory runs vs. full production runs).
- Keep human-in-the-loop: require SME signoff for major prior changes and incorporate domain review after the first calibration cycle.
- Incremental updates: treat calibration as periodic (monthly or per milestone) rather than ad-hoc; record deltas and re-evaluate impact on governance triggers.

Closing insight

Well-executed calibration translates historical evidence and expert judgment into a probabilistic model that communicates not just what the ADI is, but how confident you should be in it. A disciplined sequence—clear targets, defensible priors, rigorous posterior sampling, comprehensive forward Monte Carlo, and thorough validation—turns ADI into a transparent, auditable decision instrument that guides risk triage, resource allocation, and continuous learning.

Diagnostics and uncertainty reporting — expanded guidance

This section describes a practical, auditable diagnostics toolkit for ADI calibration and posterior reporting. The goal is to quantify how much trust to place in ADI outputs, detect model weaknesses, and prioritize data collection and model refinement where it matters most.

Bootstrapped confidence intervals and posterior credible intervals

- What they are and why both matter
  - **Bootstrap confidence intervals** measure sampling variability in empirical calibration targets by resampling observed seed events; they show how sensitive summary statistics are to which events were observed.
  - **Posterior credible intervals** come from the Bayesian posterior predictive distribution[48] and quantify uncertainty conditional on the model, priors, and data.
  - Reporting both highlights whether uncertainty is dominated by data scarcity (bootstrap) or by model / prior structure (posterior).
- How to compute and present them
  - Bootstrap: resample seed events with replacement (e.g., 1,000 replicates), recompute target statistics and downstream metrics (Footprint, ADI), and report percentile bands (50%/90%).
  - Posterior: run forward simulations across posterior parameter draws (and stochastic model noise), then compute percentile bands for the same metrics.
  - Present side-by-side bands and a short interpretive note: if posterior bands are much wider than bootstrap bands, model structure or priors are a key source of uncertainty.
- Governance action
  - Store both sets of intervals and the code/seed used to produce them for auditability.

Sensitivity analysis (identify brittle parameters)

- Purpose
  - Find which parameters (kernel decay β, kernel scale α, transfer slopes, threshold offsets, elements of E) most strongly change ADI, velocity, and reach when perturbed.
- Practical approach

[48] The Bayesian posterior predictive distribution is the probability distribution of future or unobserved data, given the data you've already observed. It integrates over uncertainty in the model parameters using the posterior distribution.

  - Local sensitivity: perturb one parameter at a time (±10–30%) and report relative change in median ADI and tail quantiles.
  - Global sensitivity: sample parameter space (Sobol, Morris, or variance-based methods) to attribute variance in outputs to parameter groups.
  - Visualizations: tornado plots for one-at-a-time impacts; heatmaps or Sobol indices for global attribution.
- Interpretation and prioritization
  - Parameters that cause large output swings with small perturbations are **high-leverage** and should be prioritized for better data, stronger priors, or conservative governance controls.
  - Document which parameters are brittle and the remediation plan (data collection, expert re-elicitation, conservative defaults).

Cross-validation and out-of-sample checks

- Why it's essential
  - Cross-validation tests whether the calibrated model generalizes beyond the events used for fitting and guards against overfitting to idiosyncratic historical shocks.
- How to run it
  - Where multiple independent seed events exist, use k-fold, leave-one-out, or time-blocked holdouts consistent with event dependence. Calibrate on the training folds, predict the held-out events, and compute predictive metrics (precision, recall, mean timing error, coverage).
  - Aggregate metrics across folds and report mean ± variance.
- Governance thresholds
  - Define acceptable performance bands in advance (e.g., median timing error ≤ X days; 90% credible interval coverage ≥ 0.85) and require remediation if thresholds are missed.

Coverage, calibration checks, and diagnostic plots

- Coverage testing
  - For each level X% (e.g., 50%, 90%), measure the empirical fraction of held-out or injected events whose observed metric (e.g., nodes affected) falls inside the predicted X% credible interval. Well-calibrated models have empirical coverage ≈ nominal coverage.
- PIT and calibration histograms
  - Use Probability Integral Transform (PIT) histograms or rank histograms to detect overdispersion (U-shaped PIT) or underdispersion (peaked PIT). These indicate the posterior predictive spreads are too narrow or too wide.
- Posterior predictive checks (PPC)
  - Generate replicated datasets from the posterior predictive distribution and compare summary statistics and distributional shapes to real data. Highlight mismatches and their likely parameter or structural causes.

Worked example — compact, reproducible scenario

- Setup
  - Observed: 40 independent seed events of type A; observed mean affected = 12 nodes (30-day window); median time to first downstream activation = 3 days.
  - Priors: kernel scale α ~ Gamma(2,1); decay β ~ Exp(0.5); logistic slope α_g ~ Normal(1,0.5); thresholds θ_j centered at the 90th percentile where available, otherwise Uniform.
- Calibration and diagnostics workflow

  1. Run MCMC and obtain posterior draws for α, β, α_g, θ_j.
  2. Posterior predictive: for 1,000 posterior draws run 200 stochastic forward simulations each → empirical distribution of ADI(H) and Footprint_i(H).
  3. Bootstrap: resample the 40 events 1,000 times to produce bootstrap bands for mean affected and timing.
  4. Cross-validate: hold out 20% of events, calibrate on 80%, predict held-out set; compute precision/recall and timing errors.
  5. Coverage check: inject 10 synthetic seeds with known coupling; count how often observed affected counts fall in the 90% predictive interval (target ≈ 90%).

- Possible outcomes and actions
  - If empirical coverage << nominal → posterior predictive too narrow; options: widen priors, add process noise, or revise model structure.
  - If sensitivity analysis shows ADI dominated by one edge in E → prioritize data collection or expert review for that edge.

Practical recommendations and governance rules

- Report distributions, not only point estimates
    - Dashboards should show medians with 50% and 90% bands and an explanatory sentence about the dominant source of uncertainty.
- Version and provenance everything
    - Track calibration dataset, priors, posterior samples, bootstrap seeds, and forward simulation seeds in a versioned registry for reproducibility and audit.
- Tiered validation cadence
    - Daily: automated synthetic injection health checks to detect pipeline regressions.
    - Monthly: backtesting on newly logged events and recalibration if performance drifts.
    - Quarterly: governance review of calibration choices, thresholds, and high-leverage parameters.
- Avoid false precision
    - Where data are weak, report wider intervals and flag decisions that would rely on narrow tails. Use conservative operational rules for escalation until uncertainty is reduced.
- Act on sensitivity results
    - Focus measurement and instrumentation efforts on parameters and edges with the largest influence on ADI variance. Resource data collection to reduce the largest sources of uncertainty.
- Maintain a calibration playbook
    - Codify targets, acceptable diagnostics thresholds, sampler settings, and reporting templates so calibration is repeatable and auditable.

Closing insight

Robust diagnostics and transparent uncertainty reporting are the governance safeguards that make ADI defensible and operational. Presenting credible intervals, sensitivity bands, cross-validation results, and coverage diagnostics side by side with central estimates lets decision makers see both signal and its reliability. That clarity reduces misinterpretation, prioritizes the right data investments, and ensures ADI informs timely, risk-sensitive governance actions rather than spurious alarms.

2.11 Dashboard and reporting recommendations

Dashboard and reporting recommendations

A purpose-built ADI dashboard translates model outputs into operational signals for governance, triage, and strategic decision making. The following guidance describes a concise set of visual tiles, interaction behaviors, metric definitions, and reporting language that together make ADI actionable and auditable for executives and practitioners. Each tile is explained (what it shows, why it matters, how to read it), with practical display and governance suggestions.

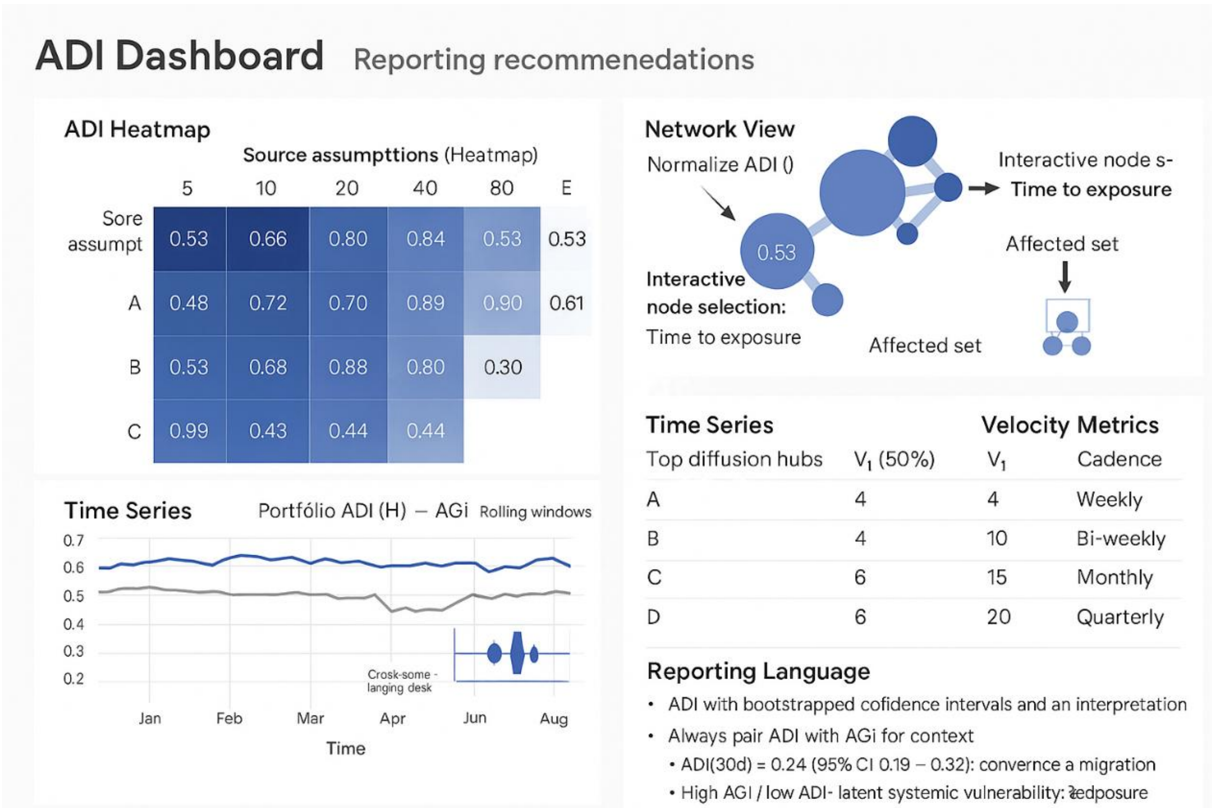

Essential tiles (visuals and interpretation)

ADI heatmap

- What it shows: rows = **source assumptions**; columns = discrete horizon steps (τ = 0..H) or cumulative footprint buckets (e.g., 0–7d, 8–30d, 31–90d). Cell color = normalized footprint or per-source contribution to ADI_i(H).
- Why it matters: quickly surfaces which sources generate early vs. late systemic exposure and enables temporal prioritization.
- How to read: bright rows early in the horizon indicate fast propagators; rows with broad brightness across columns indicate long-tail reach.
- Display tips: include tooltip with numeric Footprint_i(H), reach R_i(H), and confidence bands when hovered.

Network view (interactive)

- What it shows: graph of assumptions; **node size = ADI_i** (or Footprint_i(H)), **node color = AGI contribution** or consequence tier, **edge width = |e_{ij}|** (post-sparsification).
- Why it matters: reveals structural drivers of diffusion—hubs, bridges, and bottlenecks—and links surface where interventions (contract clauses, controls) will most reduce spread.
- Interaction features: click a node to show time-to-exposure curve, affected set list, top outbound edges with confidence scores, and example historical co-migration events. Enable filtering by consequence band, confidence tier, or edge type.
- Governance tip: display edge provenance (data vs. expert) on demand and flag high-leverage edges for review.

