Systemic Risk Channels in Digital Finance: An Integrated Taxonomy

Jörg Osterrieder[a,b,*], Lennart John Baals[a], Codruţa Mare[c]

[a] *Department of Behavioural, Management and Social Sciences, University of Twente, Enschede, The Netherlands*

[b] *University of Applied Sciences of the Grisons, Chur, Switzerland*

[c] *Department of Statistics-Forecasts-Mathematics, Faculty of Economics and Business Administration, and Interdisciplinary Centre for Data Science, Babeş-Bolyai University, Cluj-Napoca, Romania*

Abstract

Financial innovation built on blockchain infrastructure has produced a parallel digital finance ecosystem spanning decentralized protocols, centralized crypto platforms, and stablecoin settlement layers. The systemic risk transmission mechanisms embedded in this ecosystem are not captured by frameworks designed for regulated financial institutions with observable balance sheets and quarterly reporting. This paper develops an integrated taxonomy of systemic risk channels in digital finance, classified relative to established theory. Through a hybrid methodology combining systematic literature review, theoretical derivation, and crisis case evidence from twenty-five episodes between 2009 and 2026, we identify eight channels: network contagion, liquidity spirals, stablecoin peg-breaks, composability and smart-contract-cascade risk, liquidation cascades, counterparty concentration, information asymmetry, and gateway-bank risk. Composability and gateway-bank risk are genuinely novel features of programmable finance with no traditional analog; four channels are hybrids that extend established mechanisms with digital-native features including AMM-liquidity procyclicality, MEV-extraction, and bridge-exploit contagion; two are direct extensions of canonical theory. We document thirty-four cross-channel interactions that produce compounding cascade dynamics, illustrated by the 2022 sequence through Terra/Luna, Three Arrows Capital, and FTX. We develop four conjectures on cross-channel coupling, digital-native acceleration, dual-layer opacity, and gateway fragility. The taxonomy provides a mechanism-based framework for fintech researchers, digital-finance practitioners, and macroprudential regulators.

Keywords: [JEL] G01, G18, G23, G29, O33, E58

---

1. Introduction

Financial innovation enabled by blockchain technology, programmable finance, and tokenised settlement has created a parallel digital finance ecosystem whose scale and complexity demand systematic risk assessment. Fintech activity in this space spans decentralized protocol design,

*Corresponding author.

Email addresses: j.r.osterrieder@utwente.nl (Jörg Osterrieder), l.j.baals@utwente.nl (Lennart John Baals), codruta.mare@econ.ubbcluj.ro (Codruţa Mare)

URL: 0000-0002-0872-1080 (Jörg Osterrieder), 0000-0002-4920-7462 (Codruţa Mare)

centralized crypto platform operation, and stablecoin-based payment and settlement infrastructure; each of these subdomains introduces transmission mechanisms absent from the traditional banking system on which existing systemic risk frameworks were calibrated. Between 2020 and 2024, total value locked in decentralized finance (DeFi) protocols surged from under $1 billion to over $150 billion at its peak, centralized cryptocurrency exchanges processed trillions of dollars in annual trading volume, and stablecoins intermediated more than $7 trillion in on-chain transfers annually (International Monetary Fund, 2021; Financial Stability Board, 2023b). The global cryptocurrency market capitalization reached approximately $3 trillion at its November 2021 peak a magnitude that places digital finance among the world's largest asset classes and squarely within the scope of macroprudential concern. The numbers alone justify scrutiny.

The collapse of the Terra/Luna ecosystem in May 2022 destroyed approximately $45 billion in value within five days. The resulting shock then triggered a cascading chain of counterparty failures Three Arrows Capital, Celsius Network, Voyager Digital, and ultimately FTX that collectively inflicted an additional estimated $15–20 billion in direct losses, wiping out retail depositors who had been promised institutional-grade risk management (Financial Stability Board, 2023b). This was not a series of unrelated failures. It was a connected cascade in which each collapse activated multiple transmission channels simultaneously. The distress propagated at blockchain speed, traversing both transparent on-chain interactions and opaque off-chain bilateral lending agreements, and reached the traditional financial system through the banking relationships that mediate fiat-to-crypto conversion.

The systemic risk frameworks developed in response to the 2007–2009 Global Financial Crisis have proven insufficient for this environment. Measures such as CoVaR (Adrian and Brunnermeier, 2016) and Marginal Expected Shortfall (Acharya et al., 2017) were designed for regulated financial institutions with observable balance sheets, quarterly reporting obligations, and centralized clearing infrastructures. Digital finance violates each of these assumptions. DeFi protocols operate through autonomous smart contracts with no central management, no fiduciary obligations, and no identifiable legal entity. Centralized crypto platforms frequently lack audited financial statements and often operate across multiple jurisdictions without consolidated oversight (Financial Stability Board, 2023b). The foundational premise of institution-based systemic risk measurement that the relevant entities can be identified, their exposures observed, and their distress conditioned upon breaks down in an ecosystem built on pseudonymous addresses and permissionless participation. Counterparty exposures span both on-chain interactions and off-chain bilateral agreements that are visible to neither regulators nor other market participants (International Organization of Securities Commissions, 2022). The compression of crisis timelines from weeks to hours, the absence of circuit breakers, and the permissionless nature of participation further distinguish digital finance from the institutional settings for which existing tools were calibrated.

Prior taxonomic efforts cluster in two forms. Policy-oriented frameworks by the Financial Stability Board, the BIS Committee on Payments and Market Infrastructures, and IOSCO organize digital-finance risks by market segment (DeFi, stablecoins, crypto-asset service providers) or by regulatory jurisdiction (Financial Stability Board, 2023b; Bank for International Settlements, 2022; International Organization of Securities Commissions, 2022); these surveys are comprehensive but are not structured around transmission mechanism. Academic frameworks such as that of Aramonte et al. (2022) analyse the DeFi stack in depth but focus on a single subdomain. Our integrated taxonomy differs in its mechanism-based organization across all four subdomains (DeFi, centralized finance (CeFi), stablecoins, tokenised TradFi) and in its explicit mapping of cross-channel interactions, an axis absent from the policy frameworks and underdeveloped in the academic ones.

The academic literature on systemic risk in digital finance has grown rapidly but remains fragmented across subfields that rarely communicate. The field has consistently considered each channel in isolation, independent from the others, with individual studies examining: stablecoin run dynamics and their parallels to traditional bank runs (Uhlig, 2022), DeFi liquidation mechanics and the role of automated execution in cascade amplification (Perez et al., 2021), network contagion patterns in the crypto lending ecosystem (Aramonte et al., 2022), and the opacity of centralized platforms that enabled institutional fraud (Jalan and Matkovskyy, 2023). Each body of work advances understanding of a single mechanism but does not connect to the others in a unified analytical framework. Policy reports from the Financial Stability Board, the Bank for International Settlements, and the International Monetary Fund have catalogued risk factors and proposed regulatory responses (Financial Stability Board, 2022; Bank for International Settlements, 2022), but these reports organize risks by market segment or regulatory jurisdiction rather than by transmission mechanism. The gap between the granularity of individual academic studies and the breadth of policy surveys motivates the taxonomy we develop.

What remains absent from the literature and what motivated this study is an integrated framework that identifies the complete set of transmission channels through which systemic risk propagates in digital finance, classifies each channel relative to established traditional finance theory, and systematically maps the cross-channel interactions that transform individual risk events into system-wide crises. The taxonomy was constructed through a hybrid methodology combining systematic literature review with theoretical derivation and crisis case evidence. A structured search through OpenAlex identified an initial candidate set of fourteen potential channels across approximately 2,400 unique works. A composite scoring framework incorporating literature volume, citation impact, and crisis evidence narrowed the set to eight channels meeting thresholds for theoretical depth, empirical documentation, and cross-domain relevance. This paper fills the identified gap.

We identify eight distinct channels: (1) network contagion and interconnectedness, (2) liquidity spirals and fire sales, (3) stablecoin de-pegging and run dynamics, (4) composability and smart contract cascade risk, (5) leverage and liquidation cascades, (6) counterparty and concentration risk, (7) information asymmetry and opacity, and (8) fiat-crypto gateway and banking channel risk. For each channel we specify the transmission mechanism, theoretical foundations, digital-native features, and crisis evidence. The taxonomy spans all four major domains of digital finance: DeFi, CeFi, stablecoins, and the emerging tokenized traditional finance sector.

Two of the eight channels composability risk and gateway risk are genuinely novel, with no close analog in traditional finance. Composability risk arises from the permissionless stacking of DeFi protocols that creates invisible dependency chains propagating failures at blockchain speed; no traditional financial system allows arbitrary, unauthorized composition of institutional balance sheets. Gateway risk describes the systemic fragility introduced by the small number of banking institutions that mediate between traditional and digital finance; this channel exists only at the boundary between two financial systems and has no precedent in either system alone. Two channels network contagion and counterparty concentration are direct extensions of canonical theories (Allen and Gale, 2000; Duffie, 2010). The remaining four liquidity spirals, stablecoin runs, liquidation cascades, and information asymmetry are hybrids rooted in established theory but with digital-native features that substantially transform their dynamics (Brunnermeier and Pedersen, 2009; Diamond and Dybvig, 1983; Akerlof, 1970).

The hybrid channels are not merely traditional mechanisms operating in a new setting; each incorporates structural features that change the mechanism qualitatively. Liquidity spirals in DeFi involve automated market makers whose bonding curves create procyclical price dynamics absent from traditional order-book markets, and liquidation bots that compete to seize collateral through gas price auctions rather than through negotiated margin calls. The margin call, in effect, has been automated away. Stablecoin runs operate without deposit insurance or lender-of-last-resort facilities and, in the algorithmic case, involve reflexive token dynamics in which the act of redemption mechanically destroys the backing asset. Liquidation cascades are executed by permissionless bots that extract maximal value from distressed positions through transaction reordering a phenomenon known as maximal extractable value (MEV) that lacks a direct traditional analog (Daian et al., 2020). Information asymmetry in CeFi combines the traditional opacity of unregulated intermediaries with the novel transparency of on-chain activity, creating a hybrid information environment in which sophisticated actors can observe what retail participants cannot.

The cross-channel interactions documented in this paper are not hypothetical. The 2022 crisis sequence activated all eight channels in a specific temporal ordering: the Terra/UST stablecoin run triggered liquidity spirals in DeFi protocols, which activated liquidation cascades, which propagated through network contagion to Three Arrows Capital and other leveraged funds, which exposed counterparty concentration in CeFi lending markets, which revealed the information asymmetry underlying Celsius and Voyager, which culminated in the FTX collapse that activated gateway risk through the failure of Silvergate and Signature banks (Briola et al., 2023; Liu et al., 2023). Understanding these interactions requires a taxonomy that identifies each channel individually and then maps their coupling dynamics the dual contribution that this paper provides.

