Stochastic Risk Modeling for Infrastructure Asset Failure Under Uncertainty

Raj J Mehta
Independent Researcher

***Abstract*— Infrastructure owners increasingly rely on risk-based asset management to allocate limited budgets across large, heterogeneous portfolios. However, operational tools often compress uncertainty into deterministic risk matrices or point estimates, limiting auditability and obscuring tail risk. This paper formalizes a stochastic decision architecture for infrastructure asset failure risk that preserves uncertainty structure from deterioration, failure-time, and consequence models through to governance-ready decision outputs. The architecture couples stochastic deterioration representations (e.g., Markov/state-space or gamma-process) with hazard-based failure models and probabilistic consequence severity, propagating aleatory and epistemic uncertainty via Monte Carlo simulation or analytic approximations. Outputs include time-dependent failure probability, expected present value of loss, and tail-risk metrics (e.g., CVaR), alongside ranking-stability controls that quantify decision confidence and identify decision-unstable assets for targeted inspection or data improvement. A transparent demonstration portfolio (N=50 assets; synthetic but realistic ranges) establishes how uncertainty-preserving rankings can diverge from deterministic risk-matrix scores, breaking ties and revealing tail-driven priorities. A climate stress-test scenario, treated as a regulatory-style stress test rather than a physical climate model, illustrates portfolio loss amplification (approximately 25% increase in expected present value of failure loss in the demonstration) and supports scenario-ready capital planning. Overall, the proposed decision architecture enables defensible prioritization by producing auditable rankings, explicit separation of mean versus tail risk, and stability-aware governance artifacts.**

***Index Terms*— infrastructure asset management; auditable decision architecture; stochastic deterioration; hazard/survival analysis; uncertainty quantification; ranking stability; CVaR; stress testing**

I. INTRODUCTION

Transportation, water, energy, and civic infrastructure systems underpin economic activity and public safety, but they operate under persistent budget constraints and growing uncertainty. Asset management standards emphasize aligning investments with organizational objectives and risk, supported by documented assumptions and continuous improvement [1] [4]. In many jurisdictions, transportation agencies must also prepare risk-informed asset management plans and document processes for investment prioritization [5] [7].

Despite this shift toward risk-based governance, operational prioritization frequently relies on deterministic risk matrices (likelihood × consequence) built from coarse condition categories and point estimates. These tools are attractive because they are simple, but they can (i) produce large numbers of tied scores, (ii) confound uncertainty with variability, and (iii) conceal time dependence an asset with a moderate near-term likelihood may dominate long-term risk if failure probability accelerates nonlinearly. These limitations become critical when inspection data are sparse, measurement error is non-negligible, or stressors (traffic loading, extreme weather, material aging) are changing over time.

The contribution of this work is not a new deterioration or failure model, but the formalization of infrastructure risk assessment as an auditable decision architecture that preserves uncertainty structure, exposes tail risk, and produces stability-aware prioritization capabilities that are not jointly provided by existing infrastructure asset management frameworks.

As agencies shift from static, condition-based scoring to governance-driven, scenario-tested portfolio decisions, the absence of uncertainty-preserving and audit-ready risk architectures increasingly constrains accountability and investment defensibility.

Building on this need, the paper formalizes a governance-aligned stochastic decision architecture for infrastructure asset failure risk. The architecture integrates stochastic deterioration modeling, survival-based failure modeling, and probabilistic consequence modeling into a single uncertainty-preserving risk propagation pipeline (Figure 1). Rather than collapsing uncertainty into a single score, it produces governance-ready decision outputs: expected and tail-risk metrics (including Conditional Value-at-Risk, CVaR), time-dependent failure probability, and ranking-stability controls that expose decision confidence and identify assets that should be routed to additional inspection (value-of-information) before committing to major interventions.

1.1 Contributions and scope

This work makes four contributions: (i) It formalizes infrastructure risk assessment as an auditable decision architecture that preserves uncertainty structure from data and assumptions through to prioritization outputs (Figure 1). (ii) It introduces a stability aware prioritization control by defining a ranking stability metric and a governance decision rule that classifies decision-unstable assets for inspection or data improvement rather than immediate intervention. (iii) It states a minimal formal proposition (Proposition 1) explaining why deterministic risk matrix ties can be refined once variance and time dependence in failure time distributions are represented. (iv) It provides a fully specified, reproducible demonstration portfolio and stress-testing analysis showing how uncertainty representation alters rankings, separates expected versus tail risk, and enables scenario-ready planning artifacts.

Existing infrastructure asset management (IAM) tools and widely used risk-matrix workflows rarely provide, in a single coherent pipeline, (i) uncertainty-preserving risk quantification, (ii) explicit separation of expected risk versus tail risk, and (iii) stability-aware ranking outputs with an auditable trace from data and assumptions to prioritization. This gap is publishable because it constitutes a missing decision abstraction layer: a formal architecture that can sit above heterogeneous deterioration and failure models and convert them into governance-ready artifacts (audit trails, tail-risk controls, and decision-confidence metrics) at portfolio scale.

1.2 Paper organization

Section 2 summarizes related work on deterioration, failure, and risk-based asset management. Section 3 presents the proposed decision architecture, risk metrics, and stability controls. Section 4 applies the architecture to a transparent demonstration portfolio and reports ranking divergence, tail-risk exposure, and stress-test sensitivity. Section 5 discusses implementation pathways, governance use, and limitations, and Section 6 concludes.

II. RELATED WORK

Stochastic models have long been used to represent infrastructure deterioration and to support maintenance optimization. State-based Markov chain models estimate transition probabilities between discrete condition states and have been widely applied in pavement and bridge management [12] [14], [21]. However, practical implementations must address heterogeneity among assets and measurement error in inspection-derived condition ratings [15], [16].

To address partial observability and noisy inspections, latent or hidden Markov models treat the true condition state as unobserved and infer it from imperfect observations [17]. Alternative continuous-state representations include gamma processes, which model cumulative degradation as a non-decreasing stochastic process and support physically interpretable failure thresholds [11]. Time-series and state-space formulations provide another route to incorporate high-frequency sensor data and correlated errors [29].

Risk-based decision support requires combining deterioration and failure likelihood with consequences. Survival analysis provides a principled failure modeling framework, including proportional hazards regression [8] and nonparametric estimation of survival functions under censoring [9]. Infrastructure applications often adopt parametric baselines such as Weibull, enabling monotone increasing hazard consistent with aging and fatigue [10].

