

Systemic Risk Channels in Digital Finance: A Comprehensive Taxonomy

Jörg Osterrieder* Lennart John Baals[†] Codruța Mare[‡]

April 25, 2026

Abstract

The rapid growth of decentralized finance (DeFi), centralized cryptocurrency platforms (CeFi), and stablecoin ecosystems has introduced novel channels of systemic risk that existing frameworks developed for traditional, regulated financial institutions are ill-equipped to capture. This paper develops the first comprehensive taxonomy of systemic risk transmission channels across all major segments of digital finance. Through a hybrid methodology combining systematic literature review, theoretical derivation, and crisis case evidence drawn from twenty-five episodes between 2009 and 2026, we identify eight distinct channels: (1) network contagion, (2) liquidity spirals, (3) stablecoin runs, (4) composability risk, (5) liquidation cascades, (6) counterparty concentration, (7) information asymmetry, and (8) fiat-crypto gateway risk. Two channels composability risk and gateway risk are classified as genuinely novel, with no meaningful analog in traditional finance. Four are hybrids that extend established mechanisms with significant digital-native features, and two are direct extensions of canonical financial theory. We document cross-channel interactions that produce compounding cascade dynamics, as demonstrated by the 2022 crisis sequence spanning Terra/Luna, Three Arrows Capital, Celsius, and FTX. An evolutionary analysis traces the development of the systemic risk landscape from the Mt. Gox era to the present, and a forward-looking assessment examines emerging risks from tokenized real-world assets, CBDC interoperability, and AI-driven trading. The taxonomy provides a unified analytical framework for researchers and regulators seeking to understand, measure, and mitigate systemic risk in digital finance.

JEL Classification: G01, G18, G23, G28, O33

Keywords: Systemic risk, decentralized finance, stablecoins, cryptocurrency, contagion, financial stability, blockchain, digital assets, taxonomy, financial regulation

*Corresponding author. Department of Behavioural, Management and Social Sciences, University of Twente, The Netherlands; and University of Applied Sciences of the Grisons, Chur, Switzerland. Email: j.r.osterrieder@utwente.nl. ORCID: 0000-0002-0872-1080.

[†]Department of Behavioural, Management and Social Sciences, University of Twente, The Netherlands. Email: l.j.baals@utwente.nl. ORCID: to be inserted.

[‡]Department of Statistics-Forecasts-Mathematics, Faculty of Economics and Business Administration, and Interdisciplinary Centre for Data Science, Babeş-Bolyai University, Cluj-Napoca, Romania. Email: codruta.mare@econ.ubbcluj.ro. ORCID: 0000-0002-4920-7462.

Contents

1	Introduction	5
2	Background and Conceptual Foundations	9
2.1	Systemic Risk in Traditional Finance	9
2.2	The Digital Finance Landscape	12
2.3	Why Traditional Frameworks Are Insufficient	15
3	Methodology	17
3.1	Systematic Literature Search	17
3.2	Taxonomy Construction	18
3.3	Conceptual Propositions	22
3.4	Limitations and Methodological Considerations	23
4	A Taxonomy of Systemic Risk Channels in Digital Finance	25
4.1	Network Contagion and Interconnectedness	25
4.1.1	Mechanism	26
4.1.2	Theoretical Foundations	26
4.1.3	Digital-Native Features	27
4.1.4	Case Evidence	28
4.1.5	Cross-Domain Manifestation	29
4.1.6	Literature Assessment	29
4.2	Liquidity Spirals and Fire Sales	30
4.2.1	Mechanism	30
4.2.2	Theoretical Foundations	31
4.2.3	Digital-Native Features	32
4.2.4	Case Evidence	33
4.2.5	Cross-Domain Manifestation	33
4.2.6	Literature Assessment	34
4.3	Stablecoin De-Pegging and Run Dynamics	34
4.3.1	Mechanism	34
4.3.2	Theoretical Foundations	35
4.3.3	Digital-Native Features	36
4.3.4	Case Evidence	37
4.3.5	Cross-Domain Manifestation	38
4.3.6	Literature Assessment	38
4.4	Composability and Smart Contract Cascade Risk	39
4.4.1	Mechanism	39
4.4.2	Theoretical Foundations	40
4.4.3	Digital-Native Features	40
4.4.4	Case Evidence	41
4.4.5	Cross-Domain Manifestation	42
4.4.6	Literature Assessment	42
4.5	Leverage and Liquidation Cascades	43

