

Intelligent Decentralized Insurance: Analytical Perspectives on Automated Risk and Trust Systems

Wei Zhang¹, Li Zhao^{2,*}, Hao Chen³

¹ School of Computer Science and Engineering, Southeast University, No.2 Southeast University Road, Jiangning District, Nanjing 211189, P.R.China

² School of Computer Science and Technology, Huazhong University of Science and Technology, 1037 Luoyu Road, Hongshan District, Wuhan 430074, P.R.China

³ College of Computer Science and Technology, Zhejiang University, 866 Yuhangtang Road, Xihu District, Hangzhou 310058, P.R.China

* Corresponding author: lizhao@hust.edu.cn

ARTICLE INFO Received July 15, 2023 Revised September 25, 2023 Accepted November 16, 2023 Available Online December 30, 2023 DOI 10.63646/jaiaa.2023.010405 License Creative Commons Attribution 4.0 International Licence (CC BY 4.0) Publisher INATGI, United States of America Journal JAIAA - ISSN 3067-7386	Abstract Decentralized insurance combines distributed-ledger infrastructure, self-executing contracts, and machine intelligence to perform risk pooling, pricing, and claim settlement without a single coordinating intermediary. This paper develops an analytical perspective on such intelligent decentralized insurance, asking how automation reshapes the cost structure of risk transfer and how trust is produced when no central insurer stands behind the promise to pay. A four-plane reference architecture is proposed that separates participants and assets, data and oracles, automation logic, and governance assurance, and three families of automated risk architectures are examined through this lens: algorithmic mutual pools, parametric trigger contracts, and token-governed coverage protocols. The analysis is supported by a structured synthesis of peer-reviewed research published between 2016 and 2024 and by stylized quantitative models. The expense decomposition indicates that automation can plausibly compress operating loads from roughly 29 percent of premium toward 12 percent, while settlement latency falls from weeks to minutes for fully automated parametric designs. A pool-stability model shows, however, that loss dependence places a hard floor under the solvency benefits of scale, and an oracle redundancy model quantifies how trigger integrity improves geometrically with independent data sources. The paper argues that the binding constraints on intelligent decentralized insurance are no longer purely technological: they concern data integrity, contract security, governance legitimacy, and regulatory recognition. A research agenda is outlined that links actuarial risk-sharing theory, information systems design, and institutional economics. Keywords: Decentralized insurance; smart contracts; blockchain; artificial intelligence; risk pooling; oracles; trust systems; insurtech
--	---

I. INTRODUCTION

Insurance is, at its core, an institution for organizing trust among strangers. A policyholder pays a premium today in exchange for a promise whose value is realized only if an uncertain loss occurs tomorrow, and the credibility of that promise has historically been manufactured by a centralized firm: its capital, its brand, its actuaries, and the regulatory apparatus that supervises it. The financial technology movement of the past decade has demonstrated that many functions once thought to require such centralized institutions can be re-implemented as digital services built on shared infrastructure (Lee and Shin, 2018; Gomber et al., 2018). Within insurance specifically, digitalization has already reorganized distribution, underwriting, and claims workflows, and it has begun to alter which risks are insurable at all (Eling and Lehmann, 2018). The question this paper

addresses is more radical: what happens to the economics and the trust architecture of insurance when the coordinating intermediary itself is replaced by a combination of distributed ledgers, self-executing contracts, and machine intelligence?

Two technological currents make the question timely. The first is the maturation of blockchain as an organizational technology rather than merely a payments experiment. A distributed ledger provides an append-only, jointly verified record of state, which allows mutually distrusting parties to coordinate on the history of premiums paid, claims filed, and capital held without delegating bookkeeping to any one of them (Lu, 2018; Zheng et al., 2018). Smart contracts extend this shared state with shared logic: program code that executes deterministically when predefined conditions are observed, so that the rules of a risk pool become self-enforcing artifacts rather than clauses requiring interpretation and goodwill (Christidis and Devetsikiotis, 2016; Wang et al., 2019). The second current is the rapid progress of artificial intelligence in exactly the analytical tasks that insurance depends upon, including pattern-based pricing, anomaly and fraud detection, and conversational service automation (Lu, 2019a; Riikinen et al., 2018; Eling et al., 2022). Decentralized finance has already shown that lending, exchange, and derivative functions can be recomposed from these primitives into open protocols (Schär, 2021; Xu et al., 2024); insurance is a natural, if more demanding, next domain.

The demand side is equally important. The cost of operating a conventional insurer is substantial: acquisition commissions, policy administration, and adversarial claims handling together absorb a material share of every premium dollar, and these frictions are passed to policyholders or expressed as coverage exclusions. At the same time, opaque pricing and discretionary claim denial sustain a chronic deficit of consumer confidence, a problem that personalization based on behavioral data has arguably sharpened rather than relieved (Cevolini and Esposito, 2020). Decentralized designs respond to both frictions simultaneously. By automating administration they attack the expense ratio, and by making pool finances and decision rules publicly inspectable they attempt to substitute verifiable computation for institutional reputation as the basis of trust (Hawlitschek et al., 2018; Kar and Navin, 2021).

Despite a fast-growing literature, three gaps motivate the present analysis. First, technical surveys of blockchain insurance applications tend to catalog prototypes without an economic account of where automation actually saves cost and where it merely relocates it (Gatteschi et al., 2018; Dominguez Anguiano and Parte, 2023). Second, the actuarial literature on peer-to-peer risk sharing has developed rigorous allocation theory (Denuit, 2019; Denuit et al., 2022) largely in isolation from the information-systems literature on oracles, contract security, and platform governance. Third, claims about trust are often asserted rather than analyzed: replacing an insurer with code does not eliminate trust requirements but redistributes them onto data feeds, developers, auditors, and token-holding voters (Lumineau et al., 2021). This paper aims to integrate these strands into a single analytical treatment.

The contribution is threefold. The paper (1) proposes a four-plane reference architecture that makes the division of labor between ledgers, oracles, AI components, and governance mechanisms explicit; (2) compares three automated risk architectures, namely algorithmic mutual pools, parametric trigger contracts, and token-governed coverage protocols, with respect to their cost structure, settlement behavior, and residual trust assumptions; and (3) provides stylized quantitative perspectives on expense compression, settlement latency, pool stability under loss dependence, and oracle redundancy, which together identify the conditions under which intelligent decentralized insurance is economically and institutionally viable. Figure 1 introduces the reference architecture that organizes the remainder of the paper.

