

Digital-Twin-Guided Multi-Agent Resilience Finance

Connor Noble

c0613727@gmail.com

Research Article

Keywords: digital twin, agent-based modeling, climate-resilient infrastructure, model predictive control, green finance, deep uncertainty

Posted Date: September 4th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10789792/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Digital-Twin-Guided Multi-Agent Resilience Finance

Robust model predictive control for climate-stressed infrastructure networks

Original Research Article | Anonymous manuscript for peer review

Corresponding Author: Connor Noble

Abstract

Climate-resilient infrastructure finance is difficult because physical condition, hazard intensity, stakeholder behavior, and policy capacity evolve together. This study develops a digital-twin-guided multi-agent computational framework in which noisy infrastructure observations are assimilated into node-level risk forecasts, heterogeneous agents update adaptation behavior through social spillovers and financial incentives, and a robust model-predictive-control (MPC) rule reallocates a constrained green-finance budget across a dependency network. The synthetic system contains 72 infrastructure nodes and four governance regimes evaluated over 900 heavy-tailed climate-policy scenarios: a fixed baseline, digital-twin-only monitoring, finance-only reallocation using noisy forecasts, and fully coupled digital-twin plus adaptive finance control. The coupled regime lowers mean criticality-weighted service loss from 55.01% to 53.74% and equity-weighted residual risk from 18.41% to 17.59%, while increasing equilibrium adaptation from 34.25% to 36.45%. Digital-twin-only monitoring performs similarly in extreme-tail service loss, revealing an important policy boundary: better targeting can improve average and distributional outcomes without guaranteeing tail dominance when cascading failures are strong. Finance-only reallocation is less reliable because noisy state estimates mistarget scarce resources. The study contributes an interpretable computational architecture linking digital twins, agent-based adaptation, network cascades, equity constraints, and robust finance allocation. The results emphasize that climate-resilience technology is most useful when its state estimates are explicitly connected to adaptive financing rules and stress-tested under model uncertainty.

Keywords: digital twin; agent-based modeling; climate-resilient infrastructure; model predictive control; green finance; deep uncertainty

1 Introduction

1.1 Infrastructure resilience as a cyber-physical-financial system

Climate-resilient infrastructure is increasingly managed through a combination of sensing, digital models, operational adaptation, and financing. Yet these components are frequently assessed separately. A digital twin may improve situational awareness without changing where capital is allocated. A resilience fund may provide resources without a sufficiently granular state estimate. Stakeholder adaptation may respond to perceived risk, peer behavior, affordability, and policy credibility rather than to the planner's forecast. When infrastructure dependencies create cascading failures, these separations become economically consequential.

The special issue on technological and financial innovation in climate risk governance explicitly identifies digital twins, simulation, dynamic optimization, multi-agent modeling, and computational evaluation of climate-resilient infrastructure as priority methods (Song et al., 2026). This paper responds by integrating all five in one computational experiment. The framework represents physical assets as nodes in a dependency network, uses a Bayesian linear update as the digital-twin state estimator, models adaptation as an endogenous logistic fixed point with social spillovers, and allocates green finance with a risk- and equity-weighted receding-horizon rule.

The study is designed as a method-development experiment rather than a case-study forecast. No claim is made that the synthetic network represents a specific power grid, transport system, or municipality. This controlled setting allows the true hazard state to be observed by the analyst, while the policy sees only noisy measurements. It therefore enables direct evaluation of a core digital-twin proposition: whether improved state estimation translates into better decisions after behavioral and financial feedback are included.

1.2 Research questions and novelty

Four questions structure the analysis. RQ1 asks how much digital-twin state estimation reduces service disruption when finance remains fixed. RQ2 asks whether risk-weighted finance helps when the forecast used for targeting is noisy. RQ3 asks whether endogenous adaptation and network spillovers create threshold behavior that changes policy ranking. RQ4 asks whether a coupled digital-twin and finance controller improves both aggregate resilience and equity-weighted risk across a broad ensemble of climate-policy states.

The first methodological novelty is the feedback between the cyber layer and the economic layer. A digital twin is represented not as a static visualization but as a state-estimation process that updates node-level hazard forecasts. The second novelty is behavioral endogeneity: households, firms, public agencies, and finance intermediaries are abstracted into heterogeneous adaptation propensities that respond to risk signals, subsidies, and neighboring adoption. The third novelty is a constrained finance controller that assigns a minimum floor to every node while concentrating the remaining budget on high-risk, high-criticality, and high-disadvantage nodes.

A fourth contribution is the explicit treatment of policy limits. Rather than force the coupled regime to dominate every outcome, the simulation allows stronger network cascades to erode gains in the extreme tail. This matters for decision making under deep uncertainty. Robustness should mean acceptable performance over diverse plausible futures, not an implausible guarantee that one policy is optimal in every realization (Lempert et al., 2006; Marchau et al., 2019).

- H1: Digital-twin feedback improves node-risk ranking and raises adaptation where signal precision is high.
- H2: Finance-only targeting can underperform when noisy forecasts concentrate resources on the wrong nodes.
- H3: Coupled monitoring and finance reduce average service loss and equity-weighted risk relative to the baseline.
- H4: Strong dependency cascades create a tail-risk boundary in which monitoring improvements may matter as much as financial retargeting.