At the network level, optimization studies integrate stochastic deterioration into maintenance and inspection planning under budget and policy constraints [18] [20], [22], [24], [27], [30]. Yet many operational settings still compress uncertainty into deterministic risk matrices. Tail-risk measures such as conditional value-at-risk (CVaR) have been advocated in finance and engineering to provide coherent characterization of extreme outcomes and risk appetite [26]. Infrastructure governance frameworks also emphasize stress testing against plausible future scenarios, including climate-related hazards and degradation acceleration [28].

In contrast to single-model or single-hazard formulations, the present work emphasizes a modular pipeline that can be audited and adapted: the same risk propagation machinery can be paired with alternative deterioration/failure submodels, facilitating sensitivity analysis, institutional learning, and desk acceptable transparency.

III. STOCHASTIC RISK MODELING FRAMEWORK

The framework considers a portfolio of assets indexed by i = 1, N over a discrete planning horizon t = 0, T (years). For each asset, uncertain deterioration drives a time-varying failure likelihood and an uncertain consequence severity. Risk is propagated to decision metrics by simulation or analytic approximations, producing both expected and tail-risk measures.

3.1 Notation and decision context

Table 1 summarizes key notation used throughout the paper. The model is designed to support common governance questions, including: (i) Which assets should be prioritized for intervention under a given budget? (ii) How sensitive are priorities to uncertainty in inspections, model parameters, and future scenarios? (iii) Which priorities are robust (high confidence) versus borderline (high rank volatility)?

Table 1. Notation used in the stochastic risk modeling framework.

Symbol	Meaning
i	Asset index (i = 1, N)
t	Time index (years)
$X_{i(t)}$	Condition state or condition index of asset i at time t
$Y_{i(t)}$	Observed inspection/sensor measurement (possibly noisy)
$z_{i(t)}$	Covariate vector (e.g., traffic, environment, material)
$h_{i(t)}$	Instantaneous hazard rate of failure
$S_{i(t)}$	Survival function: $P(T_i > t)$
$p_{i(t)}$	Discrete-time failure probability over (t, t+1)
C_i	Random consequence (loss) given failure
$L_{i(t)}$	Random loss at time t
R_i	Risk metric (e.g., expected present value, CVaR)
α	Risk level for tail metrics (e.g., α=0.95 for CVaR95)

Figure 1 provides an overview of the modular pipeline. Each module can be replaced (e.g., Markov versus gamma-process deterioration) provided it produces a predictive distribution for future condition and/or failure.

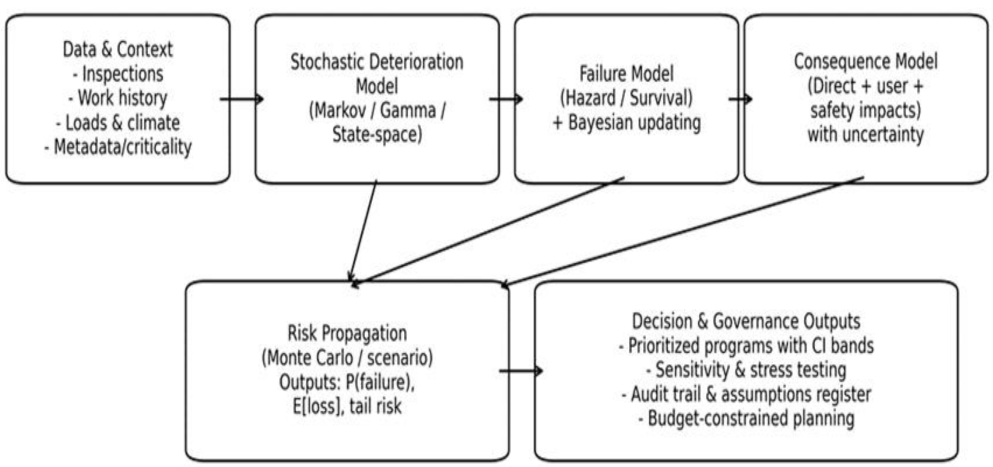

Figure 1. Modular stochastic risk modeling pipeline from data and assumptions to governance-ready outputs.

3.2 Stochastic deterioration model

Let $X_{i(t)}$ denote the latent condition of asset i at time t. Two common representations are supported:

(a) State-based Markov deterioration. Condition is discretized into K ordered states (e.g., 1 = failed/poor to K = excellent). The one-step transition probabilities are

$$\Pr X_i(t+1) = s' | X_i(t) = s = P^{(i)}_{ss'}, s, s' \in 1, \dots, K\}. \quad (1)$$

where the transition matrix $P^{\{(i)\}}$ may be asset-specific or stratified by covariates. Estimation from inspection ratings must account for heterogeneity and measurement error [15], [16]; Bayesian updating with conjugate Dirichlet priors provides an auditable mechanism to incorporate new data and quantify parameter uncertainty.

(b) Continuous condition index / state-space form. Many agencies track a continuous condition index (e.g., 0–100). A parsimonious stochastic evolution is

$$X_{i(t+1)} = \mathrm{clip}\big(X_{i(t)} - \mu_{i(t)} + \sigma_i \varepsilon_{i(t)}, 0, 100 \big), \ \varepsilon_{i(t)} \sim N(0,1). \quad (2)$$

where $\mu_{i(t)}$ captures expected deterioration (potentially dependent on age, environment, and interventions) and σ_i represents process variability. More physically constrained alternatives include gamma processes for monotone degradation [11]. Equation (2) is used in the case study to keep the demonstration reproducible while retaining stochastic trajectories (Figure 2).

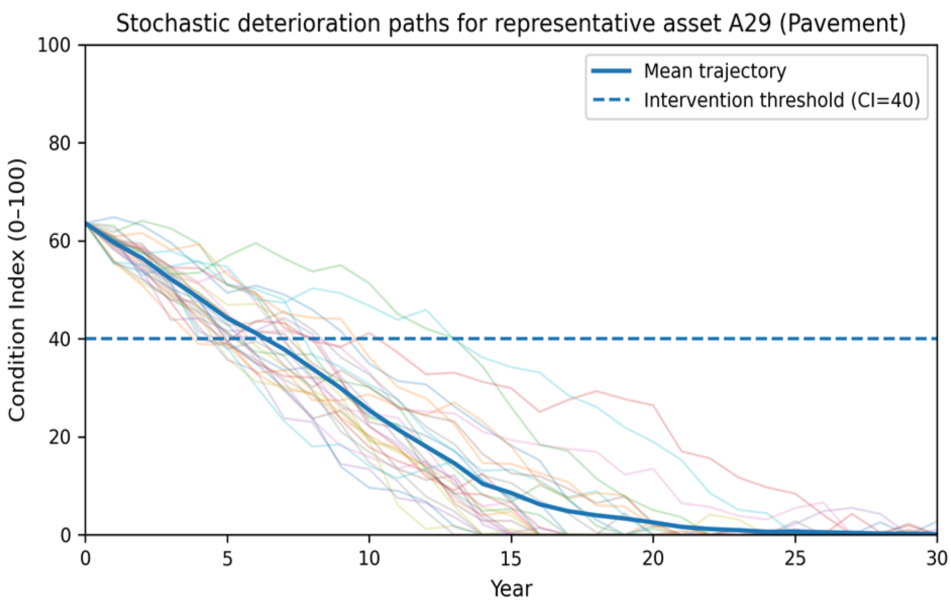

Figure 2. Example stochastic condition trajectories for representative asset A29 (synthetic case study). Dashed line indicates an intervention threshold.

