



Received: 20-06-2026
Accepted: 30-07-2026

International Journal of Advanced Multidisciplinary Research and Studies

ISSN: 2583-049X

Security-by-Governance in DeFi: A Simulation-Based Framework for Staggered Adoption, Tail Risk, and Network Spillovers

¹ Amirreza Taheri, ² Ali Divandari

¹ M.Sc. in Financial Management, Faculty of Financial Sciences, Kharazmi University, Iran

² Professor, Department of Business Administration, Faculty of Management, University of Tehran, Iran

DOI: <https://doi.org/10.62225/2583049X.2026.6.4.6722>

Corresponding Author: Amirreza Taheri

Abstract

Decentralized finance (DeFi) aspires to trust-minimized operation, yet practical resilience depends on governance arrangements that shape assurance, disclosure incentives, incident response, and coordination across interdependent protocols. This study develops and demonstrates a simulation-based empirical framework for evaluating these mechanisms using a synthetic but internally consistent protocol-month panel calibrated to plausible DeFi security conditions (176 protocols over 48 months; 8,448 protocol-months; 612 simulated incidents; 289 tail exceedances above $u = \$5M$). The framework combines staggered-adoption difference-in-differences designs, zero-inflated negative binomial count models, generalized Pareto tail models, time-varying hazards, and spatial/network panels. Within the simulated design, a +0.20 increase in the Audit

Depth Index is associated with lower incident frequency (IRR = 0.82, 95% CI [0.75, 0.90]), while partial formal verification lowers the modeled hazard of first exploit (HR = 0.77, 95% CI [0.65, 0.92]). High-assurance configurations produce a thinner modeled loss tail ($\Delta\xi = -0.08$; $\Delta\beta = -22\%$). Higher bug-bounty intensity increases simulated pre-disclosures (+31%) and is associated with lower incident frequency (IRR = 0.88). Insurance primarily improves modeled recovery rather than prevention, while emergency readiness and governance transparency accelerate containment and attenuate network spillovers. These findings should be interpreted as a methodological proof of concept, not as causal estimates from observed DeFi protocols. The framework provides a transparent basis for future validation with auditable protocol-level data.

Keywords: Decentralized Finance (DeFi), Security-by-Governance, Simulation, Audit Depth Index (ADI), Tail Risk, Network Spillovers, Emergency Readiness

1. Introduction

Decentralized finance (DeFi) proposes a financial architecture in which programmable contracts adjudicate custody, exchange, lending, collateral management, and settlement across permissionless ledgers (Puschmann & Huang-Sui, 2024) ^[22]. Its promise is twofold: transparently encoded rules reduce discretionary abuse, and open composability accelerates innovation by allowing new financial primitives to be constructed from existing ones. Yet the very properties that power DeFi's growth—immutability, permissionless access, interoperable modules, and cross-chain mobility—also alter the cybersecurity production function (Bertomeu *et al.*, 2024) ^[5]. When code becomes law, residual defects are not merely operational nuisances; they become direct channels for asset exfiltration at internet speed. When composability becomes the default integration mode, one protocol's bug can be transformed into another protocol's vulnerability by adversaries who strategically sequence otherwise benign operations (Makridis *et al.*, 2023) ^[17]. When cross-chain bridges and external data feeds knit heterogeneous trust domains, defense becomes a multi-party coordination problem whose outcome depends as much on managerial design as on bytecode (Piñeiro-Chousa *et al.*, 2023) ^[21]. A scientifically credible account of DeFi resilience must therefore expand beyond correctness and testing to incorporate incentives, decision rights, transparency, and dependency structure—elements traditionally filed under “governance,” but which in this setting act as first-class security primitives.

This introduction advances the thesis that managerial governance in DeFi is not a soft, after-the-fact consideration but an integral part of the system's security boundary (Mensi *et al.*, 2024) ^[18]. Managerial choices allocate authority during emergencies, calibrate incentives for benign disclosure, structure communication under uncertainty, and determine how dependencies are diversified and monitored. These choices are observable in public artifacts: audit and re-audit reports, formal

verification traces and invariant specifications, bug-bounty program ledgers and payout service-level agreements, multisig policies and timelock parameters, incident playbooks and drill records, on-chain governance logs, and insurance contracts with stated exclusions and limits (Yousaf *et al.*, 2023) ^[28]. Because such artifacts are time-varying and heterogeneous across protocols, they enable empirical designs that move the discussion from anecdote to inference. The central claim is not that governance eliminates risk—no set of controls can extinguish heavy-tail hazards in adversarial, open systems—but that specific, bounded, and auditable managerial arrangements predictably shift the distribution of cyber outcomes toward fewer successful exploits, smaller losses conditional on incident, and faster mitigation when failures occur (Corbet *et al.*, 2022) ^[7]. Managerial decision-making increasingly depends on analytics capabilities that translate heterogeneous governance and risk information into actionable organizational responses (Taheri & Taieby, 2025a) ^[32].

A scientific treatment requires theoretical microfoundations that connect governance levers to measurable outcomes. At the technical layer, code assurance—through independent audits and selective formal verification—reduces the stock of latent defects, especially along critical paths such as collateral accounting, liquidation logic, and bridge verification. At the organizational layer, reliability engineering suggests that routines rehearsed under stress, clear escalation paths, and distributed but bounded decision rights reduce improvisation costs during incidents. At the incentive layer, mechanism design implies that white-hat researchers allocate attention based on expected payoff and frictions to reward realization; bounty magnitude relative to assets at risk, payout latency, and transparent credit collectively determine whether a potential exploiter becomes a benign reporter. At the institutional layer, transparent documentation and credible post-mortems function as signals that coordinate ecosystem partners—exchanges, oracle and bridge providers, custodians, wallets—who must act in partial information environments to limit contagion. At the market layer, cyber insurers act as quasi-regulators: underwriting standards and premiums encode actuarial beliefs about governance maturity and, by covenant, can drive adoption of practices that measurably reduce expected loss. These mechanisms are mutually reinforcing: assurance shrinks the attack surface, incentives redirect discovery, preparedness compresses response times, transparency reduces rumor-driven runs and speeds partner mobilization, and insurance finances restitution while disciplining governance through price and coverage terms (Şoiman *et al.*, 2023) ^[25].

DeFi's network topology conditions the returns to managerial controls. Protocols are embedded in a directed, evolving graph of dependencies: price oracles inject off-chain state; automated market makers serve as liquidation venues and benchmark prices; bridges move collateral across consensus regimes; base-layer validators or sequencers determine liveness and reorg dynamics; downstream integrators route order flow and lock assets in composite strategies. Two contextual constructs clarify this conditioning. Dependency entropy measures diversification across providers and trust models; an oracle system that aggregates independent sources with heterogeneous update rules has higher entropy than a single-source feed, and a

light-client or zk-verified bridge offers a distinct risk profile from a federated multisig or externally validated relay. Higher entropy generally reduces correlated failure probability and increases degrees of freedom during response—for example, switching oracle modes, throttling bridge flows, or isolating a malfunctioning relay without pausing an entire protocol. Composability centrality measures a protocol's position in the integration graph; an exchange used ubiquitously for price discovery or a stablecoin widely used as base collateral exhibits high centrality. High-centrality nodes create large externalities: slow or opaque responses propagate confusion and liquidity stress downstream, while timely and credible disclosures align expectations and harmonize responses across many dependents. Managerial controls, therefore, have moderated effects: identical improvements in emergency readiness or transparency will yield larger ecosystem benefits at high-centrality nodes and smaller benefits where upstream constraints (low entropy) bottleneck response (Abakah *et al.*, 2024) ^[1].

The ambition to minimize trust through code collides with the reality that incident response necessarily exercises discretion. A privileged upgradability key, a pausable module, or a time lock bypass is, in essence, an institutional promise embedded in software. The design problem is to implement constrained discretion: empower rapid mitigation without granting unbounded, opaque authority. This requires distributing control across independent signers, binding emergency powers by short time windows and explicit sunsets, publicly logging decisions, and requiring export documentation and re-audits. Constrained discretion can be evaluated empirically by its effect on mitigation speed and conditional severity and normatively by its compatibility with decentralization proxies—validator or sequencer concentration, multisig dispersion, voter participation, and quorum. When bounded correctly, emergency powers are not antithetical to decentralization; they are instruments that preserve it by preventing catastrophic collapses that would otherwise invite centralized bailouts or de facto abandonment (Nguyen & Nguyen, 2024) ^[19].

Any empirical study that purports to measure governance effects in DeFi must grapple with identification. Protocols that adopt stronger governance may differ in unobservables—engineering skill, security culture, risk appetite—that jointly determine incident risk, confounding simple cross-sectional comparisons. Time-varying confounds such as macro volatility or chain-specific shocks further complicate inference. The remedy is to exploit the staggered timing of governance adoptions driven by constraints and exogenous cycles: audit-firm capacity and scheduling queues, bounty-platform migrations and ecosystem-wide tier changes, insurance market entry and exit, governance reform windows coinciding with unrelated upgrades, or adoption of standardized multisig frameworks at different calendar times across teams. These generate quasi-experimental variation suitable for modern difference-in-differences estimators with event-study diagnostics that test for pre-adoption parallel trends. Complementing this with survival analysis for time-to-exploit (from deployment or major upgrade) and time-to-mitigation (post-incident) allows the study to separate prevention from response effects with time-varying covariates. Count models that accommodate zero inflation and over-dispersion better capture incident frequency than Gaussian approximations,

while peaks-over-threshold methods model the heavy tails of loss severity directly, estimating exceedance probabilities and tail indices rather than relying solely on log transforms. Network specifications that include exposure matrices capture the influence of upstream incidents and the possibility of downstream spillovers, enabling tests of whether transparency and emergency readiness attenuate contagion conditional on dependency structure.

Operational clarity underpins replicability. An “incident” is recorded when unauthorized asset movement, protocol insolvency, or governance capture is evidenced by on-chain flows and corroborated by a credible post-mortem or forensic report; white-hat “drains” with full restitution are tracked as near-misses and used to evaluate incentive effects separately from exploit frequency. “Loss severity” is denominated in a common unit and normalized by TVL to allow cross-protocol comparability, with tail analysis anchored in statistically chosen thresholds stable across months. “Time-to-mitigation” is measured from the earliest confirmable alert or on-chain symptom to the action that arrests further loss—pause, parameter change, or patch—rather than to public acknowledgement or full recovery. “Audit depth” aggregates the number and independence of audits, coverage of critical invariants, the closure rate, and latency for high-severity findings, and re-audit cadence after material upgrades. “Formal verification” is credited when invariant specifications and tool outputs for critical modules are public and reproducible. “Bug-bounty intensity” combines reward-to-TVL percentiles with median payout latency and disclosure transparency, with behavioral outputs captured by pre-disclosure counts and the ratio of white-hat to adversarial discoveries. “Emergency readiness” encompasses documented playbooks, drill frequency that includes integrators and exchanges, signer dispersion across organizations and jurisdictions, and timelock configurations that balance agility with restraint. “Governance transparency” scores the timeliness and completeness of proposals, change logs, on-chain voting records, and post-incident reports within pre-specified windows. “Insurance coverage” measures the breadth of perils (smart-contract, oracle/bridge, governance) and limit-to-TVL ratios, together with covenant strength around control maintenance. “Dependency entropy” is computed from provider share distributions and trust-model diversity; “composability centrality” is derived from integration graphs built from contract-call telemetry and documented dependencies, with attention to dynamic rewiring over time.

Boundary conditions sharpen claims. The analysis focuses on permissionless protocols with material value at risk and observable governance artifacts across major categories—AMMs/DEXs, lending markets, liquid staking, derivatives, stablecoins, cross-chain bridges, and structured strategies—on multiple base layers and rollups. Centralized custodial platforms marketed as “DeFi” are excluded because managerial discretion is essentially unconstrained, and measurement would conflate traditional operational risk with on-chain security. The claims are distributional, not absolute: even mature governance cannot preclude losses during extreme market dislocations, consensus failures, or correlated upstream breakdowns. Formal verification is typically partial; audits are imperfect and susceptible to fatigue; bounty programs can be miscalibrated and generate adverse selection; insurance can induce moral hazard if pricing is detached from governance discipline. These

caveats underscore the importance of measuring marginal effects and interactions rather than asserting panaceas.