4.5.1	Mechanism	43
4.5.2	Theoretical Foundations	44
4.5.3	Digital-Native Features	44
4.5.4	Case Evidence	45
4.5.5	Cross-Domain Manifestation	46
4.5.6	Literature Assessment	46
4.6	Counterparty and Concentration Risk	47
4.6.1	Mechanism	47
4.6.2	Theoretical Foundations	48
4.6.3	Digital-Native Features	49
4.6.4	Case Evidence	49
4.6.5	Cross-Domain Manifestation	50
4.6.6	Literature Assessment	50
4.7	Information Asymmetry and Opacity	51
4.7.1	Mechanism	51
4.7.2	Theoretical Foundations	52
4.7.3	Digital-Native Features	53
4.7.4	Case Evidence	53
4.7.5	Cross-Domain Manifestation	54
4.7.6	Literature Assessment	55
4.8	Fiat-Crypto Gateway and Banking Channel Risk	55
4.8.1	Mechanism	55
4.8.2	Theoretical Foundations	56
4.8.3	Digital-Native Features	57
4.8.4	Case Evidence	57
4.8.5	Cross-Domain Manifestation	58
4.8.6	Literature Assessment	58
5	Cross-Channel Interactions and Amplification Mechanisms	59
5.1	Interaction Taxonomy: Amplifying, Dampening, and Sequencing	59
5.2	The Five Strongest Feedback Loops	60
5.3	Crisis Case Study: Terra/Luna 2022	61
5.4	Crisis Case Study: FTX 2022	62
5.5	Crisis Case Study: Black Thursday 2020	63
5.6	Market Structure and Cascade Acceleration	63
5.7	Cross-Domain Contagion Pathways	64
5.8	Interaction Matrix and Channel Coupling Intensity	65
5.9	Implications for Systemic Risk Assessment	65
6	Evolutionary Dynamics of Systemic Risk	66
6.1	Early Era: Exchange Failures and Bilateral Risk (2011–2017)	66
6.2	DeFi Summer and Emerging Complexity (2020–2021)	67
6.3	The Great Unraveling (2022)	68
6.4	Maturation and Institutional Entry (2023–Present)	68
6.5	Channel Evolution Trajectories	69

7	Emerging Channels and Forward-Looking Assessment	70
7.1	Tokenized Real-World Assets and TradFi Bridges	70
7.2	CBDC Interoperability Risks	70
7.3	AI-Driven Trading and Algorithmic Herding	71
7.4	Cross-Chain Composability and Layer-2 Fragmentation	72
7.5	Framework Durability	72
8	Policy Implications	73
8.1	Regulatory Gaps Highlighted by the Taxonomy	73
8.2	Macroprudential Tools for Digital Finance	74
8.3	Microprudential Considerations	75
8.4	Monitoring Architecture	75
8.5	International Coordination Challenges	76
8.6	Institutional Design Priorities	77
9	Conclusion	78
9.1	Taxonomy Contributions	78
9.2	Key Findings	79
9.3	Limitations	79
9.4	Future Research Directions	80
9.5	Closing Perspective	81
A	Literature Search Protocol	95
B	Complete Channel Interaction Matrix	98
C	Crisis Event Database	99

1 Introduction

The digital transformation of financial services has created a parallel financial system whose scale and complexity demand systematic risk assessment. 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](#)). 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 within the scope of macroprudential concern.

The collapse of the Terra/Luna ecosystem in May 2022 destroyed approximately \$45 billion in value within five days a loss confined to Terra/Luna holders and UST depositors. 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 beyond the initial Terra/Luna destruction ([Financial Stability Board, 2022a](#)). This sequence was not a series of unrelated failures but 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 total estimated losses across the 2022 cascade exceeded \$2 trillion in market capitalization, dwarfing the scale of any previous crypto market event.

The systemic risk frameworks developed in response to the 2007–2009 Global Financial Crisis have proven insufficient for this new 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 regulatory oversight ([Financial Stability Board, 2023b](#)). The foundational premise of institution-based systemic risk measurement that the relevant entities can be identified, their exposures can be observed, and their distress can be conditioned upon breaks down in an ecosystem built on pseudonymous addresses and permissionless participation.

Counterparty exposures in digital finance span both on-chain interactions which are transparent to anyone with blockchain analytics capability 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 or months to hours, the absence of circuit breakers or trading halts, and the permissionless nature of participation further distinguish digital finance from the institutional settings for which existing systemic risk tools were calibrated ([Alamsyah et al., 2024](#)). When Bear Stearns faced a liquidity crisis in March 2008, the Federal Reserve organized a weekend rescue; when Celsius Network faced an analogous crisis in June 2022, the platform froze customer withdrawals

within hours and entered bankruptcy proceedings within weeks, with no regulatory authority capable of organizing an orderly resolution.