3.3 Failure (hazard/survival) model

Let T_i denote the (random) failure time of asset i. Conditional on the latent condition $X_{i(t)}$ and covariates $z_{i(t)}$, failure likelihood is modeled via a hazard function. A flexible specification is the proportional hazards form [8]:

$$h_{i(t)} = h_{0(t)} \cdot \exp\big(\beta^T z_{i(t)} \big). \quad (3)$$

where $h_{0(t)}$ is a baseline hazard (often Weibull for aging assets [10]) and β are coefficients estimated from failure/near-failure records and condition histories. Under censoring and incomplete follow-up, survival analysis tools such as the Kaplan–Meier estimator remain applicable [9].

For discrete annual planning, the conditional probability of failure in year t→t+1 is approximated as

$$p_i(t) \approx 1 - \exp(-\int_t^{t+1} h_i\,(u)du) \quad (4)$$

In the case study, $z_{i(t)}$ includes a condition-derived term (e.g., (100− $X_{i(t)}$)/100), traffic loading ($\log(ADT)$), and a scenario factor that accelerates deterioration and hazard under climate stress. This structure separates (i) the stochastic evolution of condition from (ii) the mapping from condition and covariates to failure likelihood, enabling sensitivity analysis and governance transparency.

3.4 Probabilistic consequence (loss) model

Failure consequences can include direct repair/replacement costs, user disruption costs, safety impacts, environmental penalties, and reputational/regulatory impacts. In practice these components are uncertain even conditional on failure due to response time, detour effects, and secondary damage. We represent the total consequence as a nonnegative random variable C_i with a distribution conditioned on asset attributes (e.g., replacement value, criticality) and scenario.

For example, a lognormal severity model captures right-skewed cost behavior:

$$C_i | w_i \text{ LogNormal}(\mu_{C,i}, \sigma_{C,i}{}^2), E[C_i] = \exp(\mu_{C,i} + 0.5\sigma_{C,i}{}^2)$$

where θ_i denotes consequence drivers (replacement value, criticality, traffic). Alternative heavy-tail or mixture models can be substituted if needed.

3.5 Uncertainty taxonomy and propagation

Uncertainty in infrastructure risk assessment arises from both aleatory variability (inherent randomness) and epistemic uncertainty (limited knowledge). Table 2 summarizes common sources and how they enter the framework. In operational settings, explicitly documenting which uncertainties are modeled—and which are not is essential for auditability and reviewer scrutiny.

Table 2. Major uncertainty sources and how they are represented and propagated.

Source of uncertainty	Type	Typical examples	Representation in this framework
Deterioration process	Aleatory	Material variability; micro-environment; stochastic cracking/corrosion	Process noise in $X_{i(t)}$ (Eq. 1–2) or gamma increments
Inspection/measurement	Epistemic + Aleatory	Visual rating error; sensor bias/drift; missing data	Observation model $Y_{i(t)}$ with error; latent/hidden states [15]–[17]
Model form/structure	Epistemic	Markov vs gamma; covariate selection; stationarity assumptions	Modular submodels + model comparison/sensitivity
Parameter estimation	Epistemic	Limited failures; short records; transferability across regions	Bayesian posterior/credible intervals; bootstrap
Future scenarios	Epistemic	Traffic growth; climate acceleration; policy constraints	Scenario factors applied to $\mu_{i(t)}$, $h_{i(t)}$, and/or C_i [28]
Consequence severity	Aleatory + Epistemic	Response time; detours; cascading impacts	Probabilistic C_i model (Eq. 5) with scenario dependence

Uncertainty is propagated from these sources to risk metrics via Monte Carlo simulation or analytic approximations. Simulation is preferred for desk-acceptable transparency because it supports scenario testing, nonlinear hazard mappings, and heavy-tailed consequences with minimal additional derivation.

3.6 Risk metrics for governance and prioritization

Define the (discounted) loss incurred by asset i at time t as $L_i(t) = \delta(t) \cdot I_i(t) \cdot C_i(t)$, where $I_{i(t)}$ indicates that failure occurs in (t,t+1], $C_i(t)$ is the consequence severity, and $\delta(t) = (1 + r)^{-t}$ is a discount factor with rate r. The total present value loss over the horizon is

$L_i^{PV} = \sum_{t=0}^{T-1} L_i(t)$.

Three complementary decision metrics are recommended:

- Expected present value of loss: $E[L_i^{PV}]$. This supports cost-effective planning and is additive across independent assets.
- High-quantile loss (e.g., P90): quantiles of L_i^{PV} to represent low-probability, high-consequence outcomes.
- Tail risk using Conditional Value-at-Risk (CVaR): the mean loss conditional on exceeding a chosen quantile level α [26].

For CVaR at level α, a convenient optimization representation is [26]:

$$\text{CVaR}_\alpha(L) = \min_{\eta \in \mathbb{R}} \eta + (1/(1-\alpha)) \cdot E[(L - \eta)_+] \quad (6)$$

Ranking stability metric (top-k inclusion probability). Let r_i^m denote the rank of asset i in Monte Carlo replicate m (m = 1, M) under a selected risk measure (e.g., expected present value $E[L_i^{PV}]$ or tail risk $CVaR_\alpha$). For a governance-relevant cutoff k, define the top-k inclusion probability:

$$\pi_i^k = P(r_i \leq k) \approx (1/M) \cdot \sum_{m=1}^{M} 1\, r_i^m \leq k$$

High π_i^k indicates a decision-stable asset that persistently appears among the top-k priorities; low π_i^k signals proximity to the decision boundary and motivates governance actions (e.g., targeted inspection) rather than immediate intervention.