Reconstructing this sequence reveals a structural property of the cascade: once the first failure propagated, each subsequent collapse activated channels that had accumulated exposure observable only in retrospect. Understanding this amplification architecture, not any single failure, is what an analytical framework for digital finance must supply.

This paper provides the first integrated taxonomy of digital-finance systemic risk channels organized by transmission mechanism. As Brunnermeier and Pedersen (2009) first distinguished loss spirals, margin spirals, and predatory trading as separate mechanisms before modeling their equilibrium interactions, classification must logically precede formal modeling of individual channel dynamics. We make three principal contributions. First, the eight channels we identify are classified as novel, hybrid, or extension relative to traditional finance theory. The classification provides a structured map of where existing theory applies, where it must be extended, and where genuinely new theory is needed. Second, we document the cross-channel interactions and amplification dynamics that produce compounding cascade effects, reconstructing the temporal structure of the 2022 crisis sequence to show how all eight channels activated in a specific ordering. Third, we identify specific literature gaps within each channel, providing a structured research agenda for formal modeling of systemic risk in digital finance.

The remainder of this paper proceeds as follows. Section 2 describes the literature search protocol, channel identification process, and selection criteria. Section 3 presents the eight-channel taxonomy, devoting a subsection to each channel. Section 4 analyzes cross-channel interactions and amplification dynamics. Section 5 draws policy implications. Section 6 concludes with a summary of contributions, limitations, and a research agenda.

2. Methodology

The taxonomy was constructed through a hybrid methodology combining systematic literature review, theoretical derivation, and crisis case evidence. The approach proceeds in two stages: a structured search to identify candidate channels and assess their scholarly coverage, followed by a two-stage selection process combining quantitative composite scoring with qualitative assessment of mechanism distinctiveness and theoretical foundation.

2.1. Literature Search and Channel Identification

The literature search was conducted using OpenAlex, an open-access index covering more than 250 million scholarly works, supplemented by targeted searches of policy report repositories maintained by the Bank for International Settlements, the Financial Stability Board, the International Monetary Fund, the International Organization of Securities Commissions, and the European Central Bank [e.g., International Monetary Fund (2023)]. OpenAlex was selected for its broad coverage of working papers, preprints, and policy documents critical to a rapidly evolving field in which much relevant research appears outside traditional journal outlets.

The search began with the identification of fourteen candidate systemic risk channels derived from three complementary sources: the established systemic risk literature in traditional finance (Allen and Gale, 2000; Diamond and Dybvig, 1983; Brunnermeier and Pedersen, 2009), policy reports cataloguing digital finance risks (Financial Stability Board, 2022; International Monetary Fund, 2021), and analysis of twenty-five major crisis episodes spanning 2014–2025 across DeFi, CeFi, stablecoin, and traditional banking domains, documented in our crisis chronology (Appendix B). The fourteen candidates comprised the eight channels that ultimately survived selection network contagion, liquidity spirals, stablecoin runs, composability risk, liquidation cascades, counterparty concentration, information asymmetry, and gateway risk plus six additional candidates (oracle manipulation, regulatory contagion, governance failure, real-world asset transmission, cross-chain bridge vulnerability, and validator concentration) that were ultimately merged into surviving channels or deferred. For each candidate, three to four search queries using Boolean combinations of channel-specific terms and systemic risk vocabulary were executed, restricted to works published between 2009 and 2026.

The initial search retrieved approximately 2,400 unique works across all fourteen channels after deduplication. Title and abstract screening retained approximately 620 works (a retention rate of approximately 25%), applying inclusion criteria requiring each work to address systemic risk or contagion in digital finance and to present original analysis theoretical, empirical, or regulatory [e.g., Xu and Vadgama (2023)]. Each retained work was coded against the fourteen candidate channels, with many works assigned to multiple channels.

2.2. Composite Scoring and Channel Selection

The transition from fourteen candidates to eight selected channels was governed by a two-stage process combining quantitative composite scoring with qualitative assessment of mechanism distinctiveness and theoretical foundation. The composite scoring was applied to the full set of approximately 2,400 works retrieved from OpenAlex, prior to manual relevance screening, to provide a preliminary quantitative ranking. Each candidate was scored on three dimensions:

(i) **Literature volume** (weight: 0.35): the number of works in the channel cluster, normalized to a 0–1 scale.

(ii) **Citation impact** (weight: 0.35): the mean citation count of the ten most-cited works in the channel cluster, normalized against the highest-scoring cluster.

(iii) **Crisis evidence** (weight: 0.30): the number of distinct crisis episodes activating the channel, weighted by the logarithm of estimated losses associated with each episode.

The citation impact dimension is subject to noise from the keyword-based search: some retrieved works are topically unrelated to their assigned channel, inflating citation scores for channels with broad query terms. A relevance-filtered robustness check (available from the authors) confirms that all qualitative retention decisions remain stable under filtered scoring.

The composite score served as a structured screening input rather than an algorithmic cutoff. Channel selection proceeded in two stages. In the first stage, the composite scores provided a quantitative ranking of the fourteen candidates according to scholarly coverage and empirical grounding (Table 1). In the second stage, each candidate was assessed qualitatively against three criteria: (i) whether the channel represents a distinct transmission mechanism or is a subcase of another channel, (ii) whether established theoretical foundations exist to sustain detailed analysis, and (iii) whether the channel activated in documented crisis episodes through a mechanism not already captured by a higher-scoring channel.

The ranking is robust to alternative weighting schemes, as shown in Table 2. This two-stage process produced retention decisions that deviate from a simple rank-order cutoff in both directions. Two high-scoring candidates cross-chain bridge vulnerability (rank 1, composite 0.71) and validator concentration (rank 6, composite 0.56) were merged into parent channels because they represent specific instances of broader mechanisms already captured in the taxonomy. Bridge vulnerability is a subcase of composability risk (protocol interdependence); validator concentration is a subcase of counterparty concentration (infrastructure-layer concentration). Conversely, two lower-scoring candidates network contagion (rank 9, composite 0.38) and stablecoin runs (rank 11, composite 0.30) were retained despite their lower composite scores because each possesses foundational theoretical grounding and a distinct transmission mechanism: network contagion is anchored in the interbank contagion framework of Allen and Gale (2000), while stablecoin runs extend the bank run model of Diamond and Dybvig (1983). Excluding either would leave a theoretically important transmission pathway unanalyzed.

The remaining excluded candidates were handled as follows. Oracle manipulation was merged into composability risk as a specific instance of permissionless protocol interdependence. Governance failure was split between composability risk and counterparty concentration. Regulatory contagion was distributed between gateway risk and counterparty concentration. Real-world asset transmission was deferred to the forward-looking assessment given insufficient scale and crisis evidence (Bank for International Settlements, 2024). Splits applied two criteria: (i) the candidate channel encompassed mechanisms with materially different transmission pathways, and (ii) at least one such pathway was already covered by a higher-scoring channel. The merging and splitting decision rules are formalized in Appendix A.

We acknowledge that the boundary between retained and excluded channels involves judgment. The scoring framework structures that judgment by providing quantitative evidence on scholarly coverage and crisis activation; the qualitative stage ensures that the final taxonomy reflects mechanism distinctiveness and theoretical coherence rather than a mechanical score cutoff.

To assess the robustness of these decisions, we recomputed composite scores under three alternative weighting schemes (Table 2). Under all three schemes, the same qualitative retention decisions hold: high-scoring candidates merged as subcases (bridge vulnerability, validator

Table 1: Composite scoring results for fourteen candidate systemic risk channels. Scores are normalized to $[0, 1]$. The Retention Decision column explains the qualitative rationale applied in the second stage of channel selection. Weights: literature volume 0.35, citation impact 0.35, crisis evidence 0.30.

Rank	**Channel**	**Lit.**	**Cit.**	**Crisis**	**Comp.**	**Retention Decision**
1	Bridge Vulnerability	0.57	1.00	0.52	0.71	Merged into Composability Risk (subcase of protocol interdependence)
2	Liquidity Spirals	0.76	0.31	1.00	0.67	Retained
3	Gateway Risk	1.00	0.33	0.68	0.67	Retained
4	Composability Risk	0.57	0.58	0.74	0.63	Retained
5	Counterparty Conc.	0.57	0.21	1.00	0.58	Retained
6	Validator Conc.	0.57	0.96	0.08	0.56	Merged into Counterparty Conc. (infrastructure-layer concentration)
7	Information Asymmetry	0.57	0.18	0.62	0.45	Retained
8	Liquidation Cascades	0.57	0.22	0.34	0.38	Retained
9	Network Contagion	0.57	0.16	0.41	0.38	Retained foundational theory (Allen and Gale, 2000)
10	Governance Failure	0.57	0.23	0.21	0.35	Split: Composability Risk / Counterparty Conc.
11	Stablecoin Runs	0.63	0.03	0.25	0.30	Retained foundational theory (Diamond and Dybvig, 1983)
12	Oracle Manipulation	0.60	0.13	0.13	0.29	Merged into Composability Risk (protocol interdependence)
13	Regulatory Contagion	0.57	0.18	0.07	0.29	Distributed: Gateway Risk / Counterparty Conc.
14	RWA Transmission	0.57	0.22	0.00	0.28	Deferred insufficient crisis evidence

concentration) remain subcases regardless of weights, and lower-scoring candidates retained for theoretical reasons (network contagion, stablecoin runs) continue to possess the foundational grounding that justified their inclusion.

Table 2: Sensitivity of composite scores to alternative weighting schemes. The retention decision column indicates whether the qualitative decision changes under any scheme. No channel changes status.

Channel	**Primary (35/35/30)**	**Equal (33/33/33)**	**Crisis-dom. (25/25/50)**	**Lit.-dom. (50/25/25)**	**Decision Stable?**
Bridge Vuln.	0.71	0.70	0.65	0.67	Yes (merged)
Liquidity Spirals	0.67	0.69	0.76	0.70	Yes
Gateway Risk	0.67	0.67	0.67	0.75	Yes
Composability Risk	0.63	0.63	0.66	0.62	Yes
Counterparty Conc.	0.58	0.60	0.70	0.59	Yes
Validator Conc.	0.56	0.54	0.42	0.55	Yes (merged)
Info. Asymmetry	0.45	0.46	0.50	0.49	Yes
Liquidation Cascades	0.38	0.38	0.37	0.43	Yes
Network Contagion	0.38	0.38	0.39	0.43	Yes (theory)
Stablecoin Runs	0.30	0.30	0.28	0.38	Yes (theory)

The classification dimensions introduced in Section 1 (novel, hybrid, extension) are applied to each surviving channel, and each channel is mapped against the four domains (DeFi, CeFi, Stablecoins, Tokenized TradFi) to identify its cross-domain footprint (Brunnermeier and Pedersen, 2009; Diamond and Dybvig, 1983; Akerlof, 1970).