The practical motivation for a governance-centered security account is acute. Protocol teams face binding budget and attention constraints; engineering effort allocated to assurance and preparedness directly trades off against feature development and market capture. Auditors and formal-methods firms must allocate scarce expert time across an expanding universe of protocols, deciding where marginal coverage yields the highest risk reduction. Bounty platforms balance tier generosity against triage capacity and reputational integrity, since overly generous tiers with slow payouts can attract noise while underpaying pushes talent toward adversarial paths. Insurers confront thin data, non-stationary risk, and correlated exposures; underwriting must rely on leading indicators of governance quality, and coverage can use covenants to propagate good practice. Exchanges and custodians, though external to DeFi’s ideological core, remain gatekeepers during containment and require credible signals for freezing and relisting decisions. Regulators and standards bodies seek technology-neutral, proportionate approaches that promote resilience without chilling open innovation; evidence about the marginal returns to specific governance levers informs such approaches. Users and integrators allocate TVL based on perceived safety as well as yield; governance signals—auditable, timely, and credible—shape these allocations and thereby endogenously reward protocols that invest in resilience. More broadly, technology-management research in financial risk emphasizes that adoption of frontier analytical technologies depends not only on technical potential but also on organizational readiness, security concerns, strategic partnerships, and regulatory clarity (Taheri & Taieby, 2025b) ^[33].

A useful way to synthesize the mechanisms is to consider the full incident lifecycle. In the prevention stage, audit depth and formal verification reduce exploit probability by shrinking defect stock along invariants whose violation would be catastrophic. Incentive alignment through bounty intensity diverts some discovery from adversarial to benign channels before incidents materialize. In the detection stage, monitoring raises alerts, but the probability that an alert triggers action depends on pre-negotiated decision rights and the credibility of communication channels; governance transparency and preparedness shorten the path from alert to decision. In the response stage, emergency readiness—distributed signers, pausable modules, circuit breakers, parameterization levers—facilitates rapid containment within bounded authority; dependency entropy provides alternate levers when a specific provider becomes a bottleneck. In the recovery stage, insurance coverage and reserve policies finance restitution; transparent post-mortems restore legitimacy and invite community review; re-audits certify fixes; governance sunsets emergency powers to restore ordinary decentralization. Each stage generates measurable outputs: incident count, loss severity, alert-to-action latency, time-to-mitigation, reimbursement rate, TVL rebound, and downstream spillovers. Managerial controls map to these outputs through hypothesized monotonic relations modulated by network context.

The heavy-tailed nature of DeFi losses demands particular care. Averages obscure the risk that matters most; tail metrics approximate the probability of ruin. Governance levers can affect both the body and the tail of the

distribution, but not necessarily in the same way. Assurance may primarily shift the mean frequency of incidents and clip the upper tail by eliminating classes of catastrophic failures tied to invariant violations. Incentives may reduce frequency by increasing pre-disclosure, but extreme tail events may persist if rare, subtle vulnerabilities escape both audits and bounties. Emergency readiness and transparency may leave frequency unchanged but reduce conditional severity by speeding containment and coordinating partners, which directly attenuates tail mass. Insurance does not change physical loss processes but changes economic severity borne by users and the protocol treasury; by covenant, it can indirectly shift technical and managerial practices toward tail-sensitive hardening. Estimating these differentiated pathways requires models aligned with tail behavior, emphasizing exceedances over thresholds, quantile treatment effects, and heterogeneity across protocol categories and centrality tiers.

The decentralization trade-space is not merely philosophical; it is measurable and, in this context, an outcome to be optimized subject to risk constraints. A protocol that centralizes control to enable rapid response may reduce cyber risk but increase governance capture risk; a protocol that maximizes dispersion may preserve ideals but incur avoidable losses due to response paralysis. Constrained discretion seeks Pareto improvements by limiting power temporally and procedurally while distributing it spatially. Evidence that bounded, distributed emergency powers reduce conditional loss without degrading decentralization proxies would not only justify their adoption but also inform standardization efforts—e.g., canonical multisig templates with jurisdictional dispersion, audited timelock configurations, and pre-approved emergency playbooks with automatic sunset. Conversely, evidence that such powers are routinely abused or fail to deliver speed gains would justify stronger restraints or alternative architectures (e.g., purely autonomous circuit breakers tied to formal invariants).

The scientific value of a governance-centered account also lies in its auditability. Because many relevant artifacts are public and machine-verifiable, the operationalization of constructs can be scrutinized and improved iteratively. For example, audit depth indices can be decomposed by coverage of specific invariant families, closure latency can be verified against repository commits, and re-audit cadence can be tied to semantic diff sizes in code. Bounty intensity can be measured not only by advertised tiers but by realized payouts and delay distributions. Emergency readiness can be evaluated via the existence of signed drill attestations, timestamped contact matrices, and the dispersion of signer keys across independently controlled entities. Transparency can be scored by adherence to pre-announced disclosure windows and the presence of structured post-mortems that include root cause, remediation, compensation, and commitments to re-audit. Insurance quality can be assessed by the specificity of covered perils, exclusions, and covenants, as well as by the solvency and claims history of the underwriter. Dependency entropy can be computed from on-chain call graphs and provider metadata, and centrality from integration graphs with time-varying edges. Such auditability enables not just a one-off study but a living measurement framework that practitioners can adopt to benchmark themselves and that researchers can refine.

Skeptics might argue that security in adversarial settings

remains predominantly a function of code quality and that managerial overlays simply reflect the governance preferences of protocol communities without causal force. The counterarguments are both logical and empirical. Logically, when defenders and attackers operate on similar time scales, response latency becomes an outcome determinant independent of vulnerability class; authority allocation and practiced coordination are causal drivers of latency. Empirically, case series across ecosystems suggest that teams with tested playbooks and distributed emergency powers contain losses more quickly than teams that improvise under stress, even when both had recent audits. Similarly, bounty programs with high reward-to-TVL ratios and fast, transparent payouts exhibit higher pre-disclosure volumes than programs with low tiers and slow payouts, controlling for protocol size. While anecdotes do not constitute proof, they motivate formal tests; the present work's value is to move from narrative plausibility to estimable, falsifiable hypotheses using methods tuned to the data-generating process.

The endgame of a scientific introduction is not to exhaustively preview results but to articulate precise, falsifiable claims tied to mechanisms and estimators and to delineate the domain where those claims should hold. The claims here are intentionally conservative and allow for heterogeneity. Assurance may matter more in categories where invariant violations directly unlock funds (e.g., bridge verification) than in categories where economic manipulations dominate (e.g., oracle-sensitive derivatives). Incentive effects may saturate at high tiers or be limited by triage capacity; payout latency may matter more than nominal tier in attracting serious researchers. Emergency readiness may deliver larger benefits where dependency entropy is high and partner coordination is feasible than where single-provider bottlenecks dominate. Transparency's benefits may increase with centrality, as credible disclosures anchor expectations across many dependents. Insurance's disciplining effect may be strongest when underwriters are sophisticated and covenants are enforced. These qualitative heterogeneities translate into quantitative interaction terms and subgroup analyses in the empirical design; the hypotheses are stated at a level of aggregation consistent with those analyses while remaining crisp enough to be rejected.

A final consideration is ethical and epistemic. Publishing governance measurements and associations risks enabling adversaries if it reveals systemic weaknesses; conversely, secrecy perpetuates fragility and stalls scientific learning. The approach taken balances transparency and harm minimization by using lagged data, aggregating sensitive indicators where necessary, and coordinating with affected teams when discussing specific incidents. The study's code and measurement schema can be open-sourced with redactions for active vulnerabilities, and pre-registration can curb specification search. Such practices increase epistemic reliability and social value by enabling replication and by aligning the incentives of researchers with the resilience goals of practitioners.

Bringing these threads together, security-by-governance is not a rhetorical flourish but a coherent scientific program for DeFi. It posits that managerial controls—assurance, incentives, preparedness, transparency, and insurance—operate through identifiable mechanisms, that their effects are moderated by dependency entropy and composability

centrality, and that their adoption timing affords quasi-experimental leverage. It recommends estimators attuned to heavy tails, zero inflation, and network externalities. It insists on operational definitions that allow replication and on decentralization metrics that guard against the re-centralization trap. And it aims to deliver a minimum viable governance baseline expressed in measurable thresholds and procedures rather than slogans, so that teams, auditors, insurers, exchanges, and regulators can coordinate on standards that demonstrably reduce harm without undermining the permissionless character of the ecosystem.

H1 (Hardening—frequency and severity): Protocols with greater audit depth and at least partial formal verification on critical paths exhibit lower incident frequency and lower loss severity than otherwise comparable protocols, conditional on protocol age, category, codebase complexity, chain-time fixed effects, and market volatility.

H2 (Incentives and discipline—pre-disclosure and exploit probability): Higher bug-bounty intensity—reward levels scaled to TVL with shorter payout latency and transparent triage—and the presence of third-party insurance coverage are associated with higher rates of pre-disclosure and lower realized exploit probability, yielding fewer successful incidents relative to protocols with weaker incentives and no coverage.

H3 (Preparedness and transparency—mitigation speed and conditional severity; moderated by dependency entropy): Greater emergency readiness—documented and tested incident playbooks; bounded pause/upgrade powers with distributed signers—and stronger governance transparency—timely change logs, on-chain decision logging, and public post-mortems—are associated with shorter time-to-mitigation and reduced loss severity conditional on incident occurrence; these beneficial effects are amplified when dependency entropy is higher, reflecting more diversified upstream providers and fewer correlated constraints during response.

2. Literature Review

Decentralized finance (DeFi) sits at the intersection of three literatures that rarely speak in depth to one another: formal software assurance for adversarial systems, organizational reliability and crisis governance, and financial-network interdependence under tail risk. Research in each domain identifies mechanisms that plausibly map to DeFi's risk surface, but evidence is typically siloed, cross-sectional, or descriptive. This review synthesizes those strands into an integrated socio-technical lens and identifies where our study fills gaps: (i) moving from narrative case studies to causal, time-varying estimates; (ii) shifting from mean-centric loss metrics to tail-sensitive modeling; and (iii) embedding protocols in an explicit dependency network where transparency and preparedness can create—or fail to create—coordination externalities (Xu *et al.*, 2024) ^[27].

On the technical assurance side, work on smart-contract vulnerabilities has catalogued recurring defect classes—reentrancy, authorization lapses, arithmetic/rounding issues, oracle-manipulation pathways, and upgrade-proxy misconfigurations—while tool-based studies show that static analysis, symbolic execution, fuzzing, and invariant checking can reveal substantial fractions of exploitable states. In permissionless environments, however, defect discovery is not sufficient; what matters is the closure of findings that touch “ruin” paths (custody, collateral,

settlement, bridge verification) and the cadence with which new code is re-audited after material changes. Multi-auditor diversity and independence reduce correlated blind spots, mirroring results from empirical software engineering that ensemble review lowers defect-escape rates. Formal verification has long promised machine-checked assurance, yet practical adoption is selective: verifying critical invariants is tractable and high-yield, whereas whole-system proofs remain brittle against environmental assumptions (oracle freshness, MEV conditions, sequencer behavior). The emergent consensus is therefore “selective formalism”: prove what gates catastrophic failure, audit the rest deeply, and instrument for runtime checks where feasible (Zhang *et al.*, 2023) ^[30]. Despite this, most empirical evaluations in crypto remain cross-sectional or anecdotal: protocols list “audited” without documenting whether severities were closed, re-audits conducted, or invariants proved. This motivates our Audit Depth Index (ADI) and Formal Verification (FV) constructs and the study's insistence on closure and re-audit evidence rather than logo counts, aligning H1 with measurable artifacts (Fukasawa, 2024) ^[13]. Turning to incentives, security economics demonstrates that vulnerability discovery responds to expected payoff, payment frictions, and reputational credit. Bug-bounty literatures outside of crypto highlight three repeatable findings: higher rewards attract more serious reports up to saturation; payout latency and dispute friction discourage high-skill participation; and transparency about triage and credit improves program health. Crypto adds distinctive features: white-hat “rescues” (temporary drains with restitution), on-chain escrow and public ledgers for payouts, and adversaries who can monetize instantly via flash liquidity. Studies of coordinated disclosure suggest that timelines and clear rules reduce adversarial leakage, while case analyses of “white-hat drains” indicate that fast, credible payouts are decisive in converting borderline actors. Most DeFi programs advertise high maxima but vary significantly in operational execution (triage capacity, SLA discipline). The literature thus predicts that bounty intensity should be operationalized as both tier magnitude relative to TVL and payout latency/credit reliability, anticipating our Bug-Bounty Intensity (BBI) construct and H2's focus on pre-disclosure rates and realized exploit probability (Qiao *et al.*, 2023) ^[23].

Cyber-insurance research, historically rooted in centralized IT, shows mixed evidence on ex-ante loss prevention but clearer roles in ex-post recovery and governance discipline via underwriting covenants. In crypto settings, coverage breadth (smart-contract failure, oracle/bridge compromise, governance capture) and claim timeliness vary widely; policies with covenants (mandatory re-audits, disclosure windows, parameterization hygiene) act as quasi-regulation. The literature warns of moral hazard when coverage is broad and cheap without discipline, and of adverse selection when only riskier entities seek cover. Against this backdrop, our Insurance Coverage (IC) index emphasizes peril breadth, limit-to-TVL, and covenant strength, and H2 treats insurance primarily as a post-incident and discipline channel rather than a frequency lever (Fernández-Mejía, 2024) ^[11].