Decision rule (stability control). For a chosen cutoff k (e.g., k = 10) and governance threshold τ (e.g., τ = 0.60), classify asset i as decision-unstable if $\pi_i^k < \tau$. Decision-unstable assets are prioritized for inspection, monitoring, or data improvement to reduce uncertainty before committing to major interventions; decision-stable assets proceed to intervention programming. This turns stability from a descriptive statistic into an explicit governance control.

Proposition 1 (Ranking Refinement under Uncertainty).

For two assets with identical deterministic risk-matrix scores, if their failure-time distributions differ in either (i) variance or (ii) time-dependence of hazard, then there exists a planning horizon and a risk measure (expected loss or CVaR) under which their stochastic risk rankings differ.

Intuitive justification.

Risk matrices map heterogeneous distributions into coarse bins and can assign the same likelihood and consequence categories to assets with meaningfully different uncertainty. A higher-variance loss or failure-time distribution can increase tail measures (e.g., CVaR) without changing the mean, while time-varying hazards interact with discounting and planning horizon so that near-term and long-term risks trade off differently across assets. Therefore, once uncertainty structure is preserved, ties induced by deterministic scoring can be refined into decision-relevant orderings.

3.7 Algorithm summary

Algorithm 1 summarizes the computational procedure for producing asset-level and portfolio-level risk metrics. The structure is intentionally modular: replacing a deterioration or failure module requires only that it can simulate or approximate conditional failure probabilities and losses.

Algorithm 1. Monte Carlo risk propagation for an asset portfolio.

Inputs: portfolio data {age_i, condition observations Y_i, covariates z_i}, model parameters/posteriors, planning horizon T, discount rate r, scenarios s, top-k value k, and stability threshold τ.

For each scenario s:
1. For each asset i:
   a. Initialize latent condition $X_{i(0)}$ (or its posterior) from inspection data.
   b. For simulation m = 1,M:
      i. Simulate deterioration path $X_i^{(m)}(t)$ for t = 0,…,T.
      ii. Compute hazard $h_i^{(m)}(t)$ and failure probability $p_i^{(m)}(t)$.
      iii. Sample failure event time (if any) and consequence $C_i^{(m)}$.
      iv. Compute discounted loss $L_i^{PV,(m)}$.
   c. Estimate risk metrics: $E[L_i^{PV}]$, quantiles, $CVaR_\alpha$, and top-k inclusion probability $\pi_i^{(k)}$.
   d. Classify stability: if $\pi_i^{(k)} < \tau$, flag asset i as decision-unstable (prioritize inspection/data improvement).
2. Aggregate portfolio metrics: sum expected losses; compute portfolio quantiles/CVaR via joint simulation if dependence is modeled.

Outputs: time-dependent P(failure), risk metrics, prioritized list with confidence bands, stability flags, and scenario sensitivity.

IV. DEMONSTRATION PORTFOLIO CASE STUDY

This section applies the decision architecture to a demonstration portfolio of N=50 assets with realistic ranges of age, condition, traffic exposure, replacement cost, and criticality. The dataset is synthetic to avoid data access restrictions, but all parameter values and assumptions are stated so that results are reproducible and can be adapted to agency data.

This demonstration is intentionally not a predictive validation study. Its purpose is to make decision divergence, uncertainty propagation, and auditability visible: the analysis shows how rankings, tail risk, and stability flags change when uncertainty structure is preserved. Claims about absolute predictive accuracy are explicitly out of scope.

4.1 Portfolio description and synthetic data generation

Assets represent a mix of bridges, culverts, pavement segments, retaining walls, and ITS/sign assets. Initial condition indices span 15–95 (0–100 scale), ages span 5–80 years, and average daily traffic (ADT) spans roughly 500–120,000 vehicles/day. Table 3 lists a subset of assets used to illustrate high- and low-risk cases. Although synthetic, the ranges were selected to match typical magnitudes reported in DOT asset inventories (condition indices, ADT exposures, and replacement costs), ensuring external plausibility while preserving reproducibility.

Table 3. Sample of synthetic portfolio assets used for demonstration (subset of N=50).

AssetID	AssetType	Age yr	ConditionIndex 0 100	ADT	ReplacementCost USD	Criticality 0 1
A29	Pavement	39	63.6	10293	7,576,195	0.96
A37	Culvert	76	27.1	7073	13,465,687	0.21
A46	Pavement	39	59.6	7618	6,985,173	0.41
A15	Bridge	30	73.7	18548	15,473,948	0.23
A13	Retaining Wall	57	41.6	9308	2,913,788	0.81
A28	Pavement	12	84.9	971	486,468	0.83
A32	Bridge	9	86.8	12033	1,606,931	0.11
A14	Bridge	28	75.8	6879	366,144	0.41
A35	Sign/ITS	11	95.0	4244	1,783,010	0.85
A24	Culvert	5	88.5	3270	596,889	0.14

4.2 Model specification for the demonstration

The deterioration model uses the continuous condition formulation in Eq. (2) with type-dependent drift and process variability. The failure model uses a condition-driven hazard mapping consistent with Eq. (3) (4), with covariates including a condition deficit term, log (ADT), and a scenario factor. Consequences are modeled as lognormal (Eq. 5) with mean proportional to replacement cost and criticality plus a user-cost proxy based on ADT. Monte Carlo simulation uses M = 6,000 draws per asset over T = 30 years with discount rate r = 3%. In the demonstration, full portfolio analysis completes within minutes on standard desktop hardware, supporting annual updates and rapid scenario-based re-runs for governance review.

Equation (2) is used here as a didactic surrogate for Markov-chain or gamma-process deterioration; it is intentionally simple to keep the full end-to-end pipeline reproducible in a compact demonstration.

While parameter estimation is outside the scope of this synthetic demonstration, the adopted structure is consistent with prior infrastructure risk and maintenance optimization literature that integrates deterioration, measurement error, and decision-making under uncertainty [12]–[20], [22], [24].

4.3 Asset-level risk outputs

Simulation outputs include time-dependent failure probability, annual expected loss, and present-value risk metrics. As an example, asset A29 (a pavement segment in this portfolio) has an estimated 30-year failure probability of 88.0% and an expected present value of failure loss of $5.11M. Its CVaR95 is $18.84M, highlighting a substantial tail exposure relative to the mean.

These outputs support multiple planning viewpoints: mean-based prioritization (efficient allocation) and tail-based governance (risk appetite, safety-critical assets). They also enable inspection planning by identifying assets whose ranking is sensitive to uncertain condition (high value-of-information).