2.3. Conceptual Conjectures

The taxonomy generates four conjectures with testable implications. We deliberately label these as conjectures rather than propositions or hypotheses: each is grounded in the qualitative synthesis developed in this paper and in documented crisis evidence, but formal empirical validation, which would require structural modelling and a substantially enlarged sample of crisis episodes beyond the twenty-five documented here, is left to future work. The appendix (§Appendix A) documents the decision rules underpinning the classification on which these conjectures rest.

Conjecture 1 (Cross-Channel Amplification).. When two or more channels activate simultaneously, the aggregate systemic impact exceeds the sum of individual channel effects. Crisis episodes activating three or more channels should exhibit losses disproportionately larger than single-channel events of comparable initial shock magnitude.

Conjecture 2 (Digital-Native Acceleration).. The structural features of digital finance, automated execution, 24/7 operation, and permissionless participation, compress the timeline of systemic risk propagation relative to traditional finance by at least an order of magnitude for channels operating through on-chain mechanisms (Briola et al., 2023).

Conjecture 3 (Dual-Layer Opacity).. The coexistence of transparent on-chain and opaque off-chain network layers creates information asymmetry that amplifies contagion beyond what either a fully transparent or fully opaque system would produce, because informed agents exploit partial observability to front-run contagion (Jalan and Matkovskyy, 2023).

Conjecture 4 (Gateway Fragility).. The systemic fragility introduced by fiat-crypto gateway concentration increases with gateway bank market share, and the failure of a gateway bank serving a disproportionate share of the industry can trigger cascading effects across both digital and traditional financial systems (Galati and Capalbo, 2023).

3. A Taxonomy of Systemic Risk Channels

The taxonomy identifies eight channels through which systemic risk transmits in digital finance. Each channel is analyzed through a consistent framework: a formal definition, the transmission mechanism and its theoretical foundations, key crisis evidence, and the digital-native features that distinguish the channel from its traditional counterpart. Table 3 provides a consolidated summary, with rows grouped by classification: extensions of canonical theory first, hybrids next, genuinely novel channels last. The detailed subsections that follow proceed in the same theory-proximity ordering, beginning with those closest to established theory.

3.1. Network Contagion and Interconnectedness

Network contagion refers to the transmission of financial distress through direct and indirect linkages among entities in the digital finance ecosystem, encompassing on-chain transaction graphs, shared liquidity pool participation, cross-protocol token dependencies, and bilateral CeFi lending relationships.

The theoretical foundation is well established: Allen and Gale (2000) modeled contagion through interbank claims, Gai and Kapadia (2010) demonstrated fat-tailed loss distributions in arbitrary network topologies, and Acemoglu et al. (2015) proved a phase-transition property where connectivity shifts from diversifying to amplifying above a critical threshold. Scale-free degree distributions first characterized by Barabási and Albert (1999) are now observed in DeFi transaction networks, and Battiston et al. (2012) operationalized the resulting fragility through the DebtRank measure.

The Three Arrows Capital (3AC) insolvency of June 2022 is the canonical case. The hedge fund's bilateral OTC positions with more than twenty counterparties were invisible until default, when aggregate losses exceeded \$3 billion (Jalan and Matkovskyy, 2023). The FTX collapse demonstrated contagion across both CeFi and DeFi layers simultaneously, with social media compressing the informational channel to approximately 72 hours (Cookson et al., 2023). The digital-native dimension is threefold: on-chain observability lets sophisticated actors front-run contagion; pseudonymous participation obscures the true exposure network; and composability creates mechanical contagion edges encoded in smart contract logic, where a failure at one protocol automatically impairs another without human intervention (Schär, 2021; Daian et al., 2020).

Literature gap. Existing contagion models assume identifiable counterparties and observable exposure matrices; they do not handle pseudonymous addresses or the bilateral OTC opacity that defines crypto lending networks.

3.2. Liquidity Spirals and Fire Sales

Liquidity spirals occur when declining asset prices trigger margin calls or collateral liquidations, forcing asset sales that further depress prices and tighten funding conditions in a self-reinforcing feedback loop amplified in digital finance by automated market makers and hard-coded liquidation thresholds.

Table 3: Summary of the eight systemic risk channels in digital finance.

Channel	Class.	Definition	Theory Base	Key Evidence	Digital-Native Feature
Network Contagion	Ext.	Transmission of distress through direct/indirect linkages among entities	Allen & Gale (2000); Acemoglu et al. (2015)	3AC cascade ($3B+ losses, 20+ creditors)	Dual-layer on-chain/off-chain topology
Counterparty Concentration	Ext.	Disproportionate share of critical functions concentrated in few entities	Duffie (2010); too-big-to-fail literature	FTX collapse ($8.7B customer losses)	Multi-function commingling; no resolution framework
Liquidity Spirals	Hybrid	Self-reinforcing feedback between declining prices, margin calls, and forced sales	Brunnermeier & Pedersen (2009)	Black Thursday 2020 ($8.3M MakerDAO bad debt)	AMM bonding curves create procyclical pricing
Stablecoin Runs	Hybrid	Coordinated redemption causing loss of peg, with systemic spillovers	Diamond & Dybvig (1983)	Terra/UST collapse ($45B destroyed)	Reflexive algorithmic death spiral; no deposit insurance
Liquidation Cascades	Hybrid	Automated collateral seizure triggers chain reaction of further liquidations	Brunnermeier & Pedersen (2009); margin literature	Terra/Anchor cascade; Curve near-miss 2023	Permissionless bots; MEV extraction; 12-second execution
Information Asymmetry	Hybrid	Gap between insider knowledge and depositor/counterparty observability	Akerlof (1970); Stiglitz & Weiss (1981)	FTX fraud; Celsius/Voyager; QuadrigaCX	Hybrid transparency-opacity across DeFi/CeFi
Composability Risk	Novel	Permissionless protocol nesting creates dependency chains propagating failures	None (genuinely novel)	Euler Finance 2023 ($197M); bridge hacks ($1.7B+)	Unbounded, permissionless dependency depth
Gateway Risk	Novel	Fragility from concentrated institutions bridging traditional and digital finance	None (genuinely novel)	Silvergate, SVB, Signature failures (Mar 2023)	Bidirectional contagion; asymmetric speed

The well-known Brunnermeier and Pedersen mechanism loss spirals, margin spirals, and fire-sale externalities operates with particular intensity in DeFi, where leverage ratios adjust automatically rather than through discretionary risk management (Brunnermeier and Pedersen, 2009; Shleifer and Vishny, 2011). Lending protocols define each borrowing position by a loan-to-value ratio embedded directly in on-chain code. When a collateral asset's price breaches the liquidation threshold, any external party can seize the collateral and sell it on AMMs, whose constant-product pricing functions mechanically push prices lower in proportion to sale size (Klages-Mundt and Minca, 2022).

Black Thursday, March 12, 2020, provides the canonical case. The COVID-19 panic triggered a roughly 50% single-day decline in Bitcoin. On MakerDAO, thousands of positions fell below their 150% threshold, but extreme network congestion prevented liquidation bots from bidding, producing zero-bid auctions and $8.3 million in protocol bad debt (Klages-Mundt and Minca, 2022). The Terra/Luna collapse of May 2022 exhibited spiral dynamics at vastly larger scale: as UST de-pegged, the reflexive mint-and-burn mechanism created selling pressure on LUNA, destroying approximately $45 billion in market capitalization within five days (Briola et al., 2023).

Four digital-native features distinguish this channel: on-chain liquidation eliminates the time buffer of traditional margin calls (twelve seconds on Ethereum versus days); AMMs provide continuous but procyclical liquidity with deterministic price impact; liquidity providers can withdraw in real time, reducing depth during stress; and on-chain transparency allows all participants to anticipate and front-run cascades (Daian et al., 2020; Perez et al., 2021).

Literature gap. Brunnermeier and Pedersen's funding-and-market liquidity framework predates AMM bonding curves and MEV-extraction; a formal model integrating constant-product price impact with priority-gas-auction dynamics is not yet available.

3.3. Stablecoin De-Pegging and Run Dynamics

The destruction of $45 billion in five days during the Terra/UST collapse of May 2022 was not a black swan it was the predictable outcome of a mechanism that Diamond and Dybvig (1983) would have recognized instantly. What made it unprecedented was the speed and the reflexive architecture that turned redemption itself into the weapon.

Stablecoin runs are coordinated redemption events in which holders rapidly convert their stablecoin holdings to other assets, causing the stablecoin to lose its peg, with systemic consequences arising from stablecoins' central role as the settlement medium across digital finance.

For reserve-backed stablecoins such as Tether (USDT) and USD Coin (USDC), the run mechanism parallels the classic bank run analyzed by Diamond and Dybvig (1983): a shock to confidence triggers self-fulfilling redemption waves. In a striking historical parallel, Gorton and Zhang (2023) connect stablecoin fragility to the nineteenth-century era of wildcat banking, while Ma et al. (2025) show that concentrated arbitrage structures create a tradeoff between secondary-market stability and run vulnerability. For algorithmic stablecoins, the mechanism is reflexive: in the Terra/Luna architecture, UST redemptions created selling pressure on LUNA, which depressed LUNA's price, reducing UST's backing and incentivizing further redemptions in a death spiral with no natural floor (Uhlig, 2022; Liu et al., 2023). This reflexive dynamic introduces a class of run dynamics with no traditional analog.

The systemic significance arises from the settlement-layer function: stablecoins denominate the majority of DeFi lending, AMM liquidity, and derivatives positions (Makarov and Schoar, 2022). The Terra/UST collapse of May 2022 destroyed approximately $45 billion within five days and triggered the sequential failures of 3AC, Celsius, Voyager, and FTX (Liu et al., 2023). The SVB/USDC de-peg of March 2023 demonstrated that reserve-backed stablecoins face run risk originating in the traditional banking system: USDC de-pegged to $0.87 when Silicon Valley Bank failed, cascading through DeFi protocols using USDC as collateral (Galati and Capalbo, 2023; Gregory et al., 2024).