Time series — portfolio ADI vs AGI

- What it shows: temporal plot of portfolio ADI(H) alongside AGI or portfolio consequence metric, with rolling windows (e.g., 7/30/90d) and a cross-correlation inset.
- Why it matters: demonstrates leading/lagging relationships—when ADI rises ahead of AGI, ADI acts as an early-warning indicator.
- Display tips: show shaded credible bands (50%/90%) for ADI, and annotate causal events (seed occurrences, major rebaselines).

Velocity table (actionable triage)

Top diffusion hubs

	V50%	V80%	Cadence
Hub 123	2.4 hr	4.7 hr	
Hub 234	2.9 hr	4.9 hr	
Hub 45	3.4 hr	6.8 hr	
Hub 1	4.1 hr	8.5 hr	Daily
Hub 108	18 hr	36 hrs	
Hub 91	48 hrs	240 hrs	Monthly

Units: hr (use days for > 72 hr)

- What it shows: ranked table of diffusion hubs with V_i(50%), V_i(80%) (time to reach 50% and 80% of that source's footprint) plus suggested cadence actions (e.g., "Immediate assessment", "Weekly monitor", "Monthly review").
- Why it matters: helps operations prioritize where to act now vs. monitor over time.
- Governance rules: set automated triggers that convert V thresholds into task items or escalations.

Advanced tiles and interactions

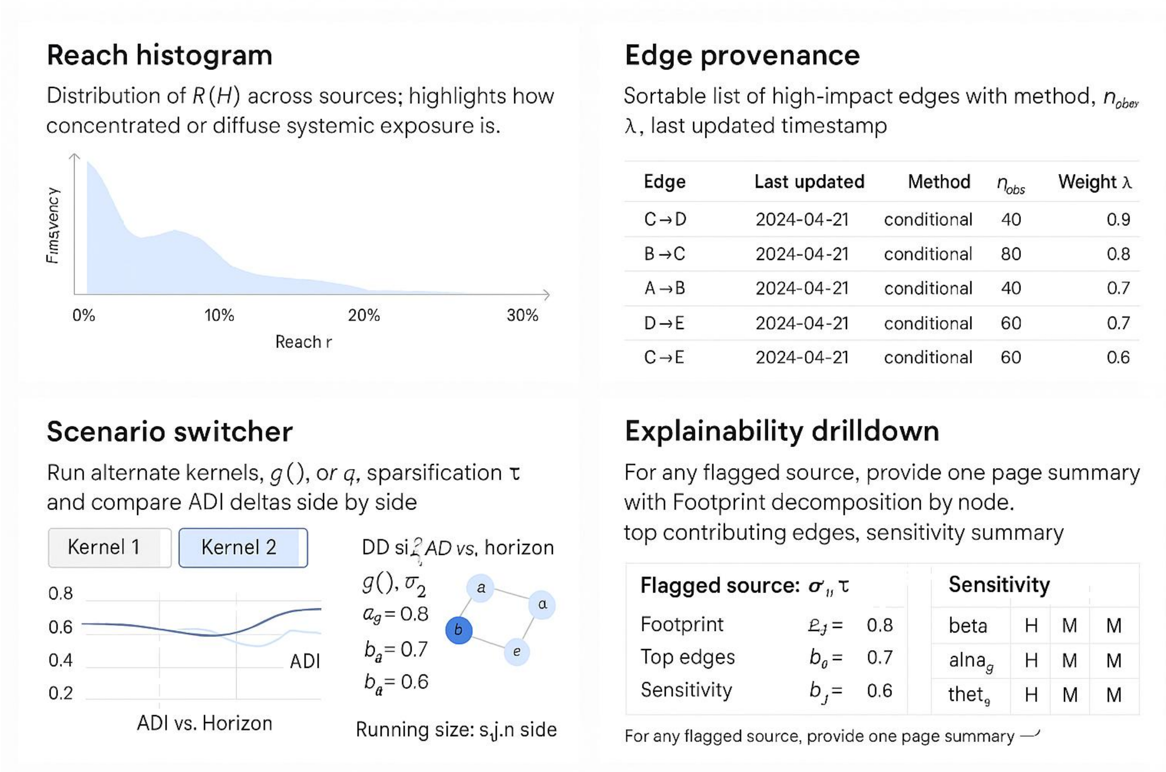

- **Reach histogram:** distribution of R_i(H) across sources; highlights how concentrated or diffuse systemic exposure is.
- **Edge provenance panel:** sortable list of high-impact edges with method, n_obs, λ weight, last updated timestamp.
- **Scenario switcher:** run alternate kernels, g(·), or sparsification τ and compare ADI deltas side-by-side.
- **Explainability drilldown**: for any flagged source, provide a one-page summary with Footprint decomposition by node, top contributing edges, and sensitivity summary.

Reporting language and framing (how to communicate results clearly)

- Pair every numeric tile with concise interpretation and confidence:

  - Example: **ADI(30d) = 0.24 (95% CI 0.18–0.32)** — "Under current assumptions and posterior uncertainty, a seeded migration typically exposes 24% of portfolio consequence within 30 days; the 95% credible interval reflects model and data uncertainty."
- Contextualize with AGI: always show AGI and explain joint regimes:
  - **High ADI / Low AGI** → latent systemic vulnerability that can escalate if left unchecked.
  - **High AGI / Low ADI** → concentrated financial or operational exposure in a few assumptions; targeted controls may be efficient.
- Use governance phrasing for actions: suggest precise next steps (e.g., "Escalate to engineering for targeted review within 48 hours," "Add to weekly monitoring list," or "Instrument event tagging for data collection").

Unified governance and operational rules (procedural)

1. Snapshot provenance and visibility (must)
   - Every dashboard snapshot displays Snapshot metadata: E version ID, estimation method for E, kernel and transfer parameters (g(·)), threshold policy, calibration date, and the posterior sample tag used for the forward Monte Carlo.
   - Tooltip or modal exposes the calibration branch, λ rules used to blend priors/data, and top data sources (tables, time windows).
2. Automated health checks and cadence (automated / scheduled)
   - Daily: automated synthetic injection health checks (small test seeds) run to verify pipeline integrity and detect unexpected regressions.
   - Weekly: operational summary (top 10 ADI sources, changes in high-leverage edges, data drift indicators).
   - Monthly: backtesting and recalibration candidate review (if new events materially change posteriors).
   - Quarterly: governance review — policy, threshold tuning, instrumentation investments, and approval of major model changes.
3. Alert thresholds, escalation gates, and SME involvement
   - Default conservative triggers (examples to adopt and tune):
     - ADI median(H) > T1 or ADI 90th percentile(H) > T2 → generate an advisory ticket. Example defaults: T1 = 0.15, T2 = 0.30 (adjust to portfolio risk appetite).
     - ADI_i(H) in top decile for the portfolio → automatic triage flag for cross-functional review.
     - ADI median > T1 AND AGI > G1 → mandatory rebaseline planning and escalation to governance board. Example: G1 = top 10% AGI.
   - SME signoff: any automatic escalation that would trigger resource reallocation, public disclosure, or contractual action requires SME confirmation within the escalation workflow.

4. Triage and monitoring rules (operational playbook)
   - Triage: immediately flag sources in the ADI top decile for a rapid assessment checklist (see Checklist section).
   - Monitoring cadence by velocity: if $V_i(50\%) < V_{thresh}$ (operational example: 7 days) increase monitoring cadence (daily review and telemetry) and assign an owner. Use $V_i(80\%)$ to decide medium-term cadence (weekly vs. bi-weekly).
   - Prioritization: weight triage priority by combined score = f(ADI_rank, AGI_tier, confidence_tier). Example: combine ranks (ADI percentile × 0.6 + AGI percentile × 0.3 + confidence_adjustment × 0.1).
5. Control design and edge remediation
   - Targeted controls should focus on high $|e_{ij}|$ edges with high ADI contribution and sufficient confidence. Controls can be software validation gates, contractual guardrails, increased testing, additional approval steps, or temporary throttles.
   - For each proposed control, require an RACI: who will implement, who will verify, and rollback criteria. Record expected effect on e_{ij} (quantified target reduction) and monitoring plan to measure efficacy.
6. Escalation triggers and rebaseline planning
   - Require rebaseline planning (operational plan to change assumptions, thresholds, or exposures) when both ADI and AGI exceed their configured escalation bands simultaneously. Rebaseline plan must include timeline, resource estimate, and impact projection using the current posterior (Monte Carlo) distribution.
   - For urgent cryptic spikes (rapid ADI rise with low data confidence), require short term containment (isolate, monitor) while collecting extra evidence; do not escalate to irreversible contractual action until SME review.
7. Audit trail and documentation (mandatory)
   - Log: every dashboard inspection, triage action, owner assignment, control action, and final decision with timestamps and actor IDs.
   - Canonical register entry for each incident must include seed event(s), propagation trace, E version and parameter set, posterior sample tag, bootstrap and posterior intervals shown, control actions taken, and outcome measures.
   - Retain artifacts for audit: code/seed used to generate Monte Carlo runs, random seeds, and visual snapshots.

Tactical templates and examples

- Example triage checklist for an ADI top-decile source:
    1. Confirm seed event and de-duplication.
    2. Inspect top 5 outbound edges and their provenance (n_obs, λ, confidence).
    3. Check V_i(50%) and V_i(80%); assign monitoring cadence.
    4. If ADI×AGI above escalation band → prepare rebaseline packet.
    5. Propose immediate mitigations for top edges; estimate expected e_{ij} reduction.
    6. Log actions and assign owner with SLAs.
- Example escalation thresholds (starting defaults, tune to org):
    - Advisory: ADI median(H) > 0.15 or ADI 90th(H) > 0.30.
    - Operational review: ADI median(H) > 0.25 or ADI_i in top 5% combined with AGI in top 10%.
    - Rebaseline escalation: ADI median(H) > 0.35 and AGI > 90th percentile.
- Example RACI for a control on edge i→j:
    - Responsible: Engineering lead for i or j.
    - Accountable: Program manager.
    - Consulted: SME who estimated e_{ij}.
    - Informed: Governance board and affected stakeholders.
    - Expected outcome: reduce e_{ij} from 0.6 to ≤0.2 within X weeks.

Confidence and conservative defaults

- Confidence tiers: tag edges and source ADI estimates as High / Medium / Low based on n_obs, posterior variance, and expert agreement. Use different operational thresholds for triggers depending on confidence (e.g., require higher ADI to escalate if confidence is Low).
- Conservative principle: when in doubt, prefer temporary containment and increased monitoring over irreversible decisions.

Implementation checklist (one-page quick start)

[ ] Embed Snapshot metadata visibly on dashboard.
[ ] Configure default alert thresholds (T1, T2, AGI band) and SME signoff flow.
[ ] Implement daily synthetic injection health check and weekly summary job.
[ ] Define V_thresh and monitoring cadences; wire automated owner assignments.
[ ] Create triage checklist and RACI template; pre-assign roles.
[ ] Build canonical register with required fields and versioning.
[ ] Define quarterly governance review agenda and decision gates.

Worked example (illustrative)

- Tile output: Heatmap shows Source A bright at 0–7d and fading later; network highlights A as a hub with large outgoing edges to cost nodes. Velocity table: V_A(50%) = 2d, V_A(80%) = 6d. Interpretation: "Source A is a fast propagator likely to affect high consequence nodes within a week—recommend immediate engineering review and temporary controls; instrument related edges for data collection."

Quick dashboard checklist

[ ] Heatmap with cumulative buckets and hover details.
[ ] Interactive network with provenance and confidence filters.
[ ] Portfolio ADI vs AGI time series with credible bands.
[ ] Velocity table with action recommendations.
[ ] Scenario switcher and explainability drilldowns.
[ ] Snapshot metadata (versions, priors, calibration date) and audit logs.