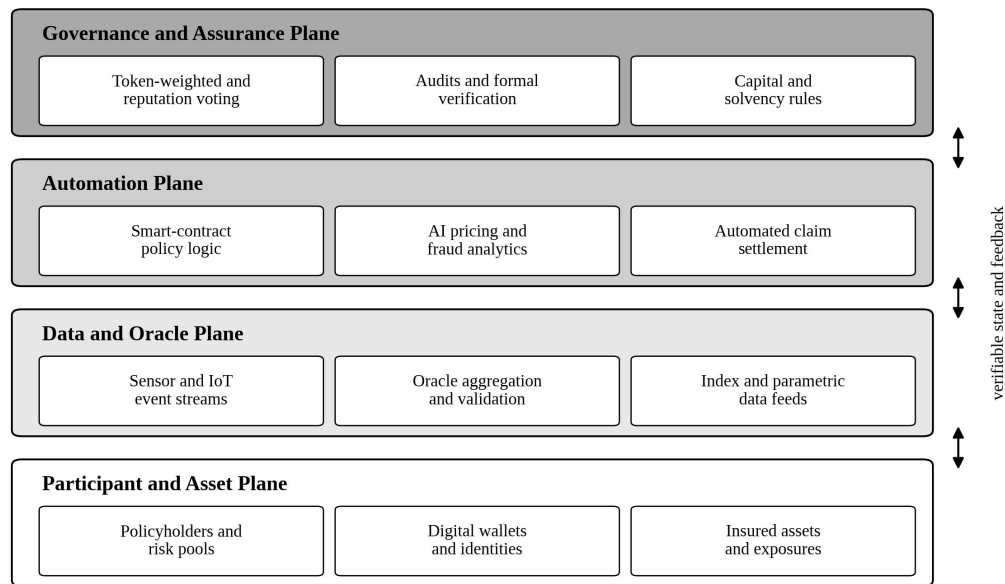


Figure 1. Four-plane reference architecture of intelligent decentralized insurance.

Figure 1 separates the system into four interacting planes. The participant and asset plane contains the policyholders, their digital identities and wallets, and the exposures being insured. The data and oracle plane transports facts about the world, including sensor streams from connected devices and curated index feeds, into machine-readable form (Lu and Xu, 2019; Xu et al., 2021). The automation plane hosts the smart contracts that encode policy terms and the AI services that price risk, screen claims, and detect anomalies (Salah et al., 2019). The governance and assurance plane supplies the residual human judgment that code cannot, through voting, auditing, formal verification, and capital rules (Beck et al., 2018). The vertical channel of verifiable state and feedback is what distinguishes this architecture from a conventional insurer with a modern IT stack: every plane reads from and writes to a record that all participants can independently check.

The paper proceeds as follows. Section II reviews the conceptual foundations in digital insurance, distributed-ledger research, and decentralized finance, and positions the contribution against prior work. Section III states the analytical framework and method. Section IV analyzes the three automated risk architectures. Section V examines the trust systems on which they depend, namely oracles, contract security, and governance. Section VI develops the quantitative perspectives. Section VII discusses strategic and regulatory implications together with the limitations of the analysis, and Section VIII concludes with a research agenda.

II. CONCEPTUAL FOUNDATIONS AND RELATED RESEARCH

The digital transformation of insurance predates distributed ledgers and provides the baseline against which decentralized designs must be judged. Research on the digitalization of the insurance value chain shows that data-rich technologies alter every link from product design to claims, and that they shift the boundary of insurability by making previously unobservable risks measurable (Eling and Lehmann, 2018). Artificial intelligence contributes along the entire chain: machine-learning models improve loss prediction and pricing granularity, natural-language systems automate service interactions, and anomaly detection strengthens fraud control (Riikinen et al., 2018; Eling et al., 2022). Text-analytic deep learning has been shown to materially improve automobile fraud detection over conventional screens, illustrating that claims integrity is increasingly an algorithmic capability rather than a purely investigative one (Wang and Xu, 2018). Sociological work cautions, however, that data-driven personalization transforms the very logic of pooling, moving insurance

from the averaging of fates toward individualized behavioral contracts, with contested distributional consequences (Cevolini and Esposito, 2020). These findings frame the first analytical theme of this paper: automation changes not only the cost of insurance operations but the social meaning of the pool.

A second foundation is the general research program on blockchain as an information-systems artifact. Early framework papers organized the field around technical, organizational, and societal levels of analysis and called for design knowledge rather than speculation (Risius and Spohrer, 2017). Systematic reviews document a rapid diffusion of blockchain applications across finance, supply chains, and government, while noting the recurring bottlenecks of scalability, interoperability, and energy cost (Casino et al., 2019; Lu, 2019b). Subsequent surveys of enterprise implementations emphasize that the value of a shared ledger is realized only when multiple parties with imperfectly aligned interests must agree on one history of events, a condition that insurance, with its insurers, reinsurers, brokers, and claimants, satisfies almost canonically (Lu, 2022; Zheng and Lu, 2022). On the contract layer, research has progressed from conceptual definitions of smart contracts (Christidis and Devetsikiotis, 2016) to architectural treatments of their lifecycle and platforms (Wang et al., 2019) and to empirical studies of the practical difficulties developers face in writing correct, secure contract code (Zou et al., 2021). Security analyses catalog recurring vulnerability classes in deployed contracts, several of which have produced large real-world losses (Atzei et al., 2017). For insurance, where the contract is the product, this body of work implies that code quality is an actuarial variable.

The third foundation is decentralized finance. Analyses of DeFi describe how lending, trading, and asset management have been rebuilt as composable open protocols in which custody, execution, and settlement occur on-chain (Schär, 2021; Chen and Bellavitis, 2020; Xu et al., 2024). Economic theory clarifies both the promise and the limits of this recomposition: decentralized consensus can reduce verification costs and weaken incumbent rents, but it introduces new distortions, and smart contracts enhance contractibility while creating fresh channels for collusion and error (Cong and He, 2019). Legal scholarship adds that the absence of identifiable intermediaries strains the assumptions of financial regulation, which traditionally attaches obligations to licensed entities rather than to autonomous code (Zetsche et al., 2020; Yeoh, 2017). Within this landscape, insurance-like DeFi services emerged to cover the ecosystem's own risks, particularly smart-contract failure, and these discretionary coverage protocols constitute the most complete field experiments in decentralized insurance to date.

At the intersection of these streams, a specifically actuarial literature on decentralized risk sharing has matured. Conditional-mean risk sharing provides a principled rule for allocating a pool's realized losses back to heterogeneous members, with attractive fairness and convex-order properties (Denuit, 2019). A systematic axiomatization of risk-sharing rules establishes which allocation schemes are actuarially defensible for peer-to-peer designs (Denuit et al., 2022), while models of multi-risk mutual aid characterize equilibrium contributions in networks that operate without an insurer's balance sheet (Abdikerimova and Feng, 2022). In parallel, the parametric insurance literature has formalized index design and the management of basis risk, the residual mismatch between an index payout and the true loss (Figueiredo et al., 2018; Lin and Kwon, 2020), with agricultural index insurance providing the longest empirical record of both the promise and the adoption frictions of trigger-based cover (Jensen and Barrett, 2017; Benami et al., 2021). Finally, organizational research on blockchain governance examines how decentralized autonomous organizations allocate decision rights through tokens and reputation, and whether such arrangements constitute a genuinely new mode of organizing collaboration or a re-skinned shareholder firm (Hassan and De Filippi, 2021; Santana and Albareda, 2022; Lumineau et al., 2021; Yermack, 2017).