2 Literature and conceptual foundation

2.1 Digital twins for climate resilience

Digital twins are increasingly proposed as decision-support systems for infrastructure exposed to climate stress. Recent reviews emphasize real-time synchronization, predictive analytics, and the integration of operational data with simulation models (Ali et al., 2025; Cavus et al., 2026; Zhu & Jin, 2025). Work on climate-neutral and resilient cities similarly frames urban digital twins as tools for connecting heterogeneous data to planning decisions (Ricciardi & Callegari, 2023). The common promise is a shift from static diagnosis to iterative sensing, forecasting, intervention, and evaluation.

However, a digital twin does not create resilience by observation alone. The economic value of a better state estimate depends on whether operational rules, budgets, and institutions can respond. This paper therefore treats the twin as an information engine inside a larger governance loop. The digital representation updates risk beliefs; those beliefs alter adaptation incentives and finance allocation; interventions change failure probability; and realized failures feed the performance metrics used for subsequent policy design.

2.2 Agent-based adaptation and insurance-finance interactions

Agent-based models are useful when adaptation depends on heterogeneous thresholds, bounded rationality, social influence, and policy incentives. Jenkins et al. (2017), for example, use an agent-based framework to study flood-risk management, homeowner protection, government investment, and insurance. Our implementation is more stylized but follows the same computational logic: behavior is generated endogenously from local states and interactions rather than imposed as a fixed aggregate response.

2.3 Robust decision making under deep uncertainty

Infrastructure decisions have long lifetimes and are exposed to uncertain hazard trajectories, technology performance, budgets, and policy regimes. Robust decision-making approaches therefore emphasize performance across many plausible futures rather than optimization for a single forecast (Lempert et al., 2006; Dittrich et al., 2016). The broader decision-making-under-deep-uncertainty literature provides tools such as robust decision making, dynamic adaptive planning, and adaptive pathways (Marchau et al., 2019). Scenario discovery also shows that multiple objectives can alter which failure regions appear policy-relevant (Shortridge & Guikema, 2016).

This study adopts the computational spirit of that literature but uses a lightweight model predictive controller rather than a full robust optimization solver. At each scenario, the controller constructs a risk score from the digital-twin forecast, vulnerability, infrastructure criticality, and disadvantage weight. A budget floor protects low-score nodes from complete exclusion, while the remaining budget is allocated nonlinearly toward higher-risk nodes. The resulting rule is interpretable, fast, and suitable for repeated ensemble testing.

The financing literature on climate-resilient infrastructure emphasizes both mainstreaming resilience into investment decisions and mobilizing additional public-private finance (OECD, 2024). Our model focuses on allocation rather than funding source. The control variable can be interpreted as grants, concessional credit, insurance subsidy, resilience bonds, or blended finance, provided the instrument reduces failure probability or raises adaptation propensity through the specified mechanism.

2.4 Research positioning

The paper's niche is the interface among digital twins, behavior, network reliability, and finance. These areas are often modeled separately; the computational architecture allows them to be stress-tested jointly and exposes where information quality, allocation concentration, and cascades interact.

3 Infrastructure digital twin and network model

3.1 Physical dependency graph

The synthetic infrastructure system contains $N = 72$ nodes embedded in a two-dimensional spatial domain. Edges connect nodes within a distance threshold, producing a sparse undirected dependency network. Nodes may be interpreted generically as substations, transport links, pumping stations, communication hubs, or other assets whose service can be disrupted by climate hazards. Each node i has vulnerability v_i , criticality c_i , and disadvantage weight d_i .

Table 1 Digital-twin and infrastructure state variables

Variable	Interpretation	Range / baseline
v_i	Physical vulnerability	0.25-0.90 synthetic
c_i	Service criticality	0.60-1.80 synthetic
d_i	Equity/disadvantage weight	0.50-2.00 synthetic
A_{ij}	Dependency adjacency	Distance-threshold graph
H_i	True hazard intensity	Heavy-tailed common + local
y_i	Observed hazard signal	True hazard + noise
m_i	Digital-twin posterior mean	Filtered node risk
u_i	Finance allocation	Budget-constrained control
a_i	Adaptation probability	Endogenous agent response

The dependency network is not intended to match a specific topology. Its purpose is to create localized clusters and cascading exposure. Isolated nodes are prevented by connecting any otherwise isolated node to its nearest neighbor. Failure propagation then depends on the fraction of failed neighbors and node vulnerability.

$$G = (V, E), A_{ij} = 1\{\text{distance}(i, j) < 0.19\} \quad (1)$$

3.2 Digital-twin observation and state estimation

A climate scenario draws a common heavy-tailed hazard component and node-specific noise. The true node hazard is clipped to a physically interpretable range and increases with structural vulnerability. The policy does not observe this state directly. It receives $y_i = H_i + \epsilon_i$, where the observation-noise variance depends on the governance regime. The digital-twin regimes use substantially lower observation noise than the baseline and finance-only regimes.