Closing insight

A disciplined dashboard turns ADI from a model artifact into a governance instrument: clear visuals, explicit uncertainty, provenance, and action-oriented language ensure decision makers see not just risk signal but also what to do about it—and why they can trust the recommendation.

Simplified Assumption Diffusion Dashboard

Linked to Assumption with AGI Parameters

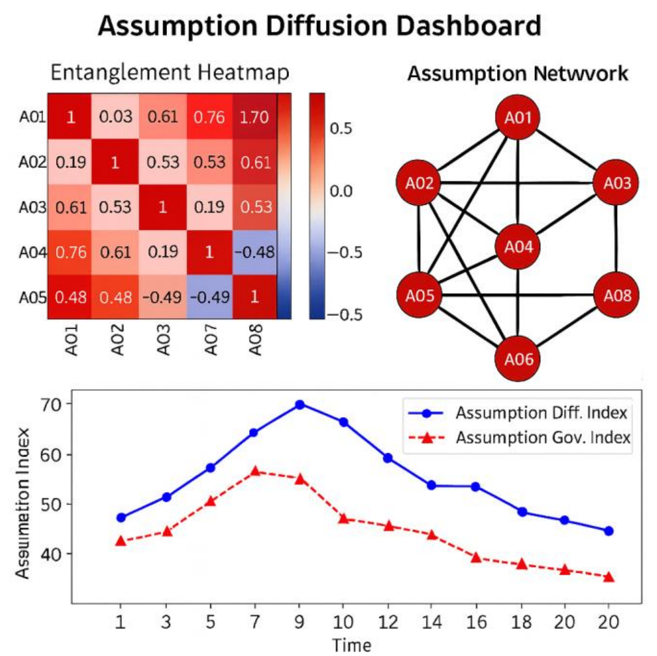

Assumption Mapping and Scores

ID	Category	BaselineValue	CurrentValue	Units	DeltaCrit	α0	λ (per month)	ω cost	ω safety	ω schedule	ω reputation	Entanglement Group
A1	EconomicCost	1,000,000	1,120,000	USD	100,000	0.85	0.015	8	0	6	2	ECON_COMMOD
A2	Performance	100	92	%prod	6	0.8	0.02	0	0	9	1	PROD
A3	Technical	0.95	0.93	availability	0.02	0.9	0.01	2	8	4	1	TECH
A4	Environmental	Permitted	Delayed	categorical	mapped dist=0.6	0.75	0.03	1	3	2	1	PERMIT
A5	Client	FundingOnTime	Delayed	categorical	mapped dist=0.8	0.7	0.04	2	1	3	6	CLIENT
A6	EconomicCost	FuelPriceIndex=1	FuelPriceIndex=2.5	index	0.5 index	0.8	0.02	6	0	2	1	ECON_COMMOD
A7	Stakeholder	Sentiment=0	Sentiment=-1	ordinal dist	0.5	0.85	0.01	1	4	3	2	STAKE
A8	Performance	CrewAvailability=1.0	0.85	fraction	0.1	0.9	0.03	0	0	7	1	HR

2.12 Validation, limitations, and governance caveats

This section makes explicit the assumptions and failure modes of the Aggregated Diffusion Index (ADI), explains why each matters, illustrates consequences with brief examples, and gives concrete mitigation and governance actions so ADI remains reliable and defensible in practice.

Dependence on the entanglement matrix E

- What this means
    - ADI propagation and footprint outputs are driven by the entanglement coefficients e_{ij}. If those coefficients are wrong, ADI will propagate shocks along incorrect paths or with incorrect intensity.
- How errors show up
    - Overstated e_{ij} on peripheral edges creates false positive systemic chains and inflated ADI. Understated e_{ij} on real hubs hides true systemic risk and produces false negatives.
- Worked example
    - If a true hub has e_true = 0.6 but your estimate ê = 0.1 (due to sparse data), ADI may under-prioritize that hub and delay mitigation actions until consequences materialize.
- Mitigations and governance
    - Tag every edge with provenance and a confidence tier; require high-leverage edges (large ADI contribution or centrality) to reach a minimum confidence before driving automatic escalations.
    - Prioritize instrumentation and data collection for edges that contribute most to ADI variance (identified via sensitivity analysis).
    - Use hybrid estimates (data + expert priors) with explicit λ rules and record the rule used for each e_{ij}.

Model specification risk: kernels and transfer functions

- What this means
  - ADI uses temporal kernels (how influence decays over time) and transfer functions g(·) (how source movements translate to target response). Mis specifying their shape or scale biases velocity measures and timing predictions.
- Why it matters
  - Velocity ($V_i(p)$) and time-to-threshold metrics are sensitive to kernel tails and transfer nonlinearity; an exponential kernel fit when real decay is heavy-tailed will understate late propagation and overstate early speed.
- Worked example
  - A model using a fast exponential kernel predicts $V_i(80\%) = 3$ days; a heavy-tailed kernel consistent with data predicts $V_i(80\%) = 12$ days — operational cadence and escalation differ materially.
- Mitigations and governance
  - Calibrate kernels and g(·) where historical propagation exists; perform prior predictive checks and posterior predictive checks.
  - Use model selection diagnostics (e.g., information criteria on fitted VAR/coupling models, or out-of-sample coverage) to distinguish plausible kernel families.
  - Report velocity with confidence bands and avoid single-number cadence triggers unless supported by narrow posteriors.

Data sparsity and expert elicitation

- What this means
  - Many domains lack abundant, clean migration logs. Where historical data are sparse or biased, expert judgment must supplement estimates.
- Why it matters
  - Expert inputs introduce subjectivity and can create overconfidence or inconsistent scaling across edges if not structured.
- Best practices for elicitation
  - Use structured protocols (Cooke's method, seed questions, calibration exercises) and translate elicited probabilities or directional judgments to the same $[-1,1]$ scale used for data estimates.
  - Capture full uncertainty ranges (min/median/max); encode these as priors rather than point estimates.
  - Weight experts by calibration performance when historic test data exist.
- Conservative uncertainty handling
  - Where data are sparse, widen priors, increase shrinkage toward zero, and surface wide credible intervals in reports and triggers. Use conservative operational defaults for escalations when confidence is low.

Complementarity with other risk analyses

- What this means
    - ADI measures propagation potential through the assumption network; it does not by itself measure tail risk of direct losses, scenario severity, or conditional portfolio loss metrics.
- How to combine analyses
    - Use ADI together with percentile metrics (e.g., 95th percentile footprint), conditional expected shortfall (CES), and scenario-based stress tests that impose extreme seeds or compound shocks.
    - Treat ADI as a leading indicator of systemic spread and use CES/portfolio loss metrics to quantify consequence magnitude and remediation cost.
- Example workflow
    - If ADI(30d) rises materially for a source, run targeted scenario stress tests seeded at that source to compute CES and prioritized remediation cost estimates.

2.13 Practical mitigations, reporting and checklist

- Real-time safeguards
    - Never rely on a single ADI point estimate for irreversible actions. Require SME confirmation for escalations that carry financial, contractual, or public impact.
    - Use multi-tier alerts: advisory (monitor), operational (assign owner), and escalation (rebaseline/contract action) driven by jointly calibrated ADI/AGI bands and confidence tiers.
- Validation and monitoring
    - Backtest ADI on historical seeds, run synthetic injection checks, and report empirical coverage of credible intervals; treat systematic under- or over-coverage as a call to revise model structure or priors.
- Documentation and auditability
    - For every decision driven by ADI, capture the dashboard snapshot ID, E version, posterior sample tag, trigger thresholds, SME approvals, and executed controls in the canonical register.
- One-page checklist for safe use

[ ] Is edge confidence high for top contributors?
[ ] Were kernels and transfer functions calibrated for the relevant seed class?
[ ] Are ADI outputs presented with 50%/90% bands and bootstrap checks?
[ ] Is an SME assigned before automated escalation?
[ ] Is the triggering rule documented and versioned?

2.14 Reproducibility Checklist and Closing summary

ADI is a practical, governance-ready signal for propagation risk that converts networked assumption relationships into a reproducible, auditable indicator of how shocks can spread through a portfolio. Its chief value is early detection: by surfacing where coupling (E), timing (kernels), and thresholds combine to produce broad reach or rapid velocity, ADI steers limited mitigation effort to the places that most reduce systemic fragility.

Key strengths

- Early warning: highlights sources that are likely to propagate before consequences (AGI) rise.
- Governance alignment: uses the same consequence scale as AGI and records versioned provenance so outputs are auditable.
- Operational focus: supports triage, cadence decisions, and targeted controls by combining ADI rank, AGI magnitude, velocity metrics, and edge confidence.

Primary limits and guardrails

- Garbage-in/garbage-out: ADI's correctness depends directly on the quality of e_{ij}, kernel and transfer specification, and input data.
- Model risk: mis specified kernels or transfer functions bias timing and magnitude predictions; present velocity and ADI with credible bands, not single values.
- Data sparsity: where logs are thin, rely on structured elicitation and conservative uncertainty (wider priors, stronger shrinkage).
- Complementarity: ADI signals diffusion potential but does not replace tail-risk, CES, or scenario stress testing.

Practical operating rules (essence)

- Surface uncertainty: always present ADI as a distribution (median + 50%/90% bands) alongside bootstrap checks.
- Prioritize measurement: use sensitivity analysis to target instrumentation for edges and parameters that dominate ADI variance.
- Govern triggers: codify snapshot metadata, escalation bands, SME sign-off, and an audit trail before any irreversible action.
- Integrate analyses: on ADI upticks, run targeted scenario and CES runs seeded at high-ADI sources to quantify consequence and cost.

How to deploy it in practice

- Implement ADI in standard data stacks: version E and calibration artifacts, automate daily health checks and synthetic injections, and precompute posterior-driven Monte Carlo snapshots for interactive dashboards.
- Operationalize playbook: embed triage checklists, velocity-based cadences, RACI templates for controls, and quarterly governance reviews into routine ops.
- Iterate: treat E and calibration as living artifacts—recalibrate when new events arrive, reassign priority to edges with persistent uncertainty, and record every decision and its outcomes.

Final perspective - Reproducibility

When governed and instrumented, ADI transforms an entangled set of assumptions into a decision-grade signal: it helps detect latent systemic vulnerabilities earlier, focus controls where they matter most, and link propagation potential to consequence-based decisions. The combination of reproducible method, explicit uncertainty, and clear governance rules preserves ADI’s early-warning advantage while preventing over-reliance on any single estimate or automated trigger.

Reproducibility checklist

For any published ADI results, archive and publish:

- Canonical CSV with assumption metadata, consequence weights, and migration histories.
- Entanglement matrix E and sparsification rules.
- Kernel, transfer, and threshold priors and final calibrated values.
- Random seeds, Monte Carlo traces, and posterior samples.
- Code notebook or script used to generate ADI outputs and dashboard figures.
- Change log and document hashes for dataset and code artifacts.

3.0 From Assumptions to Anticipation: AGI and ADI as Predictive Governance Instruments

The central focus of this paper is that assumptions are not passive metadata; they are the living, shifting foundation of every plan, estimate, and control in a Large Complex Project (LCP). This section reframes AGI (Assumption Governance Index) and ADI (Assumption Diffusion Index) as instruments that turn assumption-level activity into actionable foresight. Rather than treating these indices as ancillary analytics, they are at the heart of governance: AGI quantifies the current integrity of a project's foundation, ADI quantifies how changes to that foundation can ripple through the system, and together they convert latent uncertainty into prioritized action.