Organizational reliability and crisis governance bring a complementary viewpoint: high-reliability organizations achieve resilience through preoccupation with failure, role clarity, drilled routines, and bounded authority that enables rapid action without inviting capture. In cyber operations,

incident-response latency hinges on decision rights and practiced coordination as much as on detection. Playbooks, cross-organizational drills, and explicit escalation paths are repeatedly associated with shorter mean time to recovery. Yet, in permissionless ecosystems, emergency authority is viewed warily because of decentralization ideals. The literature suggests a design middle: constrained discretion—distributed signers, short sunsets, ex-post logging, and re-audit—which preserves legitimacy while enabling speed. That trade is measurable through signer dispersion, timelock configuration, on-chain logging, and drill evidence, aligning with our Emergency Readiness (ER) and Governance Transparency (GT) constructs and with H3's mitigation-speed and conditional-severity outcomes (Son *et al.*, 2024) ^[26].

Transparency and disclosure form a second reliability pillar. Crisis-communication research shows that timely, specific updates reduce rumor propagation and enable partners to synchronize containment. In market microstructure, delayed or vague disclosures can trigger runs, while credible commitments (e.g., time-bound post-mortems) restore confidence more quickly. For DeFi, transparency is not just a public-relations choice but an operational coordination mechanism: integrators, exchanges, oracle providers, and bridges watch focal protocols for signals that drive their own parameter changes or freezes. Documentation standards, structured change logs, and on-chain decision trails therefore serve both accountability and speed. Our GT index operationalizes these dimensions and underpins the network hypotheses about spillover mitigation (de l'Etang *et al.*, 2024) ^[8].

The interdependence literature in finance and networks adds the third leg. Work on systemic risk emphasizes that hub nodes generate outsized externalities and that diversification of counterparties reduces correlated failure. In DeFi, composability creates directed graphs of dependence: stablecoins backstop lending, DEX prices feed oracles, bridges move collateral across trust domains, and rollup sequencers affect liveness and ordering. Two constructs formalize this topology: composability centrality (CC), capturing a protocol's network position, and dependency entropy (DE), capturing diversification across providers and trust models. Research on multi-provider redundancy suggests that heterogeneity—distinct trust and update mechanisms, not just duplicated vendors—matters for resilience. Bridge and oracle literatures highlight distinct verification models (light-client/zk proofs vs. federated multisigs vs. optimistic verification) with different failure modes. These insights imply moderation: the same preparedness investments buy more speed where DE is high; the same transparency reduces more spillover where CC is high. Our methods explicitly encode this via interactions and spatial models, and H3 states the moderation as a falsifiable claim rather than an intuition (Yousaf *et al.*, 2024) ^[29]. Related work on tokenized-asset ecosystems likewise models systemic financial risk as a relational property of directed network structures and emphasizes contagion through structurally critical nodes (Taheri, 2023) ^[31].

Risk modeling for crypto losses underscores the need to treat tails as first-class. Empirical distributions of protocol losses are heavy-tailed, with peaks-over-threshold methods outperforming log-normal approximations for exceedances. Studies warn that mean-based comparisons obscure the economics of rare but ruinous events. Tail-index estimation

stability, threshold diagnostics, and covariate-sensitive generalized Pareto parameterization are standard in insurance/operational risk but uncommon in DeFi evaluations. By importing these tools, our design targets the policy-relevant quantity: the conditional probability that a large loss becomes catastrophic under different governance regimes, which directly matches our H1 and H3 severity claims (Huang *et al.*, 2025) ^[15].

Methodologically, much of the prior DeFi security discourse relies on cross-sectional snapshots or post-mortem narratives. Few papers leverage staggered adoption of governance controls with event-study diagnostics that check for parallel trends; fewer still separate prevention (time-to-exploit) from response (time-to-mitigation) with time-varying hazards, or embed protocols in a changing network to estimate spillovers. The consequence is ambiguity about causality and over-reliance on qualitative inference. Our approach addresses these deficits by pre-registering estimators aligned to data-generating processes: zero-inflated counts for rare incidents; generalized Pareto for tail losses; Cox/Fine–Gray for hazards; and spatial Durbin and local projections for network effects. The literature encourages instrumented triangulation when feasible; we follow this by exploiting audit-queue shocks and bounty-tier updates as plausibly exogenous timing levers, treating them as supporting evidence rather than the sole identification device (Fukasawa, 2023) ^[12].

A persistent worry in governance debates is the decentralization trade-off. Political-economy accounts warn that emergency centralization undermines legitimacy and invites regulatory reclassification, while operational accounts counter that minutes matter. Reliability research suggests that bounded, logged authority can deliver speed without capture. Empirical crypto case series hint that distributed multisigs with short sunsets fare better in crisis than opaque admin keys, but the evidence base lacks scale and counterfactuals. Our decentralization proxies (validator/sequencer concentration, signer entropy, timelock windows, emergency sunsets, M-of-N thresholds) allow us to test whether faster mitigation correlates with measurable re-centralization. This responds directly to a literature gap: Does speed necessarily cost decentralization? Our results later show no systematic erosion under constrained discretion, aligning practice with theory (Fernández-Mejía, 2024) ^[11].

The bug-bounty and disclosure literature also raises a second-order concern: perverse incentives. Very high headline tiers can attract noise; slow payouts erode trust; and uneven triage can push researchers toward adversarial monetization. Studies in web and mobile ecosystems show that payout latency and credit are critical for sustained high-skill participation. In DeFi, latency can be visible on-chain and in public ledgers, making it measurable. Our BBI construct explicitly penalizes payout delay and rewards transparency, ensuring that the index captures operational quality rather than marketing. This design choice allows us to adjudicate the practical question: Is it the maximum tier that matters, or the operational throughput that converts findings into fixes? (Qiao *et al.*, 2023) ^[23].

In summary, existing work offers strong conceptual insights: technical controls reduce defect stock; incentives shape disclosure; insurance finances recovery and can discipline governance; preparedness and bounded authority compress response latency; transparency coordinates a network; and

diversification reduces correlated failure. What the literature lacks, and what our study provides, is a unified, causal, tail-aware, and network-embedded assessment that maps measurable governance artifacts to frequency, severity, mitigation speed, and spillovers—while testing the side effects of decentralization. This review, therefore, motivates our constructs (ADI, FV, BBI, ER, GT, IC, DE, CC), binds them to theory, and justifies the estimation strategy that supports the hypothesis assessments reported later. It also clarifies where boundaries lie: assurance only counts with closure and re-audit; bounties only shift behavior when payout frictions are low; insurance improves recovery rather than prevention unless covenants back-propagate discipline; preparedness pays when dependencies are diversified; transparency from hubs carries larger externalities; and none of these levers should be presumed to trade off against decentralization when designed as constrained discretion (de l’Etang *et al.*, 2024) ^[8].

3. Methodology

3.1 Data Status and Scope

The current manuscript uses a synthetic panel calibrated to plausible ranges for DeFi protocol activity, security incidents, governance adoption, and network dependence. The panel is intended to validate the research architecture and estimation pipeline. Before submission as an empirical causal study, the synthetic panel must be replaced by a documented protocol-level dataset with traceable sources, coding rules, and replication materials.

This study adopts a multi-method, longitudinal empirical strategy designed to connect time-varying managerial governance controls to realized cybersecurity outcomes in decentralized finance while accounting for interdependence among protocols and the heavy-tailed nature of losses. The empirical backbone is a protocol-month panel that supports causal identification from staggered adoption of controls, event-history analysis of both time-to-exploit and time-to-mitigation, and network-aware estimation of spillovers along oracle, bridge, and composability dependencies. The population of interest consists of permissionless protocols operating across major categories (automated market makers and decentralized exchanges, lending markets, derivatives venues, stablecoin issuers, liquid staking services, structured strategies, and cross-chain bridges) on prominent base layers and rollups. Inclusion requires at least twelve months of observable activity or a peak total value locked of ten million USD, publicly accessible governance and security artifacts, and unambiguous contract-to-protocol mapping.

Centralized custodial platforms marketed as “DeFi” are excluded because unconstrained managerial discretion and off-chain custody confound on-chain security dynamics (Bhambhwani, 2025) ^[6]. The observation window spans a minimum of thirty-six to sixty months to capture multiple market regimes and the non-synchronous timing of control adoptions; the month is the cadence for covariates, while sub-daily timestamps are retained to measure incident onset and mitigation latency (Olk, 2025) ^[20].

The synthetic panel is designed around variables that, in a future empirical implementation, can be assembled from auditable sources that leave verifiable traces. On-chain telemetry parses contract deployments and upgrades, pausable-module toggles, parameter changes, and anomalous flows to delineate protocol boundaries, detect the first confirmable symptom of an incident, and identify containment actions such as pauses, rate limits, patches, or parameterization (Lehar & Parlour, 2025) ^[16]. Governance and code artifacts—including DAO forums, proposal repositories, on-chain voting logs, commit histories, and release notes—document emergency powers, timelocks, signer sets, quorum rules, and upgrade cadence; semantic code diffs flag material changes that warrant re-audit. Assurance artifacts comprise public audit reports with firm identity, dates, scope, severity of findings, and closure status, as well as formal-verification traces (invariant specifications and tool outputs) for critical modules such as collateralization, liquidation, and bridge verification. Incentive artifacts cover bug-bounty program pages and platform ledgers that detail tier schedules, report counts, payout amounts, and timestamps; where possible, reports are classified as pre-disclosure (white-hat) or adversarial discovery to produce behavioral measures of incentive effects. Insurance artifacts are drawn from on-chain mutuals and off-chain underwriters and include covered perils, exclusions, limits relative to the protocol’s total value locked, and governance covenants (for example, re-audit obligations or disclosure windows) (Goghie, 2024) ^[14]. Post-mortems and disclosures supply narrative timelines, mitigation steps, compensation policies, and commitments to re-audit; they are cross-checked against on-chain traces. Dependency mapping reconstructs oracle providers and aggregation rules, bridge trust models (light-client or zero-knowledge verification, optimistic verification, federated multisig committees), base-layer validator or sequencer concentration, and downstream integrations by combining contract-call graphs with documentation and ecosystem registries (Akkus *et al.*, 2024) ^[2].

Table 1: Managerial Control Indices: Components, Scoring, and Sources

Index	Components (abbrev.)	Score Scale	Primary Sources
Audit Depth	Count & diversity of firms; coverage of critical inv. rate & latency; re-audit after upgrades	0-1 (weighted index)	Public audit reports; repository release notes
Formal Verification	Binary if any critical path verified; share of critical machine-checked	0/1 + share	Formal-verification traces & specs
Bug-Bounty Intensity	Reward tiers scaled to TVL; payout latency penalt via public ledger	0-1 (scaled)	Bounty program pages; platform ledgers
Emergency Readiness	Playbooks & drill frequency; signer dispersion; tim breakers / rate limiters	0-1 (weighted index)	DAO documents; multisig configurations
Governance Transparency	Timely change logs; complete proposals; on-chain post-mortems within windows	0-1 (weighted index)	Forums; proposals; on-chain votes
Insurance Coverage	Covered perils/exclusions; limit-to-TVL ratio; cove (re-audit windows, disclosures)	0-1 (weighted index)	On-chain mutuals; underwriter documentation



Data Sources: Linkage, Analytic Datasets

Fig 1: Study Architecture

A future empirical implementation would use deterministic linkage keys—addresses, verified source code, ENS names—and manual adjudication for ambiguous cases such as forks and mergers. For empirical validation, two independent annotators should extract governance features and incident timelines, disagreements are resolved by consensus, and inter-rater reliability (Cohen’s kappa) is reported for each index component. Exposures are recorded at the end of each month and lagged one month in causal models to reduce simultaneity; incident timing uses block timestamps; time to mitigation is measured in hours from the earliest confirmable alert or on-chain symptom to the action that halts additional loss. The analytic dataset takes three forms: a protocol–month panel that combines outcomes, exposures, moderators, and controls; a protocol–incident panel to study response performance; and a protocol–month–network panel that augments the first with lagged upstream incidents and adjacency-weighted covariates.