Digital-native features include instantaneous on-chain redemption without queuing or regulatory intervention; a global pseudonymous holder base rendering coordination to prevent a run impractical; reflexive algorithmic mechanisms that become death-spiral channels under stress; and composability that transmits a de-peg to every protocol using the stablecoin as collateral (Kitzler et al., 2022; Lyons and Viswanath-Natraj, 2023).

Literature gap. Diamond-Dybvig captures the classical run; the reflexive death-spiral mechanism in algorithmic stablecoins, in which redemption itself destroys the backing asset, lacks a closed-form equilibrium treatment.

3.4. Composability and Smart Contract Cascade Risk

Composability risk arises from the permissionless combination and nesting of DeFi protocols, which creates chains of dependency that propagate failures across the ecosystem at blockchain speed, without possibility of human intervention or coordinated circuit-breaking.

The composability of DeFi protocols allows any smart contract to invoke any other without authorization, enabling arbitrarily complex financial products by stacking building blocks "money legos" (Schär, 2021; Kitzler et al., 2022). When an upstream protocol suffers an exploit, downstream protocols holding its tokens experience automatic balance-sheet degradation without assessment, negotiation, or forbearance. Propagation speed is limited only by block times (approximately twelve seconds on Ethereum) (Gudgeon et al., 2020). The dependency graph assembles permissionlessly, meaning no single entity has visibility into the complete set of dependencies.

This channel is genuinely novel: no traditional financial system allows arbitrary, unauthorized composition of institutional balance sheets. Traditional securitization chains involve three to five layers governed by documented legal agreements subject to rating-agency scrutiny. In DeFi, an anonymous participant can assemble an equivalent dependency chain in minutes, extend it to arbitrary depth, and govern it solely by smart contract logic whose correctness relies on informal code auditing. Consider a yield aggregator that deposits into a lending protocol that accepts a liquidity pool token as collateral a three-layer dependency chain assembled in minutes by an anonymous deployer. When the bottom layer suffers an exploit, the impairment propagates upward automatically, with no human at any layer aware the dependency existed until losses materialize. The absence of gatekeepers at any layer of the stack represents the fundamental architectural difference that makes composability risk irreducible to existing contagion frameworks (Schär, 2020). As Tolmach et al. (2021) observe, the closest analogs network contagion models assume stochastic propagation, while composability failures propagate deterministically through smart contract logic. The dependency graph is dynamic: yield aggregators and strategy vaults continuously reallocate capital based on return optimization, meaning dependencies change constantly and cannot be captured in a static snapshot.

The cross-chain dimension deserves separate emphasis. Bridge protocols which lock assets on one chain and mint synthetic representations on another create dependency links across otherwise independent blockchains. When a bridge is exploited, the synthetic tokens on the destination chain become unbacked instantaneously, but the market discovers this only as the exploit is confirmed. The Ronin bridge hack ($625M), Wormhole ($326M), Nomad ($190M), and Poly Network ($611M) together produced aggregate losses exceeding $1.7 billion in 2021–2022, establishing cross-chain composability as the highest-magnitude attack surface in DeFi (Xu and Vadgama, 2023). The Euler Finance exploit of March 2023 ($197 million) demonstrated a different facet: a single-protocol vulnerability propagating across the dependency graph to protocols whose own code was entirely secure (Xu and Vadgama, 2023).

Composability risk is the most underappreciated channel in the taxonomy: it has generated the largest aggregate losses through bridge exploits yet receives the least regulatory attention. Existing regulatory frameworks, designed for entities with identifiable counterparties and documented contractual relationships, have no mechanism for governing permissionless dependency chains whose participants may be anonymous and whose structure changes with every block. Until regulators develop tools to monitor dependency depth and cross-chain exposure in real time, composability risk will remain the ecosystem's most dangerous blind spot.

Literature gap. Composability has no traditional-finance analog and consequently no inherited theory; deterministic propagation through smart-contract-encoded dependency graphs requires a new modeling apparatus that contagion frameworks built on stochastic interbank claims do not supply.

3.5. Leverage and Liquidation Cascades

Where liquidity spirals operate through the price-funding feedback loop described above, liquidation cascades work through a different mechanism: not the gradual tightening of funding conditions, but the abrupt chain reaction of automated collateral seizure (Klages-Mundt and Minca, 2022; Perez et al., 2021). The distinction matters because the policy responses differ circuit breakers can interrupt a spiral, but they cannot prevent a cascade that executes within a single block.

When collateral prices decline past the liquidation threshold, any participant can seize collateral at a protocol-specified discount and sell it on AMMs, pushing additional positions past their thresholds. Correlated collateral holdings across protocols mean a single price decline simultaneously threatens positions across multiple venues. What distinguishes this channel is not the economics the margin-call literature is well established but the elimination of all human decision points. The threshold is hard-coded and cannot be modified without a governance vote requiring days. Autonomous bots execute with no discretion or forbearance, and the entire chain from threshold breach to market impact operates within approximately twelve seconds (Daian et al., 2020; Li et al., 2023). The margin call, in effect, has been automated away replaced by an algorithm that cannot exercise judgment and does not know the meaning of forbearance. MEV extraction amplifies the damage: front-running bots detect pending liquidation transactions in the mempool and insert trades that profit from anticipated price movements, while sandwich attacks extract additional value at the borrower's expense (Qin et al., 2022). The game-theoretic competition among liquidation bots, engaging in priority gas auctions, adds strategic interaction absent from traditional fire-sale models.

The Terra/Luna collapse triggered cascading liquidations on Anchor Protocol that amplified the death spiral, while the Curve Finance near-miss of July 2023 illustrated how a single participant's $168 million of CRV collateral distributed across Aave, Fraxlend, and Abracadabra could create synchronized cross-protocol cascade risk.

Literature gap. Margin-call models assume discretionary lender behaviour and human-mediated workouts; the strategic interaction among permissionless liquidation bots competing through priority-gas auctions, with MEV extraction superimposed, is not captured by existing fire-sale theory.

3.6. Counterparty and Concentration Risk

Counterparty concentration risk arises when a small number of entities intermediate a disproportionate share of transaction volume, custodial assets, or critical infrastructure services, such that the failure of any single entity can trigger system-wide disruption.

Despite the decentralization ethos, the ecosystem exhibits extreme concentration across exchanges, liquidity providers, oracle services, and stablecoin issuers (Bank for International Settlements, 2021). The theoretical foundation draws on Babus (2016), who predict that financial networks evolve toward concentrated structures, and Duffie (2010), who formalize slow-moving capital dynamics when intermediary capital is destroyed. The critical extension for digital finance is multi-layer concentration: FTX simultaneously served as exchange, market maker, OTC lender, and custodian, so its failure removed multiple infrastructure services at once and no institutional resolution mechanism existed to manage the fallout (Financial Stability Board, 2023b).

The Mt. Gox collapse (2014, handling approximately 70% of global Bitcoin volume) established the archetype. The FTX collapse of November 2022 proved that a decade of market evolution had not resolved the problem: within 48 hours, $8.7 billion in customer fund commingling was revealed, propagating losses to Genesis, BlockFi, and dozens of portfolio companies (Jalan and Matkovskyy, 2023). Digital-native features unconstrained growth absent licensing requirements, prohibited-in-TradFi function commingling, and regulatory attrition narrowing the gateway set ensure that concentration risk in this ecosystem is structurally self-reinforcing rather than a transient market failure.

Literature gap. Existing too-big-to-fail and SIFI frameworks were designed for entities operating under prudential supervision with separate-function requirements; multi-function commingling at the platform level (FTX simultaneously exchange, market-maker, OTC lender, and custodian) has no equivalent in regulated finance and lacks formal treatment.

3.7. Information Asymmetry and Opacity Risk

The digital finance ecosystem presents a striking paradox: it simultaneously offers more transparency than any traditional market every DeFi transaction permanently recorded on a public ledger and less transparency than any regulated exchange, with CeFi platforms operating without audited financials, regulatory filings, or obligations to disclose related-party transactions (Schär, 2021). This hybrid is genuinely disorienting from a regulatory perspective, and it is this paradox, rather than opacity alone, that makes the information channel distinctive.

The mechanism operates through three components grounded in classical theory (Akerlof, 1970; Stiglitz and Weiss, 1981): insiders produce private information about a platform's true condition; disclosure mechanisms fail to convey that information; and eventual revelation triggers confidence shocks and run dynamics. Voluntary proof-of-reserves attestations which show assets but not liabilities and can be manipulated by temporarily borrowing assets exemplify disclosure failure. The race-to-the-bottom dynamic explains why pre-crisis CeFi yields of 8–18% attracted deposits that would not have been made under full information. Transaction fee markets create additional information asymmetries that enable what Daian et al. (2020) term maximal extractable value (MEV) extraction at the expense of less-informed participants.

The QuadrigaCX collapse ($190 million, 2019) established the canonical case for information asymmetry in crypto custody (Ontario Securities Commission, 2020). The Celsius and Voyager failures of mid-2022 demonstrated it at systemic scale, with Celsius revealing $1.2 billion in undisclosed losses (Jalan and Matkovskyy, 2023). The FTX collapse combined information asymmetry with outright fraud: customer funds transferred to Alameda Research for proprietary trading, losses concealed through fabricated accounting, and FTT used as undisclosed collateral total opacity shielded by reputational capital (Jalan and Matkovskyy, 2023).

Literature gap. Akerlof-Stiglitz models the equilibrium of asymmetric information in markets with separable signals; the dual-layer regime in which on-chain transparency coexists with off-chain CeFi opacity has no canonical model and the inferred-from-MEV component has not been integrated into the run-dynamic literature.

3.8. Fiat-Crypto Gateway and Banking Channel Risk

What happens when an entire financial ecosystem depends on three banks for access to the dollar? In March 2023, the crypto industry found out.

Gateway risk refers to the systemic fragility introduced by the small number of institutions that mediate between the traditional financial system and digital finance, creating bidirectional contagion channels through which traditional bank failures can destabilize crypto markets and vice versa.

Fiat currency enters and exits crypto markets through a narrow set of gateway institutions: banks maintaining accounts for exchanges, stablecoin issuers whose reserves reside in traditional banks, regulated broker-dealers, and payment processors (Gorton and Zhang, 2023). In the crypto-to-traditional direction, a major crypto firm failure imposes direct losses on the gateway bank, potentially triggering a run. In the traditional-to-crypto direction, a gateway bank failure freezes stablecoin reserves, halts fiat settlement, and severs on-ramp infrastructure (Diop et al., 2024).