$$y_i = H_i + \epsilon_i, \epsilon_i \sim N(0, \sigma_{obs}^2) \quad (2)$$

The twin performs a one-step Bayesian linear update with prior mean 0.5 and prior variance 0.25. If $R = \sigma_{obs}^2$ denotes observation variance, posterior variance is $(1/0.25 + 1/R)^{-1}$ and posterior mean is the precision-weighted combination of prior and observation. This filter is intentionally simple so the impact of observation quality remains interpretable.

$$m_i = V_{post} (0.5/0.25 + y_i/\sigma_{obs}^2) \quad (3)$$

The digital twin is thus defined operationally by synchronization and update, not by a three-dimensional visualization. This abstraction is consistent with the decision-support role emphasized in recent digital-twin literature: the computationally relevant output is a continuously updated state estimate that can change an action rule. More complex implementations could replace the linear update with particle filtering, ensemble Kalman filtering, or learned surrogate models.

Table 2 Information and finance structure by regime

Regime	Observation noise SD	Finance allocation
Baseline	0.48	Uniform
Digital twin only	0.11	Uniform
Finance only	0.48	Risk weighted
Coupled MPC	0.11	Risk weighted + contingency

4 Multi-agent adaptation dynamics

4.1 Behavioral fixed point

Adaptation is modeled as an endogenous probability rather than a planner-controlled binary variable. The model aggregates four stakeholder types - public agencies, firms/operators, households/community organizations, and finance/insurance intermediaries - into node-specific adoption propensities. The equilibrium adoption probability depends on the perceived hazard signal, the relative finance subsidy, the mean adoption of neighboring nodes, and local vulnerability.

$$a_i = \text{logit}^{-1}[-2.2 + 1.35 S_i + 0.62 U_i + 1.05 N_i - 0.35 v_i] \quad (4)$$

Here S_i is either the node-level digital-twin forecast or a system-average signal when granular twin information is unavailable; U_i is finance allocation normalized by the cross-node average; and N_i is the average current adaptation probability among network neighbors. The model iterates this map to a fixed point using damped updating. Social spillovers create nonlinearity: once neighboring adoption becomes sufficiently high, an additional increase in incentives can move the system rapidly toward a higher-adaptation equilibrium.

The behavioral coefficients are synthetic and are not claimed as econometric estimates. Their purpose is to create bounded rational response, financial sensitivity, and social reinforcement at realistic numerical scales. The resulting adaptation rate is an endogenous intermediary outcome that helps explain why identical finance allocations can have different service-loss consequences under different information regimes.

Table 3 Multi-agent behavioral channels

Agent abstraction	Primary decision channel	Model representation
Public agency	Prioritization / grants	Risk- and criticality-weighted finance
Firms/operators	Asset protection investment	Signal and subsidy response
Households/community	Local adaptation / preparedness	Signal + social spillover
Finance/insurance	Pricing and subsidy	Allocation intensity U_i

4.2 Critical transitions and social spillovers

The adaptation map can display threshold-like behavior because the social term feeds current adoption back into the next update. A sufficiently large spillover coefficient raises the slope of the fixed-point mapping, causing small changes in perceived risk or subsidy to produce larger equilibrium changes. This mechanism is important for climate governance because incentives can have low apparent effect below a coordination threshold and large effect after adoption becomes socially self-reinforcing.




Fig. 1 Nonlinear equilibrium adaptation as the social-spillover coefficient increases. The dashed line marks a reference transition region in the synthetic model.

The curve is not interpreted as an empirically estimated tipping point. It is a structural sensitivity illustration showing why linear cost-benefit extrapolation can misstate the effect of interventions in multi-agent systems. If adaptation technologies diffuse through peer learning, supplier networks, or shared expectations, a finance policy can have indirect effects beyond the directly subsidized node.

Proposition 1: monotonic subsidy response

For a fixed signal and neighborhood state, the adoption probability is strictly increasing in normalized finance U_i because the logistic derivative is positive and the coefficient on U_i is $0.62 > 0$. This establishes local monotonicity of the direct finance-adoption channel, although equilibrium network effects can still create nonlinear aggregate responses.

5 Robust finance allocation and MPC rule

5.1 Risk-equity allocation score

The finance controller assigns each node a score equal to the product of the digital-twin risk forecast, physical vulnerability, service criticality, and an equity multiplier. The multiplicative form means a node receives high priority only when several dimensions are jointly material. This prevents a high hazard estimate at a low-criticality, low-vulnerability node from automatically dominating the budget.

$$q_i = m_i v_i c_i (0.78 + 0.32 d_i) \quad (5)$$

Finance-active regimes transform q_i through a softmax allocation. A positive floor guarantees that every node retains a minimum share of the budget, while the remaining share is targeted. The coupled controller uses slightly stronger risk discrimination because its forecast is more precise. It also represents a modest state-contingent reserve that can expand the budget when system-wide forecast risk rises. Because failure probability depends primarily on relative allocation in the baseline implementation, this reserve mainly affects interpretation and can be activated more strongly in empirical extensions.

$$u_i = B [f/N + (1-f) \exp(k q_i) / \sum_j \exp(k q_j)] \quad (6)$$

5.2 Model-predictive-control interpretation

The control is 'MPC-like' rather than a full nonlinear program: the digital twin updates the state, the finance rule recomputes allocation, agents update adaptation, and the system is re-evaluated under the new state. In a multi-period deployment the same cycle would repeat each period with observed damage and revised budgets. The current paper studies one-step ensemble control to isolate the coupling mechanism and keep the 900-scenario experiment transparent.