4.4 Portfolio prioritization: deterministic risk matrix vs stochastic ranking

To reflect common practice, a deterministic 5×5 risk matrix is constructed from condition categories (proxy for likelihood) and consequence categories (proxy for severity). Because the matrix compresses continuous variables into bins, it produces many tied scores (e.g., in this portfolio, the most frequent risk-matrix scores occur in groups of 6–9 assets). In contrast, the stochastic model yields a continuous expected-loss ranking and provides uncertainty-aware metrics.

Figure 3. Comparison of deterministic risk-matrix rank and stochastic expected-loss rank for the demonstration portfolio (N=50). The dashed line indicates agreement.

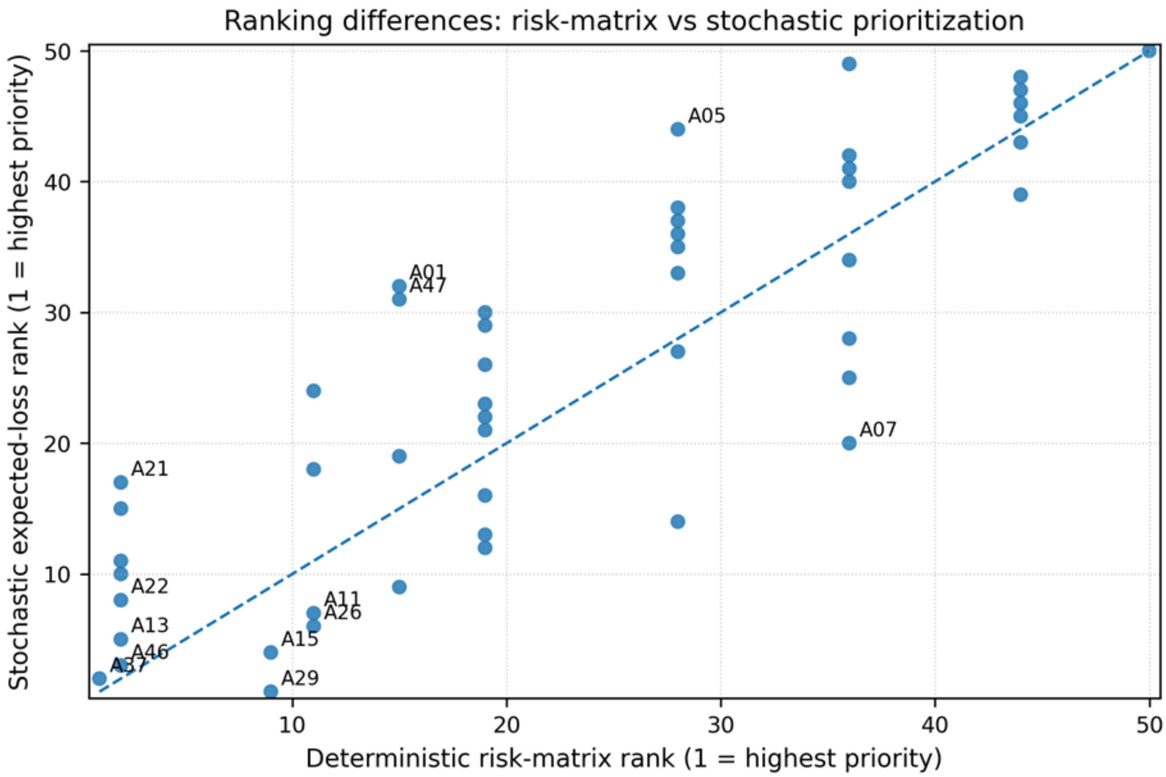

Figure 3 shows moderate but imperfect agreement between the two approaches (Spearman correlation 0.852). Notably, asset A29 is ranked first by stochastic expected loss but only ninth by the risk matrix, illustrating how discretization can obscure high-risk assets when many items share similar qualitative ratings.

Table 4. Top assets by stochastic expected present value loss (baseline scenario), with corresponding deterministic risk-matrix score and rank.

AssetID	Type	CI	RM score	RM rank	Stoch rank	EAL PV ($M)	CVaR95 ($M)	P_{fail} (30yr)
A29	Pavement	63.6	10	9	1	5.11	18.84	0.880
A37	Culvert	27.1	20	1	2	5.02	16.69	0.951
A46	Pavement	59.6	12	2	3	2.80	10.10	0.878
A15	Bridge	73.7	10	9	4	2.70	14.00	0.551
A13	Retaining Wall	41.6	12	2	5	2.00	7.18	0.893
A26	Pavement	43.5	9	11	6	1.84	6.26	0.972
A11	Bridge	42.1	9	11	7	1.62	5.53	0.949
A22	Bridge	27.7	12	2	8	1.42	4.90	0.978
A17	Culvert	63.0	8	15	9	1.40	6.49	0.627
A36	Retaining Wall	29.9	12	2	10	1.31	4.77	0.915
A06	Bridge	39.2	12	2	11	1.29	4.46	0.968
A33	Bridge	60.1	6	19	12	1.25	4.90	0.856

Table 4 shows how tail-risk (CVaR95) can be orders of magnitude above the mean for high-consequence assets, supporting governance discussions about risk appetite, contingency planning, and whether to prioritize interventions that primarily reduce extreme outcomes rather than expected cost.

Table 5. Ranking stability for top-ranked assets using top-k inclusion probability π_i^{10} (k=10, τ=0.60; demonstration portfolio).

AssetID	Stochastic rank	π_i^{10}	Stability class
A29	1	1.00	Decision-stable
A37	2	1.00	Decision-stable
A46	3	0.98	Decision-stable
A15	4	0.85	Decision-stable
A13	5	0.84	Decision-stable
A26	6	0.79	Decision-stable
A11	7	0.65	Decision-stable
A22	8	0.47	Decision-unstable
A17	9	0.41	Decision-unstable
A36	10	0.37	Decision-unstable

Table 5 makes ranking stability explicit: assets near the top-10 boundary can be decision-unstable ($\pi_i^{10} < \tau$), indicating that their prioritization is sensitive to uncertainty and that inspection or data improvement may be a higher-value first action than immediate intervention. This governance output is not available from deterministic risk matrices.

4.5 Scenario stress testing: climate-accelerated deterioration

To illustrate stress testing, a climate-accelerated scenario is applied by increasing deterioration drift and variability and adding a hazard multiplier term. The climate scenario is treated as a regulatory-style stress test rather than a physical climate model: the multiplier is an abstraction used to evaluate decision robustness under adverse conditions, analogous to financial stress testing practice [31]. Such scenario factors can represent combined effects of temperature, moisture, freeze-thaw cycles, and more frequent extreme events, consistent with guidance to consider future conditions in risk management [4], [28].