This channel is genuinely novel: it exists only at the boundary between two financial systems that were, until recently, largely separate. The closest analog is the correspondent banking literature analyzing concentrated cross-border payment functions, but even that literature assumes both sides operate under compatible regulatory regimes. Gateway institutions operate under a fundamental tension: they are regulated by traditional banking authorities yet serve an ecosystem that often lacks equivalent regulation on the other side of the bridge. This asymmetry creates a fragility that neither banking regulators nor crypto-native risk frameworks are equipped to address.

The political economy dimension compounds the structural vulnerability. Regulatory actions whether explicit (FDIC letters discouraging crypto banking relationships) or implicit (reputational pressure on banks serving crypto firms) progressively narrow the set of willing gateway institutions. Each bank that exits increases the concentration among those that remain, a self-reinforcing dynamic that Schuler et al. (2024) term regulatory attrition. By early 2023, the entire U.S. crypto industry's fiat settlement infrastructure depended on three banks: Silvergate, Silicon Valley Bank, and Signature. The fragility this created was not hypothetical for long.

The Silvergate Bank failure of March 2023 illustrates the crypto-to-traditional direction: following FTX's collapse, Silvergate experienced $8.1 billion in deposit outflows, forcing securities sales at a loss and voluntary liquidation. The SVB/USDC de-peg demonstrates the traditional-to-crypto direction: SVB's failure revealed $3.3 billion of USDC reserves at the bank, causing USDC to de-peg to $0.87 and cascading through DeFi protocols within hours (Diop et al., 2024; Galati and Capalbo, 2023). The Signature Bank closure two days later completed the triptych, eliminating both major real-time fiat settlement platforms for crypto within a single month

(Cookson et al., 2023). The asymmetric speed of contagion is particularly noteworthy: crypto-to-traditional operates at banking speed (days to weeks), while traditional-to-crypto operates at blockchain speed (hours), driven by automated stablecoin de-pegging and DeFi liquidations. No existing model of banking contagion accounts for this speed asymmetry, which makes gateway risk not merely a concentration problem but a genuinely new transmission mechanism requiring its own theoretical treatment.

Literature gap. Correspondent-banking and cross-border-payment literature treats both sides as symmetric regulated entities; bidirectional contagion at the boundary between regulated banks and an asymmetrically regulated crypto ecosystem, with order-of-magnitude speed asymmetry, has no formal counterpart and is the most under-modelled channel in the taxonomy.

4. Cross-Channel Interactions and Cascade Dynamics

Despite the fact that the field literature has consistently considered each channel in isolation, independent from the others, no major crisis in digital finance has ever activated a single risk channel that way. The eight channels interact, compound, and amplify losses far beyond what any one channel would generate independently (Acemoglu et al., 2015). The most consequential feature of the taxonomy is therefore the coupling structure: the system's true fragility lies not in any individual channel but in their interactions.

4.1. Interaction Taxonomy

We distinguish three interaction types. *Amplifying interactions* occur when activation of one channel increases the severity of another: a liquidity spiral deepens a liquidation cascade, which further deepens the spiral. *Dampening interactions* occur when one channel's activation reduces the intensity of another: rapid deleveraging from a stablecoin run may reduce available leverage for subsequent cascades. *Sequential interactions* occur when one channel's resolution triggers a different channel with a time lag: the FTX collapse activated counterparty concentration immediately, but gateway risk at Silvergate materialized weeks later (Shleifer and Vishny, 2011).

Amplifying interactions dominate the empirical record. Of the 28 unique undirected channel pairs, we identify 18 amplifying, 4 dampening, and 12 sequential interactions. Six channel pairs exhibit more than one interaction type for example, the liquidity spirals/liquidation cascades pair is both amplifying (during acute stress) and sequential (post-crisis deleveraging triggers delayed liquidations) yielding thirty-four typed interactions across the twenty-eight pairs. Amplification outnumbers dampening by more than four to one. The system-level risk is therefore superadditive: aggregate risk exceeds the sum of individual channel risks.

4.2. The Three Strongest Feedback Loops

The tightest loop operates between **liquidity spirals and liquidation cascades**. A price decline triggers automated liquidations; bots seize collateral and sell it on AMMs, depressing prices further along the bonding curve; lower prices push additional positions past thresholds, triggering the next wave. The interaction is multiplicative: the speed of automated liquidation feeds the spiral's intensity, and the spiral's depth determines the next liquidation wave's breadth. Both Black Thursday 2020 and the Terra/Luna collapse exhibited this loop as their core amplification mechanism (Klages-Mundt and Minca, 2022; Qin et al., 2022; Daian et al., 2020).

The second loop connects **counterparty concentration and network contagion**. When a failing entity occupies a central network position by virtue of exchange volume, market-making activity, and venture investments the contagion radius scales dramatically with centrality. FTX's simultaneous roles as exchange, market maker (via Alameda), and venture investor meant its failure activated contagion across exchange, OTC lending, and venture capital channels simultaneously (Jalan and Matkovskyy, 2023). As Acemoglu et al. (2015) demonstrate, highly connected nodes can flip a system from shock-absorbing to shock-amplifying.

The third loop connects **stablecoin runs and gateway risk**, operating at the boundary between digital and traditional finance. When a gateway bank holding stablecoin reserves fails, the stablecoin de-pegs and propagates through DeFi and CeFi via stablecoins' settlement-layer function. Conversely, a stablecoin run forcing rapid reserve liquidation can impair the bank's balance sheet. The SVB/USDC episode activated both directions within a single week: SVB's failure caused USDC to de-peg to $0.87, and the resulting uncertainty contributed to the run on Signature Bank (Gorton and Zhang, 2023; Galati and Capalbo, 2023).

4.3. The 2022 Cascade Sequence

The 2022 crisis sequence spanning Terra/Luna (May), Three Arrows Capital (June), Celsius and Voyager (June–July), and FTX (November) activated all eight channels in temporal succession and provides the most comprehensive empirical test of cross-channel dynamics. The Terra/UST stablecoin run triggered liquidity spirals across DeFi as the Luna Foundation Guard sold $3.5 billion in Bitcoin reserves. Liquidation cascades on Anchor amplified the death spiral. From there, the contagion propagated through network, counterparty, and information channels in rapid succession each failure revealing the next layer of hidden exposure before culminating in the FTX collapse that activated gateway risk through Silvergate's deterioration (Liu et al., 2023; Briola et al., 2023). For the hundreds of thousands of retail users whose funds were frozen at Celsius, Voyager, and FTX, the distinction between channels was academic they experienced a single, compounding loss of trust in the entire ecosystem.

The escalating severity Black Thursday (three channels, $8 million protocol losses), Terra/Luna (five channels, $45 billion destroyed), FTX (four channels, $8.7 billion customer losses) provides suggestive evidence of a *cascade pattern*. Across the twenty-five documented episodes, channel count and loss severity display an approximately monotonic association, with every episode that breached three simultaneous channels producing losses at least an order of magnitude larger than those that did not. We flag this pattern as a candidate regularity that formal empirical testing on an enlarged sample should validate, rather than as an estimated threshold.

Figure 1 visualizes this escalation, mapping each major crisis episode to the specific channels it activated.

4.4. Interaction Matrix

Table 4 presents the full 8×8 undirected interaction matrix. Coupling strength is rated as strong $(+++)$, moderate $(++)$, or weak $(+)$ based on crisis evidence.

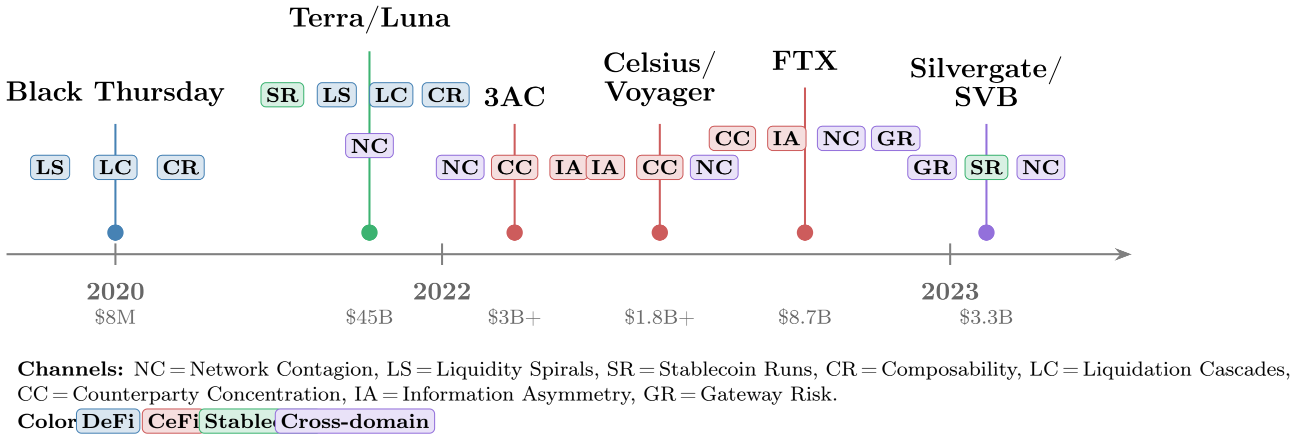

Figure 1: Crisis timeline showing channel activation patterns, 2020–2023. Each event maps to the systemic risk channels activated during the episode. Loss figures represent direct documented losses. Badge colors indicate the primary domain of each channel's activation in that episode.

Table 4: Channel interaction matrix. NC = Network Contagion, LS = Liquidity Spirals, SR = Stablecoin Runs, CR = Composability Risk, LC = Liquidation Cascades, CC = Counterparty Concentration, IA = Information Asymmetry, GR = Gateway Risk.

	NC	LS	SR	CR	LC	CC	IA	GR
NC		++	+	+	+	+++	++	++
LS	++		++	++	+++	+		+
SR	+	++		++		+	+	+++
CR	+	++	++		+++	++		
LC	+	+++		+++		+		
CC	+++	+	+	++	+		+++	++
IA	++		+			+++		+
GR	++	+	+++			++	+	

Legend. +++ Strong coupling: activation of one channel reliably triggers or materially amplifies the other; documented in multiple crisis episodes. ++ Moderate coupling: well-documented conditional linkage; amplification occurs under stress but is not automatic. + Weak coupling: plausible indirect or episodic connection. Blank: negligible or no meaningful interaction pathway.