Algorithm 1 Digital-twin-guided finance control

1. Observe noisy infrastructure state y and update the digital-twin posterior mean m .
2. Calculate risk-equity score q_i for each node.
3. Allocate budget using a floor plus softmax targeting.
4. Iterate multi-agent adoption to a fixed point using risk, subsidy, and neighbor adoption.
5. Compute node failure probabilities and propagate cascading failures.
6. Store service loss, cascade fraction, equity-weighted residual risk, adoption, and allocation concentration.

5.3 Failure and cascade model

Initial node failure probability increases with true hazard and structural vulnerability and decreases with adaptation and resilience finance. Because finance is scarce, the allocation term is nonlinear in each node's share relative to the network mean. A second interaction lets targeted finance be more effective at high-hazard nodes, representing the idea that pre-positioned protection is most valuable where exposure is currently material.

$$P(F_i=1) = \text{logit}^{-1}\{-2.20 + 2.25 H_i + 1.25 v_i - 0.90 a_i - 0.55 \sqrt{U_i} - 0.58 U_i H_i\} \quad (7)$$

After initial failures are sampled, the model applies four cascade rounds. In each round, a functioning node faces an additional failure probability proportional to the share of failed neighbors and its vulnerability. This mechanism represents dependency loss without committing to a sector-specific engineering flow model. It is strong enough to create systemic tail events and weak enough to retain meaningful policy differences.

$$P(\text{cascade}_i) = \min[0.40 n_{\text{fail}_i} (0.70 + v_i), 0.65] \quad (8)$$

Criticality-weighted service loss is the fraction of total network criticality associated with failed nodes. Unlike a raw node count, this measure gives greater consequence to failure of important assets while remaining bounded between zero and one. Equity-weighted residual risk similarly weights ex-ante residual failure probability by the disadvantage index.

$$\text{ServiceLoss} = \sum_i F_i c_i / \sum_i c_i \quad (9)$$

6 Experimental design and reproducibility

The experiment uses 900 independent climate-policy scenarios with random seed 20260823. The common hazard component follows a Student-t distribution with five degrees of freedom and is combined with node-level Gaussian variation and vulnerability. Structural node attributes and the dependency graph are fixed across scenarios, while hazard, observations, initial failures, and cascades are redrawn. Common latent hazard states are shared across governance regimes to support paired comparisons.


Table 4 Baseline experimental settings

Parameter	Baseline value	Role
Infrastructure nodes	72	Network scale
Monte Carlo scenarios	900	Ensemble robustness
Baseline budget	8 synthetic units	Common finance envelope
Twin observation SD	0.11	High-quality state estimate
Noisy observation SD	0.48	Low-quality state estimate
Cascade rounds	4	Dependency propagation
Cascade coefficient	0.40	Neighbor-failure sensitivity
Adoption damping	0.65 old / 0.35 new	Fixed-point stability

Outcomes are summarized by mean service loss, 95% service-loss quantile, 95% expected shortfall, cascade fraction, equity-weighted residual risk, mean adaptation rate, and allocation coefficient of variation. The latter measures how concentrated finance becomes under risk targeting. A uniform allocation has essentially zero coefficient of variation, while finance-active regimes produce substantial concentration.

Generative-AI disclosure. OpenAI ChatGPT was used during manuscript preparation to assist with drafting, computational code development, and language refinement. All mathematical statements, simulation specifications, numerical outputs, references, and interpretations were reviewed for internal consistency, and responsibility for the final manuscript rests with the human author(s).

7 System architecture



Feedback loop: observations update the twin, behavior responds, and finance is re-optimized.

Fig. 2 Digital-twin, multi-agent, and finance-control feedback architecture.

The architecture separates five functions: physical state, digital estimation, behavior, finance control, and outcome evaluation. The separation is methodologically useful because it identifies where errors arise. Observation noise is a digital-twin problem; misaligned incentives are a behavioral-finance problem; cascade amplification is a network problem. A single aggregate resilience score would conceal these distinctions.

The controller uses only information available through its regime. The finance-only regime therefore demonstrates a key failure mode: an adaptive algorithm can be sophisticated but still allocate poorly when its state estimate is noisy. Conversely, the digital-twin-only regime can observe well but lacks the reallocation mechanism needed to exploit high-resolution differences in risk.

Algorithm 2 Ensemble evaluation

1. Generate one latent hazard vector H shared by all four regimes.
2. For each regime, generate observation y with regime-specific noise and update the twin state.
3. Compute finance allocation, behavioral equilibrium, initial failures, and four cascade rounds.
4. Record the seven outcome metrics.
5. Repeat for 900 scenarios and compare distributions rather than a single forecast.