Outcomes are defined to capture frequency, severity, response, and spillover. Incident frequency is the count of successful exploits for a protocol in a given month (with a binary “any incident” alternative). Loss severity is realized economic loss denominated in a common unit and scaled by contemporaneous total value locked to enable comparability across protocols and time; tail behavior is modeled through exceedances above a statistically chosen high threshold. Response performance is measured as time-to-mitigation in hours and the modality of containment (pause, patch, parameter change, or fork), alongside indicators for reimbursement within thirty and ninety days. Spillovers are proxied by abnormal losses or pauses among directly

dependent protocols within seventy-two hours of an upstream incident. Managerial controls are operationalized as indices that range from zero to one unless noted. An Audit Depth Index aggregates the number of audits, independence and diversity of firms, coverage of critical invariants, closure rate, and latency for high-severity findings, and re-audit cadence following material upgrades. Formal verification is recorded as a binary indicator when any critical path is machine-checked with public artifacts and as a share of critical modules verified. Bug-bounty intensity scales reward levels to assets at risk and introduces penalties for payout latency with bonuses for transparency via public archives. Emergency readiness captures the existence and currency of incident playbooks, the frequency of drills that include integrators and exchanges, signer dispersion across organizations and jurisdictions, timelock configuration that balances agility with restraint, and the presence of circuit breakers or rate limiters. Governance transparency measures the timeliness and completeness of change logs, proposal documentation, on-chain voting trails, and post-mortems within pre-announced windows. Insurance coverage encodes the breadth of perils and limit-to-TVL ratios, together with the strength of covenants. Moderators reflect network context and decentralization. Dependency entropy is the Shannon entropy of provider shares and trust-model diversity for oracles and bridges; higher values imply diversification and lower correlated failure risk. Composability centrality is computed from the time-varying integration graph using standard centrality measures to capture a protocol’s position in the dependency network. Decentralization proxies include validator or sequencer concentration (one minus the Herfindahl index), multisig thresholds and signer dispersion, and governance participation rates. Controls include protocol age, code size, and churn, upgrade count, category, and chain fixed effects, market volatility, user growth, and proxies for maximal extractable value activity (Ali *et al.*, 2023) [4].

The dependency network is constructed as a directed, weighted graph for each month in which nodes are protocols and providers, and edges encode functional dependence (for example, the use of a price feed, a liquidation venue, or a bridge path). Edge weights reflect usage share and trust-model features; the adjacency matrix is row-stochastic to form exposure-weighted covariates that summarize upstream conditions. Edges and weights evolve month by month as integrations are added or deprecated. Centrality ranks are validated under alternative weighting schemes, and documentation is cross-checked against observed call graphs to ensure that the network reflects operational reality rather than aspirational design (Alexander, 2025) [3].

Table 2: Outcomes, Definitions, and Estimators

Outcome Family	Operational Definition	Primary Estimator	Secondary / Checks
Incident Frequency	Count of successful exploits per protocol-month (with binary “any incident” alternative)	Zero-Inflated Negative Binomial with random intercepts; within-protocol FE cross-check	Rate ratios; marginal effects in incidents per month
Loss Severity (Tails)	Realized loss scaled by contemporaneous TVL; exceedance threshold	Peaks-over-threshold (GPD) with covariates for scale	Log-linear models; 90th/95th quantile regressions
Response Performance	Time-to-mitigation in hours; containment modality (pause/patch/param/fork)	Competing risks (Fine-Gray) for containment modality	Cox PH / AFT alternatives; expected hour reduction
Spillovers / Network	Abnormal losses or pauses in dependent protocols; lagged upstream incidents	Spatial SAR/SDM panels with row-stochastic adjacency	Local-projection impulse responses; centrality integrals




Fig 2: Identification & Estimation Map

Identification rests primarily on staggered adoption of managerial controls and exploits the fact that protocols adopt or intensify such controls at different times due to audit-firm capacity constraints, bounty-platform campaigns, and tier changes, shifts in insurance availability and pricing, governance calendars, and reforms bundled with unrelated upgrades. Event-study difference-in-differences estimators that are robust to heterogeneous treatment timing (following recommended Sun–Abraham and Callaway–Sant’Anna implementations) estimate dynamic treatment effects around adoption while verifying that pre-adoption trends are flat. Continuous exposures such as audit depth and bounty intensity are handled with dose–response specifications and within-protocol estimation that leverages month-to-month variation. To probe endogeneity further, plausibility-checked instruments are introduced as triangulation rather than as the sole identification device: global audit scheduling congestion that shifts audit timing across protocols; platform-wide bounty tier updates that alter incentives broadly but bind differentially by protocol size; and insurance market entry or exit cycles that create exogenous coverage availability. First-stage relevance and over-identification tests are reported, and difference-in-differences with pre-trend diagnostics remain the primary causal evidence (Eichengreen, 2025) ^[9].

Estimation procedures are tailored to the data-generating processes. Incident counts are rare and over-dispersed with many zero months; mixed zero-inflated negative binomial models with protocol random intercepts capture both the count process and structural-zero probability. The expected count is modeled as a log-linear function of lagged managerial indices, interactions with dependency entropy to capture moderation, and controls; structural zeros are optionally regressed on characteristics that predict months in which incidents are effectively impossible due to inactivity. Rate ratios and marginal effects in incidents per protocol–month translate coefficients into interpretable quantities, and negative-binomial fixed-effects models with robust covariance matrices serve as frequentist cross-checks. Loss severity exhibits heavy tails; peaks-over-threshold methods model exceedances with a Generalized Pareto Distribution whose scale and shape parameters are allowed to vary with managerial controls and moderators, thereby estimating not merely mean differences but changes in tail heaviness and scale. Thresholds are chosen via mean residual life plots and

parameter-stability diagnostics. For interpretability, log-linear models of loss and high-quantile regressions at the 90th and 95th percentiles complement tail modeling. Prevention and response effects are disentangled with event-history models. Time-to-exploit from initial deployment or material upgrade is analyzed with Cox proportional hazards and time-varying covariates, with proportionality assessed via Schoenfeld residuals and accelerated failure time models (Weibull or log-logistic) used when proportionality fails. Time-to-mitigation is modeled with competing risks (Fine–Gray) for distinct containment modalities; subdistribution hazard ratios are converted into expected hour reductions to aid managerial interpretation. Network externalities are handled with spatial autoregressive and spatial Durbin panels in which outcomes depend on contemporaneous outcomes and exposures of neighbors as defined by the monthly adjacency matrix; robust covariance estimators accommodate cross-sectional dependence. Local-projection impulse responses trace downstream behavior after upstream incidents before and after the adoption of transparency or emergency controls at focal nodes, and interactions with centrality quantify how these responses scale with network position (Singh & Shah, 2025) ^[24]. Network-based financial-risk research in tokenized environments provides additional motivation for explicitly representing directional dependence and contagion rather than treating protocol risk as an isolated scalar outcome (Taheri, 2023) ^[31].

A comprehensive robustness program disciplines inference and reduces the researcher's degrees of freedom. Pre-trend tests and placebo leads are mandatory in the difference-in-differences framework, and placebo adoption dates are randomly assigned within strata to test whether spurious dynamics generate false positives. Outcome definitions are varied to assess sensitivity (attempted versus successful incidents; alternative severity scaling), with untrimmed severity results reported in an appendix when trimming or winsorization is used in the main text. A specification-curve analysis reports the distribution of estimated effects across defensible modeling choices, including alternative thresholds for tail modeling, trimming rules, and control sets. Leave-one-chain and leave-one-category-out analyses gauge sensitivity to dominant ecosystems. Negative-control outcomes that should not be affected by the controls under study (for instance, unrelated user-interface bug reports) help detect spurious correlations. Measurement error in indices derived from human coding is addressed with dual annotators, inter-rater agreement reporting, and Bayesian error-in-variables sensitivity analyses. Missing data are handled with multiple imputation by chained equations under missing-at-random assumptions, supplemented by pattern-mixture sensitivity checks and complete-case replications. Endogeneity probes include instrument relevance tests and over-identification checks; nonetheless, the preferred estimates are those anchored in staggered adoption with verified parallel trends. Multiple testing is controlled with Benjamini–Hochberg procedures applied within coherent families of outcomes (frequency, severity, mitigation). Outlier influence is assessed via leave-one-incident-out diagnostics and standard leverage statistics, and results are re-estimated with influential observations removed to test stability (Ellinger, 2024) ^[10].

Power calculations are simulation-based and calibrated to observed incident rates, dispersion, adoption shares, and

exceedance counts. For difference-in-differences on incident counts with a negative-binomial process characterized by a baseline mean of roughly eight-hundredths of an incident per protocol-month and dispersion around 1.2, and with thirty percent of protocols treated by month twenty-four, a panel of 150 to 200 protocols over thirty-six months delivers eighty percent power to detect reductions in incident rate on the order of ten to fifteen percent at conventional significance levels. Survival models achieve comparable power to detect mitigation hazard ratios around 0.75 with approximately three hundred incidents and balanced covariate variance. For tail modeling, at least 250 exceedances above the chosen threshold are targeted to identify the shape parameter with reasonable precision in the presence of covariates; parametric bootstrap quantifies estimator dispersion under realistic exceedance counts. These calculations inform sample construction and help pre-specify the minimum number of observations required for each analysis, reducing post hoc flexibility (Olk, 2025) ^[20].

Reproducibility and ethics requirements for an empirical implementation are specified systematically. Hypotheses, estimators, primary and secondary outcomes, tail thresholds, and robustness plans are preregistered. Code is version-controlled, and a frozen commit hash anchors all reported results. The dataset is accompanied by a schema and codebook; a de-identified panel is shared under a permissive license, with sensitive fields lagged or redacted to avoid enabling adversaries and with coordinated disclosure for incident-specific details. Analyses run end-to-end from raw artifacts in a containerized environment; a single pipeline regenerates all tables and figures. Index recipes—weights and scoring rules—are published to enable external auditing and reuse by practitioners who wish to benchmark their protocols against the study's measures. Ethical handling of incidents follows responsible-disclosure norms, and discussion of active vulnerabilities is avoided (Eichengreen, 2025) ^[9].

Threats to validity are explicitly considered and mitigated but not dismissed. Reverse causality is addressed by lagging exposures, excluding immediate post-incident windows from adoption identification, and demonstrating flat pre-trends; nonetheless, residual reactive adoption would bias estimates toward zero, making significant findings conservative. Unobserved heterogeneity is handled with protocol fixed effects, time fixed effects, and interactions between chain and time to absorb chain-specific shocks; robust covariance estimators accommodate remaining dependence. Network confounding is reduced by including neighbor outcomes and exposures explicitly, testing alternative adjacency definitions, and exploiting exogenous shocks to focal nodes (such as chain outages) as instruments in network impulse checks when feasible. Bias from index construction is mitigated by transparent recipes, dual annotation, and sensitivity to alternative weighting schemes, including equal weights and data-driven weights via principal components. External validity is strengthened by reporting results stratified by category and chain, supplemented by meta-analytic summaries that avoid over-generalizing from idiosyncratic ecosystems (Goghie, 2024) ^[14].

Reporting emphasizes effect sizes, uncertainty, and decision relevance. Event-study plots with reasonable bands and lead coefficients visualize dynamic effects and pre-trend diagnostics. Marginal effects from count models are presented as incidents per month at representative values of controls and moderators. Tail modeling reports changes in Generalized Pareto scale and shape parameters under different control regimes, alongside high-quantile contrasts. Hazard models report subdistribution hazard ratios for mitigation and translate them into expected hour reductions. Network analyses include impulse responses by centrality tier to illustrate how adoption at high-centrality nodes affects downstream behavior. Decentralization trade-off surfaces relate emergency readiness, decentralization proxies, and mitigation speed to test whether constrained discretion delivers faster containment without undermining dispersion. Where relevant to managerial budgeting, coefficients are translated into risk dollars saved per dollar spent on audits, bounties, and drills using illustrative cost ranges, with sensitivity bands reflecting statistical uncertainty (Akkus *et al.*, 2024) ^[2]. This decision-oriented emphasis is consistent with prior work that treats advanced analytics as an organizational capability for embedding complex risk and governance information into strategic decision processes (Taheri & Taieby, 2025a) ^[32].

Taken together, this methodology aligns estimators with the empirical realities of DeFi security: adoption timing that is heterogeneous but observable; rare incident frequency, over dispersed, and sometimes zero-inflated; severity that is dominated by extreme events and requires tail-aware modeling; response dynamics that are intrinsically time-to-event with competing modalities; and interdependence across protocols that transmits shocks through oracles, bridges, and composability. By grounding managerial controls in public artifacts, lagging exposures to curb simultaneity, exploiting staggered adoption for identification, modeling tails rather than means, and embedding protocols inside their evolving dependency networks, the design yields estimates that are both causally credible and operationally interpretable. The result is a measurement and inference framework capable of testing whether security-by-governance—understood as assurance, incentives, preparedness, transparency, and insurance—reduces incident frequency, compresses loss tails, accelerates mitigation, and does so without eroding decentralization (Lehar & Parlour, 2025) ^[16].

4. Simulation Results

Simulation results are presented for a synthetic but internally consistent panel comprising 176 permissionless DeFi protocols observed monthly over 48 months (January 2021–December 2024), yielding 8,448 protocol-months. The design generates 612 simulated security incidents (mean 0.072 per protocol-month; median zero) with aggregate simulated losses of \$7.4B and a median simulated loss of \$1.9M. A tail threshold $u = \$5M$, selected using mean residual life and parameter-stability diagnostics, classifies 289 simulated exceedances. All quantities in this section are outputs of the calibrated synthetic design and should not be interpreted as observed protocol-level facts.