5. Policy Implications

The taxonomy reveals an uneven match between the channel structure of digital finance systemic risk and existing regulatory frameworks. Channels most effectively addressed are those with traditional-finance analogs: the FSB's 2023 recommendations target counterparty concentration and information asymmetry through licensing, asset segregation, and disclosure requirements; the Bank of England's systemic stablecoin framework targets run dynamics through reserve backing and wind-down planning (Financial Stability Board, 2023b; Bank of England, 2023; Financial Stability Board, 2023c). These initiatives address the channels that dominated the 2022 crisis sequence but leave significant gaps in channels structurally embedded in DeFi architecture. Some channels are well addressed. Others are barely visible to regulators. The human cost of these gaps was borne disproportionately by retail participants individual depositors at Celsius and FTX who had no access to the counterparty information that would have revealed their exposure.

Three gaps stand out. First, composability risk the most digitally native channel is largely unaddressed: no current proposal anywhere in the world requires protocols to map, disclose, or limit their dependency exposures, and DeFi lending protocols set their own liquidation parameters without regulatory input or stress-testing requirements (International Organization of Securities Commissions, 2022). Second, gateway risk operates at the boundary between two regulatory perimeters; neither crypto-specific regulators nor banking supervisors fully capture the bidirectional contagion dynamics, as the SVB/USDC and Silvergate episodes made concrete (International Monetary Fund and Financial Stability Board, 2023). Third, liquidation cascade parameters thresholds, collateral ratios, liquidation bonuses are set by protocol governance without framework for calibration against historical tail scenarios.

Because systemic risk exceeds the sum of individual channels (Section 4), channel-specific regulation is necessary but insufficient. The regulatory architecture must also address cross-channel interactions. For regulators and policymakers reading this taxonomy, we identify four priorities that follow directly from the cross-channel analysis:

- **Cross-channel stress testing.** The taxonomy provides the scenario architecture: a Terra-type scenario tests stablecoin run, liquidity spiral, and liquidation cascade channels simultaneously; an FTX-type scenario tests counterparty concentration, information asymmetry, and network contagion (Financial Stability Board, 2023a).
- **Gateway bank supervision.** Coordinated monitoring across crypto regulators and banking supervisors to track concentration of crypto deposits at individual banks, stablecoin reserve custody arrangements, and fiat settlement volumes through gateway institutions (International Monetary Fund, 2023).
- **Stablecoin reserve transparency.** Moving beyond periodic audits to real-time reserve attestation that captures both assets and liabilities, with mandatory disclosure of custodial arrangements and settlement timeframes. Algorithmic designs warrant either prohibition or substantially higher regulatory scrutiny (Gorton and Zhang, 2023).
- **Composability risk monitoring.** Developing on-chain analytics infrastructure to map protocol dependency graphs in real time, identify concentrated dependency exposures, and provide early warning of cascading failure pathways leveraging the full transparency of DeFi data that is available in principle but underexploited in practice (Bank for International Settlements, 2021).

The monitoring challenge is heterogeneous across channels. Liquidity spirals, liquidation cascades, and composability risk operate primarily on-chain and are fully observable through blockchain analytics; the challenge is analytical capacity rather than data availability. Counterparty concentration, information asymmetry, and network contagion in the CeFi layer involve off-chain exposures that blockchain analytics cannot capture, requiring mandatory reporting analogous to trade repositories for OTC derivatives (Financial Stability Board, 2023b). Gateway risk demands monitoring that spans both perimeters information that exists but is fragmented across multiple agencies and jurisdictions. Building a consolidated monitoring framework integrating on-chain analytics, off-chain regulatory reporting, and cross-perimeter gateway data into unified systemic risk surveillance is the most urgent institutional design priority (International Monetary Fund, 2023; Bank for International Settlements, 2021).

6. Conclusion

This paper develops an integrated taxonomy of systemic risk transmission channels in digital finance, identifying eight channels organized by transmission mechanism and classified relative to traditional finance theory. The taxonomy makes three contributions. First, the identification and classification of the channel set: two genuinely novel channels (composability risk and gateway-bank risk), four hybrids (liquidity spirals, stablecoin peg-breaks, liquidation cascades, information asymmetry), and two extensions (network contagion, counterparty concentration), provides a structured map of where existing theory applies, where it must be extended, and where genuinely new theory is needed. Second, the cross-channel interaction analysis documents thirty-four pairwise interactions, identifies five strongly coupled channel pairs, and surfaces a suggestive cascade pattern above which the system transitions from shock-absorbing to shock-amplifying, a regularity with direct implications for regulatory design that warrants formal empirical testing. Third, the channel-by-channel identification of literature gaps provides a structured research agenda for formal modeling of systemic risk in digital finance (Allen and Gale, 2000; Acemoglu et al., 2015; Diamond and Dybvig, 1983). Each contribution builds on the previous one.

Limitations

The analysis is subject to several limitations. The composite scoring methodology involves weighting choices that, while robust to sensitivity analysis, reflect analytical judgment. The scoring pipeline operates on the full OpenAlex retrieval rather than the manually screened subset, meaning citation impact scores are influenced by off-topic papers from the keyword-based search; a relevance-filtered robustness check confirms the qualitative selection is invariant. The breadth of the taxonomy necessarily limits the depth devoted to any single channel; subsequent focused studies building on this foundation can develop formal models and empirical tests for individual channels. The taxonomy is grounded in crisis evidence through early 2025; emerging channels such as tokenized real-world asset transmission and AI-driven algorithmic herding may alter the channel structure as the ecosystem matures. The crisis evidence dimension treats episodes uniformly across shock types; channel-specific responses to DeFi-only versus CeFi-only versus banking-channel shocks would require a stratified analysis beyond the current chronology.

In addition, the data collection pipeline introduces three structural features that influence the scoring results. The per-channel retrieval limit of 200 works, combined with sorting by citation count, means that the literature volume dimension reflects only the most-cited subset of the literature for each channel. For the ten channels where the retrieval limit was binding, the literature volume sub-score is identical (0.57), reducing the discriminating power of this dimension. Additionally, the citation-count sorting systematically favors older, established work over recent publications with fewer accumulated citations, creating a temporal bias in the literature sample. Finally, one candidate channel real-world asset transmission has zero crisis events in the chronology database, resulting in a permanent zero score on the crisis evidence dimension (30% of the composite weight). This structural disadvantage reflects the genuine absence of documented crises in tokenized RWA markets as of the analysis cutoff, but it means the scoring framework cannot distinguish between channels that are low-risk and channels where risk has simply not yet materialized.

The OpenAlex API does not guarantee stable result ordering for queries with tied citation counts, and result sets may shift between executions due to database updates. While the search parameters and queries are fully documented and the scoring code is deterministic given fixed inputs, exact reproducibility of the literature corpus depends on the API state at the time of retrieval. We mitigate this by archiving the complete retrieved corpus as a static dataset.

Reporting standards.. The methodology combines a systematic literature retrieval with theoretical derivation and crisis-case evidence. It does not claim PRISMA-2020 compliance: there is no registered protocol, no PRISMA flow diagram, no formal inter-rater reliability assessment, and no per-study risk-of-bias evaluation. We deliberately frame the design as a hybrid synthesis appropriate to a taxonomy contribution rather than as a meta-analysis. Readers seeking strict PRISMA compliance should treat this paper as a structured narrative review with explicit retrieval and scoring procedures, and consult the limitations enumerated above for the specific deviations from PRISMA reporting standards.

Finally, the crisis evidence dimension is influenced by event multiplicity: the Terra/Luna collapse activates five channels simultaneously, contributing disproportionate weight to those channels. The log-loss weighting partially mitigates this by compressing the loss range, but cross-channel crisis activation remains an inherent feature of interconnected financial events rather than a methodological artifact. Similarly, the literature volume counts papers in all channels to which they were assigned, meaning that broadly relevant works contribute to multiple channels' volume scores. This multi-channel assignment is intentional it reflects the genuine cross-cutting nature of systemic risk scholarship but it means that channel volume scores are not independent.

Five directions for future research emerge. First, the development of formal theoretical models for composability risk and gateway risk the two novel channels that currently lack any formal treatment would fill the most significant gaps. A model of cascading failure in permissionlessly composed systems would provide the theoretical foundation for stress-testing DeFi protocols; a model of bidirectional gateway contagion would inform institutional design for the fiat-crypto boundary (Xu and Vadgama, 2023; Gorton and Zhang, 2023). Second, the construction of quantitative systemic risk measures calibrated to digital finance extending CoVaR and SRISK or developing new measures tailored to on-chain data would enable real-time monitoring (Qin et al., 2022). Third, two emerging technological developments will reshape the channel structure: the growth of tokenized real-world assets, which creates novel contagion pathways between DeFi composability risk and traditional asset markets through tokenized collateral; the proliferation of CBDCs, which introduces interoperability risks extending the gateway and stablecoin run channels; and the deployment of AI-driven trading strategies, which may produce algorithmic herding that amplifies liquidity spirals through correlated machine-generated behavior (Bank for International Settlements, 2024; International Monetary Fund, 2024). Fourth, empirical validation of the four conjectures, particularly the cross-channel amplification and cascade pattern regularities, through on-chain data analysis and cross-crisis comparison would test the taxonomy's core claims. Most urgently, the design of macroprudential tools specifically adapted to the multi-channel, cross-domain architecture documented here constitutes the critical next step for policy research (Financial Stability Board, 2023a; International Monetary Fund, 2023).

The central insight of this taxonomy is that digital finance both creates genuinely novel systemic risk channels and accelerates the dynamics of channels inherited from traditional finance. The 2022 crisis sequence made this visible. As digital finance continues to mature and its connections to the traditional system deepen through tokenization and institutional adoption, the need for a rigorous, mechanism-based framework for systemic risk assessment will only intensify. The taxonomy we offer here is a first step; the formal models, empirical tests, and regulatory tools it calls for are now urgent.

Appendix A. Decision Rules and Classification Protocol

Because channel classification was executed primarily by the first author, with review and validation by the co-authors, we document here the decision rules and inclusion-exclusion criteria applied at each stage of the taxonomy's construction. The aim is to make the coding decisions reproducible in principle, so that subsequent work can verify or challenge each judgement.

Channel identification (fourteen candidate set)

A candidate transmission channel was retained for scoring if it satisfied all three of the following: (i) it appears in the peer-reviewed systemic risk literature in traditional finance, in the digital finance literature, or in a policy report from the FSB, BIS, IOSCO, IMF, or ECB; (ii) it can be described by a transmission mechanism with identifiable shock source, propagation medium, and affected agents; (iii) it has at least one documented activation in the crisis chronology (§Appendix B). Mechanisms that failed (iii) were retained as candidates only if their theoretical foundations were strong enough to justify inclusion for completeness; the real-world-asset transmission channel was the sole such case and is deferred to the forward-looking research agenda.