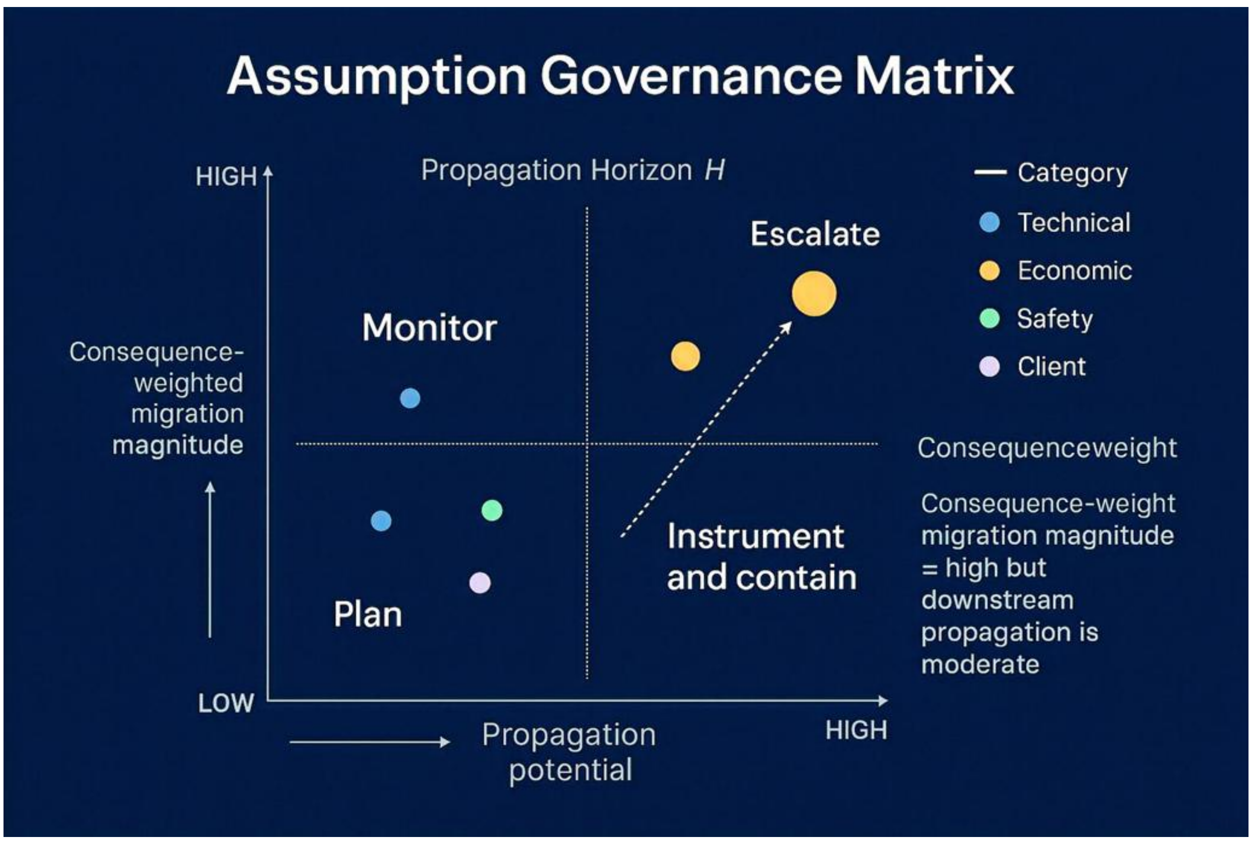

3.1 AGI as a predictive signal for project performance

AGI is predictive because it elevates what is normally invisible—assumption migration—into a normalized, auditable signal that presages realized losses and governance events. Where conventional LCP metrics (earned value, cost variance, schedule indices, discrete risk registers) report what has already moved, AGI reports the fragility of the model that generates those outcomes. An AGI constructed from per-assumption migration Mi, consequence weight Wi, and time-aware confidence Ci does three things that conventional indicators do not:

- it makes aging and shock sensitivity explicit (confidence decay and event multipliers),
- it centers materiality by weighting migrations by governance consequence, and
- it signals entanglement through co-migration and cluster amplification diagnostics rather than assuming independence.

Practically, this means AGI will often rise before a cost overrun or rebaseline appears in EVM. A cluster of small migrations across high-weight assumptions, or a single migration in a high-weight, low-confidence item, will raise AGI and prompt governance review long before realized cost or schedule metrics show the impact. The governance value is concrete: earlier, evidence-traceable rebaseline decisions; more targeted contingency allocation; and auditable escalation trails that link decisions to the state of assumptions rather than to ex-post explanations.

3.2 Why ADI is essential and how it strengthens AGI's predictive power

ADI complements AGI by answering a different operational question: given an assumption change now, what will be the systemic exposure, how fast will it unfold, and which assumptions will act as hubs? AGI answers "how big is the current problem?" ADI answers "how will this problem move and where should I act?" This forward-looking orientation is crucial because many governance failures arise not from an isolated assumption breach but from its propagation through entangled[49] dependencies.

ADI's value propositions are threefold. First, it converts directional coupling (the entanglement matrix E) into consequence-weighted footprints, giving governance a ranked list of sources by systemic potential rather than raw incidence. Second, by returning velocity metrics (time to reach p% of a footprint) ADI maps propagation into operational cadences—telling an owner whether a source needs immediate containment or routine monitoring. Third, ADI produces reach metrics that reveal latent hubs whose instantaneous AGI contribution may be modest but whose long-term systemic footprint is large.

---

[49] This is analogous to the entanglement of quantum systems and recognized in Quantum Project Management.

It is important to stress that ADI must be interpreted alongside AGI. ADI without AGI risks overplaying propagation from low-consequence nodes; AGI without ADI risks underplaying the future pathways by which current migrations will become material. Together they produce a two-axis decision surface (magnitude × propagation potential) that supports precise triage, evidence-backed escalation, and efficient measurement investment.

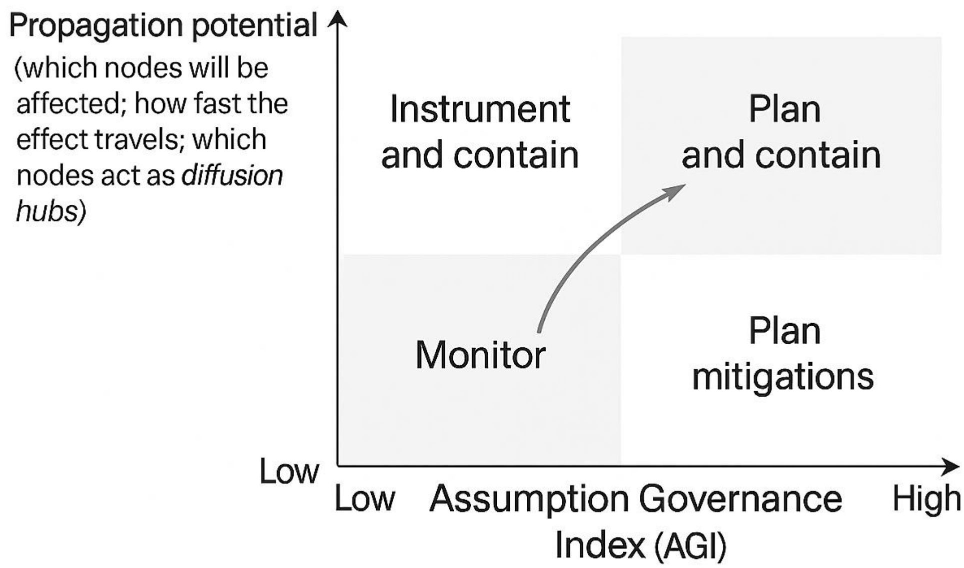

The two-axis decision surface (magnitude × propagation) is crucial for assumption management in quantum project management.

3.3 Capturing complexity, emergence, uncertainty and entanglement

Quantum Project Management (QPM) provides the philosophical and practical lens that justifies AGI/ADI design choices. QPM observes that LCPs exhibit entanglement (tight, sometimes asymmetric coupling among elements), nonlinearity (small changes generating outsized effects), measurement feedback (observation changes the system), and emergence (new modes that cannot be reduced to component behavior). AGI and ADI operationalize QPM in three direct ways:

- **Entanglement is encoded.** AGI uses co-migration diagnostics and PCA-based cluster amplification so that correlated migrations multiply governance attention rather than being averaged away. ADI formalizes directional coupling with a signed matrix E and propagates influence iteratively, capturing asymmetry and pathways that simple correlation misses.

- **Nonlinearity and emergence are recognized and gated**. Both indices include nonlinear components (cluster amplification factors, superlinear transfer functions) and calibrated thresholds (Tc) so that emergent systemic states trigger scaled governance actions rather than linear, one-size responses.
- **Measurement feedback is modeled.** AGI/ADI log updates, apply observation-amplification multipliers, and track second-order migrations that follow measurement events. This makes the system self-aware: the act of measuring and intervening becomes part of the modeled dynamics so governance can distinguish real change from measurement-driven artifacts.

Handling uncertainty is a first-class concern. AGI includes time decay and event multipliers to express aging and shock sensitivity; ADI is calibrated through a Bayesian pipeline with priors, MCMC/posterior sampling, Monte Carlo forward simulations, and bootstrap checks. The result is distributional outputs—median trajectories with 50% and 90% bands—that show both central tendency and reliability. Presenting both bootstrap (sampling variability) and posterior predictive bands (model + prior uncertainty) is a governance best practice: it shows whether uncertainty arises from data scarcity, model form, or parameter ambiguity.

3.4 From migration detection to improved project outcomes

The operational payoff of routinely monitoring AGI and ADI is measurable. First, early detection of fragility reduces surprise: governance can schedule rebaseline planning or contingency allocation with greater lead time and evidence. Second, ADI-driven triage focuses scarce intervention capacity where it reduces systemic risk most—fast hubs for containment, broad hubs for data collection and remediation. Third, prioritized instrumentation reduces model uncertainty over time, meaning the same dashboard evolves from expert-heavy to data-driven and more precise.

These improvements are not hypothetical. When assumptions that drive procurement schedules, commodity costs, or client funding are monitored as first-class governance artifacts, organizations see fewer emergency procurements, fewer unplanned contingency draws, and shorter decision cycles for rebaselining. Moreover, because AGI and ADI are designed for auditability (versioned E, snapshot metadata, posterior samples, logged SME approvals), every governance action can be traced to the evidence and model state that motivated it—essential for sponsor confidence and external review.

3.5 Implementation guidance and governance implications

Realizing the predictive value of AGI/ADI requires disciplined operational practices. First, treat E (the entanglement map) as a living, versioned governance artifact: record provenance, confidence tiers, estimation method, and shrinkage rules. Second, use sensitivity and posterior diagnostics to prioritize instrumentation: focus on edges and parameters that dominate ADI variance. Third, codify joint AGI/ADI escalation rules that require SME confirmation for irreversible actions and that map velocity thresholds to operational cadences (immediate containment, daily monitoring, weekly review).

Calibration and validation must be institutionalized. Calibrate kernels and transfer functions against historical propagation events where possible; run synthetic injection experiments to test coverage; apply ROC analysis to tune AGI bands for desired true/false positive tradeoffs. Present both central estimates and uncertainty bands in dashboards and require that governance decisions reference snapshot metadata (E version, posterior tag, kernel/transfer parameters, Monte Carlo seed).

Finally, recognize the cultural dimension: assumption management requires explicit practices (regular register updates, peer review of consequence weights, mandatory evidence attachments for migrations above thresholds) and incentives that reward good registration and penalize gaming. When organizational culture values clear provenance, timely updates, and peer review, AGI/ADI transition from analytical curiosities to operational control levers that materially reduce systemic fragility.