8 Results

8.1 Risk-weighted allocation in the infrastructure network




Fig. 3 Representative coupled-regime network. Node size is proportional to adaptive finance allocation; color encodes the digital-twin risk forecast.

The representative network shows that allocation is spatially heterogeneous even though the model has no pre-defined high-risk district. Clusters emerge from the joint realization of hazard, vulnerability, criticality, and disadvantage. Larger nodes are not necessarily the highest-hazard nodes because the controller also accounts for service importance and equity weighting. This multidimensional prioritization is the principal difference from a hazard-only rule.

The allocation floor is important for robustness. Pure softmax targeting can starve low-score nodes and increase failure outside the identified hotspot, especially when network dependencies allow those failures to cascade back into the protected cluster. The floor therefore acts as a diversification constraint: it limits overconfidence in the digital twin and prevents the controller from treating low estimated risk as zero risk.

8.2 Main service-loss outcomes




Fig. 4 Mean criticality-weighted service loss across governance regimes.

Table 5 Main ensemble results

Regime	Mean service loss	VaR95	ES95	Cascade fraction	Equity loss
Baseline	55.01%	79.13%	84.53%	55.01%	18.41%
Digital twin only	54.00%	77.63%	83.54%	53.94%	17.64%
Finance only	54.77%	79.07%	84.22%	55.00%	18.50%
Coupled MPC	53.74%	78.79%	83.60%	54.00%	17.59%

The coupled controller reduces mean service loss from 55.01% to 53.74%, a relative reduction of 2.3%. Digital-twin-only monitoring reaches 54.00%, showing that much of the average resilience gain arises from better state information and the adaptation behavior it enables. Finance-only targeting reaches 54.77% and performs less reliably because resource concentration is based on noisier node forecasts.

8.3 Cascading failure and tail-risk boundary




Fig. 5 Distribution of the fraction of infrastructure nodes failed after cascade propagation.

The cascade distributions reveal a more nuanced result than the mean service-loss bars. The coupled and digital-twin-only regimes have very similar extreme-tail performance. In the baseline output, coupled 95% expected shortfall is 83.60% versus 83.54% for digital-twin-only. This difference is negligible relative to scenario variability and should not be interpreted as tail superiority. The correct inference is that targeted finance improves average and equity outcomes, while strong dependency cascades can erase part of the advantage in the most severe states.

This boundary is substantively important. A climate-resilience policy should not be judged only by whether it dominates a benchmark at every quantile. If a cascade mechanism is sufficiently strong, resilience may require structural network investment - redundancy, islanding, backup capacity, or dependency redesign - rather than only better sensing and marginal finance reallocation. The model therefore identifies when the policy problem changes category.

Proposition 2: cascade limitation

If marginal finance reduces initial failure probability but does not alter the dependency matrix, then as the conditional probability of neighbor-induced failure approaches one, differences in initial failure rates contribute progressively less to terminal cascade outcomes. Thus, targeting can improve average performance without guaranteeing dominance in the far tail when network contagion is sufficiently strong.

8.4 Adaptation behavior




Fig. 6 Distribution of equilibrium adaptation rates across the 900-scenario ensemble.

Average equilibrium adaptation rises from 34.25% in the baseline to 36.43% under digital-twin-only monitoring and 36.45% under the coupled regime. The similarity of the latter two values reflects the behavioral structure: accurate node-level signals are the dominant driver of adoption, while finance affects adoption through a smaller subsidy coefficient. This result highlights an important distinction between behavioral and engineering resilience. Finance can reallocate physical protection even when it produces only a modest additional change in aggregate adoption.

Finance-only targeting leaves adaptation close to baseline because the signal provided to agents is system-average and the targeted subsidy is based on noisier forecasts. Some nodes receive large allocations, but those allocations do not consistently align with high true hazard. This is an example of algorithmic precision without informational precision: an allocation rule can be highly differentiated while still being poorly targeted.

Table 6 Behavioral and allocation diagnostics

Regime	Adaptation rate	Allocation CV	Interpretation
Baseline	34.25%	0.00	Uniform
Digital twin only	36.43%	0.00	Uniform
Finance only	34.35%	0.54	Targeted
Coupled MPC	36.45%	0.54	Targeted

8.5 Finance intensity and distributional effects




Fig. 7 Sensitivity of service and equity-weighted losses to the finance-intensity multiplier around the coupled baseline.

The finance-intensity sensitivity experiment shows diminishing marginal gains. Increasing the budget multiplier lowers both service loss and equity-weighted residual risk, but the slope flattens as the multiplier rises. This is consistent with two mechanisms in the model. First, adaptation probability is bounded by the logistic response. Second, concentrated finance cannot eliminate network-induced failures arriving through dependent neighbors.

The equity effect is not automatic. The controller includes disadvantage in q_i , so higher-disadvantage nodes receive additional weight conditional on hazard, vulnerability, and criticality. Removing the disadvantage term in an ablation shifts allocation toward high-criticality assets and modestly worsens the equity metric while slightly improving some engineering outcomes. This illustrates why multi-objective climate governance should make distributive priorities explicit instead of assuming they will emerge from efficiency optimization.

$$q_i = risk_i \times vulnerability_i \times criticality_i \times equity_multiplier_i \quad (10)$$

In policy use, the equity multiplier could be calibrated through participatory governance, statutory service obligations, or explicit minimum-resilience standards rather than chosen solely by the modeler.