The academic literature on systemic risk in digital finance has grown rapidly but remains fragmented across subfields that rarely communicate. Individual studies examine specific channels in isolation: stablecoin run dynamics and their parallels to traditional bank runs (Uhlig, 2022), DeFi liquidation mechanics and the role of automated execution in cascade amplification (Gudgeon et al., 2020), network contagion patterns in the crypto lending ecosystem, and the opacity of centralized platforms that enabled institutional fraud. Each body of work advances understanding of a single mechanism but does not connect to the others in a unified analytical framework.

The network contagion literature has developed sophisticated models of how shocks propagate through crypto lending networks, but these models do not account for the simultaneous activation of liquidity spirals in AMM pools or the reflexive dynamics of stablecoin de-pegging (Aramonte et al., 2022; Allen et al., 2022). The stablecoin literature has produced rigorous analyses of run dynamics, but these analyses do not integrate the role of gateway institutions whose banking relationships determine whether stablecoin reserves can be accessed during stress. The information asymmetry literature has documented the opacity of CeFi platforms, but has not connected this opacity to the specific contagion channels through which undisclosed exposures propagate losses (Huan and Renn, 2025). The result is a literature that understands individual trees but has not mapped the forest.

Policy reports from the Financial Stability Board, the Bank for International Settlements, the International Monetary Fund, and the European Central Bank have catalogued risk factors and proposed regulatory responses (Financial Stability Board, 2022a; Bank for International Settlements, 2022). These reports identify overlapping risk categories typically organized by market segment or regulatory jurisdiction without providing a unified analytical framework organized by transmission mechanism. The FSB’s 2022 assessment, for example, catalogs risks under headings such as “operational vulnerabilities,” “market integrity,” and “financial stability,” cutting across the transmission channels rather than mapping them. The gap between the granularity of individual academic studies and the breadth of policy surveys motivates the taxonomy we develop.

What remains absent, to our knowledge, 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 (Briola et al., 2023). This paper fills that gap by developing the first comprehensive taxonomy of systemic risk transmission channels across all major segments of the digital finance ecosystem.

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. Each channel is analyzed through a consistent lens specifying its transmission mechanism, theoretical foundations, digital-native features, historical crisis evidence, cross-domain manifestation, and literature gaps (Diamond and Dybvig, 1983). 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; to our knowledge, 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. 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 (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 (Daian et al., 2020). 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 in Anchor and other lending platforms, which propagated through the network contagion channel 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; Alamsyah et al., 2024). Understanding these interactions requires a taxonomy that identifies each channel individually and then maps their coupling dynamics the dual contribution that this paper provides.

The taxonomy was constructed through a hybrid methodology combining systematic literature review with theoretical derivation and crisis case evidence. We conducted a structured search through OpenAlex that identified an initial candidate set of fourteen potential channels across approximately 2,400 unique works after deduplication (International Monetary Fund, 2021; Financial Stability Board, 2023b). 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. Six excluded candidates oracle manipulation, regulatory contagion, governance failure, real-world asset transmission, cross-chain bridge vulnerability, and validator concentration were merged into surviving channels as subtypes or deferred to the forward-looking assessment rather than discarded (Xu and Vadgama, 2023). The methodology ensures that the taxonomy is evidence-based rather than impressionistic, grounded in the actual distribution of scholarly attention and crisis experience across the field.

The timeliness of this taxonomy reflects a critical window in the development of digital finance. The 2022 crisis sequence provided the first comprehensive stress test of the digital finance ecosystem, generating crisis evidence across all eight channels that earlier periods lacked (Bank for International Settlements, 2022). Simultaneously, the growing integration of digital and traditional finance through stablecoin banking relationships, Bitcoin ETFs, tokenized treasuries, and institutional custody services means that the next major digital finance crisis will propagate into the traditional financial system through channels that were barely relevant in 2022. A taxonomy that maps these channels before the next crisis provides the conceptual infrastructure that regulators and researchers need to anticipate, rather than merely react to, systemic events.

The scope of the taxonomy spans four domains: DeFi (decentralized exchanges, lending protocols, yield aggregators, and derivatives platforms), CeFi (centralized exchanges, custodians, and crypto lending platforms), stablecoins and central bank digital currencies, and the emerging tokenized traditional finance sector. The organizing principle is transmission mechanism rather than domain, asset type, or chronological emergence. Each channel describes a distinct causal pathway through which distress propagates, and the cross-channel interaction analysis documents how the activation of one channel triggers or amplifies others a compounding dynamic that the 2022 crisis sequence illustrated with devastating clarity.