Table I summarizes how the present paper draws these strands together. Relative to prior surveys of blockchain insurance (Gatteschi et al., 2018; Kar and Navin, 2021; Dominguez Anguiano and Parte, 2023), the distinctive move is analytical: each architecture is treated as a bundle of automation choices with measurable consequences for expense, latency, solvency, and residual trust, rather than as a use case to be described.

Table I. Research streams underpinning intelligent decentralized insurance and their analytical roles.

Research Stream	Representative Findings	Role in This Paper
Digital insurance and AI analytics	Digitalization reshapes the full value chain and the boundary of insurability; AI improves pricing, service, and fraud detection	Baseline cost structure and the intelligence components of the automation plane
Blockchain and smart-contract systems	Shared ledgers coordinate distrusting parties; contract code is powerful but error-prone and hard to patch	Feasibility and security constraints on self-executing policy logic
Decentralized finance and its economics	Open protocols lower verification costs but create new governance and regulatory problems	Institutional template and cautionary evidence for coverage protocols
Actuarial risk-sharing theory	Axiomatic allocation rules and mutual-aid equilibria for pools without an insurer balance sheet	Formal basis for the algorithmic mutual architecture and the stability model
Parametric and index insurance	Trigger design determines basis risk; adoption depends on trust in the index	Foundation of the parametric architecture and the oracle integrity model
DAO and blockchain governance	Token voting redistributes, rather than removes, discretion and power	Analysis of the governance and assurance plane

The synthesis in Table I also exposes the research gap this paper addresses. The actuarial stream presumes reliable data and honest computation; the systems stream presumes that someone has solved the economics of pooling; and the governance stream rarely engages either. An integrated treatment that traces a premium through all four planes of Figure 1, identifying at each step what is automated, what is trusted, and what it costs, is still missing. Providing such a treatment is the purpose of the sections that follow.

III. ANALYTICAL FRAMEWORK AND METHOD

The analysis rests on three complementary lenses. The first is a transaction-cost lens: an insurance arrangement is viewed as a chain of verification, coordination, and enforcement activities, and each architecture is assessed according to which activities are performed by code, by data infrastructure, by collective human judgment, or not at all. Blockchain economics suggests that distributed verification is cheapest where the facts to be verified are natively digital, and most expensive where physical-world states must be attested (Cong and He, 2019). The second is an information-economics lens, focusing on how each design handles adverse selection, moral hazard, and fraud, the classic asymmetries that conventional insurers manage through underwriting and claims investigation, and that automated systems must manage through data, incentives, and algorithmic screening (Eling et al., 2022; Wang and Xu, 2018). The third is a socio-technical trust lens: following research on trust-free systems, decentralization is treated as a relocation of trust, from an institution to a configuration of artifacts and communities, whose weakest element determines the credibility of the whole (Hawlitschek et al., 2018; Lumineau et al., 2021).

Methodologically, the paper combines a structured literature synthesis with stylized quantitative modeling. The synthesis covers peer-reviewed work published between 2016 and 2024 in the information systems, actuarial science, finance, and computer science literatures, identified through keyword searches on decentralized insurance, peer-to-peer insurance, parametric insurance, smart contracts, blockchain oracles, and decentralized autonomous organizations, and screened for analytical rather than purely promotional content. The quantitative component consists of four deliberately simple models: an expense-decomposition comparison across operating models, a settlement-latency comparison across process designs, a normal-approximation model of pool shortfall probability under correlated losses, and a binomial model of oracle trigger integrity under redundancy. The parameter values are stylized: they are chosen to lie within ranges reported in the cited literature and public industry disclosures, and the figures should be read as analytical illustrations of mechanisms rather than as estimates for any particular firm or protocol. This choice reflects the state of the field, in which audited, comparable operating data for decentralized insurance remain scarce, and it keeps the

logic of each argument fully transparent and reproducible.

Three architecture families serve as the units of analysis. An algorithmic mutual pool is a group of participants who contribute to a common fund whose allocation rules are encoded in contract logic, in the spirit of the actuarial peer-to-peer literature (Denuit et al., 2022; Abdikerimova and Feng, 2022). A parametric trigger contract pays a predefined amount when an observable index crosses a threshold, automating settlement end to end (Lin and Kwon, 2020). A token-governed coverage protocol is a standing, open-membership capital pool whose underwriting and claims decisions are made by token-weighted or reputation-weighted voting (Santana and Albareda, 2022). The three families are ideal types; deployed systems mix their elements, but the typology isolates the distinct automation and trust choices each embodies.

IV. AUTOMATED RISK ARCHITECTURES

This section analyzes the three architecture families in turn, attending to how each automates the canonical insurance functions of pooling, pricing, and settlement, and to the asymmetric-information problems each design inherits or creates.

Algorithmic mutual pools recover the oldest organizational form in insurance, the mutual, and rebuild it on programmable infrastructure. Members contribute premiums to a transparent fund; realized losses are allocated back to members according to a rule fixed in advance; and any surplus is returned, donated, or rolled forward rather than retained as profit. The actuarial foundations are now well developed: conditional-mean risk sharing allocates the pool's aggregate loss to members in proportion to their conditional expected contribution to it, and satisfies desirable fairness and risk-reduction properties even for heterogeneous members (Denuit, 2019). The axiomatic treatment of risk-sharing rules identifies which allocation schemes preserve actuarial fairness, monotonicity, and willingness to participate (Denuit et al., 2022), while equilibrium models of mutual aid show how multi-risk networks can price participation without any external capital provider (Abdikerimova and Feng, 2022). What the ledger adds to this mathematics is enforceability and observability: contributions, balances, and allocations are computed by code that every member can inspect, which directly addresses the suspicion of discretionary claim denial that burdens conventional insurers (Hawlitschek et al., 2018). The design also realigns incentives, since unclaimed funds do not become the operator's profit; the residual information problems are peer-level moral hazard and the entry of bad risks, which the architecture must counter with data-driven screening at admission and AI-based anomaly detection at claim time (Riikkinen et al., 2018).