8.6 Efficiency-equity frontier




Fig. 8 Joint comparison of mean service loss and equity-weighted residual risk.

The coupled and digital-twin-only regimes occupy the lower-left region of the efficiency-equity plot, while finance-only is close to the baseline on aggregate service loss and slightly worse on equity-weighted residual risk. This outcome is a warning against interpreting adaptive allocation as inherently beneficial. When allocation is conditional on a noisy forecast, concentration can increase vulnerability at nodes that the model mistakenly classifies as low risk.

A robust controller should therefore combine targeting with safeguards. The floor used here is one such safeguard. Others include caps on concentration, minimum service standards, uncertainty penalties, or distributionally robust optimization. The choice depends on institutional objectives. A public infrastructure agency may prioritize continuity and equity constraints, while a private operator may place greater weight on expected cost and asset criticality.

Table 7 Design safeguards suggested by the computational results

Policy design feature	Benefit	Failure mode addressed
Allocation floor	Maintains minimum protection	Forecast misclassification
Equity multiplier	Protects disadvantaged nodes	Efficiency-only concentration
Twin calibration audit	Improves state ranking	Overconfident sensing
Cascade stress test	Reveals dependency risk	Local optimization blindness

8.7 Robustness to policy and observation uncertainty




Fig. 9 Distribution of service loss across the ensemble. Boxplots emphasize regime overlap and uncertainty.

The boxplots reinforce that uncertainty remains large after policy improvement. None of the regimes collapses the service-loss distribution into a narrow band. The main value of the coupled architecture is a shift in expected and equity-weighted performance, not elimination of climate variability. This distinction is consistent with robust decision-making principles: the objective is to reduce vulnerability across plausible futures, not to produce false certainty.




Fig. 10 Sensitivity of normalized policy variance to digital-twin observation noise.

As observation noise rises, policy variance increases because risk ranks change more across repeated measurements. Coupling reduces the growth rate of this variance in the sensitivity design by combining the risk estimate with vulnerability, criticality, and allocation-floor constraints. In a real system, a similar diagnostic could inform whether a proposed sensor network is precise enough to support automated financial triggers.

9 Ablation studies

Ablation analysis asks which model components are necessary for the observed performance. Four conceptual ablations are considered. Removing node-level digital-twin information replaces S_i with a system-average signal. Removing adaptive finance forces uniform u_i . Removing social spillovers sets the neighbor-adoption coefficient to zero. Removing the equity multiplier sets d_i to a constant. The first two correspond directly to the named governance regimes; the latter two are structural sensitivity checks.

Table 8 Component ablations and interpretation

Ablation	Expected effect	Observed interpretation
No node-level twin	Worse risk ranking	Finance-only close to baseline
No adaptive finance	Less targeted protection	Twin-only improves information but not allocation
No social spillover	Flatter adoption response	Fewer threshold effects
No equity weight	Lower protection of disadvantaged nodes	Efficiency-equity frontier shifts

The finance-only result is especially informative because it shows that the same adaptive allocation formula can be weak or counterproductive when the forecast is poor. This is exactly the kind of coupling issue the special issue seeks to illuminate: technology and finance should not be assessed as independent checkboxes. A financial innovation that depends on data quality inherits the failure modes of the data pipeline.

The social-spillover ablation changes not only average adoption but also the sensitivity of adaptation to subsidy. With no spillover, the behavioral response becomes closer to a collection of independent logistic decisions. Under stronger spillover, localized financing can diffuse indirectly through peer effects, creating larger but less linear system responses. This underscores the value of multi-agent simulation over representative-agent aggregation when policy diffusion is central to the research question.

10 Policy interpretation

10.1 When should digital twins trigger finance?

The model suggests that a digital-twin signal should trigger adaptive finance only after three conditions are satisfied. First, the state estimator should demonstrate calibration and ranking accuracy under out-of-sample or stress conditions. Second, the finance rule should include diversification safeguards so uncertainty cannot concentrate resources excessively. Third, the infrastructure dependency model should be stressed to determine whether local protection can plausibly reduce systemic loss or whether network redesign is required.

The coupling is most valuable in the middle region of the risk distribution: risk is high enough that differentiation matters, but not so extreme that cascades dominate every intervention. In low-risk states, uniform maintenance may be sufficient. In extreme states, emergency capacity and redundancy may matter more than marginal targeting. This state-dependent interpretation is more useful than a blanket claim that digital twins or green finance are universally effective.

10.2 Implications for climate-resilient infrastructure finance

The control variable is instrument-neutral, but the computational logic maps naturally to real finance. A resilience bond could lower capital cost for high-score projects; a public grant program could allocate matching funds based on updated risk; an insurer could subsidize verified adaptation; or a development bank could prioritize loans where climate vulnerability and social disadvantage overlap. What matters is that the instrument has a measurable mechanism linking allocation to reduced failure probability or higher adaptation.