The paper makes three principal contributions. First, we provide the first systematic taxonomy of eight channels organized by transmission mechanism and 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 (Briola et al., 2023). Third, we identify specific literature gaps within each channel, providing a structured research agenda for formal modeling of systemic risk in digital finance. The taxonomy provides the conceptual architecture within which formal models can be developed; the formalization of individual channel dynamics and their interactions constitutes a primary avenue for future work identified in Section 9.

We emphasize that this taxonomy is not merely descriptive. By grounding each channel in established economic theory where applicable and identifying precisely where that theory must be extended, we provide a platform for the next generation of formal systemic risk models. The two novel channels represent, to our knowledge, genuine theoretical lacunae that the economics profession has not yet addressed. The six hybrid and extension channels reveal specific assumptions in canonical models human decision-making in liquidation, sequential service in bank runs, bilateral exposure observability in contagion networks that digital finance systematically violates (Financial Stability Board, 2022a). Identifying these violated assumptions is the first step toward developing systemic risk models calibrated for a financial

system that operates at the speed of code.

The remainder of this paper proceeds as follows. Section 2 presents the conceptual foundations by reviewing systemic risk theory in traditional finance (Section 2.1), mapping the institutional architecture of digital finance (Section 2.2), and identifying the structural features that render existing frameworks insufficient (Section 2.3). Section 3 describes the literature search protocol, channel identification process, and selection criteria (Aramonte et al., 2022; Allen et al., 2022). Section 4 presents the eight-channel taxonomy, devoting a subsection to each channel. Section 5 analyzes cross-channel interactions and amplification dynamics, with particular attention to the 2022 cascade sequence. Section 6 provides an evolutionary perspective on systemic risk events from the Mt. Gox era through the present. Section 7 assesses emerging channels including tokenized real-world assets, CBDC interoperability, and AI-driven trading. Section 8 draws policy implications from the taxonomy. Section 9 concludes with a summary of contributions, limitations, and a research agenda for the field.

2 Background and Conceptual Foundations

2.1 Systemic Risk in Traditional Finance

Systemic risk refers to the risk that the failure of a single financial institution, market segment, or payment system triggers a chain of failures that impairs the functioning of the financial system as a whole, with material adverse consequences for the real economy. The European Central Bank defines systemic risk as “a risk of financial instability so widespread that it impairs the functioning of a financial system to the point where economic growth and welfare suffer materially” (European Central Bank, 2019). The Bank for International Settlements and the Financial Stability Board adopt parallel definitions that emphasize the endogenous amplification of shocks through interconnections, common exposures, and information externalities (Bank for International Settlements, 2022; Financial Stability Board, 2018). The distinguishing feature of systemic risk, as opposed to the idiosyncratic risk of individual firm failure, is that the social cost of distress exceeds the private cost borne by the failing institution a negative externality that provides the economic rationale for macroprudential regulation, deposit insurance, and lender-of-last-resort facilities (Adrian and Brunnermeier, 2016).

The modern academic literature on systemic risk coalesces around four canonical transmission channels that provide the theoretical foundation for the digital finance taxonomy we develop. First, Diamond and Dybvig (1983) formalized the bank run as a coordination failure among depositors, demonstrating that maturity transformation the financing of illiquid long-term assets with liquid short-term liabilities creates a fragile equilibrium in which rational depositors may run if they believe others will do the same. Deposit insurance and lender-of-last-resort facilities eliminate the bank-run equilibrium in theory, but practical limitations became apparent during the 2007–2008 crisis when runs migrated to uninsured shadow-banking instruments such as asset-backed commercial paper, repo markets, and money market mutual funds (Gorton and Metrick, 2012). The key insight for digital finance is that any system performing maturity transformation without credible backstops

inherits the coordination failure, regardless of whether the “bank” is a traditional depository institution, a money market fund, or a CeFi lending platform (Allen and Gale, 2000).

Second, Allen and Gale (2000) developed a model of financial contagion through interbank networks, showing that the structure of the network determines whether shocks are absorbed or amplified. Complete networks with diversified bilateral exposures exhibit resilience, while incomplete networks with concentrated linkages are vulnerable to cascading failure. Acemoglu et al. (2015) extended this framework to demonstrate that the relationship between connectivity and stability is non-monotonic: below a critical threshold, additional connections improve shock absorption by distributing losses across more counterparties; above the threshold, the same connections become channels for contagion that amplify shocks system-wide. Elliott et al. (2014) complemented this analysis by showing that cross-holdings of financial assets create integration effects that can either diversify or concentrate risk depending on the degree of portfolio overlap and the magnitude of the initial shock (Upper, 2011).