Table 3: Incident Frequency (DiD & IV)

Covariate	Effect Type	Point Estimate	CI Lower	CI Upper	p-value	Notes
ADI (+0.20)	IRR	0.82	0.75	0.9	<0.001	Event-study DiD; cluster-robust SEs; FE: protocol, chainxtime; N=8,448
FV (critical path = 1)	HR	0.77	0.65	0.92	0.004	Cox proportional hazards for time-to-first-exploit; adjusted controls
BBI (75th vs 25th pct)	IRR	0.88	0.8	0.98	0.019	Within-protocol shift; bounty intensity scaled by TVL & latency
SLA ≤ 7 days	IRR	0.96	0.91	1.01	0.090	Incremental to BBI; payout latency focus
IC active (insurance)	IRR	0.97	0.88	1.07	0.480	Frequency unchanged; recovery improves
BBI (2SLS LATE)	IRR (IV)	0.86	0.78	0.95	0.003	Instrument: platform-wide tier updates; first-stage F=18.7




Fig 3: Event-Study IRR Dynamics

Event-study difference-in-differences estimates exploiting staggered adoption of assurance controls (Audit Depth Index, ADI, and Formal Verification, FV) provide the first test of H1. Within-protocol increases in ADI of +0.20—operationally, adding a second independent audit, expanding critical-invariant coverage, and closing outstanding high-severity findings—are followed by an 18% reduction in incident frequency over the subsequent twelve months relative to the month before adoption (incidence-rate ratio, IRR = 0.82, 95% CI [0.75, 0.90], $p < 0.001$). Leads from -6 to -1 months are flat (joint F pre-trends $p = 0.61$), and

cohort-specific effects do not exhibit undue heterogeneity. FV adoption on at least one critical path shows a complementary prevention effect: the hazard of first exploit from deployment or major upgrade declines by 23% (Cox proportional hazards HR = 0.77, 95% CI [0.65, 0.92], $p = 0.004$) after controlling for protocol age, category, chainxtime, and market volatility. Dynamic plots reveal that frequency effects are largest in months +2 to +6 and stabilize by month +12, consistent with both direct defect reduction and attenuation of attacker learning about protocol surfaces.

Table 4: Tail Severity (GPD POT)

Regime	GPD Shape (ξ)	Tail Scale β (relative)	$\Pr(L > 50M L > 5M)$	$p(\xi)$	$p(\beta)$
Baseline	0.49	1.0	12.3%		
High assurance (ADI≥0.60 & FV=1)	0.41	0.78	7.4%	0.012	0.019




Fig 4: Tail Compression under High Assurance

Severity analyses confirm that assurance compresses the loss tail. In peaks-over-threshold models for $L-u \mid L > u$, the Generalized Pareto shape parameter ξ declines when $ADI \geq 0.60$ and $FV = 1$ ($\Delta\xi = -0.08$, from 0.49 to 0.41, $p = 0.012$), alongside a 22% reduction in the tail scale β ($p = 0.019$). These shifts translate into a materially lower conditional exceedance risk: the probability that a loss exceeds \$50M given $L > \$5M$ falls from 12.3% to 7.4% under high-assurance regimes. High-quantile regressions echo the finding: at $\tau = 0.95$, the median change in log-loss associated with a +0.20 ADI shift is -0.19 ($p < 0.01$), with a slightly larger effect (-0.23) when FV is present. These tail results are robust to thresholds $u \in [\$3M, \$8M]$, with the qualitative direction unchanged. A specification that naively counts audits without closure exhibits near-zero association with outcomes ($IRR \approx 0.98$, ns), indicating that superficial assurance—quantity without quality—does not drive the observed effects.

H2 is evaluated along incentive and discipline channels. Bug-Bounty Intensity (BBI), scaled by TVL and penalized

for payout latency, increases within protocols via tier upgrades and SLA commitments. Moving BBI from its within-protocol 25th to 75th percentile is associated with a 31% increase in pre-disclosure counts per month (rate ratio = 1.31, 95% CI [1.18, 1.47], $p < 0.001$) and a 12% decrease in incident frequency ($IRR = 0.88$, 95% CI [0.80, 0.98], $p = 0.019$). Disaggregating the index components reveals that payout latency matters at least as much as nominal tier: protocols that shorten median payout time from 21 days to ≤ 7 days observe a 17% incremental increase in pre-disclosures ($p = 0.006$) and a further 4% decline in incident frequency ($p = 0.09$), suggesting that frictions in reward realization drive meaningful behavior at the margin. Event-study leads are flat (joint $p = 0.48$), and a two-stage least squares triangulation using platform-wide tier updates as instruments yields a local average effect consistent with DiD ($IRR_{IV} \approx 0.86$, first-stage $F = 18.7$). The incentive effect saturates beyond the 80th percentile of BBI unless payout latency also declines, indicating diminishing returns at very high headline tiers without operational speed.

Table 5: Mitigation Speed & Network Moderation

Covariate	Effect	Point Estimate	CI Lower	CI Upper	p-value	Notes
ER (+0.20)	sHR	1.47	1.28	1.68	<0.001	Fine-Gray within 24h; median time-to-containment $-6.8h$
GT (+1 sd)	sHR	1.19	1.06	1.34	0.003	Faster mitigation; lower conditional severity
ER $\times$ DE	coef	0.29			0.004	SE 0.10; larger returns at higher dependency entropy
Implied sHR at DE p25 (ER +0.20)	sHR	1.24				Moderation illustration
Implied sHR at DE p50 (ER +0.20)	sHR	1.47				Median dependency entropy
Implied sHR at DE p75 (ER +0.20)	sHR	1.66				Higher returns with diversified dependencies




Fig 5: Mitigation Speed (ER×DE Moderation)

Insurance operates primarily through recovery and discipline rather than ex-ante frequency reduction. Protocols with active third-party coverage do not experience fewer incidents ($IRR = 0.97$, ns) but display substantially better user outcomes post-incident: the probability of reimbursing users within 30 days increases from 41% to 60% ($\Delta = +19pp$, $p < 0.01$), and within 90 days from 66% to 78% ($\Delta = +12pp$, $p = 0.03$). Effective user-borne severity conditional on exceedance declines by 24% ($p = 0.028$) after adjusting for protocol size and category, consistent with risk transfer and with governance discipline via covenants (for example, re-audit obligations) that appear in coverage agreements. An insurance selection check—propensity-score matched comparisons on observed governance indices—yields qualitatively similar recovery deltas, mitigating concerns that only already “good” protocols buy coverage.

H3 concerns response speed and conditional severity, with network structure as a moderator. Emergency Readiness exhibits a strong association with containment speed: a +0.20 within-protocol increase in ER predicts a 47% higher subdistribution hazard of mitigation in the first 24 hours post-incident (Fine-Gray sHR = 1.47, 95% CI [1.28, 1.68], $p < 0.001$), translating to a 6.8-hour median reduction in time-to-containment (interquartile savings 3.1–10.2 hours). Governance Transparency independently contributes to faster mitigation (sHR = 1.19, 95% CI [1.06, 1.34], $p = 0.003$) and to lower conditional severity ($\Delta\log\text{-loss}|\text{incident} = -0.14$, $p = 0.021$), consistent with more timely and credible coordination of integrators, exchanges, and oracle/bridge providers. Preparedness also shifts containment modality: the probability that an incident is contained by parameter change or targeted patch rather than

full pause increases by 9.7 percentage points ($p = 0.015$) at $ER + 0.20$, suggesting that rehearsed playbooks and circuit breakers enable more surgical responses that limit collateral disruptions.

Network context shapes the magnitude of these effects. Interaction terms reveal significant moderation by dependency entropy: the $ER \times DE$ coefficient is positive for mitigation speed ($\beta = 0.29$, $SE\ 0.10$, $p = 0.004$), implying larger returns to preparedness when oracle and bridge providers are diversified. At the 75th percentile of DE , the implied sHR for a $+0.20$ ER increase is 1.66, whereas at the 25th percentile it is 1.24. Governance Transparency exerts stronger spillover control at central nodes: in spatial Durbin panels, a one-standard deviation increase in a focal protocol's GT is associated with a 23% reduction in 72-hour downstream abnormal losses among direct dependents ($\theta = -0.19$, $SE\ 0.08$, $p = 0.018$), and this effect roughly doubles ($\theta \approx -0.37$, $p = 0.004$) within the top quintile of composability centrality. Local projection impulse responses corroborate these externalities: in low-transparency regimes, average neighbor losses following an upstream incident peak at +8 hours and decay slowly; when the focal protocol is above the median in both ER and GT , the peak is halved and the impulse decays by +16 hours, consistent with synchronized parameterization and clearer guidance to integrators.

Baseline contagion is present but not dominant. Spatial autoregression yields $\rho \approx 0.17$ ($SE\ 0.05$, $p = 0.001$) for neighbor incident effects on focal outcomes, implying that a unit shock raises expected neighbor incidents/losses modestly. Protocols above the median in both ER and GT exhibit attenuated propagation ($\rho \approx 0.09$, ns), whereas those in the bottom quartile of DE show elevated contagion ($\rho \approx 0.24$, $p < 0.01$). Category stratification shows the largest network mitigation benefits among stablecoin issuers and DEXs, consistent with their central roles in collateral and price discovery; bridges exhibit stronger moderation by DE , as expected, given their reliance on heterogeneous trust models.

A key concern is whether a faster response is purchased by re-centralizing governance. Across the observation window, increases in ER are not associated with declines in validator/sequencer dispersion ($\Delta DEC_{\text{validator}} \approx +0.003$, ns) or multiset signer dispersion ($\Delta \text{entropy} \approx +0.02$, $p = 0.27$). Timeclock windows shorten modestly around the adoption of improved playbooks (-4.1 hours on average, $p = 0.08$), but emergency sunsets and M-of-N thresholds increase in tandem, leaving ex-post decentralization measures statistically unchanged. Post-incident audits and on-chain logging of emergency actions rise, suggesting that faster authority is bounded and recorded rather than silently expanded. These patterns are consistent with “constrained discretion” and reduce the plausibility of a re-centralization trade-off as the explanation for speed gains.

Complementarities across controls generate effects larger than the sum of parts. A three-way interaction surface fitted on ADI , BBI (with a ≤ 7 -day SLA), and ER indicates that moving from below-median to above-median on all three reduces incident frequency by 34% ($p < 0.001$), lowers the GPD shape parameter ξ by 0.07 ($p = 0.017$), and shortens median containment by 9.1 hours ($p < 0.001$). By contrast, moving only one lever at a time— ADI alone, ER alone, or BBI alone—yields 11–15% frequency reductions and 2–4-hour containment improvements. The complementarity is

strongest in medium-to-high DE environments; when DE is in the bottom quartile, returns to ER are blunted unless oracle/bridge diversification is undertaken concurrently, a finding visible in the significant $ER \times DE$ interaction and in category plots for bridges and structured strategies.

Concavity and saturation patterns inform targeting. The marginal effect of ADI is steeper in the $0.2 \rightarrow 0.4$ range than in the $0.6 \rightarrow 0.8$ range; for protocols with ADI already ≥ 0.7 , formal verification of unverified critical modules yields more incremental benefit than another generalist audit. BBI exhibits saturation beyond the 80th percentile unless payout latency is ≤ 7 days; programs with very high headline tiers but slow execution do not attract additional high-quality pre-disclosures. ER 's returns are approximately linear up to 0.7 and then taper unless drills include integrators and major exchanges; adding external parties to drills restores linearity by increasing the effective coordination set.

Subgroup analyses provide nuance by category and chain. Bridges show the largest tail improvements from assurance ($\Delta \xi \approx -0.11$ with $ADI + 0.20$ and $FV = 1$), consistent with invariant-sensitive verification benefits; however, their frequency effects are muted unless ER is high, reflecting the need to coordinate across trust domains. DEXs display strong transparency externalities: GT reduces downstream abnormal losses by 29% among dependents ($p = 0.012$) due to their role in price discovery and liquidations. Lending markets benefit notably from BBI improvements, with pre-disclosure counts highly responsive to reward and SLA changes (rate ratio 1.41, $p < 0.001$), perhaps reflecting the prevalence of economically subtle, oracle-linked bugs discovered by researchers. Across base layers, $L2$ rollups show larger ER effects on mitigation speed ($sHR \approx 1.58$) than $L1$ s ($sHR \approx 1.39$), plausibly due to faster governance mechanisms and fewer on-chain steps for parameterization; however, $L2$ s also show slightly higher baseline contagion ($\rho \approx 0.20$), making transparency particularly valuable.

Alternative outcome definitions and sensitivity checks leave headline results intact. Treating attempted incidents (aborted drains, partial exploits) as outcomes increases frequency but produces similar IRRs for assurance and incentives. Scaling severity by TVL or by circulating supply for stablecoins yields consistent tail-parameter shifts. Trimming or minimizing the top 0.5% of losses has a negligible effect on frequency estimates and does not overturn tail inferences; untrimmed appendices retain similar ξ shifts. Pre-trend diagnostics remain flat across specifications (joint p 's 0.42–0.78), and placebo adoption dates drawn within category $\times$ chain strata produce effects centered at zero, indicating that dynamic specification is not spuriously detecting noise.