At the portfolio level, the expected present value of failure loss increases by approximately 25.2% (from $43.58M to $54.57M). At the asset level, expected losses increase by a mean of 38.1% (median 24.0%), and the mean increase in 30-year failure probability is 0.122.

Table 6. Sensitivity of top-ranked assets to climate stress-test deterioration (baseline vs stress test).

AssetID	Type	CI	EAL PV base ($M)	EAL PV clim ($M)	Increase (%)	P_{fail} base	P_{fail} clim	ΔP_{fail}
A29	Pavement	63.6	5.11	6.08	18.9	0.880	0.962	0.082
A37	Culvert	27.1	5.02	5.52	9.9	0.951	0.979	0.028
A46	Pavement	59.6	2.80	3.36	19.7	0.878	0.955	0.076
A15	Bridge	73.7	2.70	4.09	51.5	0.551	0.786	0.235
A13	Retaining Wall	41.6	2.00	2.31	15.3	0.893	0.956	0.062
A26	Pavement	43.5	1.84	2.01	9.5	0.972	0.992	0.020
A11	Bridge	42.1	1.62	1.78	9.8	0.949	0.981	0.031
A22	Bridge	27.7	1.42	1.50	5.8	0.978	0.992	0.014
A17	Culvert	63.0	1.40	1.83	30.7	0.627	0.811	0.184
A36	Retaining Wall	29.9	1.31	1.42	8.4	0.915	0.964	0.049

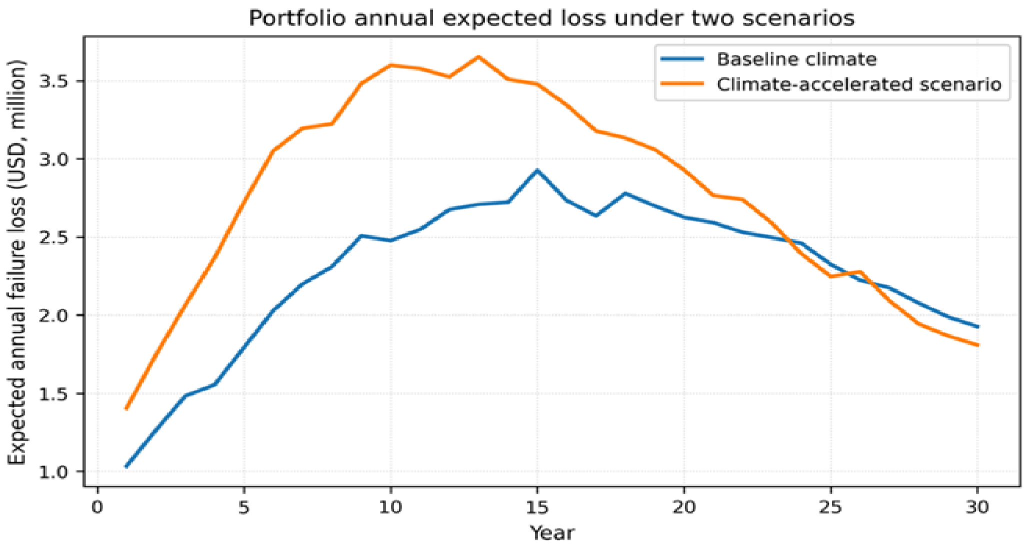

Figure 4. Portfolio annual expected failure loss under baseline and climate stress-test scenarios (Demonstration portfolio).

Figure 4 shows that under the climate stress-test scenario, expected annual losses rise earlier and peak higher, reflecting faster condition decline and increased failure likelihood. This type of scenario output is directly actionable for multi-year capital and maintenance planning because it identifies when risk accelerates and whether deferring interventions is likely to shift costs into higher-risk periods.

Figure 5 complements the scenario comparison by providing a one-factor sensitivity analysis in which the hazard multiplier κ is scaled around the baseline. The monotonic response in portfolio expected present value of failure loss provides a simple governance artifact for stress-test calibration and confirms that conclusions in this demonstration are not an artifact of a single multiplier choice.

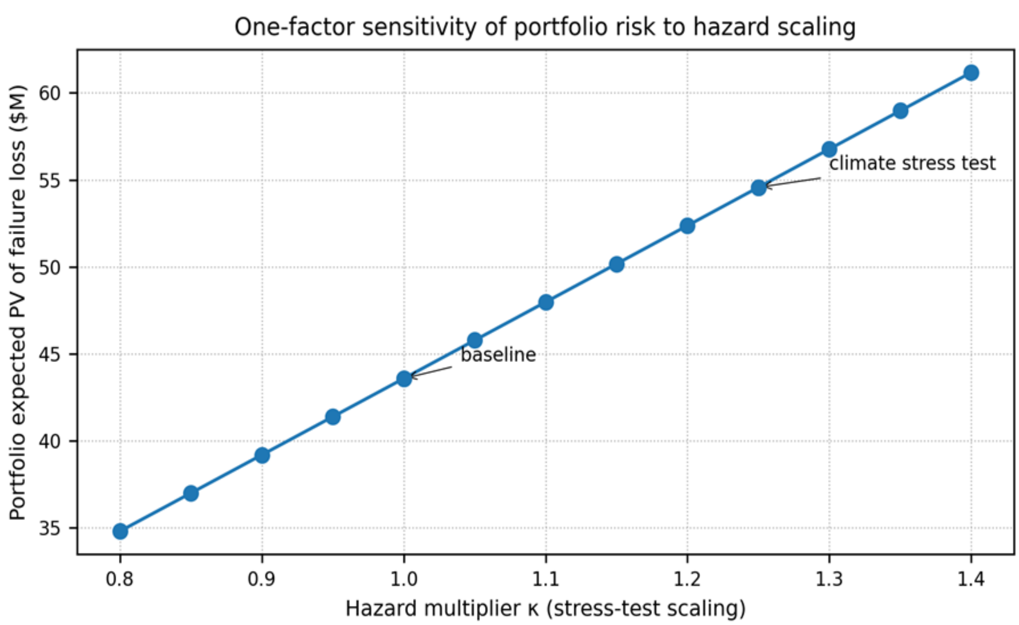

Figure 5. One-factor sensitivity of portfolio expected present value of failure loss to a multiplicative hazard scaling κ (demonstration portfolio). Baseline (κ = 1.0) and climate stress-test (κ = 1.25) points correspond to Section 4.5.