The network contagion literature has since expanded to analyze core-periphery structures, scale-free topologies, and multiplex networks in which institutions are connected through multiple types of exposure simultaneously (Duffie, 2010). Gai and Kapadia (2010) demonstrated that financial networks are robust yet fragile able to absorb small shocks but vulnerable to cascading failure when initial losses exceed a tipping point determined by network structure. Haldane and May (2011) drew on ecological network models to argue that modern banking systems exhibit the same vulnerability as ecosystems: homogeneity and tight coupling increase efficiency under normal conditions but amplify fragility during stress. Battiston et al. (2012) operationalized this insight through the DebtRank measure, which captures the feedback between node distress and network-wide loss propagation, providing a tool for identifying systemically important institutions based on network centrality rather than balance-sheet size alone. Glasserman and Young (2016) provided an authoritative survey of the mechanisms through which contagion propagates in financial networks, establishing the canonical framework for subsequent empirical and theoretical work. A consistent finding is that both the topology of the network and the size distribution of exposures determine whether an initial shock is contained locally or amplifies into a system-wide event. Core-periphery structures, in which a small number of highly connected institutions intermediate most transactions, are particularly fragile because the failure of a core institution propagates losses to all peripheral counterparties simultaneously (Financial Stability Board, 2018). This finding has direct implications for digital finance, where a small number of centralized exchanges, stablecoin issuers, and market makers occupy analogous core positions.

Third, Brunnermeier and Pedersen (2009) formalized the interaction between market liquidity and funding liquidity in a model of mutually reinforcing spirals. When asset prices decline, the value of collateral held by leveraged intermediaries falls, triggering margin calls. Intermediaries forced to sell assets into declining markets depress prices further, which tightens funding conditions and triggers additional margin calls. The loss spiral and margin spiral operate as a positive feedback loop that can transform a modest initial shock into a severe liquidity crisis. Shleifer and Vishny (2011) provided complementary analysis demonstrating that forced sellers impose fire-sale externalities on other market participants by depressing asset prices below fundamental values. Gromb and Vayanos (2010) extended this analysis to show how margin constraints reduce the arbitrage capacity of financial intermediaries,

widening the gap between market prices and fundamental values during stress.

The liquidity spiral mechanism proved central to the 2007–2008 crisis, when declining values of mortgage-backed securities triggered margin calls on leveraged financial institutions holding those securities as collateral, forcing fire sales that depressed prices further and tightened credit conditions system-wide. [Brunnermeier \(2009\)](#) documented how losses in the US mortgage market, which were large in absolute terms but modest relative to the \$8 trillion in stock market wealth ultimately destroyed, amplified through these feedback mechanisms into the most severe financial crisis since the Great Depression. The crisis demonstrated that the severity of a systemic event depends not on the size of the initial shock but on the amplification dynamics embedded in the system’s institutional structure.

Fourth, the development of quantitative systemic risk measures provided tools for identifying systemically important institutions and monitoring aggregate vulnerability. [Adrian and Brunnermeier \(2016\)](#) introduced CoVaR, measuring the value-at-risk of the financial system conditional on an individual institution being in distress, shifting the focus from the riskiness of individual institutions to the externalities each imposes on the system. [Acharya et al. \(2017\)](#) developed complementary measures Marginal Expected Shortfall and SRISK that capture the expected capital shortfall of a firm during a systemic crisis. These tools have been widely adopted for regulatory stress testing, the designation of globally systemically important financial institutions (G-SIBs), and the calibration of macroprudential capital surcharges ([European Central Bank, 2019](#)).

The fire-sale mechanism has been extensively studied because it creates a direct link between individual institutional distress and market-wide price dislocations. When multiple institutions hold the same asset and one is forced to sell, the price decline impairs the balance sheets of all other holders, potentially triggering further forced sales ([Gromb and Vayanos, 2010](#); [Shleifer and Vishny, 2011](#)). This “fire-sale externality” is a canonical example of the negative externality that justifies macroprudential regulation: the selling institution bears only its own losses, not the losses it imposes on other market participants through price depression. The mechanism operates with particular force when markets are illiquid and when many institutions are leveraged against the same collateral.