The OECD (2024) emphasizes the need to mainstream climate resilience into infrastructure finance. The present model adds a computational layer: mainstreaming should include how risk information is produced, how uncertainty is represented, and how updated risk changes funding decisions over time.

11 Computational scalability and implementation

The baseline system is intentionally small enough for transparent experimentation but large enough to exhibit network effects. Each scenario requires one vector Bayesian update, one softmax allocation, an iterative behavioral fixed point, Bernoulli sampling of initial failures, and four sparse matrix-vector cascade operations. Complexity scales approximately linearly with the number of Monte Carlo scenarios and, for sparse networks, near-linearly with the number of edges.

For city-scale deployment, the main computational challenge is not the controller itself but the fidelity of the digital twin and hazard simulator. High-resolution hydrologic, power-flow, transport, or building models can be orders of magnitude more expensive than the economic policy layer. Surrogate modeling, emulation, and multi-fidelity simulation therefore become important extensions. The economic controller can remain unchanged if the surrogate exposes calibrated node-level risk distributions.

A second implementation issue is time synchronization. Real sensor streams arrive at different frequencies and may fail during the hazard they are intended to monitor. A production digital twin should therefore carry data latency and missingness as state variables, not treat them as generic Gaussian noise. Third, financial execution has delay: a grant or loan approved after a warning may not affect physical vulnerability before the event. A multi-period MPC version should explicitly model lead times and irreversible investments.

11.1 Reproducibility sequence

Algorithm 3 Minimal replication

1. Set seed 20260823; generate 72 node positions, dependency edges, vulnerability, criticality, and disadvantage.
2. For 900 scenarios, draw common and local hazards; share latent hazards across regimes.
3. Generate regime-specific observations, posterior means, finance allocations, adaptation fixed points, failures, and cascades.
4. Calculate service loss, VaR95, ES95, cascade fraction, equity loss, adaptation, and allocation concentration.
5. Generate sensitivity curves for finance intensity, observation noise, and social spillovers.

12 Limitations and research agenda

The first limitation is the absence of engineering calibration. The dependency network is generic, and failure probabilities are synthetic. A sector-specific study would need validated fragility curves, physical flow constraints, restoration times, and realistic asset interdependencies. The present results therefore demonstrate mechanisms rather than estimate the resilience return of a real investment program.

Second, the digital twin is represented by a one-step linear Bayesian update. Real twins may integrate physics-based models, machine learning, remote sensing, and streaming telemetry. Those systems can exhibit bias, concept drift, cyber risk, and correlated sensor failure. Third, the agent model aggregates heterogeneous stakeholders into a common functional form. Empirical behavioral calibration could reveal different risk perceptions, liquidity constraints, and trust in public warnings.

Fourth, the finance controller is a transparent heuristic rather than a solved multi-stage stochastic program. This is a deliberate choice for interpretability and computational speed, but future work should compare robust MPC, distributionally robust optimization, reinforcement learning, and multi-objective evolutionary methods under identical scenarios. Fifth, equity is represented through exogenous node weights rather than endogenous household welfare or political economy.

Finally, the study uses synthetic data and should not be presented as evidence about any named jurisdiction. Its appropriate role is methodological: it demonstrates how to build and audit a coupled digital-twin-finance simulator before empirical deployment. The next research step is calibration to a specific infrastructure sector with transparent data provenance and a pre-specified validation protocol.

13 Discussion and conclusion

The computational experiments show that climate-resilient infrastructure governance benefits from treating digital monitoring, behavior, finance, and network dependence as one system. High-quality digital-twin information raises adaptation and improves average service continuity. Risk-weighted finance can add value when that information is accurate, but the finance-only regime demonstrates that differentiated allocation is not inherently superior to a uniform rule. Poor state estimation can make a sophisticated controller confidently wrong.

The coupled regime produces the best mean service-loss and equity-weighted outcomes in the baseline experiment, but it does not dominate digital-twin-only monitoring in every extreme-tail statistic. This is not a defect in the method; it is a policy insight. When cascading dependencies dominate, the problem shifts from allocating marginal protection to redesigning the network itself. Robust governance therefore requires a portfolio of sensing, targeted finance, redundancy, and recovery capacity.

For Computational Economics, the study contributes a modular test bed combining Bayesian state estimation, multi-agent fixed points, nonlinear failure probabilities, network cascades, and adaptive financial control. Each module is interpretable and replaceable. The architecture is suitable for empirical extension to power, transport, water, communications, or urban infrastructure, provided the hazard and engineering layers are recalibrated.

The broader implication is that technological innovation and financial innovation should be evaluated jointly. A digital twin has limited governance value if its forecasts cannot change resource allocation. A financial innovation has limited resilience value if it responds to noisy or stale risk information. Computational coupling makes those dependencies visible and allows policy designers to identify the conditions under which integration improves resilience, equity, or neither.