Glossary

- **AGI — Assumption Governance Index** A normalized, governance-ready KPI that aggregates per-assumption migration, consequence weights, and time-aware confidence to summarize the materiality of assumption migration across a project or portfolio.
- **ADI — Assumption Diffusion Index** A network-based index that quantifies how changes to one or more assumptions propagate through an entanglement matrix, reporting spread, velocity, and reach of propagated exposures.
- **Acrit — Critical Tolerance (DeltaCrit)** A per-assumption threshold that defines the migration magnitude considered materially significant for governance; used to normalize numeric migrations.
- **Affected Indicator ($a_j(t)$)** A binary function that flags when propagated influence on node j at time t exceeds its node-specific threshold, marking that assumption as "affected."
- **Amplification factor (B)** A calibrated scalar applied to cluster- or event-driven signals to represent nonlinear escalation when co-migration or entanglement thresholds are breached.
- **Canonical register** The single authoritative inventory of assumptions used for governance: contains IDs, baseline/current values, Acrit, confidence, weights, entanglement metadata and evidence links.
- **Category Governance Index (CGI_k)** A category-level AGI: the AGI-style aggregation computed only for assumptions in category k (e.g., Technical, Economic, Client).
- **Cluster amplification ($I_c(t)$ or $T_c(t)$)** A nonlinear gating function applied to principal-component-derived cluster scores that increases AGI/CGI contributions when a cluster's co-migration exceeds its sensitivity threshold.
- **Confidence decay ($C_i(t)$)** Time- and event-driven reduction of baseline confidence for assumption i, typically modeled exponentially and modulated by event multipliers.
- **Consequence weight (W_i)** The normalized aggregate weight for an assumption derived from component scores (e.g., cost, safety, schedule, reputation) that scales its governance importance.

- **Copula** A statistical device used to impose or model dependence structures (tail and rank dependence) among simulated assumption migration processes.
- **Coverage (calibration diagnostic)** The empirical frequency with which observed outcomes fall inside predicted credible intervals; used to assess calibration quality.
- **Decay kernel / Temporal kernel (K(Δt))** A time-dependent function that governs how influence attenuates (or persists) across propagation lags; examples include exponential, power-law, boxcar, and delayed kernels.
- **DP_i(t) — Diffusion Potential (instantaneous)** A scalar measuring the immediate propagation potential of a state change at node i, typically computed as the product of the source signal and the sum of absolute outbound entanglement coefficients.
- **Entanglement / Entanglement matrix (E = [e_ij])** A directed, signed matrix of coefficients e_ij ∈ [-1,1] encoding the strength and sign of influence from assumption i to assumption j; the structural backbone of ADI propagation.
- **Evidence link** A documented reference (URL, file hash or attachment index) stored in the canonical register that substantiates an assumption update or migration claim.
- **Event-driven amplification multiplier (φ_i(t) or similar)** A per-assumption multiplier ≥1 that increases sensitivity (speeds confidence decay or raises migration impact) following disruptive events.
- **Footprint_i(H)** The cumulative, consequence-weighted exposure (probability-weighted) that a single seeded source i generates over propagation horizon H.
- **Granger-style causality** A time-series testing approach used to infer directional predictability (i → j) and to support estimation of directed entanglement coefficients where temporal data permit.
- **H — Propagation horizon** The number of discrete time steps used to evaluate forward propagation from seeded source events when computing footprints and ADI(H).
- **Kernel tail behavior** The long-lag profile of a temporal kernel (e.g., exponential decays quickly; power law has a heavy tail) that influences velocity and reach of diffusion.

- **M_i(t) — Migration metric / Migration score** The per-assumption normalized migration measure in [0,1] representing the magnitude of change from baseline (numeric normalized by Acrit or categorical mapped to an ordinal distance).
- **Monte Carlo (forward simulation)** Repeated stochastic simulation across posterior parameter draws (and process noise) to produce distributions for ADI, Footprints, velocity, and reach for uncertainty quantification.
- **NUTS / HMC / MCMC (samplers)** Classes of Bayesian sampling algorithms (No-U-Turn Sampler / Hamiltonian Monte Carlo / Markov Chain Monte Carlo) recommended for posterior estimation during ADI calibration.
- **Observation-amplification factor (q_i(t))** A parameter or multiplier that models how measurement or reporting actions alter subsequent migration dynamics (measurement feedback).
- **p_i(t) / influence score** The dynamic propagated influence state of node i at time t in the discrete iterative ADI update rule.
- **PC / Principal Component (PCA)** An orthogonal latent vector extracted from a weighted covariance of migration series used to identify dominant co-migration clusters for amplification diagnostics.
- **Posterior predictive checks (PPC)** Diagnostics comparing replicated datasets drawn from posterior predictive distributions to observed data, used to identify model misspecification.
- **Pruning / sparsification threshold (T)** A governance-defined cutoff used to zero-out weak entanglement coefficients (|e_ij| < T) to improve interpretability and computational tractability.
- **QPM — Quantum Project Management** Theoretical framing that treats large complex projects as entangled, emergent systems where measurement feedback, nonlinearity, and coupling must be explicitly managed.
- **R — Correlation matrix (or R matrix)** The matrix of pairwise correlations (e.g., Pearson) computed between assumption migration time series; used for co-migration mapping and cluster discovery.
- **Reach metric (R_i(H))** The fraction of portfolio consequence weight potentially affected by source i over horizon H; highlights systemic breadth even when instantaneous impact is small.

- **Rebaseline** A formal governance action to update project baselines (scope, cost, schedule) when underlying assumptions or their aggregated signals (e.g., AGI) indicate material drift.
- **Receiver Operating Characteristic (ROC)** A calibration analysis tool plotting true positive rate vs false positive rate across thresholds; recommended for tuning AGI bands and governance detection trade-offs.
- **Sparsity / matrix hygiene** Post-processing practices (clipping, sparsification, directionality resolution, normalization, confidence tagging) applied to E to ensure robustness and auditability.
- **Sensitivity analysis** Systematic perturbation of model parameters (kernels, thresholds, e_ij, transfer slopes) to identify brittle or high-leverage inputs driving ADI variance.
- **Snapshot metadata** Versioned record attached to each dashboard or analytic snapshot including E version ID, calibration date, posterior tag, parameter settings and provenance for audit trails.
- **Thresholds (θ_j or ϑ_j)** Node-specific cutoffs used to convert propagated continuous influence into binary affected flags; selected by governance to balance sensitivity and false alarm rates.
- **Triage / Velocity-based cadence (V_i(p))** Operational rules mapping velocity metrics (time to reach p% of source footprint) to monitoring cadences and escalation actions.
- **Transfer function (g(·))** The node-level mapping that transforms a raw influence score into the effective signal transmitted across edges; can be linear, saturating, rectified, or piecewise.
- **UFE — Unallocated Future Expense** A traditional cost-reserve concept referenced as part of the context for AGI's relationship to cost uncertainty and contingency planning.
- **Weighted covariance (E_w)** The covariance matrix of migration series pre- or post-weighting by consequence scores used to prioritize governance-relevant co-migration structure for PCA.
- **Workaround / Edge-level intervention** Targeted, localized governance action (e.g., temporary hold, microlearning, contractual guardrail) applied to reduce propagation along high-impact edges or hubs.

- **W_tot (portfolio total weight)** The sum of all per-assumption consequence weights used to normalize ADI and footprint measures to portfolio scale.
- **Weighted moments (μ_wM, σ_wM)** Mean and dispersion statistics calculated with consequence weights applied, used to prioritize shifts in high-impact assumptions for governance attention.
- **Zero-trust / conservative defaults** Governance principle to prefer containment, wider uncertainty bands, and SME confirmation when data or edge confidence is low.

AI Disclosure

AI tools were used to support specific technical and editorial tasks for this manuscript. Between September and November 2025 the author used Microsoft Copilot (text/code assistance), DALL·E (image generation), and Grammarly (language editing) to generate, clean, and format the symbolic equations appearing in the AGI and ADI sections and associated appendices; to assist background literature review and clarify terminology; and to format tables, lists, and illustrative figures where appropriate. All AI outputs were iteratively prompted, reviewed, and validated by the author, who confirms the accuracy of the equations, figures, and definitions and accepts full responsibility for the integrity and interpretation of the work. All DALL·E images are original, and no third-party copyrighted material were uploaded to any AI service.

References

Blue Skies Consulting. (2025). *Project governance: The art of measuring effectively*. https://blueskiesconsulting.com/project-governance-the-art-of-measuring-effectively/

Committee of Sponsoring Organizations of the Treadway Commission. (2017). *Enterprise risk management: Integrating with strategy and performance*. https://www.coso.org/Documents/2017-COSO-ERM-Integrating-with-Strategy-and-Performance-Executive-Summary.pdf https://www.academia.edu/42188027/Enterprise_Risk_Management_Integrating_with_Strategy_and_Performance

Crouhy, M., Galai, D., & Mark, R. (2014). *The essentials of risk management* (2nd ed.). McGraw-Hill Education. https://archive.org/details/essentialsofrisk0000crou_n6p4/page/n5/mode/2up

Data Processing Pipelines — EITC. http://www.eitc.org/research-opportunities/new-media-and-new-digital-economy/data-science-and-analytics/workflows-and-pipelines-in-ai-and-ml/data-processing-pipelines

Data Science & AI Glossary 2024: Abaka AI. https://www.abaka.ai/glossary

Efron, B., & Tibshirani, R. J. (1994). *An introduction to the bootstrap*. CRC Press.

Hollmann, J. K. (2016). *Project risk quantification: A guide to cost and schedule risk management*. AACE International. https://decisions-books.com/PRQ

How To Install R | DigitalOcean. https://www.digitalocean.com/community/tutorial_collections/how-to-install-r

International Organization for Standardization. (2018). *ISO 31000: Risk management—Guidelines*. ISO.

ITS Project Management Office. (2024). *8 key components of project governance*. https://pmo.its.uconn.edu/2024/12/23/pmo-project-governance/

Jolliffe, I. T., & Cadima, J. (2016). Principal component analysis: A review and recent developments. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374(2065), 20150202. https://doi.org/10.1098/rsta.2015.0202 https://pubmed.ncbi.nlm.nih.gov/26953178/

Logikheim. (2024). *Making assumptions work for your project: A structured approach for project controls*. https://www.logikheim.com.au/post/making-assumptions-work-for-your-project-a-structured-approach-for-project-controls

NASA. (n.d.). *Appendix G: Cost risk and uncertainty methodologies*. https://www.nasa.gov/wp-content/uploads/2020/11/ceh_appg.pdf?emrc=691098d98ebe7

Nasution, B., Ritonga, W., Doyan, A., Pandara, P. D., Alfaris, L., & Siagian, R. C. (2023). Temporal Trends and Spatial Relationships of Radioactive Isotopes (I-131, Cs-134, and Cs-137) in Response to Nuclear Events: A Comprehensive Analysis Using Time Series Graphs, Regression, and Multivariate Techniques. https://doi.org/10.55981/aij.2023.1306

Nguyen, N. (2023). *Governance and controls in complex projects*. https://pgcs.org.au/wp-content/uploads/2024/07/2023_Nam_Nguyen.pdf

Prieto, B. (2019). *Assumption, risk driver and constraint tracking*; National Academy of Construction https://www.naocon.org/wp-content/uploads/Assumption-Risk-Driver-and-Constraint-Tracking.pdf
https://www.researchgate.net/publication/340949729_Assumption_Risk_Driver_and_Constraint_Tracking_Key_Points#fullTextFileContent

Prieto, B. (2016). Management of assumption infatuation in large complex projects. *PM World Journal*, April. https://pmworldlibrary.net/wp-content/uploads/2016/04/pmwj45-Apr2016-Prieto-Management-of-Assumption-Infatuation-in-Large-Complex-Projects.pdf
https://www.researchgate.net/publication/299979232_Management_of_Assumption_Infatuation_in_Large_Complex_Projects#fullTextFileContent

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-Pireto-Managing-Uncertainty-in-Large-Complex-Projects.pdf
https://www.researchgate.net/publication/397300009_Managing_Uncertainty_in_Large_Complex_Projects_1

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
https://www.researchgate.net/publication/379694597_Measurement_of_Complexity_in_Large_Complex_Projects#fullTextFileContent

Prieto, R. (2024a). Quantum project management. *PM World Journal*, January.
https://pmworldlibrary.net/wp-content/uploads/2024/01/pmwj137-Jan2024-Prieto-Quantum-Project-Management-.pdf
https://www.researchgate.net/publication/377219948_Quantum_Project_Management#fullTextFileContent