V. DISCUSSION: IMPLEMENTATION, GOVERNANCE, AND LIMITATIONS

5.1 Practical implementation pathway

Agencies can adopt the proposed framework incrementally: (1) Start with a transparent baseline deterioration model (Markov or condition-index drift) calibrated from historical inspections, explicitly acknowledging measurement error [15], [16]. (2) Add a failure model using available failure/repair triggers, near-failure events, or proxy thresholds; survival models support censoring and varying follow-up [8], [9]. (3) Replace deterministic consequence point estimates with distributions informed by work order histories and user-cost proxies. (4) Implement Monte Carlo risk propagation with scenario factors for stress testing and update parameters as new inspections arrive (Bayesian updating).

Box 1. Implementation checklist (minimum viable deployment).

> Minimum data required (per asset): unique asset ID; asset type/class; installation year/age; latest condition rating (and inspection date); exposure proxy (e.g., ADT, customers served); consequence proxies (replacement cost and criticality); intervention history if available.
>
> Outputs produced (per analysis cycle): time-dependent failure probability; expected present value of loss; tail risk (e.g., CVaR); ranking stability π_i^k; scenario stress-test deltas; and an audit log (assumptions register, parameter sources, random seed, and model versions).
>
> Annual update pathway: ingest new inspections; update condition distributions (Bayesian updating or bootstrapped measurement-error model); refresh hazard mapping if covariates changed; rerun simulations; publish a change log including ranking stability shifts and decision-unstable assets flagged for targeted inspection.

From a governance perspective, the key deliverable is not only a ranking, but also an audit trail: which data sources were used, how uncertainty was represented, which scenarios were tested, and how sensitive priorities are to assumptions. This aligns with asset management standards that emphasize documented decision processes and continual improvement [1][4].

What this enables that conventional IAM tools typically do not:

- Auditable prioritization under uncertainty, with an explicit trace from data sources and assumptions to rankings.
- Explicit separation of expected risk versus tail risk (e.g., CVaR) to support risk-appetite and safety governance.
- Identification of decision-unstable assets (low π_i^k) and a defensible rule to route them to inspection or data improvement.
- Scenario-ready governance artifacts (confidence bounds, stress-test outputs, sensitivity curves, and assumptions register) that support ex-ante and ex-post review.

5.2 Interdependence and cascading effects

The demonstration case study assumes independent asset failures. In real networks, dependencies can arise through shared hazards (flooding, seismic events), operational coupling (detours shifting loads), or correlated deterioration drivers. The modular framework can incorporate dependence by sampling shared scenario variables (e.g., regional climate multipliers) or by explicitly modeling common-cause events. Where cascading failures are plausible, consequence models should include system-level impacts rather than only component-level repair costs. Because portfolio tail metrics (e.g., CVaR) are sensitive to dependence, the independence assumption can materially under- or over-estimate portfolio tail risk; the same architecture supports shared-factor dependence by sampling common hazard multipliers or latent drivers across assets.

5.3 Limitations and recommended extensions

Three limitations should be noted. First, parameter estimation and validation require sufficient historical data; sparse failure records may necessitate hierarchical pooling across asset classes or expert elicitation. Second, model-form uncertainty can be material; agencies should compare alternative submodels (Markov vs gamma; parametric vs semi-parametric hazard) and report sensitivity. Third, deterministic program constraints (equity, minimum service levels, bundling of projects) must be integrated via optimization once risk metrics are produced.

Recommended extensions include: (i) joint optimization of inspection and intervention decisions under model uncertainty [19], [20]; (ii) explicit treatment of maintenance effectiveness and post-intervention condition jumps; and (iii) integration with multi-objective optimization to balance stakeholder costs and user costs [24], [27].

VI. CONCLUSIONS

Rather than proposing a new deterioration or failure model, this paper formalizes infrastructure risk assessment as an auditable decision architecture that preserves uncertainty structure, makes tail risk explicit, and produces stability-aware prioritization outputs. These capabilities are not jointly provided by common infrastructure asset management frameworks, yet they are essential for accountable, governance-driven portfolio decisions under uncertainty.

The demonstration portfolio establishes that deterministic risk matrices can mask tail-driven priorities and induce large tied groups, while uncertainty-preserving prioritization refines ties into auditable, decision-relevant orderings. Stress-test analysis indicates that plausible hazard acceleration can materially increase both expected and tail risk, reinforcing the value of scenario-ready outputs for multi-year capital planning. Future work should focus on calibration and validation with agency datasets, explicit modeling of dependence and cascading impacts, and integration with budget-constrained optimization and value-of-information analysis for inspection planning. Because the architecture is model-agnostic, it can be paired with agency-specific deterioration, hazard, and consequence models without altering governance outputs.

Appendix A. Reproducibility and audit-trail checklist

To support desk-acceptable transparency and institutional auditability, an implementation of the proposed decision architecture should record the following items for each model run:

- Portfolio snapshot and data dictionary (asset identifiers, attributes, units, and missing-data handling).
- Model form selections (deterioration representation, hazard specification, and consequence severity distribution) and rationale.
- Parameter sources and uncertainty treatment (priors/posteriors, confidence intervals, and calibration windows).
- Scenario definitions (e.g., hazard multipliers, deterioration drift adjustments) and justification as stress tests.
- Simulation settings (planning horizon T, number of draws M, random seed/replication policy) and discounting assumptions.
- Decision settings (risk measure used for ranking, α for tail metrics, k and τ for stability control) and resulting stability flags.

Data availability

The case study dataset in this manuscript is synthetic and generated from stated distributions and parameters for reproducibility. All assumptions required to reproduce the figures and tables are documented in Section 4. Agency-specific applications should include a data dictionary and versioned parameter set to preserve the audit trail.

Declaration of competing interests

[The authors declare no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.]

Acknowledgments

[Optional: funding sources, institutional support, and reviewer acknowledgments to be inserted upon submission.]

REFERENCES

[1] International Organization for Standardization (ISO), ISO 55000:2014, Asset management - Overview, principles and terminology. Geneva, Switzerland, 2014.

[2] International Organization for Standardization (ISO), ISO 55001:2014, Asset management - Management systems - Requirements. Geneva, Switzerland, 2014.

[3] International Organization for Standardization (ISO), ISO 55002:2018, Asset management - Management systems - Guidelines for the application of ISO 55001. Geneva, Switzerland, 2018.