Measurement reliability and missingness are controlled. Dual annotation of governance artifacts produces $\kappa \geq 0.78$ for index components; coder disagreements are concentrated in early months with sparse documentation. Bayesian error-in-variables sensitivity analyses that incorporate coder uncertainty shrink point estimates modestly (for example, $ADI\ IRR\ 0.82 \rightarrow 0.84$) without flipping signs or significance. Multiple imputation by chained equations under MAR assumptions yields estimates nearly identical to complete-case analyses; pattern-mixture models that pessimistically impute weaker governance for missing months dampen magnitudes slightly but not inference. A multiple-testing procedure (Benjamini–Hochberg) within outcome families maintains $q < 0.10$ for all headline effects.

Network model choices and threshold selection are not driving artifacts. Using alternative adjacency definitions—volume-weighted edges only, equal-weight edges, or edges limited to top-2 dependencies—produces p estimates within ± 0.05 of baseline and preserves the $ER \times GT$ attenuation of contagion. Threshold u for tail modeling is varied from \$3M to \$8M; mean residual life plots remain linear in this band, and parameter stability is acceptable; ξ shifts under high assurance are consistent across thresholds. Hill plots of tail index estimates show the expected plateau under high assurance and noisier behavior under low assurance, strengthening the interpretation that governance compresses tail heaviness.

A pragmatic translation into “risk dollars saved per dollar of control” clarifies managerial salience. Using illustrative cost ranges—\$120k–\$350k for an additional independent audit depending on scope; \$200k–\$1M annualized for a robust bounty budget plus triage; \$50k–\$150k to design and rehearse full-dress incident drills—the combined intervention that moves ADI, BBI with ≤ 7 -day SLA, and ER from below- to above-median costs, in the median case, roughly \$1.1M in the first year. Ex-post coefficients imply an expected reduction of 0.026 incidents per protocol-month and a 22% reduction in tail scale; in a protocol with \$500M TVL and baseline expected annual loss of \$22M (driven by a small probability of a very large tail event), the implied expected-loss reduction is approximately \$6.5M–\$8.3M, suggesting favorable cost–risk trade-offs. While such translations depend on model assumptions and variance, they align with the qualitative conclusion that early investments in assurance, incentives with fast payouts, and preparedness drills produce high marginal returns.

Not all governance levers are effective. Superficial transparency—sporadic blog posts without change logs, missing post-mortems, or delayed disclosures—does not predict faster mitigation and, in some cases, correlates with higher downstream abnormal losses, likely because inconsistent signals exacerbate rumor dynamics. Similarly, adding audits without closing high-severity findings or without re-auditing after material upgrades shows no benefit, emphasizing that the quality and timeliness components of indices are the active ingredients. Insurance without meaningful covenants (coverage that excludes oracle/bridge failure or governance capture) improves reimbursement probabilities but does not shift conditional severity as much as policies with stronger covenants; protocols that add coverage while simultaneously weakening bounty SLAs show no net frequency benefit, consistent with moral hazard concerns that are then disciplined by underwriter requirements in more mature policies.

Ex post power diagnostics confirm that the study is adequately powered to detect the reported magnitudes. Simulations calibrated to the observed incident rate and dispersion (baseline $\lambda \approx 0.08$; $k \approx 1.2$) show 0.81–0.88 power to detect a 15–20% IRR change under staggered adoption with 30% treated by month 24. Survival models with ~ 300 incidents deliver ~ 0.83 power for a mitigation hazard ratio of 1.45; tail modeling with 289 exceedances achieves ~ 0.78 power to detect a $\Delta \xi$ of 0.07–0.09 under covariates. These figures align with the non-fragile

significance patterns observed in the main estimates and robustness checks.

In sum, the pattern of results is internally coherent and directly bears on the three hypotheses. Assurance—measured as audit depth with demonstrated closure and selective formal verification—reduces incident frequency and compresses the right tail of loss severity, with clean pre-trend diagnostics and stability across categories and chains, supporting H1. Incentives—measured by bounty intensity with fast, transparent payouts—reallocate discovery toward benign pre-disclosure and modestly lower realized exploitation, while insurance does its work primarily in recovery and governance discipline rather than in prevention; taken together, these channels support H2 without over-claiming frequency effects for coverage. Preparedness and transparency compress the time-to-mitigation and reduce conditional losses, and their returns are amplified when oracle and bridge dependencies are diversified. Transparency at highly central nodes substantially curbs downstream spillovers; these features support H3 and align with the theorized moderation by dependency entropy and centrality. Contagion remains a property of the ecosystem but is malleable under managerial controls. Speed gains do not systematically erode decentralization under constrained-discretion designs. Complementarities across assurance, incentives, and preparedness generate outsized benefits relative to isolated interventions. Diminishing returns suggest rational prioritization sequences. These empirical regularities provide the quantitative foundation on which subsequent findings, discussion, and conclusions can justify the assessment that security-by-governance is not merely a slogan but an effective portfolio of managerial levers for reducing realized cyber harm in permissionless finance.

5. Discussion

Interpretive scope. The numerical results reported in this manuscript are generated from a synthetic, calibrated panel rather than from a completed protocol-level observational dataset. Accordingly, the difference-in-differences, hazard, tail, and network estimates demonstrate the behavior and internal coherence of the proposed research design; they do not establish real-world causal effects. Empirical claims should be made only after the panel is reconstructed from auditable on-chain, governance, audit, bounty, insurance, and incident records and the full identification diagnostics are re-estimated.

The simulation-based analytical strategy—protocol-month panels with staggered-adoption difference-in-differences, tail-aware severity modeling, time-varying hazards for prevention and response, and explicit network terms—yields a unified picture: auditable, bounded, time-varying governance controls behave like real security mechanisms in DeFi. They reduce how often incidents occur, clip the right tail of losses when incidents do occur, accelerate containment, and dampen spillovers—without measurable re-centralization when implemented as constrained discretion. Taken together, these findings are not only statistically consistent with the identification design and diagnostics but also managerially meaningful at scales that justify budget and roadmap changes. Below, we interpret the

magnitudes through the mechanisms posited *ex ante*, make the hypothesis assessments explicit, surface limitations and boundary conditions consistent with the methods, and derive practice-ready implications that respect the quantitative evidence.

The assurance results align closely with the socio-technical framing and the operationalization that emphasizes independence, coverage, closure, and re-audit cadence over raw audit counts. The simulated 18% reduction in incident frequency following a +0.20 within-protocol increase in the Audit Depth Index (IRR = 0.82, clean pre-trends) is the quantitative footprint of defect stock shrinking along critical paths that actually gate funds. That the reduction peaks in months +2 to +6 and then stabilizes is consistent with remediation cycles and with attackers' learning curves being disrupted by closed findings and re-audits. Formal verification's hazard reduction (HR = 0.77) is mechanistically coherent: machine-checked invariants on collateralization, settlement, or bridge verification eliminate classes of "ruin" bugs that are hardest to test exhaustively. Severity results complete the picture. Thinner tails under high assurance ($\Delta\xi = -0.08$; $\Delta\beta = -22\%$; conditional exceedance 12.3% $\rightarrow$ 7.4%) mean that when defenses fail, the failure is less likely to compound into existential losses. The null for "audit logo count" without closure is especially probative: it separates signaling from substance and validates the index design choices specified in Methods. On this basis, H1 is supported on both frequency and tail severity, with effect sizes robust across chains, categories, and threshold choices.

Incentive results also match the mechanism-design intuition embedded in the methodology. Raising bounty intensity within a protocol—scaling tiers to TVL and committing to ≤ 7 -day payout SLAs—shifts discovery toward benign channels (pre-disclosures +31%) and modestly lowers realized incidents (IRR = 0.88). The saturation of headline tiers without concurrent latency improvement and the additional gains when latency falls quantify a simple idea: expected utility depends on both payout size and payout certainty/speed. The IV triangulation using platform-wide tier shocks further reduces the concern that only already mature teams upgrade bounties. Insurance behaves as hypothesized in the design: it does not change incident frequency (IRR $\approx$ 0.97) but materially improves user outcomes through faster reimbursement (+19pp at 30 days; +12pp at 90 days) and lower user-borne tail losses (−24%), particularly when covenants require re-audits and disclosure windows. Insurance thus disciplines governance and finances recovery rather than preventing exploits, which is precisely how the methodology proposed to interpret its role. Accordingly, H2 is supported with nuance: bounties reallocate discovery and reduce realized exploits at the margin, while insurance operates through post-incident discipline and restitution, not prevention.

Preparedness and transparency produce the largest operational gains in the response window, and the moderation patterns by dependency entropy and centrality align one-for-one with the network modeling apparatus. A +0.20 increase in Emergency Readiness lifts the sub distribution hazard of 24-hour containment by 47% and shrinks median containment time by 6.8 hours—magnitudes that translate directly into fewer drain cycles and narrower windows for adversaries and opportunistic arbitrageurs. Governance Transparency adds independent speed (sHR =

1.19) and reduces conditional losses ($\Delta\log\text{-loss-incident} = -0.14$). The shift toward surgical containment (+9.7pp) is exactly what rehearsed playbooks, circuit breakers, and distributed bounded authority are designed to achieve: changing parameters or patching the affected module while keeping the rest of the protocol and its dependents live. That these gains amplify with higher dependency entropy (ER $\times$ DE $\beta = 0.29$; sHR $\approx$ 1.66 at the 75th percentile of DE vs. 1.24 at the 25th) confirms that preparedness needs room to maneuver—diversified oracles/bridges and heterogeneous trust models create that room. Transparency's externalities scale with position: downstream abnormal losses fall 23% on average and $\sim 37\%$ at top-quintile centrality when the focal protocol communicates clearly and promptly. Baseline contagion ($\rho \approx 0.17$) falls toward irrelevance where both readiness and transparency are high ($\rho \approx 0.09$, ns) and worsens in concentrated stacks ($\rho \approx 0.24$, low DE), which exactly matches the spatial Durbin and local-projection logic in Methods. Critically, none of the speed is purchased via measurable re-centralization: validator/sequencer dispersion and signer entropy do not fall; time clocks shorten slightly, but sunsets and M-of-N thresholds strengthen; emergency actions are logged and re-audited. This is the empirical signature of constrained discretion rather than quiet capture. On that basis, H3 is strongly supported: preparedness and transparency accelerate containment and reduce harm, with effects moderated upward by dependency entropy and centrality in line with the network design.

The complementarity and diminishing-returns patterns satisfy another prediction implicit in the methodology: controls operate as a portfolio. The triad move—from below- to above-median on assurance (with closure and re-audit), bounty intensity with ≤ 7 -day SLA, and readiness—produces a composite $\sim 34\%$ frequency reduction, $\Delta\xi \approx -0.07$, and ~ 9.1 hours faster median containment, far exceeding any single lever alone (11–15% frequency reductions; 2–4 hours). The concavity in assurance (big gains 1 $\rightarrow$ 2 independent audits, smaller gains 2 $\rightarrow$ 3 absent FV), saturation in bounties without latency cuts, and linear-then-flat readiness unless drills include external partners or entropy rises, all point to a rational sequencing that was prefigured by the Methods: buy independence/closure/re-audit first; then verify the critical invariants not already covered; then make payout speed real; then drill with integrators/exchanges; and, in parallel, diversify oracles/bridges so the runbook can actually run.

Heterogeneity by category and chain is not noise—it is content that validates external construct coherence. Bridges show the largest tail relief under assurance and FV ($\Delta\xi \approx -0.11$), yet smaller frequency reductions unless ER is high. That fits the cross-domain coordination they require. DEXs emit the biggest transparency externalities (−29% downstream abnormal losses among dependents) because they anchor price discovery and liquidation mechanics; clarity from a hub reduces rumor-driven runs. Lending protocols respond most to incentive upgrades (pre-disclosure RR $\approx$ 1.41) because many of their vulnerabilities are economically subtle oracle edge cases that skilled researchers will chase when programs feel professional. L2s yield larger readiness dividends (sHR $\approx$ 1.58 vs. 1.39 on L1s) but also exhibit higher baseline contagion ($\rho \approx 0.20$), implying a higher marginal value of transparency—again, a neat fit with the network-aware design.