Prieto, R. (2024b). Quantum project management book announced. *PM World Journal*, January.
https://pmworldlibrary.net/wp-content/uploads/2025/01/250110-Prieto-Quantum-Project-Management-book-announced.pdf

Prieto, R. (2024c) “Quantum Project Management: A monograph on a new theory for management of large complex projects”; ISBN: 978-1-304-08165-0, December 2024 https://www.researchgate.net/publication/387334326_Quantum_Project_Management_A_monograph_on_a_new_theory_for_management_of_large_complex_projects#fullTextFileContent

Project Management Institute. (2017). *A guide to the project management body of knowledge (PMBOK® Guide)* (6th ed.). Project Management Institute. https://drive.google.com/file/d/1Ln91Ies6LZ78W2aRQhW68dJ9vRX1oPba/view

ProjectManager.com. . (2025). *Project risk analysis: Quantitative & qualitative techniques*. https://www.projectmanager.com/training/how-to-analyze-risks-project

ProjectPractical. (2024). *Project assumptions: 30 examples and how to manage them*. https://www.projectpractical.com/project-assumptions-examples-and-how-to-manage-them/

ROC And AUC Curves In ML Made Simple & How To Tutorial. https://spotintelligence.com/2024/06/17/roc-auc-curve-in-machine-learning/

Stoltz, S. (2025). Evaluating project governance: Metrics and KPIs for success. *Weblog*. https://www.shaunstoltz.com/2025/08/07/evaluating-project-governance-metrics-and-kpis-for-success/#google_vignette

Taskford. (2024). *Project governance for large-scale and complex projects*. https://taskford.com/en/blog/project-governance

WSDOT. (n.d.). *Project risk analysis model (PRAM) user's guide*. Washington State Department of Transportation. https://wsdot.wa.gov/publications/fulltext/CEVP/PRAM_Users_Guide.pdf

Appendices[50]

- Appendix A: Structuring the Assumption Register for Governance Integrity
- Appendix B: Worked Example: AGI Computation in Practice
- Appendix C: Monte Carlo Validation and Calibration Protocol
- Appendix D: Governance Dashboard and Audit Architecture
- Appendix E: Implementation Roadmap and Research Agenda
- Appendix F: Assumption Mapping and Scores (n = 8); AGI Signal Ranking and Escalation Triggers
- Appendix G Reproducibility Checklist

[50] Extensive appended information has been included in the paper rather than making it separately available. The intent is to facilitate deeper examination by serious practitioners.

Appendix A — Structuring the Assumption Register for Governance Integrity

The assumption register is the foundational data artifact upon which the Assumption Governance Index (AGI) and all related metrics are computed. Its structure must support traceability, auditability, and analytical rigor. In large complex projects (LCPs), assumptions are made across multiple domains—technical, environmental, stakeholder, economic, client, and productivity—and often by diverse actors with varying levels of authority and insight. As Prieto (2016) emphasized in "Management of Assumption Infatuation in Large Complex Projects," assumptions are frequently made unconsciously and without coordination, leading to compounded uncertainty and systemic fragility[51].

To counteract this, the assumption register must be designed not merely as a passive repository but as an active governance instrument. It must support:

- Granular tracking of assumption values over time
- Categorization for domain-specific governance
- Confidence decay modeling to reflect aging of assumptions
- Consequence weighting to prioritize governance attention
- Entanglement tagging to identify correlated assumptions
- Evidence linkage to support audit and decision justification

Recommended Register Schema

Each row in the register represents a single assumption. The following fields are required:

- Field Name
- Description
- ID
- Unique identifier (e.g., A-001)
- Category
- Domain classification (e.g., Technical, Economic/Cost)
- Baseline Value
- Initial value or state of the assumption
- Baseline Timestamp
- Date the baseline was established
- Current Value
- Most recent assessed value
- Units

[51] Prieto, R. (2024) "Quantum Project Management." PM World Journal feature paper, January

- Measurement units (e.g., USD, %, ordinal label)
- DeltaCrit (Δcrit) (critical tolerance)
- Critical tolerance threshold for material migration
- Baseline Confidence (α0)
- Initial confidence in the assumption (0–1 scale)
- Lambda (λ) Volatility parameter for confidence decay
- Omega (ω) Cost, Omega (ω) Safety, Omega (ω) Schedule, Omega(ω) Reputation Consequence scores (0–10 scale)
- Owner/Responsible party for assumption management
- Last Update
- Date of most recent update
- Evidence Link - URL or reference to supporting documentation
- Update Notes - Free-text rationale for update
- Qualitative Map Key- Mapping key for categorical assumptions
- Entanglement Group - Optional tag for correlated assumptions
- Archived - Boolean flag for retired assumptions

This schema supports deterministic computation of AGI and CGI[52] metrics, facilitates peer review of consequence scores, and enables automated alerts when assumptions breach migration thresholds.

References

Prieto, R. (2016) "Management of Assumption Infatuation in Large Complex Projects." PM World Journal archive, April 2016.

Prieto, R. (2024) "Quantum Project Management." PM World Journal feature paper, January

[52] The Category Governance Index (CGI_k) breaks AGI down by assumption category — such as Technical, Economic, Stakeholder, Environmental — to show which domains are driving governance risk.

Appendix B — Worked Example: AGI Computation in Practice

To illustrate the AGI methodology, consider a representative project with eight assumptions spanning multiple categories. This example demonstrates how per-assumption migration scores (Mi), consequence weights (Wi), confidence decay factors (Ci), and aggregate metrics are computed.

Assumption Set												
ID	**Category**	**BaselineValue**	**CurrentValue**	**Units**	**DeltaCrit**	**α0**	**λ (per month)**	**ω cost**	**ω safety**	**ω schedule**	**ω reputation**	**Entanglement Group**
A1	EconomicCost	1,000,000	1,120,000	USD	100,000	0.85	0.015	8	0	6	2	ECON_COMMOD
A2	Performance	100	92	%prod	6	0.8	0.02	0	0	9	1	PROD
A3	Technical	0.95	0.93	availability	0.02	0.9	0.01	2	8	4	1	TECH
A4	Environmental	Permitted	Delayed	categorical	mapped dist=0.6	0.75	0.03	1	3	2	1	PERMIT
A5	Client	FundingOnTime	Delayed	categorical	mapped dist=0.8	0.7	0.04	2	1	3	6	CLIENT
A6	EconomicCost	FuelPriceIndex=1	FuelPriceIndex=2.5	index	0.5 index	0.8	0.02	6	0	2	1	ECON_COMMOD
A7	Stakeholder	Sentiment=0	Sentiment=-1	ordinal dist	0.5	0.85	0.01	1	4	3	2	STAKE
A8	Performance	CrewAvailability=1.0	0.85	fraction	0.1	0.9	0.03	0	0	7	1	HR

Step-by-Step Computation

1. Migration Scores (Mi): Calculated as absolute deviation from baseline normalized by Δcrit. Categorical assumptions use mapped distances (e.g., 0.6 for A4).

2. Consequence Weights (Wi): Sum ω components and normalize across all assumptions. For example, A1 has ω_total = 16, and Wi = 16 / 87[53] ≈ 0.1839.

3. Confidence Decay (Ci): Apply exponential decay: $Ci = \alpha 0 * \exp(-\lambda * \Delta t)$. For A1, with $\Delta t = 9$ months, Ci ≈ 0.7438.

4. Contribution to AGI: Compute T_i[54] = Wi * Ci * Mi. Sum across all assumptions and multiply by 100 to get AGI.

5. Category Indices (CGI_k): Aggregate T_i within each category and normalize by category weight sum.

6. Statistical Metrics: Compute mean migration (μ_M), weighted mean (μ_wM), standard deviation (σ_M, σ_wM), and identify top contributors.

---

[53] Sum of ω across all eight assumptions

[54] T_i is the contribution term for assumption *i* in the Assumption Governance Index (AGI) framework. It quantifies how much that individual assumption contributes to the overall governance risk score at a given point in time.

Results

- AGI = 61.4 → Amber band: program-level review triggered
- CGI_Technical = 82.2 → High exposure
- CGI_Economic = 71.7 → Elevated risk
- Top contributors: A3 (Technical), A1 (Economic/Cost), A2 (Performance)

This example demonstrates how AGI surfaces systemic risk from assumption migration and supports targeted governance actions.

Appendix C — Monte Carlo Validation and Calibration Protocol

To validate the AGI methodology and calibrate parameters (e.g., λ, β, τ_c), a Monte Carlo simulation framework is recommended. This approach enables testing under controlled conditions and supports statistical inference about AGI's predictive power.

Simulation Design

- **Assumption Population:** Generate synthetic assumptions with realistic distributions for baseline values, $\Delta crit$, $\alpha 0$, λ, and ω components.
- **Migration Processes:** Model Vt using stochastic processes:
  - Geometric Brownian Motion for prices
  - Ornstein-Uhlenbeck for productivity
  - Markov chains for categorical states
- **Correlation Structures:** Impose entanglement via copulas or factor models (e.g., commodity shocks affecting multiple assumptions).
- **Event Injection:** Simulate disruptive events (e.g., permit denial) and observe AGI response.

Metrics and Calibration

- **Detection Performance**: Measure true positive rate, false positive rate, and lead time for AGI-triggered governance actions.
- **Threshold Calibration:** Use ROC curves to set AGI bands and $\Delta\mu_{crit}$ values.
- **Parameter Sensitivity:** Vary λ, β, τ_c to assess robustness.

Implementation

- Use Python (NumPy, Pandas) or R for simulation
- Archive inputs, outputs, and random seeds for reproducibility
- Document scenario definitions and calibration results

This validation protocol ensures AGI is not only theoretically sound but empirically effective in real-world governance contexts.

Appendix D — Governance Dashboard and Audit Architecture

To operationalize AGI, a governance dashboard and audit package must be developed. These tools support real-time monitoring, stakeholder engagement, and regulatory compliance.

Dashboard Components

- **Top-Line Metrics:**
  - AGI with color band and trend arrow
  - Sparkline of AGI over time
  - Category CGIs with delta indicators
- **Diagnostic Panels:**
  - Top contributors table (ID, Mi, Wi, Ci, T_i)
  - Entanglement heatmap (correlation matrix)

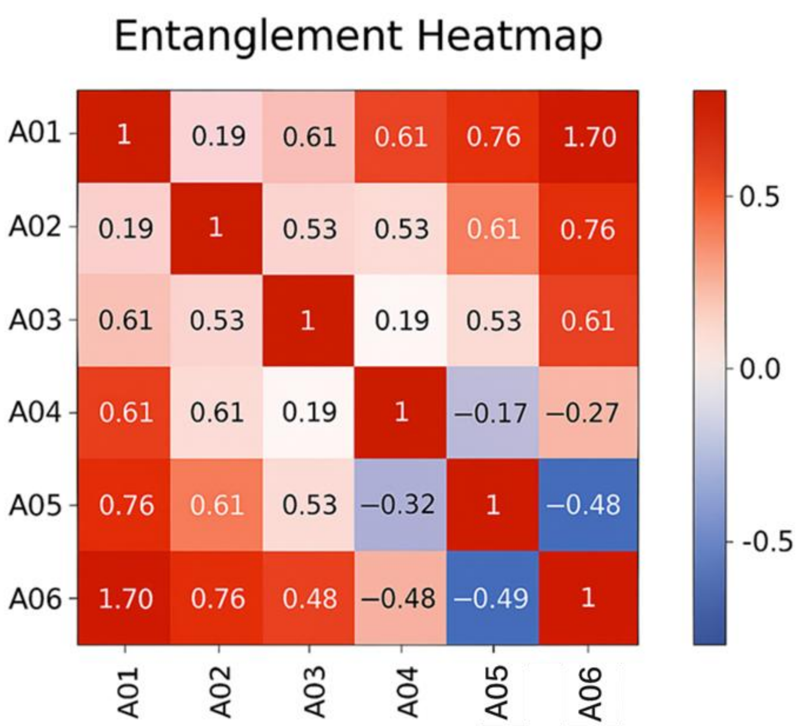

- Statistical metrics (μ_M, σ_M, skewness, kurtosis)[55]
- **Drill-Through Views:**
    - Assumption detail pages with update history, evidence, and action logs
- **Alerts and Workflows:**
    - Automated triggers for AGI band changes, CGI spikes, and Mi breaches
    - Role-based routing to PMO, sponsors, and risk managers