Retention vs merging (eight final channels)

A candidate was retained as a stand-alone channel if it had a distinct transmission mechanism not already captured by a higher-scoring retained channel, even if its composite score was below the rank-order cutoff. Bridge vulnerability and validator concentration were merged into parent channels because, although their composite scores were high, their mechanisms are specific instances of composability risk (protocol interdependence) and counterparty concentration (infrastructure-layer concentration), respectively. Oracle manipulation was merged into composability risk for the same reason. Governance failure was split between composability risk (when failure propagates through protocol interdependence) and counterparty concentration (when a single multi-sig key-holder or DAO treasury controls outsized resources). Regulatory contagion was distributed between gateway-bank risk (bank-side regulatory withdrawal) and counterparty concentration (crypto-firm-side regulatory enforcement). These merging rules are applied consistently across all episodes in the chronology.

Classification as novel, hybrid, or extension

A channel was classified as *novel* if no traditional-finance mechanism shares both its shock source and its propagation medium. Composability risk qualifies because permissionless protocol nesting has no traditional analog in the securitisation or correspondent-banking literatures. Gateway-bank risk qualifies because bidirectional contagion between regulated banks and an unregulated crypto ecosystem has no pre-2014 precedent. A channel was classified as *extension* if an established traditional-finance model captures its dynamics without substantial modification; network contagion (Allen and Gale, 2000) and counterparty concentration (Duffie, 2010) meet this bar. The remaining four channels were classified as *hybrid*: each is grounded in an established theoretical framework (Brunnermeier–Pedersen 2009 for liquidity spirals and liquidation cascades, Diamond–Dybvig 1983 for stablecoin peg-breaks, Akerlof 1970 and Stiglitz–Weiss 1981 for information asymmetry), but the digital-native implementation introduces features (automated market makers, reflexive token designs, MEV-extraction, hybrid on-chain/off-chain transparency) that require adaptation of the base model.

Episode-to-channel coding rules

For each of the twenty-five episodes in the chronology (§Appendix B), channels were coded as activated if at least one of the following obtains: (i) peer-reviewed or policy-report analysis of the episode explicitly attributes losses to the channel; (ii) contemporaneous on-chain or accounting evidence shows the channel's transmission mechanism was operative; (iii) losses in the episode cannot be accounted for without invoking the channel's mechanism. Ambiguous cases were coded conservatively: where two channels could plausibly explain the same loss, the channel with the strongest independent documentation was recorded, and the alternative was noted as an interaction partner rather than an independent activation. This conservatism biases the channel-count distribution downward and makes the cross-channel cascade pattern (Section 4.3) a lower bound on the true coupling structure.

Borderline decisions explicitly flagged

Three episodes involved judgement calls worth naming. (i) The SushiSwap vampire attack (2020-09) was coded as governance failure plus composability risk, but a reasonable alternative coding would include liquidity spirals; we excluded the latter because no asset price trajectory was reported in the documented analyses. (ii) The Ronin bridge hack (2022-03) was coded as a bridge-vulnerability instance of composability risk rather than as a counterparty concentration event, despite Axie Infinity's high bridge market share, because the proximate loss mechanism was smart-contract exploit rather than single-entity failure. (iii) The SVB/USDC episode (2023-03) activates both stablecoin peg-break and gateway-bank risk; we coded both, which slightly inflates the channel count relative to a more parsimonious coding.

Future work with multiple fully independent coders can revisit these decisions; the documentation here is the qualitative analogue of an inter-rater reliability procedure, and we regard it as the minimum defensible standard for a taxonomy led by a primary coder with co-author validation in an emerging literature.

Appendix B. Crisis Chronology

Table B.5 lists the twenty-five documented crisis episodes used in the composite scoring and cross-channel analysis. Channel abbreviations: NC Network Contagion; LS Liquidity Spirals; SR Stablecoin Runs; CR Composability Risk (including bridge, oracle, and governance subcases); LC Liquidation Cascades; CC Counterparty Concentration (including validator subcase); IA Information Asymmetry; GR Gateway Risk (including regulatory-contagion subcase); GF Governance Failure (split between CR and CC per the appendix classification protocol).

Declarations

Availability of data and materials

The literature search metadata retrieved from OpenAlex, the composite scoring data, the channel-classification coding, and the crisis chronology supporting the analyses are deposited at Zenodo (DOI to be inserted at acceptance stage; currently available from the corresponding author at j.r.osterrieder@utwente.nl prior to deposit confirmation). The Python source code used to execute the literature search, compute composite scores, and audit the bibliography is available at the corresponding author's institutional repository and will be archived at Zenodo alongside the data.

Table B.5: Documented systemic-risk episodes in digital finance, 2014-2025. Losses in USD where quantifiable.

#	Date	Episode	Domain	Channels	Losses
1	2014-02	Mt. Gox Collapse	CeFi	CC, GR, IA	$460M
2	2016-06	The DAO Hack	DeFi	CR, GF	$60M
3	2016-08	Bitfinex Hack	CeFi	CC, IA	$72M
4	2017-07	BTC-e Seizure and Shutdown	CeFi	CC, GR	n.d.
5	2019-02	QuadrigaCX Collapse	CeFi	CC, GR, IA	$190M
6	2020-03	Black Thursday (COVID crash)	DeFi/CeFi	GR, LC, LS, NC	n.d.
7	2020-09	SushiSwap Vampire Attack	DeFi	CR, GF, LS	n.d.
8	2021-06	Iron Finance / TITAN Collapse	DeFi	CR, LS, SR	$2.0B
9	2021-08	Poly Network Hack	DeFi	CR	$611M
10	2022-02	Wormhole Bridge Hack	DeFi	CR, LS	$326M
11	2022-03	Ronin Bridge Hack	DeFi/CeFi	CC, CR	$625M
12	2022-05	Terra/Luna Collapse	Stablecoins/DeFi	CC, CR, LS, NC, SR	$45.0B
13	2022-06	Three Arrows Capital Insolvency	CeFi	CC, IA, LS, NC	$3.5B
14	2022-06/07	Celsius and Voyager Failures	CeFi	CC, GR, IA, LS	$5.4B
15	2022-08	Nomad Bridge Hack	DeFi	CR	$190M
16	2022-10	Mango Markets Exploit	DeFi	CR, LC	$114M
17	2022-11	FTX Collapse	CeFi	CC, GR, IA, LS, NC	$8.7B
18	2023-03	Euler Finance Exploit	DeFi	CR, LC, LS	$197M
19	2023-03	SVB Failure and USDC De-Peg	Stablecoins/TradFi	GR, LS, NC, SR	n.d.
20	2023-07	Curve Finance Pool Exploit	DeFi	CR, LC, LS	$62M
21	2023-07	Multichain Bridge Collapse	DeFi	CC, CR, GF	$126M
22	2024-05	DMM Bitcoin Exchange Hack	CeFi	CC, GR, IA	$305M
23	2024-07	WazirX Multi-Sig Exploit	CeFi	CC, CR, GR	$235M
24	2025-02	Bybit $1.5B Hack	CeFi	CC, CR, IA, LS	$1.5B
25	2025-03	Hyperliquid Whale Manipulation Events	DeFi	CR, LC, LS	$15M

Notes. "n.d." indicates losses not directly quantifiable (market-wide drawdowns or contagion effects not attributable to a single entity). Channel codings follow the decision rules documented in Appendix A. Governance failure (GF) episodes are retained as a distinct code in this chronology for traceability even though the main taxonomy splits GF between Composability Risk and Counterparty Concentration.

Competing interests

The authors declare that they have no competing interests.

Funding

This document is based upon work from COST Action CA19130, supported by COST (European Cooperation in Science and Technology). COST is a funding agency for research and innovation networks. Their Actions help connect research initiatives across Europe and enable scientists to grow their ideas by sharing them with their peers. This boosts their research, career, and innovation. The collaboration with the COST Action CA21163 Text functional and other high-dimensional data in econometrics: New models, methods, applications is acknowledged.

Financial support by the Swiss National Science Foundation within the project Mathematics and Fintech - the next revolution in the digital transformation of the Finance industry (IZCNZ0-174853) is gratefully acknowledged. The authors are also grateful for financial support from the Swiss National Science Foundation under the grants IZSEZ0-211195 (Anomaly and Fraud Detection in Blockchain Networks), 205487, and 190703, and within the project Narrative Digital Finance: a tale of structural breaks, bubbles & market narratives (IZCOZ0-213370). The authors acknowledge funding from the European Union's Horizon 2020 research and innovation program FIN-TECH: A Financial supervision and Technology compliance training program under the grant agreement No 825215 (Topic: ICT-35-2018, Type of action: CSA).

The authors gratefully acknowledge the support of the Marie Skłodowska-Curie Actions under the European Union's Horizon Europe research and innovation program for the Industrial Doctoral Network on Digital Finance, acronym: DIGITAL, Project No. 101119635.

This work is further supported by Innosuisse, the Swiss Innovation Agency, under Innovation Project 133.672 IP-SBM “A Specialized and Secure AI Orchestrator for Swiss Financial Compliance,” a research collaboration between the University of Applied Sciences of the Grisons (FHGR) and WeCanGroup SA.

C. Mare gratefully acknowledges support from Babeş-Bolyai University (contract PFE-550-UBB) and from a grant of the Ministry of Research, Innovation and Digitization of Romania, CNCS-UEFISCDI, project number PN-IV-P2-2.1-TE-2023-1317, within PNCDI IV.

The funding bodies had no role in the design of the study; in the collection, analysis, and interpretation of data; or in writing the manuscript.

Authors' contributions

J.O. (Jörg Osterrieder): Conceptualization, Methodology, Investigation, Formal analysis, Writing – Original draft, Writing – Review & editing, Supervision, Project administration, Funding acquisition. **L.J.B.** (Lennart John Baals): Validation, Writing – Review & editing. **C.M.** (Codruța Mare): Conceptualization, Validation, Funding acquisition, Writing – Review & editing. All authors read and approved the final manuscript.

Ethics approval and consent to participate

Not applicable. This study did not involve human participants, human data, or animal subjects; it is based on published literature, publicly available policy reports, and documented crisis episodes.

Acknowledgements

The authors are grateful to working group members and management committee members of the COST (Cooperation in Science and Technology) Action CA19130 Fintech and Artificial Intelligence in Finance. This European network was established in 2018 and now encompasses more than 430 researchers from 51 countries internationally. The authors also thank members of the MSCA Industrial Doctoral network on Digital Finance, founded in 2024 and now encompassing more than 140 members from more than 15 countries globally. The collaboration with the COST Action CA21163 Text, functional and other high-dimensional data in econometrics: new models, methods, applications is acknowledged.