A fifth strand, though less often framed as a standalone channel in the traditional literature, underpins all four pillars: information asymmetry. [Akerlof \(1970\)](#) demonstrated that asymmetric information between buyers and sellers can cause market breakdown through adverse selection, while [Stiglitz and Weiss \(1981\)](#) showed that information asymmetry in credit markets leads to rationing rather than price-clearing equilibria. In the systemic risk context, opacity about counterparty exposures amplifies contagion because market participants cannot distinguish solvent from insolvent counterparties during stress, triggering precautionary withdrawal that deepens liquidity crises ([Allen and Gale, 2000](#)). The 2007–2008 crisis demonstrated this mechanism when uncertainty about which institutions held toxic mortgage-backed securities froze the interbank lending market entirely, converting a solvency problem at specific institutions into a system-wide liquidity crisis.

These five pillars bank runs, network contagion, liquidity spirals, conditional risk measurement, and information asymmetry interact in practice and rarely operate in isolation. The 2007–2008 crisis combined information asymmetry about mortgage-backed security exposures with network contagion through the interbank lending market, liquidity spirals in structured credit products, and runs on shadow-banking vehicles including money market

funds and repo markets (Gorton and Metrick, 2012; Elliott et al., 2014; Upper, 2011). The crisis revealed that cross-channel interaction produces compounding dynamics more severe than any single channel operating alone an insight that carries directly into our digital finance taxonomy, where the same compounding dynamic operates at vastly greater speed.

Each pillar identifies a transmission mechanism with analogs in digital finance, but the structural properties of blockchain-based systems automated execution, permissionless participation, and 24/7 global operation alter the dynamics of each mechanism in ways the original models did not contemplate (Diamond and Dybvig, 1983; Acemoglu et al., 2015; Brunnermeier and Pedersen, 2009). Diamond-Dybvig assumes deposit insurance and lender-of-last-resort facilities absent from crypto markets. Allen-Gale assumes observable bilateral exposures, an impossibility in pseudonymous markets. Brunnermeier-Pedersen assumes human decision-making timelines that smart contracts compress to seconds. Adrian-Brunnermeier requires institution-level return data that most crypto entities do not produce. The taxonomy developed in Section 4 maps each digital finance channel to its traditional analog where one exists and identifies the specific assumptions that digital finance violates.

2.2 The Digital Finance Landscape

Digital finance encompasses a heterogeneous set of institutions, protocols, and instruments that use blockchain technology and cryptographic techniques to provide financial services. For the purposes of this taxonomy, we organize the landscape into four domains, each with distinct institutional structures, governance arrangements, and risk profiles (Werner et al., 2022; Xu and Vadgama, 2023; Auer et al., 2023). Understanding the architecture of each domain is a prerequisite for analyzing the systemic risk channels that operate within and across them, because the structural features of each domain determine which channels are most salient and how they interact with channels in other domains.

The first domain, decentralized finance (DeFi), consists of financial protocols deployed as smart contracts on programmable blockchains, predominantly Ethereum. Automated market makers (AMMs) such as Uniswap and Curve Finance replace traditional order-book exchanges with liquidity pools governed by deterministic pricing functions: liquidity providers deposit token pairs into a pool, and the pool algorithmically adjusts the exchange rate based on the ratio of reserves (Werner et al., 2022; Mohan, 2022). The constant-product formula that Uniswap pioneered ensures continuous liquidity at all price levels but creates procyclical dynamics during stress, as large trades move prices mechanically and liquidity providers face impermanent loss when prices diverge from their initial deposit ratio (Bank for International Settlements, 2021b). By the peak of “DeFi Summer” in 2021, total value locked across DeFi protocols exceeded \$150 billion.

Lending protocols such as Aave and Compound allow users to deposit collateral and borrow other assets, with loan terms collateralization ratios, interest rates, and liquidation thresholds determined entirely by smart contract parameters rather than bilateral negotiation (Xu and Vadgama, 2023). Overcollateralization requirements (typically 125–150% of the borrowed amount) substitute for credit assessment, and liquidation is executed automatically by permissionless bots when collateral values fall below the threshold. Yield aggregators such as Yearn Finance automate capital deployment across multiple protocols to maximize returns, creating additional layers of protocol interdependence. The defining

structural feature of DeFi is composability: any protocol can call the functions of any other protocol without permission, enabling the construction of complex financial products by stacking simple building blocks a property often described as “money legos” (Schär, 2021). This composability drives innovation but creates systemic dependencies that no single actor designs, monitors, or controls.