Skewness	Kurtosis
Definition Skewness measures the asymmetry of the distribution of assumption migrations (the Delta or migration metric M_i). Positive skewness means a long right tail (occasional large upward migrations); negative skewness means a long left tail (occasional large downward migrations).	**Definition** Kurtosis measures tail heaviness and peak sharpness of the migration distribution relative to a normal distribution. Excess kurtosis (kurtosis − 3) is commonly used to compare to the normal baseline.
Sample formula $$g_1 = \frac{\frac{1}{n}\sum_{j=1}^{n} (x_j - \bar{x})^3}{(\frac{1}{n}\sum_{j=1}^{n} (x_j - \bar{x})^2)^{3/2}}$$	**Sample formula (excess kurtosis)** $$g_2 = \frac{\frac{1}{n}\sum_{j=1}^{n} (x_j - \bar{x})^4}{(\frac{1}{n}\sum_{j=1}^{n} (x_j - \bar{x})^2)^2} - 3$$
Interpretation in assumption migration • **Positive skew**: most migrations are small but rare events produce large increases in migration (sudden escalations in risk or loss of confidence). These assumptions need monitoring for tail events and may drive conservative governance actions. • **Negative skew**: most migrations are small but rare events produce large decreases (rapid recovery or overcorrection). These may reflect episodic fixes, corrective actions, or data artifacts.	**Interpretation in assumption migration** • **High (positive) excess kurtosis**: heavy tails and more frequent extreme migrations than expected. Indicates elevated tail risk and that standard deviation understates true risk. Governance should emphasize stress tests, event-driven controls, and contingency plans. • **Low (negative) excess kurtosis)**: light tails and fewer extreme migrations. Distribution concentrated around the mean; routine monitoring and trend detection may suffice.

[55] Use skewness and kurtosis together: skewness flags direction of asymmetry; kurtosis signals tail risk magnitude. An assumption with positive skew and high kurtosis is a high-priority tail-risk driver.

Audit Package Specification

- **Archived Run File:** Includes register snapshot, computed metrics, and parameter settings
- **Evidence Index:** Links to supporting documents with checksums
- **Change Log:** Tracks edits with timestamps and rationales
- **Governance Actions Log:** Records decisions, approvals, and outcomes
- **Reproducibility Script:** Enables re-running AGI computation from archived inputs

Security and Compliance

- Role-based access controls
- Immutable archival with retention policies
- Tamper-evident evidence storage
- Exportable audit trail for external review

This architecture ensures AGI is not only analytically rigorous but operationally robust, supporting transparent, accountable governance in complex project environments.

References

• Nguyen, N. (2023). “Governance and Controls in Complex Projects.”

Appendix E: Implementation Roadmap and Research Agenda

Implementation Roadmap

This roadmap outlines a phased strategy for operationalizing the ADI–AGI framework across construction safety governance, skilled trades investment, and hyperscale project delivery. Each phase is designed to build traceability, accelerate feedback loops, and embed assumption intelligence into decision cycles.

Phase 1: Foundation and Tooling (Months 0–3)

ADI–AGI Framework Objectives	
Objective	**Description**
Deploy ADI–AGI Framework	Integrate temporal kernel logic and entanglement coefficients into existing incident analysis platforms. Validate source footprint matrices against historical case data.
Standardize Inputs and Definitions	Finalize assumption taxonomies, normalize horizon buckets (1d, 7d, 30d, 90d), and align AGI weighting schema with governance priorities.
Build Modular Visuals and Appendices	Embed the figures and table that into governance documentation. Develop training decks for field teams and executive briefings.

The following figure illustrates the ADI diffusion schematic, showing how assumptions propagate via entanglement coefficients and temporal kernels.

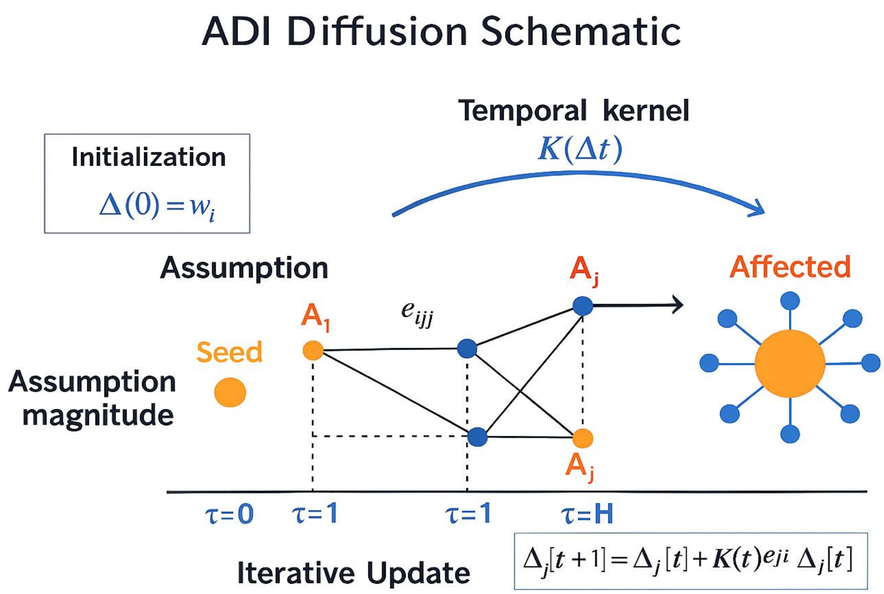

This next figure presents the ADI heatmap across horizon buckets, highlighting high-impact sources.

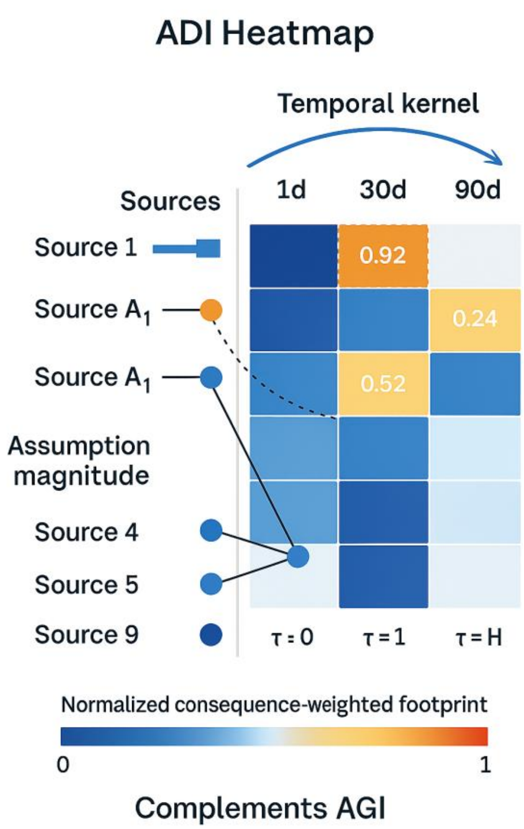

The following table maps assumption categories to normalized footprint scores and are visualized in an ADI heatmap tile and used to drive triage, velocity table actions, and explainability drilldowns..

Assumption Categories and Normalized Footprint Scores

Assumption Category	Normalized Footprint (ADI_i(H))	Interpretation
Source A	0.89	Fast propagator with broad systemic reach
Source B	0.88	High footprint, moderate velocity
Source C	0.44	Localized impact, slower propagation
Source D	0.43	Low footprint, limited downstream exposure
Source E	0.94	Dominant systemic driver; priority for triage

Phase 2: Pilot and Feedback Loop (Months 4–6)

ADI–AGI Pilot and Refinement Objectives

Objective	Description
Run Controlled Pilots	Apply ADI–AGI overlays to 3–5 active projects across regions. Track assumption propagation and governance response timing.
Surface Lessons Learned	Extract granular insights from entanglement networks and time series comparisons. Identify lagging governance signals and high-impact assumption clusters.
Refine SDRI[56] and ROI Models	Link assumption diffusion to skilled trades investment outcomes. Update escalation frameworks based on temporal footprint patterns.

The following figure visualizes entanglement networks, mapping node size to ADI_i and color to AGI contribution.

[56] SDRI = Skilled Trades Development Return Index (shorthand for the skilled-trades / workforce development ROI index discussed in the roadmap)

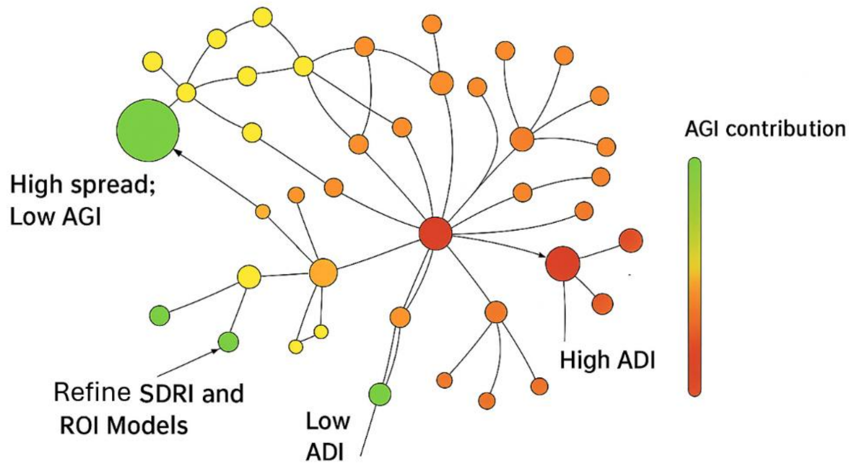

Visualizing entanglement networks, mapping node size to ADI_i and color to AGI contribution.

The following figure compares ADI(H) and AGI over behavior time, with a cross-correlation insert showing governance lag.

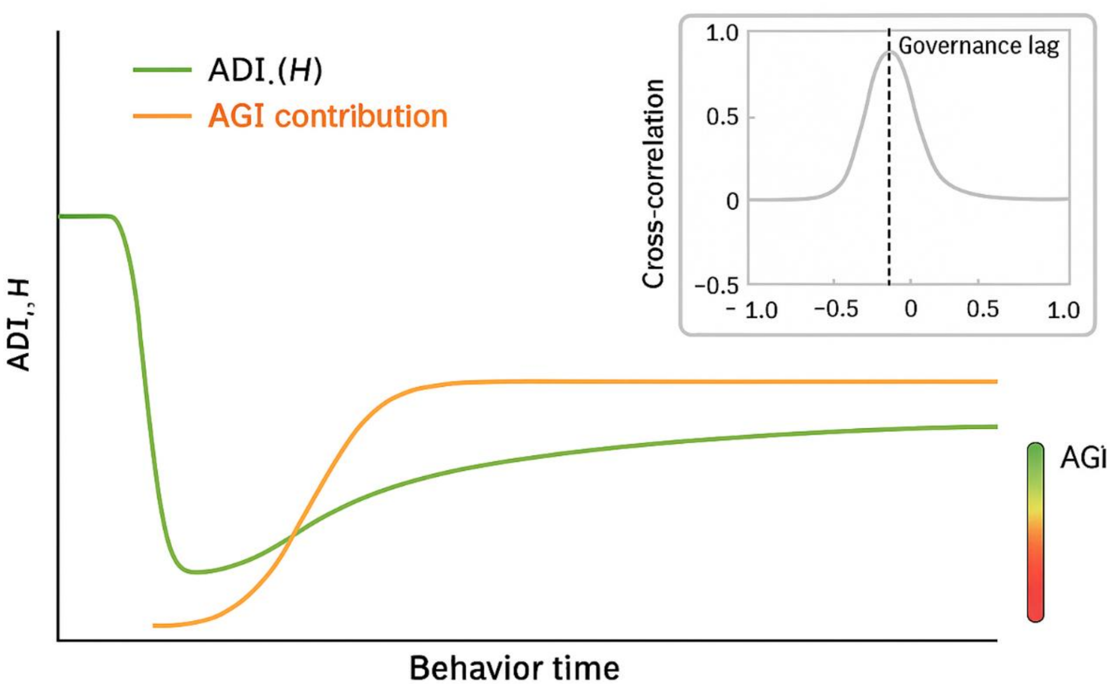

Mapping node size to ADI_i and color to AGI contribution

Phase 3: Scale and Governance Integration (Months 7–12)

Strategic Objectives for ADI–AGI Integration

Objective	**Description**
Embed into Governance Cycles	Use AGI overlays to inform quarterly safety reviews and capital planning. Align assumption tracking with audit and compliance workflows.
Expand to Sectoral Benchmarks	Compare ADI–AGI profiles across data center, infrastructure, and vertical construction sectors. Publish anonymized benchmarking overlays.
Develop Feedback-Driven Training Modules	Create role-specific training based on assumption entanglement and governance lag. Integrate visuals and time series insights into onboarding and refresher programs.