Declaration of AI Assistance

The authors used large language model assistants during drafting, editing, and proofreading of this manuscript. The authors take full responsibility for the content, analyses, and conclusions presented.

References

Acemoglu, D., Ozdaglar, A., Tahbaz-Salehi, A., 2015. Systemic risk and stability in financial networks. American Economic Review 105, 564–608.

Acharya, V.V., Pedersen, L.H., Philippon, T., Richardson, M., 2017. Measuring systemic risk. Review of Financial Studies 30, 2–47. doi:10.1093/rfs/hhw088.

Adrian, T., Brunnermeier, M.K., 2016. CoVaR. American Economic Review 106, 1705–1741. doi:10.1257/aer.20120555.

Akerlof, G.A., 1970. The market for "lemons": Quality uncertainty and the market mechanism. Quarterly Journal of Economics 84, 488–500.

Allen, F., Gale, D., 2000. Financial contagion. Journal of Political Economy 108, 1–33.

Aramonte, S., Huang, W., Schrimpf, A., 2022. DeFi Leverage. Quarterly Review Article BIS Quarterly Review, December 2022. Bank for International Settlements.

Babus, A., 2016. The formation of financial networks. RAND Journal of Economics 47, 239–272. doi:10.1111/1756-2171.12126.

Bank for International Settlements, 2021. DeFi Risks and the Decentralisation Illusion. Quarterly Review Article BIS Quarterly Review, December 2021. Bank for International Settlements. URL: https://www.bis.org/publ/qtrpdf/r_qt2112b.htm. authors: Sirio Aramonte, Wenqian Huang, Andreas Schrimpf.

Bank for International Settlements, 2022. Crypto Risk and Macro-Financial Stability: A Literature Review. Working Paper BIS Working Paper No. 1036. Bank for International Settlements. URL: https://www.bis.org/publ/work1036.htm.

Bank for International Settlements, 2024. Tokenisation in the Context of Money and Other Assets: Concepts and Implications for Central Banks. BIS Papers BIS Papers No. 148. Bank for International Settlements. URL: https://www.bis.org/publ/bppdf/bispap148.htm.

Bank of England, 2023. Regulatory Regime for Systemic Payment Systems Using Stablecoins and Related Service Providers. Discussion Paper DP 2023/4. Bank of England. URL: https://www.bankofengland.co.uk/paper/2023/dp/regulatory-regime-for-systemic-payment-systems-using-stablecoins-and-related-service-pr

Barabási, A.L., Albert, R., 1999. Emergence of scaling in random networks. Science 286, 509–512. doi:10.1126/science.286.5439.509.

Battiston, S., Puliga, M., Kaushik, R., Tasca, P., Caldarelli, G., 2012. DebtRank: Too central to fail? financial networks, the FED and systemic risk. Scientific Reports 2, 541. doi:10.1038/srep00541.

Briola, A., Vidal-Tomás, D., Wang, Y., Aste, T., 2023. Anatomy of a stablecoin's failure: The terra-luna case. Finance Research Letters 51, 103358.

Brunnermeier, M.K., Pedersen, L.H., 2009. Market liquidity and funding liquidity. Review of Financial Studies 22, 2201–2238.

Cookson, J.A., Fox, C., Gil-Bazo, J., Imbet, J.F., Schiller, C., 2023. Social media as a bank run catalyst doi:10.2139/ssrn.4422754.

Daian, P., Goldfeder, S., Kell, T., Li, Y., Zhao, X., Bentov, I., Breidenbach, L., Juels, A., 2020. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability, in: IEEE Symposium on Security and Privacy, pp. 910–927.

Diamond, D.W., Dybvig, P.H., 1983. Bank runs, deposit insurance, and liquidity. Journal of Political Economy 91, 401–419.

Diop, P.O., Chevallier, J., Sanhaji, B., 2024. Collapse of silicon valley bank and usdc depegging: A machine learning experiment doi:10.3390/fintech3040030.

Duffie, D., 2010. Presidential address: Asset price dynamics with slow-moving capital. Journal of Finance 65, 1237–1267.

Financial Stability Board, 2022. Assessment of Risks to Financial Stability from Crypto-Assets. Report. Financial Stability Board. URL: https://www.fsb.org/2022/02/assessment-of-risks-to-financial-stability-from-crypto-assets/.

Financial Stability Board, 2023a. Global Regulatory Framework for Crypto-Asset Activities. Report. Financial Stability Board. URL: https://www.fsb.org/2023/07/fsb-global-regulatory-framework-for-crypto-asset-activities/.

Financial Stability Board, 2023b. High-Level Recommendations for the Regulation, Supervision and Oversight of Crypto-Asset Activities and Markets. Final Report. Financial Stability Board. URL: https://www.fsb.org/2023/07/high-level-recommendations-for-the-regulation-supervision-and-oversight-of-crypto-asset

Financial Stability Board, 2023c. High-Level Recommendations for the Regulation, Supervision and Oversight of Global Stablecoin Arrangements: Final Report. Final Report. Financial Stability Board. URL: https://www.fsb.org/2023/07/high-level-recommendations-for-the-regulation-supervision-and-oversight-of-global-stabl

Gai, P., Kapadia, S., 2010. Contagion in financial networks. Proceedings of the Royal Society A 466, 2401–2423. doi:10.1098/rspa.2009.0410.

Galati, L., Capalbo, F., 2023. Silicon valley bank bankruptcy and stablecoins stability doi:10.1016/j.irfa.2023.103001.

Gorton, G.B., Zhang, J.Y., 2023. Taming wildcat stablecoins. University of Chicago Law Review 90, 909–971.

Gregory, G., Alessio, C., Vito, L., Patrice, S., 2024. Break a peg! a study of stablecoin co-instability doi:10.1016/j.irfa.2024.103608.

Gudgeon, L., Pérez, D., Harz, D., Livshits, B., Gervais, A., 2020. The decentralized financial crisis doi:10.1109/cvcbt50464.2020.00005.

International Monetary Fund, 2021. The Crypto Ecosystem and Financial Stability Challenges. GFSR Chapter Global Financial Stability Report, October 2021, Chapter 2. International Monetary Fund. URL: https://www.imf.org/en/Publications/GFSR/Issues/2021/10/12/global-financial-stability-report-october-2021.

International Monetary Fund, 2023. Elements of Effective Policies for Crypto Assets. IMF Policy Paper Policy Paper No. 2023/004. International Monetary Fund.

International Monetary Fund, 2024. Central Bank Digital Currency: Progress, Plans, and Risks. Staff Discussion Note. International Monetary Fund. URL: https://www.imf.org/en/Publications/Staff-Discussion-Notes/Issues/2024/02/29/Central-Bank-Digital-Currency-Progress-Plans-and-Risks-545364.

International Monetary Fund and Financial Stability Board, 2023. IMF-FSB Synthesis Paper: Policies for Crypto-Assets. Synthesis Paper. International Monetary Fund and Financial Stability Board. URL: https://www.imf.org/en/Publications/Policy-Papers/Issues/2023/09/07/IMF-FSB-Synthesis-Paper-Policies-for-Crypto-Assets-538692.

International Organization of Securities Commissions, 2022. IOSCO Decentralized Finance Report. Report. IOSCO. URL: https://www.iosco.org/library/pubdocs/pdf/IOSCOPD699.pdf.

Jalan, A., Matkovskyy, R., 2023. Systemic risks in the cryptocurrency market: Evidence from the ftx collapse doi:10.1016/j.frl.2023.103670.

Kitzler, S., Victor, F., Saggese, P., Haslhofer, B., 2022. Disentangling decentralized finance (defi) compositions doi:10.1145/3532857.

Klages-Mundt, A., Minca, A., 2022. While stability lasts: A stochastic model of noncustodial stablecoins. Mathematical Finance 32, 943–981. doi:10.1111/mafi.12357.

Li, Z., Li, J., He, Z., Luo, X., Wang, T., Ni, X., Yang, W., Chen, X., Chen, T., 2023. Demystifying defi mev activities in flashbots bundle doi:10.1145/3576915.3616590.

Liu, J., Makarov, I., Schoar, A., 2023. Anatomy of a run: The terra luna crash doi:10.3386/w31160.

Lyons, R.K., Viswanath-Natraj, G., 2023. What keeps stablecoins stable? Journal of International Money and Finance 131, 102777. doi:10.1016/j.jimonfin.2022.102777.

Ma, Y., Zeng, Y., Zhang, A.L., 2025. Stablecoin runs and the centralization of arbitrage doi:10.3386/w33882.

Makarov, I., Schoar, A., 2022. Cryptocurrencies and decentralized finance (defi) doi:10.1353/eca.2022.0014.

Ontario Securities Commission, 2020. QuadrigaCX: A Review by Staff of the Ontario Securities Commission. Technical Report. Ontario Securities Commission. URL: https://www.osc.ca/quadrigacxreport.

Perez, D., Werner, S.M., Xu, J., Livshits, B., 2021. Liquidations: DeFi on a knife-edge, in: Financial Cryptography and Data Security (FC 2021). doi:10.1007/978-3-662-64331-0_24.

Qin, K., Zhou, L., Livshits, B., Gervais, A., 2022. Quantifying blockchain extractable value: How dark is the forest? IEEE Symposium on Security and Privacy .

Schuler, K., Cloots, A.S., Schär, F., 2024. On defi and on-chain cefi: How (not) to regulate decentralized finance doi:10.1093/jfr/fjad014.

Schär, F., 2020. Decentralized finance: On blockchain- and smart contract-based financial markets doi:10.2139/ssrn.3571335.

Schär, F., 2021. Decentralized finance: On blockchain- and smart contract-based financial markets doi:10.20955/r.103.153-74.

Shleifer, A., Vishny, R., 2011. Fire sales in finance and macroeconomics. Journal of Economic Perspectives 25, 29–48.

Stiglitz, J.E., Weiss, A., 1981. Credit rationing in markets with imperfect information. American Economic Review 71, 393–410.

Tolmach, P., Li, Y., Lin, S.W., Liu, Y., 2021. Formal analysis of composable DeFi protocols, in: Financial Cryptography Workshops. doi:10.1007/978-3-662-63958-0_13.

Uhlig, H., 2022. A Luna-tic Stablecoin Crash. Working Paper 30256. National Bureau of Economic Research.

Xu, J., Vadgama, N., 2023. Sok: Decentralized finance (defi) attacks. arXiv preprint .