The second domain, centralized finance (CeFi), comprises centralized cryptocurrency exchanges (Binance, Coinbase, the defunct FTX), custodial lending platforms (the defunct Celsius Network, BlockFi, Voyager Digital), and over-the-counter trading desks (the defunct Alameda Research, Genesis Trading). CeFi platforms intermediate the majority of cryptocurrency trading volume and serve as the primary interface between retail users and the digital asset ecosystem (Auer et al., 2023). Unlike DeFi protocols, CeFi platforms operate as corporate entities with centralized management, proprietary technology stacks, and critically custodial control over user assets. The opacity of CeFi balance sheets, combined with the absence of mandatory audit, disclosure, or capital adequacy requirements in most jurisdictions, creates information asymmetries that have repeatedly enabled fraud and risk mismanagement (Ontario Securities Commission, 2020; Allen et al., 2022). The FTX collapse revealed that customer deposits had been commingled with proprietary trading funds, a practice that would have been detected immediately under traditional financial regulation.

The third domain, stablecoins, encompasses digital tokens designed to maintain a stable value relative to a reference currency, typically the US dollar. Stablecoins serve as the primary unit of account, medium of exchange, and settlement asset across both DeFi and CeFi, making their stability a precondition for ecosystem functioning (Financial Stability Board, 2022b). Three architectural categories exist, each with distinct stability mechanisms and failure modes. Fiat-backed stablecoins (Tether’s USDT, Circle’s USDC) maintain reserves of cash, cash equivalents, and short-term securities in traditional financial institutions; their stability depends on the adequacy and accessibility of these reserves. Crypto-collateralized stablecoins (MakerDAO’s DAI) are backed by overcollateralized deposits of volatile crypto assets held in smart contracts; their stability depends on liquidation mechanisms functioning correctly during market stress. Algorithmic stablecoins (the defunct TerraUSD) attempted to maintain their peg through automated supply adjustments linked to a companion volatile token, without maintaining full reserves a design that Uhlig (2022) showed is inherently fragile because redemptions create reflexive selling pressure on the companion token.

The aggregate market capitalization of stablecoins exceeded \$160 billion at its peak in early 2022, and daily on-chain transfer volumes routinely exceeded those of major traditional payment networks (European Central Bank, 2022). The Terra/UST collapse in May 2022 destroyed approximately \$18 billion in stablecoin value within days and demonstrated the catastrophic reflexivity of algorithmic designs: as UST lost its peg, the protocol minted LUNA to absorb redemptions, LUNA’s price collapsed under the supply pressure, which further eroded confidence in UST, creating a death spiral that no mechanism could arrest (Briola et al., 2023). When USDC briefly de-pegged to \$0.87 following the Silicon Valley Bank failure in March 2023, the disruption cascaded through hundreds of DeFi protocols that relied on USDC as collateral or pricing benchmark, illustrating that stablecoin stability depends not only on algorithmic design or reserve management but also on the soundness of the traditional banking institutions that custody those reserves (Galati and Capalbo, 2023).

The fourth domain, tokenized traditional finance, represents the nascent intersection

of blockchain technology with conventional financial assets. Tokenization refers to the issuance of blockchain-based representations of real-world assets, including government bonds, money market fund shares, real estate, and equities (Bank for International Settlements, 2024). While still small in scale relative to native crypto assets, the tokenized asset market has grown rapidly, with tokenized US Treasuries alone exceeding \$1 billion by late 2024. Major financial institutions including BlackRock, Franklin Templeton, and JPMorgan have launched tokenized fund products, signaling growing integration between blockchain infrastructure and traditional capital markets. The domain also encompasses institutional access products such as Bitcoin and Ethereum exchange-traded funds, regulated broker-dealers providing crypto custody, and the banking relationships that mediate fiat-to-crypto conversion (Bank for International Settlements, 2024; Auer et al., 2023; Allen et al., 2022). These gateway institutions create the bidirectional linkages between traditional and digital finance that the gateway risk channel, analyzed in Section 4.8, describes.

Six structural features distinguish digital finance from traditional finance and fundamentally alter the dynamics of systemic risk transmission. First, composability allows protocols to combine without permission, creating dependency chains that are largely invisible to individual participants and extremely difficult to circuit-break during stress (Alamsyah et al., 2024). Second, pseudonymity obscures the identities and aggregate exposures of market participants, preventing the kind of counterparty due diligence that limits concentration risk in regulated markets. Third, continuous 24/7 global operation eliminates the cooling-off periods that overnight market closures and weekend trading halts provide in traditional finance, compressing the timeline of crisis dynamics from weeks to hours (Aramonte et al., 2022). Fourth, atomic settlement the property that complex multi-step transactions either execute completely or not at all within a single blockchain block enables both unprecedented efficiency and novel attack vectors including flash loan exploits and sandwich attacks (Xu and Vadgama, 2023).