Research Agenda

This agenda identifies priority areas for advancing the ADI–AGI framework and its application to safety, governance, and skilled trades development. Each item is designed to deepen analytic rigor and enhance field-level usability.

1. Temporal Kernel Calibration

- Compare exponential vs power-law decay functions across assumption types.
- Investigate kernel sensitivity to behavior time and governance response lag.
- Validate kernel shape against observed escalation timing in incident logs.

2. Entanglement Coefficient Mapping

- Develop automated methods for estimating e_{ij} from incident narratives and project metadata.
- Explore sparsification techniques for large-scale assumption networks.
- Test stability of entanglement maps across project types and regions.

3. AGI Weighting Schema Optimization

- Test alternative weighting models (e.g., risk-adjusted, role-specific, compliance-driven).
- Validate AGI scores against real-world governance actions and outcomes.
- Align AGI thresholds with escalation protocols and audit triggers.

4. Cross-Sector Benchmarking

- Expand footprint matrices to include infrastructure, manufacturing, and energy sectors.
- Identify common assumption clusters and governance bottlenecks.
- Develop sector-specific overlays for executive briefings.

Sectoral ADI–AGI Comparison and Governance Recommendations

Sector	Median ADI_i	Median AGI_i	Dominant Assumption Clusters	Recommended Governance Actions
Energy Infrastructure	0.78	0.65	EconomicCost, Environmental	Escalation protocols for commodity volatility and permitting delays
Commercial Construction	0.52	0.71	Client, Schedule, Stakeholder	Tighten client funding reviews and stakeholder sentiment tracking
Data Centers	0.61	0.68	Technical, Performance, HR	Prioritize crew availability and technical uptime in training modules
Transportation Projects	0.74	0.59	Safety, Schedule, Environmental	Integrate safety protocol triggers and schedule drift alerts
Public Works	0.66	0.72	Reputation, Stakeholder, Permit	Enhance reputation risk mapping and stakeholder escalation logic

5. Skilled Trades ROI Linkage

- Quantify impact of assumption propagation on training investment effectiveness.
- Model feedback loops between ADI(H) patterns and workforce development metrics.
- Integrate assumption intelligence into skilled trades curriculum design.

6. Human Factors and Decision Timing

- Study how assumption diffusion affects decision latency and escalation behavior.
- Integrate behavioral economics into AGI signal interpretation.
- Develop role-specific dashboards that surface assumption lag and governance urgency.

Appendix F Assumption Mapping and Scores (n = 8); AGI Signal Ranking and Escalation Triggers

Assumption Mapping and Scores (n = 8)

ID	Category	ADI_i	AGI_i	T_i (months)	Entanglement Group	Notes
A1	EconomicCost	0.72	0.68	9	ECON_COMMOD	Amplified via commodity shock
A2	Performance	0.51	0.62	6	PROD	Crew availability decline
A3	Technical	0.48	0.57	12	TECH	Low decay, stable footprint
A4	Environmental	0.39	0.44	8	PERMIT	Permitting delay
A5	Client	0.63	0.71	6	CLIENT	Funding delay, high impact
A6	EconomicCost	0.66	0.61	9	ECON_COMMOD	Amplified via commodity shock
A7	Stakeholder	0.42	0.49	7	STAKE	Sentiment decline
A8	Performance	0.55	0.60	10	HR	Crew availability drop

AGI Signal Ranking and Escalation Triggers

Rank	Assumption ID	AGI_i	Amplification Triggered	Escalation Tier
1	A5	0.71	No	Tier 1
2	A1	0.68	Yes	Tier 1
3	A6	0.61	Yes	Tier 2
4	A2	0.62	No	Tier 2
5	A8	0.60	No	Tier 2
6	A3	0.57	No	Tier 3
7	A7	0.49	No	Tier 3
8	A4	0.44	No	Tier 3

Appendix G Reproducibility Checklist

Purpose: enable an auditable, repeatable re-run of AGI/ADI computations and the review of any governance decision that uses those outputs.

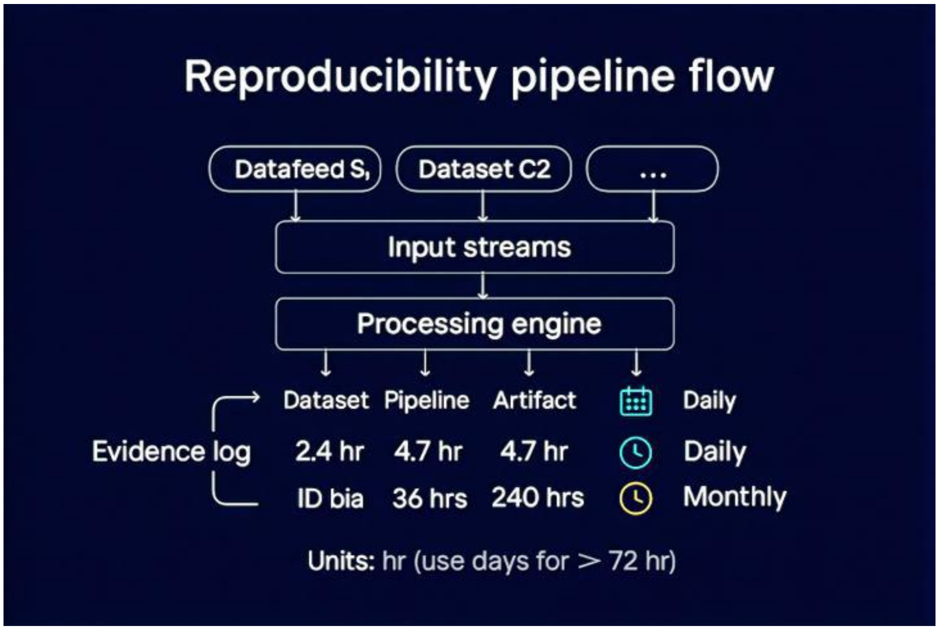

A. Snapshot metadata (must)

[ ] **Snapshot ID** (format: YYYYMMDD_TAG_V#)
[ ] **Snapshot timestamp** (UTC)
[ ] **Owner / Responsible** (name / role / contact)
[ ] **Reviewer / Approver** (name / role / signature or electronic approval)
[ ] **Purpose statement** (one-line: scope, use-case, decision supported)
[ ] **Snapshot summary** (one-paragraph: inputs changed since prior snapshot, main findings, recommended governance action)

B. Canonical data and register (must)

[ ] **Canonical register export** (single-file CSV/Parquet); include file hash (SHA256)
[ ] **Register schema** (published schema file listing required fields and allowed value ranges)
[ ] **Field-level provenance log** (for every changed record: prior value, new value, who changed it, date/time, rationale, evidence link)
[ ] **Raw source references** (list of upstream sources, extract queries or file IDs used to build register)
[ ] **Input validation report** (unit checks, null/missing checks, range checks, Δcrit plausibility) with pass/fail flags and remediation notes

[ ] **Evidence attachments index** (links or checksums for documents referenced by register entries)

C. Code, pipeline, and environment (must)

[ ] **Code repository reference** (URL, branch, commit hash/tag)
[ ] **Entry-point script & invocation** (exact command line, config file path)
[ ] **Environment manifest** (OS, language runtimes, package list with pinned versions; container image digest if used)
[ ] **Deterministic run instructions** (working directory, environment variables, and exact invocation to reproduce run)
[ ] **Random seeds** (explicit seed values used for all stochastic components)
[ ] **ETL artifacts** (transformed intermediate snapshot with file hash; note any manual steps)

D. Parameters, artifacts and matrices (must)

[ ] **Parameter registry file(s)** used (kernels, transfer function parameters, thresholds ϑ_j, decay rates λ_i, amplification B, smoothing) with version and hash
[ ] **Weighting & criticality rules** (documented formulae and any manual overrides with approver)
[ ] **Entanglement matrix E** (versioned file used for this snapshot; estimation method and confidence tier per edge) with hash
[ ] **Post-processing log for E** (mapping functions, clipping, sparsification threshold T, normalization applied; before/after statistics)
[ ] **Model code modules** (functions computing Mi, Ci, AGI, CGI_k, ADI footprints, velocity/reach) and unit tests referenced

E. Calibration, Monte Carlo and validation (must if used)

[ ] **Calibration record** (priors, sampler algorithm, chains, warmup, draws; commit/hash for calibration code)
[ ] **Posterior artifacts** (posterior sample bundle with identifier and hash)
[ ] **Monte Carlo configuration** (number of posterior draws, replicates per draw, seeds, compute batch IDs)
[ ] **Validation datasets** (held-out events or historical cases with extraction rules)
[ ] **Validation results** (ROC, precision/recall, time-to-detection delta, coverage diagnostics) with pass/fail against pre-defined acceptance criteria
[ ] **Synthetic injection tests** (injection definitions, expected outcomes, results, coverage statistics)
[ ] **Sensitivity analysis report** (parameters swept, method, ranked sensitivities)

F. Outputs, dashboard and auditables (must)

[ ] **Dashboard snapshot** (image + underlying data file; hash) and snapshot ID shown on dashboard
[ ] **Top contributors table** (CSV) and entanglement heatmap (file) used for decision briefing
[ ] **Exported AGI/ADI results** (median and credible bands; raw posterior-forward outputs accessible)
[ ] **Explainability drill (per flagged source)**: footprint decomposition, top edges, sensitivity summary included in archive

G. Governance trail and approvals (must)

[ ] **Decision-rule mapping** (AGI/ADI bands → governance actions; workflow IDs) used for this snapshot
[ ] **Change log of parameter/matrix edits** (who, when, why) with approver signatures prior to run for any non-default changes
[ ] **Action log** (any actions triggered by this snapshot: owner, task ID, SLA, closure evidence)
[ ] **Archive bundle location & hash** (where full snapshot bundle is stored; access controls listed)
[ ] **Retention policy** (how long artifacts retained and who may access)

H. Minimal reproducible run checklist (smoke test)

[ ] Export canonical register; record hash
[ ] Tag code commit; capture environment manifest
[ ] Record seeds and parameter file used
[ ] Run ETL to produce intermediate snapshot; record hash
[ ] Execute entry-point script; capture AGI/ADI outputs and hashes
[ ] Run validation smoke tests (e.g., synthetic seed check) and record results
[ ] Archive bundle and publish snapshot metadata (ID, hashes, access)

About the Author

Bob Prieto

Chairman & CEO
Strategic Program Management LLC
Jupiter, Florida, USA

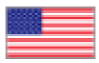

Bob Prieto is Chairman & CEO of Strategic Program Management LLC focused on strengthening engineering and construction organizations and improving capital efficiency in large capital construction programs. Previously, Bob was a senior vice president of Fluor, focused on the development, delivery, and 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.