Parametric trigger contracts automate the settlement function most completely. Coverage is written not on the loss itself but on an index correlated with it, such as cumulative rainfall, wind speed at landfall, flight departure delay, or seismic intensity, and the contract pays a predetermined schedule when the index crosses a threshold. Because the payout condition is objectively observable, a smart contract can verify it against an oracle feed and disburse funds within minutes, removing loss adjustment from the process entirely (Lin and Kwon, 2020). The economics of this design are dominated by basis risk: a well-designed index pays when the insured is actually harmed and withholds when they are not, and probabilistic index construction methods exist to quantify and minimize the mismatch (Figueiredo et al., 2018). Decades of experience with agricultural index insurance show both the appeal of the design for thin-margin, hard-to-adjust risks and the adoption frictions that arise when prospective buyers do not understand or do not trust the index (Jensen and Barrett, 2017). Remote sensing and crop modeling have materially improved index quality, tightening the link between trigger and loss and expanding the set of insurable perils (Benami et al., 2021). For the present analysis, the parametric architecture is the cleanest demonstration that automation transforms claims from an adversarial negotiation into a deterministic computation, at the price of relocating all residual uncertainty into index design and oracle integrity.

Token-governed coverage protocols generalize both designs into standing, open-membership underwriting organizations. Capital providers stake value into a shared pool and receive governance tokens; cover buyers pay premiums for protection against specified events, classically the technical failure of a named smart contract;

and claims are adjudicated by member vote rather than by a claims department. Organization theory recognizes these decentralized autonomous organizations as a distinctive attempt to coordinate through algorithmically enforced rules plus token-weighted collective choice (Hassan and De Filippi, 2021; Santana and Albareda, 2022). The governance literature, however, counsels realism. Blockchain governance redistributes discretion rather than eliminating it: proposal power, vote concentration, and off-chain influence re-create familiar principal-agent problems in new clothing (Lumineau et al., 2021; Beck et al., 2018), and analyses of ledger-based corporate governance anticipate both sharper accountability and novel manipulation risks (Yermack, 2017). Claim voting in particular embeds a conflict of interest, since the voters are also the residual claimants on the pool, an arrangement that works only if reputation, staking penalties, and appeal mechanisms make honest adjudication incentive-compatible. The protocol architecture therefore offers the broadest functional scope of the three families at the cost of the heaviest governance burden.

Figure 2 places the three families against the conventional baseline on the two dimensions where automation bites most directly: the operating expense embedded in each premium and the latency of claim settlement. The values are stylized in the sense of Section III, but their ordering and rough magnitudes follow from the mechanism analysis above and from the cost evidence reviewed in the digital-insurance literature (Eling and Lehmann, 2018; Kar and Navin, 2021).

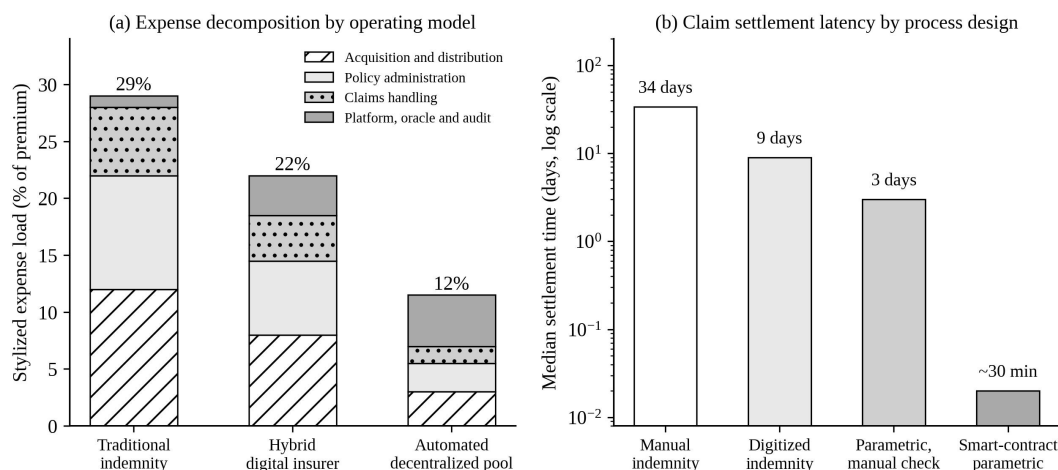


Figure 2. Stylized operating economics of automation: (a) expense decomposition by operating model; (b) claim settlement latency by process design.

Panel (a) of Figure 2 decomposes a stylized expense load for three operating models. A traditional indemnity insurer carrying acquisition, administration, claims handling, and platform costs of roughly 29 percent of premium compares with a digitally optimized hybrid at about 22 percent, in which automation thins administration but distribution and regulatory overhead persists. The automated decentralized pool reaches approximately 12 percent, but the composition matters as much as the total: conventional acquisition and administration nearly vanish, while a new block of platform, oracle, and audit expense appears, since transaction fees, redundant data feeds, and recurring security reviews are the price of removing the intermediary (Zhou et al., 2020; Zou et al., 2021). Automation does not abolish operating cost; it converts labor-intensive coordination cost into capital- and assurance-intensive infrastructure cost. Panel (b) shows the corresponding latency ladder. Manual indemnity settlement measured in weeks contracts to days under digitized workflows, while a parametric contract with automated verification settles in minutes, a five-hundred-fold improvement that changes what insurance is useful for: a payout that arrives during the disruption, rather than after it, functions as liquidity insurance and not merely as eventual reimbursement (Lin and Kwon, 2020).

The comparison also clarifies where each architecture is fragile. The algorithmic mutual depends on

admission screening and peer honesty; the parametric contract depends on index quality and oracle integrity; the coverage protocol depends on governance legitimacy and the security of a large standing treasury. These dependencies are not implementation details but the load-bearing trust assumptions of each design, and they are the subject of the next section.

V. TRUST SYSTEMS: ORACLES, CONTRACT SECURITY, AND GOVERNANCE

If decentralization removed trust, decentralized insurance would be simple. In reality it re-engineers trust into three subsystems: the data layer that tells contracts what happened, the code layer that executes the promise, and the governance layer that handles everything the code did not anticipate. Each subsystem has a distinct failure profile, and the credibility of the whole is bounded by the weakest of the three (Hawlitschek et al., 2018).

The oracle subsystem is the most distinctive. A smart contract cannot observe the world; it can only read values written to the ledger by oracles, the services that attest to off-chain facts such as rainfall, flight status, or asset prices. The oracle problem is that this attestation reintroduces exactly the third-party trust the architecture sought to remove: a deterministic contract fed by a corruptible oracle is a deterministic transmitter of corruption (Caldarelli, 2020). The systems literature responds with decentralized oracle networks, in which multiple independent sources report, outliers are filtered, aggregates are computed, and reporters post collateral that is forfeited for provable misreporting; comparative reviews document the design space of voting-based, reputation-based, and stake-based aggregation and its open weaknesses, including source correlation and data-origin authentication (Al-Breiki et al., 2020). For insurance, oracle design interacts directly with underwriting: an index is only as insurable as its feed is manipulation-resistant, which favors instrument-grade sources such as official meteorological networks and satellite products over scrapeable web data (Benami et al., 2021). The Internet-of-Things extension of the data plane raises the stakes further, since sensor-fed policies inherit the well-documented security limitations of constrained devices, and ledger-anchored device identity and integrity schemes become part of the actuarial perimeter (Lu and Xu, 2019; Xu et al., 2021).