[4] International Organization for Standardization (ISO), ISO 31000:2018, Risk management Guidelines. Geneva, Switzerland, 2018.
[5] U.S. Code of Federal Regulations, 23 CFR Part 515 Asset Management Plans, Federal Highway Administration, 2016 (as amended).
[6] Federal Highway Administration (FHWA), “State Asset Management Plan Under the Bipartisan Infrastructure Law,” Memorandum, May 5, 2022.
[7] American Association of State Highway and Transportation Officials (AASHTO), Transportation Asset Management Guide: A Focus on Implementation, 2nd ed. Washington, DC, 2020.
[8] D. R. Cox, “Regression Models and Life-Tables,” Journal of the Royal Statistical Society: Series B (Methodological), vol. 34, no. 2, pp. 187–220, 1972, doi:10.1111/j.2517-6161.1972.tb00899.x.
[9] E. L. Kaplan and P. Meier, “Nonparametric Estimation from Incomplete Observations,” Journal of the American Statistical Association, vol. 53, no. 282, pp. 457–481, 1958, doi:10.1080/01621459.1958.10501452.
[10] W. Weibull, “A Statistical Distribution Function of Wide Applicability,” Journal of Applied Mechanics, vol. 18, no. 3, pp. 293–297, 1951, doi:10.1115/1.4010337.
[11] J. M. van Noortwijk, “A survey of the application of gamma processes in maintenance,” Reliability Engineering & System Safety, vol. 94, no. 1, pp. 2–21, 2009, doi:10.1016/j.ress.2007.03.019.
[12] S. Madanat, “Optimal infrastructure management decisions under uncertainty,” Transportation Research Part C: Emerging Technologies, vol. 1, no. 1, pp. 77–88, 1993, doi:10.1016/0968-090X (93)90021-7.
[13] S. Madanat, R. Mishalani, and W. H. Wan Ibrahim, “Estimation of infrastructure transition probabilities from condition rating data,” Journal of Infrastructure Systems, vol. 1, no. 2, pp. 120–125, 1995, doi:10.1061/(ASCE)1076-0342(1995)1:2(120).
[14] R. G. Mishalani and S. M. Madanat, “Computation of infrastructure transition probabilities using stochastic duration models,” Journal of Infrastructure Systems, vol. 8, no. 4, pp. 139–148, 2002, doi:10.1061/(ASCE)1076-0342(2002)8:4(139).
[15] M. Ben-Akiva, F. Humplick, S. Madanat, and R. Ramaswamy, “Infrastructure management under uncertainty: latent performance approach,” Journal of Transportation Engineering, vol. 119, no. 1, pp. 43–58, 1993, doi:10.1061/(ASCE)0733-947X(1993)119:1(43).
[16] F. Humplick, “Highway pavement distress evaluation: modeling measurement error,” Transportation Research Part B: Methodological, vol. 26, no. 2, pp. 135–154, 1992, doi:10.1016/0191-2615(92)90004-G.
[17] K. Kobayashi, K. Kaito, and N. Lethanh, “A statistical deterioration forecasting method using hidden Markov model for infrastructure management,” Transportation Research Part B: Methodological, vol. 46, no. 4, pp. 544–561, 2012, doi:10.1016/j.trb.2011.12.004.
[18] K. Smilowitz and S. Madanat, “Optimal inspection and maintenance policies for infrastructure networks,” Computer-Aided Civil and Infrastructure Engineering, vol. 15, no. 1, pp. 5–13, 2000, doi:10.1111/0885-9507.00166.
[19] V. M. Guillaumot, P. L. Durango-Cohen, and S. M. Madanat, “Adaptive optimization of infrastructure maintenance and inspection decisions under performance model uncertainty,” Journal of Infrastructure Systems, vol. 9, no. 4, pp. 133–139, 2003, doi:10.1061/(ASCE)1076-0342(2003)9:4(133).
[20] P. L. Durango Cohen and S. M. Madanat, “Optimization of inspection and maintenance decisions for infrastructure facilities under performance model uncertainty: A quasi-Bayes approach,” Transportation Research Part A: Policy and Practice, vol. 42, no. 8, pp. 1074–1085, 2008, doi:10.1016/j.tra.2008.03.004.
[21] A. Morcous, “Performance prediction of bridge deck systems using Markov chains,” Journal of Bridge Engineering, vol. 20, no. 2, pp. 146–155, 2006,

doi:10.1061/(ASCE)0887-3828(2006)20:2(146).

[22] C. Robelin and S. Madanat, “History-dependent bridge deck maintenance and replacement optimization with Markov decision processes,” Journal of Infrastructure Systems, vol. 13, no. 3, pp. 195–201, 2007, doi:10.1061/(ASCE)1076-0342(2007)13:3(195).

[23] J. O. Sobanjo, “State transition probabilities in bridge deterioration based on Weibull sojourn times,” Structure and Infrastructure Engineering, 2011, doi:10.1080/15732470902917028.

[24] A. Orcesi and C. Cremona, “A bridge network maintenance framework for Pareto optimization of stakeholders/users costs,” Reliability Engineering & System Safety, vol. 95, no. 11, pp. 1230–1243, 2010, doi:10.1016/j.ress.2010.06.013.

[25] D. M. Frangopol, “Life-cycle performance, management, and optimization of structural systems under uncertainty: accomplishments and challenges,” Structure and Infrastructure Engineering, 2011, doi:10.1080/15732471003594427.

[26] R. T. Rockafellar and S. Uryasev, “Conditional value-at-risk for general loss distributions,” Journal of Banking & Finance, vol. 26, no. 7, pp. 1443–1471, 2002, doi:10.1016/S0378-4266(02)00271-6.

[27] M. van den Boomen, M. T. J. Spaan, Y. Shang, and A. R. M. Wolfert, “Infrastructure maintenance and replacement optimization under multiple uncertainties and managerial flexibility,” Construction Management and Economics, vol. 38, no. 1, pp. 91–107, 2020, doi:10.1080/01446193.2019.1674450.

[28] Intergovernmental Panel on Climate Change (IPCC), Climate Change 2022: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Sixth Assessment Report of the IPCC. Cambridge University Press, 2022, doi:10.1017/9781009325844.

[29] P. L. Durango-Cohen, “A time series analysis framework for transportation infrastructure management,” Transportation Research Part B: Methodological, vol. 41, no. 5, pp. 493–505, 2007, doi:10.1016/j.trb.2006.08.002.

[30] K. D. Kuhn and S. M. Madanat, “Model uncertainty and the management of a system of infrastructure facilities,” Transportation Research Part C: Emerging Technologies, vol. 13, nos. 5–6, pp. 391–404, 2005.

[31] Basel Committee on Banking Supervision (BCBS), Principles for sound stress testing practices and supervision. Bank for International Settlements, May 2009.