References

- Ali, E., Mansour, A., Abdelkader, E. M., Elshaboury, N., & Zayed, T. (2025). Digital twin for climate resilience: Transforming smart cities for a sustainable future. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-G-2025, 139-145. <https://doi.org/10.5194/isprs-archives-XLVIII-G-2025-139-2025>
- Bhave, A. G., Conway, D., Dessai, S., & Stainforth, D. A. (2017). Barriers and opportunities for robust decision making approaches to support climate change adaptation in the developing world. *Climate Risk Management*, 14, 1-10. <https://doi.org/10.1016/j.crm.2016.09.004>
- Cavus, M., Jiang, J., Allahham, A., Ramesh Rao, A. G., Scullion, E., Malamud, B. D., Sun, H., & Ng, W. P. (2026). Digital twins for hazard-resilient power grids: A systematic review and roadmap. *Renewable and Sustainable Energy Reviews*, 235, 116947. <https://doi.org/10.1016/j.rser.2026.116947>
- Dittrich, R., Wreford, A., & Moran, D. (2016). A survey of decision-making approaches for climate change adaptation: Are robust methods the way forward? *Ecological Economics*, 122, 79-89. <https://doi.org/10.1016/j.ecolecon.2015.12.006>
- Groves, D. G., Molina-Perez, E., Bloom, E., & Fischbach, J. R. (2019). Robust decision making (RDM): Application to water planning and climate policy. In V. Marchau et al. (Eds.), *Decision making under deep uncertainty*. Springer. https://doi.org/10.1007/978-3-030-05252-2_7
- Jenkins, K., Surminski, S., Hall, J., & Crick, F. (2017). Assessing surface water flood risk and management strategies under future climate change: Insights from an agent-based model. *Science of the Total Environment*, 595, 159-168. <https://doi.org/10.1016/j.scitotenv.2017.03.242>
- Lempert, R. J. (2019). Robust decision making (RDM). In V. Marchau et al. (Eds.), *Decision making under deep uncertainty* (pp. 23-51). Springer.
- Lempert, R. J., Groves, D. G., Popper, S. W., & Bankes, S. C. (2006). A general, analytic method for generating robust strategies and narrative scenarios. *Management Science*, 52(4), 514-528.

References (continued)

Marchau, V. A. W. J., Walker, W. E., Bloemen, P. J. T. M., & Popper, S. W. (Eds.). (2019). Decision making under deep uncertainty: From theory to practice. Springer. <https://doi.org/10.1007/978-3-030-05252-2>

OECD. (2024). G20/OECD report on approaches for financing and investing in climate-resilient infrastructure. OECD Publishing.

Reis, J., & Shortridge, J. (2020). Impact of uncertainty parameter distribution on robust decision making outcomes for climate change adaptation under deep uncertainty. *Risk Analysis*, 40(3), 494-511. <https://doi.org/10.1111/risa.13405>

Ricciardi, G., & Callegari, G. (2023). Digital twins for climate-neutral and resilient cities: State of the art and future development as tools to support urban decision-making. In *Technological imagination in the green and digital transition*. Springer. https://doi.org/10.1007/978-3-031-29515-7_55

Shortridge, J. E., & Guikema, S. D. (2016). Scenario discovery with multiple criteria: An evaluation of the robust decision-making framework for climate change adaptation. *Risk Analysis*, 36(12), 2298-2312. <https://doi.org/10.1111/risa.12582>

Song, M., Pirtea, M. G., & Budria, S. (2026). Call for papers: Technological and financial innovation in climate risk governance: Opportunities and challenges. *Computational Economics*, 68(3), 1889-1893. <https://doi.org/10.1007/s10614-026-11388-7>

Zhu, M., & Jin, J. (2025). Data-driven urban digital twins and critical infrastructure under climate change: A review of frameworks and applications. *Urban Planning*, 10, Article 10109. <https://doi.org/10.17645/up.10109>

Appendix A: Interpretation of figures

All figures are generated from the synthetic computational model or from explicitly labeled sensitivity transformations around the baseline. No external published figure is reproduced. The network visualization is a representative simulated scenario and is not a geographic map of a real city or infrastructure operator.

Statements and Declarations

Funding

The authors declare that no funds, grants, or other support were received during the preparation of this manuscript.

Competing interests

The authors have no relevant financial or non-financial interests to disclose.

Author contributions

The submitting author is responsible for study conception, model design, computational implementation, analysis, interpretation, and manuscript preparation. If additional authors are added before submission, this statement should be updated to reflect their actual contributions.

Data availability

No external empirical dataset was used. All reported numerical results are generated from synthetic infrastructure, hazard, behavior, and policy states using the distributions, dimensions, random seed, and algorithms documented in the manuscript. These specifications are sufficient to reconstruct the computational experiment.

Code availability

The analysis was implemented in Python using standard numerical, graph, and plotting libraries. The complete simulation source used to create the reported tables and figures is retained with the manuscript preparation files and can be submitted as supplementary material.

Ethics approval and consent

Not applicable. The study uses no human participants, animals, personal data, or confidential infrastructure records.

Generative-AI use

OpenAI ChatGPT was used to assist with drafting, computational code development, and language refinement. The human author(s) reviewed the mathematical content, simulations, references, numerical outputs, and final text and accept responsibility for the manuscript.