The contract-security subsystem determines whether the encoded promise will execute as intended under adversarial pressure. Taxonomies of smart-contract vulnerabilities, from reentrancy and unchecked external calls to mishandled exceptions and dependence on manipulable ledger state, were established early and remain depressingly current, with several classes implicated in losses of historic size (Atzei et al., 2017). Empirical studies of practitioners confirm the underlying difficulty: contract code is immutable once deployed, testing tools are immature relative to the stakes, and the gap between intended and implemented semantics is hard to close (Zou et al., 2021). For an insurance protocol the implication is stark, because the contract is simultaneously the policy wording, the claims department, and the custodian of the reserve fund. Defense in depth therefore combines independent audits, formal verification of critical invariants, capped exposure during ramp-up, circuit-breaker logic, and, increasingly, AI-assisted code analysis and runtime anomaly monitoring, an application area where the convergence of blockchain and machine intelligence is especially productive (Salah et al., 2019; Zhang and Lu, 2021). Residual code risk can itself be insured, which is precisely the niche in which token-governed coverage protocols first found product-market fit.

The governance subsystem absorbs the irreducible remainder: ambiguous claims, oracle disputes, parameter changes, and crisis response. Research on decentralized autonomous organizations shows that effective designs layer fast algorithmic rules over slower collective-choice procedures, with escalation paths between them (Hassan and De Filippi, 2021; Santana and Albareda, 2022). The blockchain-governance literature identifies the recurring pathologies, namely voter apathy, whale dominance, proposal capture, and the off-chain concentration of effective control in core developer teams (Beck et al., 2018; Lumineau et al., 2021; Yermack, 2017). Institutional economics adds a constructive frame: ledgers, smart contracts, and token votes are institutional technologies, alternative devices for producing the rules of cooperation, and their comparative efficiency depends on the transaction characteristics of the activity being governed (Allen et al., 2020). Claims adjudication, with its mixture of verifiable facts and contestable judgment, is close to the boundary where pure

code governance stops being efficient, which explains the hybrid human-plus-algorithm adjudication observed in every surviving protocol. Table II consolidates the three subsystems, their characteristic failure modes, and the countermeasures the literature supports.

Table II. Trust subsystems of decentralized insurance: failure modes and countermeasures.

Trust Subsystem	Characteristic Failure Modes	Documented Countermeasures	Residual Exposure
Oracle and data layer	Source manipulation; correlated or stale feeds; data-origin spoofing; sensor compromise	Multi-source aggregation with outlier filtering; staked reporting with penalties; instrument-grade indices; device identity anchoring	Correlated source failure during extreme events, when payouts matter most
Smart-contract layer	Reentrancy and call-handling bugs; logic divergent from intent; immutable defects; treasury drain exploits	Independent audits; formal verification of invariants; exposure caps and circuit breakers; AI-assisted code and runtime monitoring	Novel vulnerability classes and economic exploits not covered by past audits
Governance layer	Vote concentration; claim-vote conflicts of interest; proposal capture; emergency-power abuse	Reputation-weighted and quadratic voting; staking penalties for dishonest adjudication; time-locks and appeal tiers; transparency of treasury	Legitimacy crises when code outcomes contradict community expectations

Two cross-cutting observations follow from Table II. First, the countermeasure columns are populated largely by redundancy and collateral, which cost money; this is the assurance expense that reappeared in the decomposition of Figure 2 and confirms that trust in decentralized systems is produced, not free. Second, every residual exposure clusters in the tail: oracles fail together in catastrophes, exploits target the largest treasuries, and governance fractures under stress. Decentralized insurance therefore faces its hardest tests exactly when conventional insurance does, a point the quantitative analysis of the next section makes precise for the pooling function itself.

VI. QUANTITATIVE PERSPECTIVES ON POOL STABILITY AND TRIGGER INTEGRITY

The qualitative analysis identified expense, latency, dependence, and data integrity as the decisive quantities. This section examines the latter two with stylized models whose assumptions are stated fully, so that the mechanisms, rather than the particular numbers, carry the argument.

Consider first the stability of an automated mutual pool. Let each of n members face a loss with mean one unit and standard deviation two units over the coverage period, and let pairwise loss correlation be a constant ρ , capturing the event risk that strikes many members at once. Suppose the pool charges the actuarially fair premium plus a solvency buffer of twenty percent of expected losses, and approximate the aggregate loss as normal. The probability that realized losses exhaust premiums and buffer, the shortfall probability, then depends on n and ρ in the manner shown in panel (a) of Figure 3. With independent losses, pooling works exactly as the law of large numbers promises: shortfall probability falls from roughly one in three for a twenty-member pool to a negligible level beyond about a thousand members. With ρ equal to 0.05, however, the curve flattens near a thirty-three percent floor, and with ρ equal to 0.15 near forty percent, no matter how large the pool grows, because the dispersion of the average loss converges not to zero but to the systematic component that diversification cannot remove. The qualitative lesson is standard actuarial science, but its institutional implication for decentralized designs is sharp: a peer pool without access to external risk transfer is structurally confined to weakly dependent risks, and any ambition to cover catastrophe-correlated perils requires either ledger-native reinsurance layers or buffers far larger than communities will plausibly fund (Denuit et al., 2022; Abdikerimova and Feng, 2022).

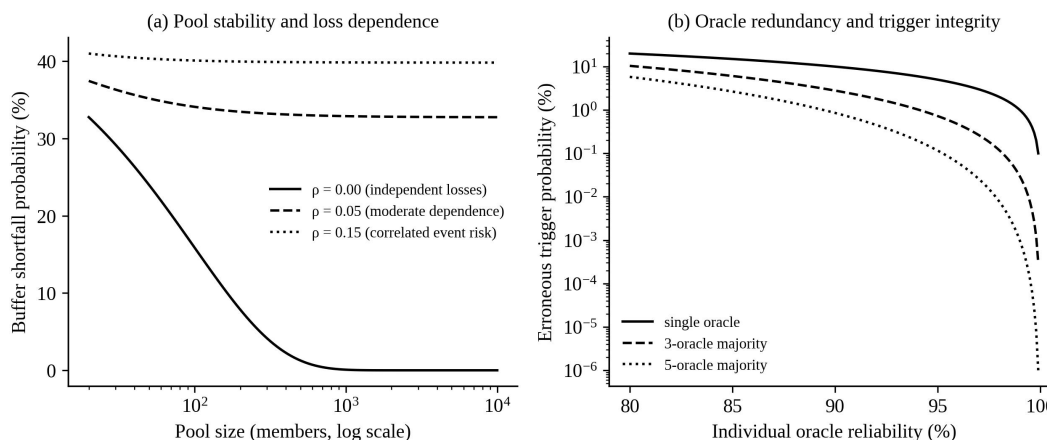


Figure 3. Stylized models of systemic constraints: (a) pool shortfall probability under loss dependence; (b) erroneous trigger probability under oracle redundancy.