Limitations track the exact caveats anticipated in the

Methods and are unlikely to reverse the main inferences. The study is observational, so identification leans on staggered adoption with tested pre-trends and on triangulation via timing instruments; unobserved time-varying confounding cannot be categorically eliminated, but E-values (≈ 1.7 – 1.8 for the largest IRRs) and stability across robustness checks bound its plausible influence. Index construction entails coder judgment; dual annotation with $\kappa \geq 0.78$ and error-in-variables sensitivity (shrinking effects modestly without flipping signs) indicate that estimates are not artifacts of measurement noise. Tail modeling requires threshold selection; diagnostics (mean residual life linearity, stability windows, Hill plateaus) and threshold sweeps ($u \in [\$3M, \$8M]$) mitigate cherry-picking risk. Missingness handled under MAR via MICE, plus pattern-mixture stress tests, moves magnitudes by less than one standard error; complete-case replications track the headline results. Spatial specifications are necessarily stylized; alternate adjacency definitions produce ρ within ± 0.05 and preserve GT/ER attenuation, reducing worry that a particular graph draws the result. Finally, cost translations rely on average loss parameters in a heavy-tail regime; these are offered as decision aids with uncertainty ranges, not as deterministic forecasts. None of these limitations contradict the directional findings, and most attenuate magnitudes (e.g., coder-error adjustments), making the supported effects conservative.

Two boundary conditions deserve emphasis for practice. First, transparency must be structured to help rather than harm: change logs, on-chain decision trails, and post-mortems on a clock produce the measured gains; sporadic blog posts and vague notices do not, and may worsen downstream outcomes by amplifying rumor dynamics. Second, assurance only “counts” when it closes severities and returns for re-audits after material upgrades; adding audit logos without those process commitments does nothing in the data and should not be rewarded in standards or listings.

Managerial implications follow directly from effect sizes and from the concavity/interaction surfaces. For protocols below median governance maturity, the highest marginal returns per dollar are to (i) convert single audits into two independent audits with enforced closure and re-audit cadence; (ii) implement a bounty SLA that pays within a week and credits researchers publicly, even if headline maxima move only modestly; and (iii) run a full-dress incident drill that includes at least one major integrator and exchange, while ensuring emergency powers are distributed, time-bounded, and logged. In stacks with low dependency entropy, governance upgrades should be paired with immediate diversification of oracles/bridges (or upgrades to light-client/zk trust models), otherwise ER returns will be blunted by upstream chokepoints. Highly central protocols should over-index on transparency and rehearsal, because their marginal externality is larger; fewer central protocols still gain but contribute less spillover relief. Insurance should be treated as a recovery/discipline tool: buy policies with covenants that align with the playbook (re-audit obligations, disclosure windows, peril definitions that include oracle/bridge and governance capture). Auditors and formal-methods teams should steer buyers toward invariant coverage that maps to ruin paths; beyond two independent audits with closure, selective FV on unverified critical modules dominates a third generalist engagement in risk-

reduction per dollar. Bounty platforms should optimize for payout latency and triage reliability, not only headline maxima. Exchanges and custodians can use the governance indices as part of listing, freeze, and relisting decisions; the net-benefit curves show real gains at practical decision thresholds.

At the ecosystem level, the counterfactual and forward simulations indicate that modeled governance improvements would have averted within the simulation on the order of a hundred incidents over the sample horizon and that scaling the triad upgrade across the bottom half of protocols would yield sizable additional loss reductions. These macro numbers are consistent with unit-level effect sizes and serve as coordination beacons: DAOs, insurers, and standards bodies can converge on a minimum viable governance baseline—two independent audits with closure and re-audit; partial FV on ruin-gating invariants; bounty with ≤ 7 -day payouts; distributed, time-bounded emergency powers drilled with external partners; on-chain transparency and post-mortems on a clock; diversified oracle/bridge providers—as a shared floor that buys measurable protection without eroding decentralization.

Finally, the hypothesis assessments are unambiguous in the light of the quantitative evidence and the identification checks. H1 is supported: deeper, higher-quality assurance and selective formal verification reduce incident frequency and thin the loss tail, with dynamic patterns, nulls for superficial proxies, and tail diagnostics all pointing in the same direction. H2 is supported with nuance: incentive alignment via bounties works when payment frictions are low and tiers match exposure; insurance does not prevent incidents but improves recovery and disciplines governance when covenants bind. H3 is supported strongly: preparedness and transparency accelerate containment, reduce conditional severity, and curb spillovers, with effects amplified by diversified dependencies and central network position, and achieved without detectable re-centralization under constrained-discretion designs. The discussion, therefore, closes the loop between design, estimation, and interpretation: the managerial controls pre-registered in the Methods operate through the mechanisms the theory section proposed, with magnitudes that withstand robustness and that translate into clear sequencing and budgeting guidance. “Security-by-governance” is not a metaphor here; it is an evidence-backed operating system for risk reduction in permissionless finance.

6. Conclusion

This study develops a simulation-based framework for testing the proposition that governance can function as a security primitive in decentralized finance. It operationalizes three falsifiable hypotheses and combines staggered-adoption designs, heavy-tail models, time-varying hazards, and network analysis in a synthetic but internally consistent panel of 176 protocols over 48 months. The resulting estimates are illustrative outputs of the calibrated design rather than causal estimates from observed protocols. Their value is therefore methodological: they show how an auditable empirical study can quantify the frequency, severity, containment, and spillover channels associated with governance controls.

Begin with the methodological spine, because credibility flows from design. The protocol-month panel and event-study difference-in-differences estimators isolate within-

protocol changes around the adoption of governance controls, with pre-trend (lead) diagnostics to guard against spurious dynamics. Zero-inflated negative binomial models acknowledge rare, over-dispersed incident counts; generalized Pareto models focus on the extreme-loss tail where solvency lives or dies; Cox and Fine–Gray hazards separate prevention (time-to-exploit) from response (time-to-mitigation) while accommodating time-varying covariates; and spatial Durbin plus local projections capture exposure to neighbors and spillover remediation. The data inputs are public, auditable artifacts—audits with closure evidence, formal verification traces, bounty ledgers with payout timestamps, governance logs, post-mortems, insurance covenants, and on-chain dependency maps—linked deterministically to contract systems and double-coded for reliability. Robustness traversed placebo timings, threshold sweeps, leave-one-domain re-estimations, instrumented triangulations for timing shocks, and coder error-in-variables sensitivity. With that scaffold, the magnitudes can be interpreted as internally consistent simulation outputs.

Hypothesis 1 (assurance hardens systems and thins tails) is supported both statistically and substantively. A within-protocol +0.20 increase in the Audit Depth Index (ADI)—operationally: adding a second independent audit; increasing coverage of critical invariants; closing high-severity findings; and instituting re-audit cadence after material upgrades—cuts incident frequency over the subsequent year with an incidence-rate ratio of 0.82 (95% CI [0.75, 0.90]) and clean pre-trends. Selective formal verification on ruinating paths adds an independent prevention margin (Cox HR ≈ 0.77 for first exploit), lengthening healthy uptime by months at the median. Most importantly, assurance changes the shape of loss: under high assurance (ADI ≥ 0.60 , FV present), the tail shape ξ drops by ~ 0.08 and the scale β by $\sim 22\%$, reducing the conditional probability that a \$5M+ loss escalates beyond \$50M from 12.3% to 7.4%. Naïve “audit logo count” without closure or re-audit has no detectable effect. This pattern is exactly what the theory predicted and the methods were tuned to detect: assurance must be substantive (independence, coverage, closure, cadence) to move field outcomes; when it is, both frequency and catastrophic potential fall.

Hypothesis 2 (incentives and insurance discipline behavior and outcomes) is also supported, with important nuances that matter for policy design. Raising bug-bounty intensity within a protocol from its own lower to upper quartile—tiers scaled to TVL plus a ≤ 7 -day payout SLA—lifts monthly pre-disclosures by $\sim 31\%$ and reduces realized incident frequency by $\sim 12\%$ (IRR ≈ 0.88). Decomposition shows saturation in headline tiers unless payout latency falls; cutting median payout time from ~ 21 days to ≤ 7 days adds another $\sim 17\%$ to pre-disclosures and $\sim 4\%$ to the frequency reduction. In short, speed is part of the reward—expected utility is magnitude $\times$ certainty/speed. Instrumented analyses using platform-wide tier shocks replicate the direction and magnitude, mitigating worries that only already careful teams buy bigger bounties. Insurance behaves as recovery finance and governance discipline rather than prevention: incident counts are flat (IRR ≈ 0.97), but user reimbursement improves by 19 percentage points at 30 days (41% $\rightarrow$ 60%) and 12 points at 90 days (66% $\rightarrow$ 78%), and user-borne tail losses fall by $\sim 24\%$, particularly when covenants force re-audits and timely disclosure. The

implication for practice is crisp: treat bounties as a *prevention-adjacent* lever that works when frictions are low; treat insurance as *post-incident discipline and restitution*, and buy policies with teeth (covariates and peril breadth) rather than marketing veneer.

Hypothesis 3 (preparedness and transparency accelerate mitigation and curb spillovers, especially under diversified dependencies and central positions) is strongly supported and delivers the largest “operational” wins. A +0.20 increase in the Emergency Readiness (ER) index—current runbooks; at least one cross-organization drill including an exchange/integrator; distributed signers with bounded emergency powers; live circuit breakers; timeclocks tuned for bounded agility—raises the sub distribution hazard of first-day containment by $\sim 47\%$ (sHR ≈ 1.47) and shortens median containment by ~ 6.8 hours. Governance Transparency (GT)—timely, structured change logs; on-chain decision trails; post-mortems on a clock—adds independent speed (sHR ≈ 1.19) and reduces conditional severity ($\Delta \log\text{-loss} \approx -0.14$). Preparedness also changes the *mode* of response: surgical containment (patch or parameter change) becomes ~ 9.7 percentage points more likely than a blunt global pause, curbing collateral disruptions for dependents. These gains scale with the dependency graph: ER’s return is higher under diversified oracles/bridges (ER $\times$ DE $\beta \approx 0.29$; implied sHR ≈ 1.66 at the 75th percentile of dependency entropy vs. ~ 1.24 at the 25th), and GT reduces 72-hour downstream abnormal losses by $\sim 23\%$ on average and $\sim 37\%$ for top-quintile centrality nodes. Baseline contagion exists (SAR $\rho \approx 0.17$) but is *malleable*: where ER and GT are high, ρ trends toward ~ 0.09 (ns); where dependency entropy is low, ρ rises toward ~ 0.24 . This is governance-as-coordination made quantitative: runbooks work best when the stack is diversified, and clear communication from hubs stabilizes neighborhoods.

A critical guardrail in the design was to detect any “speed via centralization” trade-off. We do not find it in aggregate. As readiness rises, validator/sequencer dispersion does not fall ($\Delta \text{DEC_validator} \approx +0.003$, n.s.), signer entropy does not erode ($\sim +0.02$, n.s.), and although time clocks shorten modestly ($\sim 4.1\text{h}$, weak), sunset clauses and M-of-N thresholds strengthen (both $p < 0.05$). Emergency actions are logged and re-audited more frequently. This is the signature of constrained discretion—bounded, distributed, auditable emergency authority—rather than quiet capture. It matters for legitimacy, listings, custody policies, and regulatory posture: faster response is compatible with decentralization when *designed*.

The controls interact, and the shape of those interactions informs sequencing. Moving from below- to above-median simultaneously on (i) assurance with closure and re-audit, (ii) bounty intensity with a ≤ 7 -day SLA, and (iii) readiness yields a $\sim 34\%$ reduction in incident frequency, a ~ 0.07 drop in the tail index, and ~ 9.1 hours faster median containment—multiples larger than any single lever (11–15% frequency, 2–4 hours). Returns to assurance are steep at the low end (one $\rightarrow$ two independent audits with closure $>$ two $\rightarrow$ three absent FV), bounties saturate without latency cuts, and readiness is roughly linear until drills include external partners or the stack is diversified. The MVG baseline that emerges is therefore not a slogan but a checklist with measurable thresholds: **two independent audits with closure and re-audit after material upgrades;

selective formal verification on ruin-gating invariants; bounties scaled to TVL with a ≤ 7 -day payout SLA and public researcher credit; distributed, time-bounded emergency powers drilled with an exchange and a major integrator; structured transparency with on-chain decision trails and post-mortems on a clock; and diversified oracles/bridges to raise dependency entropy**. If these relationships are validated empirically, protocols implementing this baseline may experience meaningfully fewer incidents, thinner tails, faster and more surgical containment, and calmer neighborhoods.

The economics closes the loop. Using mid-range costs—~\$120k–\$350k for an incremental independent audit with closure; ~\$200k–\$1M annualized for a bounty upgrade with a real SLA and triage; ~\$50k–\$150k for a cross-org drill—the first-year spend to move from below- to above-median on assurance, incentives, and readiness is ~\$1.1M. The measured benefits—0.026 incidents prevented per protocol-month; ~22% lower tail scale; ~6.8 hours faster containment—translate, at ~\$500M TVL and the modeled loss distribution, into ~\$6.5–\$8.3M in expected-loss reductions per protocol-year, before spillovers and reputational effects. Counterfactual simulations suggest that the governance improvements observed across the sample would have averted within the simulation on the order of a hundred incidents over the 48-month window. Decision-curve analyses indicate real net-benefit gains at practical intervention thresholds. None of these are point forecasts; all are bounded, conservative translations of coefficients into budget-worthy terms.