Panel (b) of Figure 3 turns to trigger integrity. Model each oracle as reporting the true state of the index with probability r and an erroneous state otherwise, independently across oracles, and let the contract fire on the majority report. A single oracle of ninety-five percent reliability then produces an erroneous trigger, a false payout or a false denial, in five percent of evaluations, an error rate no underwriter could absorb. Three-oracle majority voting compresses the rate to roughly 0.7 percent and five-oracle majority to about 0.1 percent, and the improvement steepens as individual reliability rises, reaching orders-of-magnitude gains at r equal to 99 percent. The binomial arithmetic explains why production oracle networks emphasize both the number and, critically, the independence of sources: the redundancy dividend in panel (b) is computed under independence and evaporates when reporters share an upstream data origin, which is precisely the correlated failure mode flagged in Table II (Al-Breiki et al., 2020; Caldarelli, 2020). Read together, the two panels make one argument: dependence is the common enemy, of solvency in the loss dimension and of integrity in the data dimension, and pricing it is the central actuarial task that decentralization adds to insurance.

Table III consolidates the quantitative perspectives by profiling the three architectures against the conventional baseline on the dimensions developed in Figures 2 and 3. The entries for expense and latency restate the stylized values of Figure 2; the dependence and oracle columns translate the models of Figure 3 into the qualitative exposure each architecture carries.

Table III. Analytical profile of automated risk architectures (stylized values from the models of Section VI).

Dimension	Traditional Indemnity Insurer	Algorithmic Mutual Pool	Parametric Trigger Contract	Token-Governed Coverage Protocol
Stylized expense load	About 29% of premium	About 12-15%; assurance costs replace administration	About 12%; oracle redundancy is the main overhead	About 13-18%; audit and governance costs dominate
Median settlement latency	Weeks (manual adjustment)	Days (automated rules, human appeals)	Minutes (fully automated trigger)	Days (claim vote and time-locks)
Exposure to loss dependence	Managed via reinsurance and capital regulation	High; floor effect of Figure 3(a) binds without reinsurance access	Moderate; event severity capped by parametric schedule	High for correlated technical failures across covered protocols
Oracle and data exposure	Low; internal adjustment substitutes for oracles	Moderate; admission and claim screening need verified data	Decisive; trigger integrity is the product, per Figure 3(b)	High; both cover triggers and governance votes consume feeds
Primary residual	Institution, brand, and	Peer honesty and	Index design and oracle	Token-holder governance

trust	regulator	allocation-rule code	independence	and treasury security
-------	-----------	----------------------	--------------	-----------------------

The profile in Table III resists a single ranking. The parametric architecture wins on latency and expense but stakes everything on the data plane; the mutual wins on incentive alignment but cannot scale into dependent risks alone; the coverage protocol wins on scope but carries the heaviest governance and security burden. Architectural choice in decentralized insurance is thus a mapping from risk characteristics, namely verifiability, dependence, and severity, to the trust subsystem one is best equipped to harden, which is an actuarial decision as much as an engineering one.

VII. STRATEGIC IMPLICATIONS, REGULATION, AND LIMITATIONS

For incumbent insurers, the analysis implies a selective rather than existential response. The expense decomposition shows that the contestable margin lies in administration and claims handling, exactly the functions automation absorbs first, while incumbent advantages, namely proprietary loss histories, regulatory licenses, and reinsurance access, sit in the layers that decentralized designs find hardest to replicate, namely dependence management and institutional legitimacy. The adoption literature accordingly finds insurers experimenting along a gradient: ledger-based reconciliation among known counterparties, then parametric products with automated settlement, and only rarely open-pool participation (Kar and Navin, 2021; Dominguez Anguiano and Parte, 2023). The trajectory mirrors fintech generally, where incumbents internalize the technologies of challengers faster than challengers internalize the institutions of incumbents (Gomber et al., 2018; Lee and Shin, 2018). For decentralized-native ventures, the complementary implication is that survival depends on occupying risks the incumbent structure serves poorly, namely digitally verifiable, weakly dependent, fast-settlement exposures, and on importing institutional credibility through audits, disclosed capital rules, and, where available, regulatory recognition.

Regulation is the decisive external variable. Insurance regulation attaches solvency, conduct, and reporting obligations to licensed legal persons, and a protocol governed by dispersed token holders fits that template badly (Zetsche et al., 2020). Early regulatory analyses of blockchain finance anticipated exactly this mismatch and proposed functional rather than entity-based supervision, regulating the activity of risk transfer wherever it occurs (Yeoh, 2017). The institutional-economics perspective sharpens the stakes: if ledgers and token governance are genuinely new institutional technologies, then premature application of entity-based rules may suppress efficient organizational variety, while regulatory abstention invites the consumer harms insurance law exists to prevent (Allen et al., 2020). Plausible middle paths visible in the literature include licensed wrappers around decentralized pools, supervisory access to on-chain records as a substitute for periodic reporting, and minimum standards for oracle redundancy and contract audit in any product marketed as insurance. The transparency of the ledger is an underexploited regulatory asset: a supervisor that can verify reserves and claim flows continuously needs fewer attestations, not more (Casino et al., 2019).

The analysis carries limitations that bound its claims. The quantitative models are stylized by design; their parameters illustrate mechanisms within literature-reported ranges and do not estimate any deployed system, and the normal and binomial approximations understate tail behavior precisely where decentralized systems are most fragile. The literature synthesized is young and survivorship-biased, since failed protocols publish less than successful ones. The architecture typology abstracts from hybrid designs that mix peer pooling with external capital. And the treatment of regulation is necessarily generic across jurisdictions whose positions are evolving rapidly. These limitations define, rather than diminish, the research agenda of the concluding section.

VIII. CONCLUSION AND FUTURE RESEARCH

This paper has analyzed intelligent decentralized insurance as a re-engineering of both the cost structure and the trust structure of risk transfer. Through a four-plane reference architecture, a comparison of three automated risk families, and stylized models of expense, latency, pool stability, and oracle integrity, the analysis

yields three conclusions. First, automation genuinely compresses the operating wedge between premiums paid and losses returned, but it does so by converting coordination labor into assurance infrastructure, so the relevant question for any design is whether oracle redundancy, audits, and governance cost less than the administration they replace. Second, the law of large numbers is not repealed by decentralization: loss dependence places a floor under pool risk that no community scale can remove, which makes ledger-native reinsurance and dependence-aware admission rules the critical missing layer of the current ecosystem. Third, trust is conserved rather than eliminated; it migrates from institutions to data feeds, code, and collective choice, and the credibility of any decentralized insurance system equals that of its weakest subsystem.

Four research directions follow directly. Actuarial research should extend risk-sharing theory to pools with endogenous membership and on-chain reinsurance layers, pricing the dependence floor identified in Section VI (Denuit et al., 2022). Information-systems research should treat oracle networks as insurable infrastructure, developing reliability metrics and independence tests that underwriting can consume (Al-Breiki et al., 2020; Caldarelli, 2020). Governance research should evaluate, with field data, which adjudication mechanisms sustain honest claim voting under conflict of interest, connecting DAO theory to claims outcomes (Santana and Albareda, 2022; Lumineau et al., 2021). And the convergence of machine intelligence with ledger infrastructure deserves study as a joint system, in which AI components price, monitor, and defend the contracts that in turn discipline the data the AI consumes (Salah et al., 2019; Zhang and Lu, 2021; Lu, 2019a). If these strands mature together, decentralized insurance can move from an ecosystem insuring its own technology toward an institution that widens access to trustworthy risk transfer, which has always been the social purpose of insurance.

REFERENCES

- Abdikerimova, S., & Feng, R. (2022). Peer-to-peer multi-risk insurance and mutual aid. *European Journal of Operational Research*, 299(2), 735-749. <https://doi.org/10.1016/j.ejor.2021.09.017>
- Al-Breiki, H., Rehman, M. H. U., Salah, K., & Svetinovic, D. (2020). Trustworthy blockchain oracles: Review, comparison, and open research challenges. *IEEE Access*, 8, 85675-85685. <https://doi.org/10.1109/ACCESS.2020.2992698>
- Allen, D. W. E., Berg, C., Markey-Towler, B., Novak, M., & Potts, J. (2020). Blockchain and the evolution of institutional technologies: Implications for innovation policy. *Research Policy*, 49(1), 103865. <https://doi.org/10.1016/j.respol.2019.103865>
- Atzei, N., Bartoletti, M., & Cimoli, T. (2017). A survey of attacks on Ethereum smart contracts (SoK). In *Principles of Security and Trust, Lecture Notes in Computer Science*, 10204, 164-186. Springer. https://doi.org/10.1007/978-3-662-54455-6_8
- Beck, R., Mueller-Bloch, C., & King, J. L. (2018). Governance in the blockchain economy: A framework and research agenda. *Journal of the Association for Information Systems*, 19(10), 1020-1034. <https://doi.org/10.17705/1jais.00518>
- Benami, E., Jin, Z., Carter, M. R., Ghosh, A., Hijmans, R. J., Hobbs, A., Kenduiwo, B., & Lobell, D. B. (2021). Uniting remote sensing, crop modelling and economics for agricultural risk management. *Nature Reviews Earth & Environment*, 2(2), 140-159. <https://doi.org/10.1038/s43017-020-00122-y>
- Caldarelli, G. (2020). Understanding the blockchain oracle problem: A call for action. *Information*, 11(11), 509. <https://doi.org/10.3390/info11110509>
- Casino, F., Dasaklis, T. K., & Patsakis, C. (2019). A systematic literature review of blockchain-based applications: Current status, classification and open issues. *Telematics and Informatics*, 36, 55-81. <https://doi.org/10.1016/j.tele.2018.11.006>
- Cevolini, A., & Esposito, E. (2020). From pool to profile: Social consequences of algorithmic prediction in insurance. *Big Data & Society*, 7(2). <https://doi.org/10.1177/2053951720939228>
- Chen, Y., & Bellavitis, C. (2020). Blockchain disruption and decentralized finance: The rise of decentralized business models. *Journal of Business Venturing Insights*, 13, e00151. <https://doi.org/10.1016/j.jbvi.2019.e00151>
- Chen, Y., Lu, Y., Bulysheva, L., & Kataev, M. Y. (2024). Applications of blockchain in Industry 4.0: A review. *Information Systems Frontiers*, 26(5), 1715-1729. <https://doi.org/10.1007/s10796-022-10248-7>
- Christidis, K., & Devetsikiotis, M. (2016). Blockchains and smart contracts for the Internet of Things. *IEEE Access*, 4, 2292-2303. <https://doi.org/10.1109/ACCESS.2016.2566339>
- Cong, L. W., & He, Z. (2019). Blockchain disruption and smart contracts. *Review of Financial Studies*, 32(5), 1754-1797. <https://doi.org/10.1093/rfs/hhz007>
- Denuit, M. (2019). Size-biased transform and conditional mean risk sharing, with application to P2P insurance and tontines. *ASTIN Bulletin*, 49(3), 591-617. <https://doi.org/10.1017/asb.2019.24>
- Denuit, M., Dhaene, J., & Robert, C. Y. (2022). Risk-sharing rules and their properties, with applications to peer-to-peer insurance. *ISSN: 3067-7386 © 2023 INATGI (Institute of Advanced Technology and Green Innovation). Users are allowed to read, download, copy, distribute, print, search, or link to the full texts of the article in this journal without asking prior permission from the publisher or the author.*
- See: <https://inatgi.in/index.php/jaiaa/index> for more information. <https://doi.org/10.63646/jaiaa.2023.010405>