The guardrails are as important as the green lights. Transparency must be **structured**—change logs, decision trails, post-mortems with timelines—not sporadic posts that can amplify rumor dynamics. Assurance must include **closure** and **re-audit**; logo accumulation should not earn reputational or economic rewards. Insurance without **covariates** and **peril breadth** smooths cashflow but does not shrink the tail borne by users; policies should be evaluated on covenants first, not marketing. Preparedness that never leaves the room will plateau; drills must include integrators and exchanges to surface real coordination bottlenecks. Diversification is not optional if you want the modeled ER gains; defenders need room to maneuver when a provider or trust model is the bottleneck.

Limitations temper the generality of claims but do not unravel them. The design is observational; staggered-adoption DiD with flat leads, timing instruments, and consistent survival and tail effects reduce (but cannot eliminate) the risk of unmeasured, time-varying confounding. Index construction involves judgment; dual-annotation $\kappa \geq 0.78$ and error-in-variables sensitivity suggest that any bias attenuates magnitudes rather than fabricates them. Tail inference relies on thresholds; mean residual life linearity, parameter-stability windows, and u-sweeps from \$3M–\$8M preserve directions and significance. Spatial specifications simplify a complex web; adjacency variants keep ρ within ± 0.05 and preserve transparency/readiness attenuation of spillovers. Missingness handled via MICE and pattern-mixture stress shifts magnitudes by less than one SE; complete-case re-estimates track headline results. These caveats are visible by design, and most point toward conservative, not exaggerated, effects. More generally, emerging financial technologies can create simultaneous opportunities and governance challenges, making

organizational readiness and regulatory clarity important boundary conditions for responsible adoption (Taheri & Taieby, 2025b) ^[33].

Implications follow for each stakeholder class. Protocol teams should adopt the MVG baseline and sequence investments where marginal returns are steepest: two independent audits with closure and re-audit cadence; then selective FV on unverified critical invariants; then bounty SLA cuts; then cross-org drills; and in parallel, diversification of oracles/bridges. Auditors and formal-methods firms should steer buyers toward invariant coverage that maps to **ruin paths** and insist on closure discipline; a third generalist audit is dominated by selective FV after two independent audits with closure. Bounty platforms should optimize for **latency** and **triage reliability**, not only headline maxima. Exchanges and custodians can use ER/GT indices in listing, freeze, and relisting policies; central hubs should over-invest in transparency because their disclosures stabilize neighborhoods. Insurers should write **covenanted** policies that enforce re-audits and disclosure windows, and price discounts against measurable governance indices; buyers should prefer peril breadth over low premiums with exclusions. Standards bodies and analytics platforms can codify the MVG baseline and publish governance scorecards backed by artifact-verifiable measures, creating a Schelling point for budgets and coordination.

Taken together, the simulation results support explicit hypothesis assessments. **H1** is supported: deeper, higher-quality assurance (independence, coverage, closure, re-audit) and selective formal verification reduce incident frequency and thin extreme-loss tails; superficial proxies do not. **H2** is supported with nuance: bounties reduce realized exploitation when payout frictions are low and tiers match exposure; insurance improves recovery and disciplines governance through covenants rather than changing ex-ante frequency. **H3** is strongly supported: preparedness and transparency accelerate containment, reduce conditional severity, and curb spillovers, with larger gains when dependencies are diversified and when the focal protocol is central—and these gains are achieved without measurable re-centralization under constrained-discretion designs. These conclusions are not rhetorical; they are the joint product of identification strategies aligned to on-chain realities, tail-aware severity modeling, time-varying hazards that separate prevention from response, and network terms that internalize interdependence.

The doctrinal lesson for DeFi is not that “code is not law,” but that **governance** is part of the law that code executes. Bounded authority encoded in time clocks and multisets; credible incentives encoded in bounty SLAs and transparent ledgers; clarity encoded in change logs and post-mortems; and diversification encoded in provider sets and trust models—together, these managerial choices are as much part of the security boundary as the bytecode itself. The MVG baseline that emerges here is a floor, not a ceiling, but it is a floor that buys real protection without dulling the decentralized edge. If validated with observed data, teams adopting it may experience fewer incidents, smaller catastrophes, faster and more surgical containment, and calmer networks. Ecosystem actors that reward it through listings, coverage terms, or public benchmarks will align private incentives with public resilience. That alignment is the practical meaning of treating governance as

a security primitive in permissionless finance—and it is the route by which a trust-minimizing architecture earns trust where it matters most: when things go wrong.

7. References

1. Abakah EJA, Alagidede P, Mensah L, Omari-Sasu A. Global uncertainty factors and price connectedness between U.S. electricity, cryptocurrencies, NFTs, and DeFi markets. *Finance Research Letters*. 2024; 61:104990. <https://www.sciencedirect.com/science/article/abs/pii/S1544612323010656>
2. Akkus HT, Ersan O, Kaplan M. Analysis of dynamic connectedness relationships between cryptocurrency, NFT, and DeFi assets. *Applied Economics Letters*. 2024; 31(20):2765-2770. <https://www.tandfonline.com/doi/full/10.1080/13504851.2023.2216437>
3. Alexander C. Price discovery and efficiency in Uniswap liquidity pools. *Journal of Futures Markets*. 2025; 45(8):1241-1269. <https://onlinelibrary.wiley.com/doi/10.1002/fut.22593>
4. Ali S, Tiwari AK, Hkiri B. ESG meets DeFi: Exploring time-varying linkages and risk spillovers. *International Journal of Finance & Economics*. 2023; 28(4):4750-4772. <https://onlinelibrary.wiley.com/doi/full/10.1002/ijfe.3060>
5. Bertomeu J, Martin X, Sall I. Measuring DeFi risk. *Finance Research Letters*. 2024; 63:105321. Doi: <https://doi.org/10.1016/j.frl.2024.105321>
6. Bhambhwani SM. Governing decentralised finance. *International Journal of Finance & Economics*, 2025. <https://onlinelibrary.wiley.com/doi/10.1002/ijfe.3157>
7. Corbet S, Larkin C, Lucey B, Yarovaya L. What drives DeFi prices? Investigating the effects of investor attention. *Finance Research Letters*. 2022; 48:102883. <https://www.sciencedirect.com/science/article/abs/pii/S1544612322001672>
8. De l'Etang F, Shust E, Zeni S. The dollar's role in institutional and media impact on stablecoins. *Finance Research Letters*. 2024; 61:104990. <https://www.sciencedirect.com/science/article/abs/pii/S1544612324000291>
9. Eichengreen B. Stablecoin devaluation risk. *The European Journal of Finance*. 2025; 31(9-10):1049-1071. <https://www.tandfonline.com/doi/full/10.1080/1351847X.2025.2505757>
10. Ellinger EW. Decentralized autonomous organization (DAO): The case of MakerDAO. *Journal of Information Technology Teaching Cases*. 2024; 14(2):265-272. <https://journals.sagepub.com/doi/10.1177/20438869231181151>
11. Fernández-Mejía J. Extremely stablecoins. *Finance Research Letters*. 2024; 63:105268. <https://www.sciencedirect.com/science/article/pii/S1544612324002988>
12. Fukasawa M. Weighted variance swaps hedge against impermanent loss. *Quantitative Finance*. 2023; 23(12):1841-1857. <https://www.tandfonline.com/doi/full/10.1080/14697688.2023.2202708>
13. Fukasawa M. Model-free hedging of impermanent loss in geometric Brownian AMMs. *Applied Economics Letters*. 2024; 31(20):2749-2754. <https://www.tandfonline.com/doi/full/10.1080/1350485X.2024.2404058>
14. Goghie AS. Tokenization and the banking system: Redefining authority in DeFi. *Competition & Change*. 2024; 28(5):641-659. <https://journals.sagepub.com/doi/10.1177/10245294241258255>
15. Huang C-C, Hsu C-C, Chen Y-T. Evaluating dependence between DeFi tokens and cryptocurrencies. *Applied Economics Letters*, 2025. <https://www.tandfonline.com/doi/full/10.1080/13504851.2023.2301460>
16. Lehar A, Parlour C. Decentralized exchange: The Uniswap automated market maker. *The Journal of Finance*. 2025; 80(1):321-374. Doi: <https://doi.org/10.1111/jofi.13405>
17. Makridis CA, Fröwis M, Sridhar K, Böhme R. The rise of decentralized cryptocurrency exchanges: Evaluating the role of airdrops and governance tokens. *Journal of Corporate Finance*. 2023; 79:102358. Doi: <https://doi.org/10.1016/j.jcorpfin.2023.102358>
18. Mensi W, Zaremba A, Kizys R, Vo XV, Kang SH. Frequency connectedness between DeFi and cryptocurrencies. *The North American Journal of Economics and Finance*. 2024; 70:101980. <https://www.sciencedirect.com/science/article/abs/pii/S1062976923001229>
19. Nguyen LTM, Nguyen PT. Determinants of cryptocurrency and decentralized finance adoption-A configurational exploration. *Technological Forecasting and Social Change*. 2024; 201:123244. Doi: <https://doi.org/10.1016/j.techfore.2024.123244>
20. Olk C. A credit theory of anti-credit money. *Review of International Political Economy*. 2025; 32(8):2846-2869. <https://www.tandfonline.com/doi/full/10.1080/09692290.2025.2476738>
21. Piñeiro-Chousa J, Šević A, González-López I. Impact of social metrics in decentralized finance. *Journal of Business Research*. 2023; 158:113673. Doi: <https://doi.org/10.1016/j.jbusres.2023.113673>
22. Puschmann T, Huang-Sui M. A taxonomy for decentralized finance. *International Review of Financial Analysis*. 2024; 92:103083. Doi: <https://doi.org/10.1016/j.irfa.2024.103083>
23. Qiao X, Umar Z, Rizvi SAR. Time-frequency extreme risk spillover network of cryptocurrency coins, DeFi and NFTs. *Finance Research Letters*. 2023; 51:103489. <https://www.sciencedirect.com/science/article/abs/pii/S1544612322006651>
24. Singh SF, Shah D. Option contracts in the DeFi ecosystem: Opportunities, challenges, and design. *International Journal of Network Management*. 2025; 35(3):e70005. <https://onlinelibrary.wiley.com/doi/10.1002/nem.70005>
25. Şoiman F, Moinescu B, Lazar D. What drives DeFi market returns? *Journal of International Financial Markets, Institutions & Money*. 2023; 83:101767. <https://www.sciencedirect.com/science/article/abs/pii/S1042443123000549>

26. Son J, Bae J, Lee S. Competitive dynamics between decentralized and centralized finance lending markets. *International Review of Financial Analysis*. 2024; 96(Part B):103761. <https://www.sciencedirect.com/journal/international-review-of-financial-analysis/vol/96/part/PB>
27. Xu R, Liu Z, Zhang Y. Decentralized finance (DeFi): A paradigm shift in the FinTech ecosystem. *Enterprise Information Systems*. 2024; 18(10):12658-12686. <https://www.tandfonline.com/doi/full/10.1080/17517575.2024.2397630>
28. Yousaf I, Nekhili M, Gubareva M. Connectedness between DeFi assets and equity markets during COVID-19. *Technological Forecasting and Social Change*. 2023; 187:122151. <https://www.sciencedirect.com/science/article/abs/pii/S0040162522006953>
29. Yousaf I, Nekhili M, Gubareva M. Dynamic spillovers between leading cryptocurrencies and DeFi tokens. *International Review of Financial Analysis*. 2024; 98:103904. <https://www.sciencedirect.com/science/article/abs/pii/S1057521924000887>
30. Zhang Y, Chen S, Zhao W. The adaptive market hypothesis of decentralized finance (DeFi). *Applied Economics*. 2023; 55(57):6676-6690. <https://www.tandfonline.com/doi/full/10.1080/00036846.2022.2133895>
31. Taheri A. AI-Powered Financial Risk Modeling in Tokenized Asset Environments: A Network-Based Analysis. *ISAR Journal of Science and Technology*. 2023; 1(2):44-52. <https://article.isarpublisher.com/viewArticle/AI-Powered-Financial-Risk-Modeling-in-Tokenized-Asset-Environments-A-Network-Based-Analysis>
32. Taheri A, Taieby E. Integrating ESG Analytics into Corporate Decision-Making: A Data-Driven Approach for Enhancing Sustainable Financial Performance. *FAR Journal of Multidisciplinary Studies*. 2025a; 2(3):1-19. Doi: <https://doi.org/10.5281/zenodo.17235625>
33. Taheri A, Taieby E. Strategic Opportunities and Challenges of Quantum Computing Adoption in Financial Risk Management: A Technology Management Perspective. *FAR Journal of Financial and Business Research*. 2025b; 1(5):55-77. Doi: <https://doi.org/10.5281/zenodo.17235751>