- Journal of Risk and Insurance, 89(3), 615-667. <https://doi.org/10.1111/jori.12385>
- Dominguez Anguiano, T., & Parte, L. (2023). The state of art, opportunities and challenges of blockchain in the insurance industry: A systematic literature review. *Management Review Quarterly*, 74, 1097-1118. <https://doi.org/10.1007/s11301-023-00328-6>
- Eling, M., & Lehmann, M. (2018). The impact of digitalization on the insurance value chain and the insurability of risks. *The Geneva Papers on Risk and Insurance - Issues and Practice*, 43(3), 359-396. <https://doi.org/10.1057/s41288-017-0073-0>
- Eling, M., Nuessle, D., & Staubli, J. (2022). The impact of artificial intelligence along the insurance value chain and on the insurability of risks. *The Geneva Papers on Risk and Insurance - Issues and Practice*, 47(2), 205-241. <https://doi.org/10.1057/s41288-020-00201-7>
- Figueiredo, R., Martina, M. L. V., Stephenson, D. B., & Youngman, B. D. (2018). A probabilistic paradigm for the parametric insurance of natural hazards. *Risk Analysis*, 38(11), 2400-2414. <https://doi.org/10.1111/risa.13122>
- Gatteschi, V., Lamberti, F., Demartini, C., Pranteda, C., & Santamaria, V. (2018). Blockchain and smart contracts for insurance: Is the technology mature enough? *Future Internet*, 10(2), 20. <https://doi.org/10.3390/fi10020020>
- Gomber, P., Kauffman, R. J., Parker, C., & Weber, B. W. (2018). On the fintech revolution: Interpreting the forces of innovation, disruption, and transformation in financial services. *Journal of Management Information Systems*, 35(1), 220-265. <https://doi.org/10.1080/07421222.2018.1440766>
- Hassan, S., & De Filippi, P. (2021). Decentralized autonomous organization. *Internet Policy Review*, 10(2). <https://doi.org/10.14763/2021.2.1556>
- Hawlicschek, F., Notheisen, B., & Teubner, T. (2018). The limits of trust-free systems: A literature review on blockchain technology and trust in the sharing economy. *Electronic Commerce Research and Applications*, 29, 50-63. <https://doi.org/10.1016/j.elerap.2018.03.005>
- Jensen, N., & Barrett, C. (2017). Agricultural index insurance for development. *Applied Economic Perspectives and Policy*, 39(2), 199-219. <https://doi.org/10.1093/aep/ppw022>
- Kar, A. K., & Navin, L. (2021). Diffusion of blockchain in insurance industry: An analysis through the review of academic and trade literature. *Telematics and Informatics*, 58, 101532. <https://doi.org/10.1016/j.tele.2020.101532>
- Lee, I., & Shin, Y. J. (2018). Fintech: Ecosystem, business models, investment decisions, and challenges. *Business Horizons*, 61(1), 35-46. <https://doi.org/10.1016/j.bushor.2017.09.003>
- Lin, X., & Kwon, W. J. (2020). Application of parametric insurance in principle-compliant and innovative ways. *Risk Management and Insurance Review*, 23(2), 121-150. <https://doi.org/10.1111/rmir.12146>
- Lu, Y. (2018). Blockchain and the related issues: A review of current research topics. *Journal of Management Analytics*, 5(4), 231-255. <https://doi.org/10.1080/23270012.2018.1516523>
- Lu, Y. (2019a). Artificial intelligence: A survey on evolution, models, applications and future trends. *Journal of Management Analytics*, 6(1), 1-29. <https://doi.org/10.1080/23270012.2019.1570365>
- Lu, Y. (2019b). The blockchain: State-of-the-art and research challenges. *Journal of Industrial Information Integration*, 15, 80-90. <https://doi.org/10.1016/j.jii.2019.04.002>
- Lu, Y. (2022). Implementing blockchain in information systems: A review. *Enterprise Information Systems*, 16(12), 1876-1907. <https://doi.org/10.1080/17517575.2021.2008513>
- Lu, Y., & Xu, L. D. (2019). Internet of Things (IoT) cybersecurity research: A review of current research topics. *IEEE Internet of Things Journal*, 6(2), 2103-2115. <https://doi.org/10.1109/JIOT.2018.2869847>
- Lumineau, F., Wang, W., & Schilke, O. (2021). Blockchain governance - A new way of organizing collaborations? *Organization Science*, 32(2), 500-521. <https://doi.org/10.1287/orsc.2020.1379>
- Riikkinen, M., Saarijarvi, H., Sarlin, P., & Lahteenmaki, I. (2018). Using artificial intelligence to create value in insurance. *International Journal of Bank Marketing*, 36(6), 1145-1168. <https://doi.org/10.1108/IJBM-01-2017-0015>
- Risius, M., & Spohrer, K. (2017). A blockchain research framework: What we (don't) know, where we go from here, and how we will get there. *Business & Information Systems Engineering*, 59(6), 385-409. <https://doi.org/10.1007/s12599-017-0506-0>
- Salah, K., Rehman, M. H. U., Nizamuddin, N., & Al-Fuqaha, A. (2019). Blockchain for AI: Review and open research challenges. *IEEE Access*, 7, 10127-10149. <https://doi.org/10.1109/ACCESS.2018.2890507>
- Santana, C., & Albareda, L. (2022). Blockchain and the emergence of Decentralized Autonomous Organizations (DAOs): An integrative model and research agenda. *Technological Forecasting and Social Change*, 182, 121806. <https://doi.org/10.1016/j.techfore.2022.121806>
- Schär, F. (2021). Decentralized finance: On blockchain- and smart contract-based financial markets. *Federal Reserve Bank of St. Louis Review*, 103(2), 153-174. <https://doi.org/10.20955/r.103.153-74>
- Wang, S., Ouyang, L., Yuan, Y., Ni, X., Han, X., & Wang, F.-Y. (2019). Blockchain-enabled smart contracts: Architecture, applications, and future trends. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 49(11), 2266-2277.

<https://doi.org/10.1109/TSMC.2019.2895123>

- Wang, Y., & Xu, W. (2018). Leveraging deep learning with LDA-based text analytics to detect automobile insurance fraud. *Decision Support Systems*, 105, 87-95. <https://doi.org/10.1016/j.dss.2017.11.001>
- Xu, L. D., Lu, Y., & Li, L. (2021). Embedding blockchain technology into IoT for security: A survey. *IEEE Internet of Things Journal*, 8(13), 10452-10473. <https://doi.org/10.1109/JIOT.2021.3060508>
- Xu, R., Zhu, J., Yang, L., Lu, Y., & Xu, L. D. (2024). Decentralized finance (DeFi): A paradigm shift in the FinTech. *Enterprise Information Systems*, 18(9), 2397630. <https://doi.org/10.1080/17517575.2024.2397630>
- Yeoh, P. (2017). Regulatory issues in blockchain technology. *Journal of Financial Regulation and Compliance*, 25(2), 196-208. <https://doi.org/10.1108/JFRC-08-2016-0068>
- Yermack, D. (2017). Corporate governance and blockchains. *Review of Finance*, 21(1), 7-31. <https://doi.org/10.1093/rof/rfw074>
- Zetzsche, D. A., Arner, D. W., & Buckley, R. P. (2020). Decentralized finance. *Journal of Financial Regulation*, 6(2), 172-203. <https://doi.org/10.1093/jfr/fjaa010>
- Zhang, C., & Lu, Y. (2021). Study on artificial intelligence: The state of the art and future prospects. *Journal of Industrial Information Integration*, 23, 100224. <https://doi.org/10.1016/j.jii.2021.100224>
- Zheng, X. R., & Lu, Y. (2022). Blockchain technology: Recent research and future trend. *Enterprise Information Systems*, 16(12), 1939895. <https://doi.org/10.1080/17517575.2021.1939895>
- Zheng, Z., Xie, S., Dai, H.-N., Chen, X., & Wang, H. (2018). Blockchain challenges and opportunities: A survey. *International Journal of Web and Grid Services*, 14(4), 352-375. <https://doi.org/10.1504/IJWGS.2018.095647>
- Zhou, Q., Huang, H., Zheng, Z., & Bian, J. (2020). Solutions to scalability of blockchain: A survey. *IEEE Access*, 8, 16440-16455. <https://doi.org/10.1109/ACCESS.2020.2967218>
- Zou, W., Lo, D., Kochhar, P. S., Le, X.-B. D., Xia, X., Feng, Y., Chen, Z., & Xu, B. (2021). Smart contract development: Challenges and opportunities. *IEEE Transactions on Software Engineering*, 47(10), 2084-2106. <https://doi.org/10.1109/TSE.2019.2942301>