Fifth, algorithmic governance replaces human decision-making with smart contract logic for critical functions including liquidation, interest rate adjustment, and collateral valuation, removing the discretion and forbearance that human intermediaries can exercise during stress events (Gudgeon et al., 2020). When a traditional bank faces a stressed borrower, the loan officer can negotiate forbearance, restructure the debt, or escalate to resolution authorities; when a DeFi lending protocol faces an undercollateralized position, the smart contract executes liquidation automatically and irrevocably. Sixth, permissionless participation means that anyone can deploy a protocol, provide liquidity, or take leveraged positions without authorization, credit assessment, or identity verification, creating an open system in which the barriers to both innovation and risk-taking are minimal.

These six features interact to produce a risk environment qualitatively different from traditional finance. The combination of composability and permissionless participation creates emergent complexity that no single actor designs or controls. Algorithmic governance and atomic settlement compress the timeline of crisis propagation below any human reaction threshold. Pseudonymity and 24/7 operation prevent the coordinated regulatory intervention weekend emergency negotiations, overnight capital raising, temporary trading halts that has historically contained traditional financial crises (Adamyk et al., 2025). Moreover, despite the decentralization rhetoric, empirical evidence reveals significant concentration at critical infrastructure layers: a small number of validators process the majority of transactions,

a handful of exchanges dominate trading volume, and two stablecoin issuers control over 80% of the stablecoin market (Ozili, 2022). This concentration recreates the too-big-to-fail dynamics that decentralization was intended to eliminate, while the absence of regulatory oversight means these concentrated entities operate without the capital requirements, stress testing, or resolution frameworks that constrain their traditional counterparts.

2.3 Why Traditional Frameworks Are Insufficient

The structural features of digital finance identified in the preceding subsection create specific analytical challenges that existing systemic risk frameworks were not designed to address. The gap between the assumptions embedded in traditional models and the realities of digital finance motivates the need for a new taxonomy built from the mechanisms themselves rather than retrofitted from existing categories. Network contagion models in the tradition of Allen and Gale (2000) and Acemoglu et al. (2015) assume that bilateral exposures between institutions are known or at least knowable to regulators and, with appropriate disclosure requirements, to counterparties. Digital finance fundamentally violates this assumption (Auer et al., 2023). On-chain DeFi interactions are transparent and can be traced through blockchain analytics, but they are pseudonymous: an address supplying \$100 million in liquidity to a lending protocol may be controlled by a regulated institution, a sanctioned entity, or an automated smart contract acting on behalf of unknown depositors. Off-chain CeFi exposures bilateral loans, OTC trading positions, and custodial relationships are almost entirely opaque, as the insolvencies of Three Arrows Capital, Celsius, and FTX revealed when counterparties discovered exposures they had not known existed (Jalan and Matkovskyy, 2023). No existing contagion model integrates the partially observable on-chain network with the fully opaque off-chain network into a unified analytical framework.

The composability of DeFi protocols creates dependency structures with no analog in traditional finance. When a yield aggregator deposits funds into a lending protocol that uses those funds as collateral to borrow from a second protocol, which in turn supplies liquidity to an automated market maker, the resulting chain of dependencies spans four protocol layers (Werner et al., 2022). A vulnerability at any layer propagates mechanically to all dependent layers at blockchain speed, with no possibility of human intervention, negotiated forbearance, or coordinated resolution. Traditional fire-sale and contagion models assume that transmission involves human decision-making portfolio managers deciding to sell, risk committees deciding to cut exposure, central banks deciding to intervene (Xu and Vadgama, 2023; Financial Stability Board, 2023b). DeFi transmission is algorithmic: smart contracts execute liquidations, rebalance pools, and adjust collateral requirements automatically, compressing the timeline of crisis propagation from days or weeks to minutes or seconds.

Oracle dependencies introduce a category of single-point-of-failure risk absent from traditional finance. DeFi protocols that require external price information virtually all lending protocols and derivatives platforms depend on oracle networks such as Chainlink to deliver accurate and timely price feeds (Xu and Vadgama, 2023). A compromised, manipulated, or delayed oracle feed can trigger incorrect liquidations, enable economic exploits, or cause cascading failures across all protocols that consume the same feed. Traditional financial markets also depend on price information, but the diversity of data sources, the regulatory oversight of benchmark administrators following the LIBOR scandal, and the involvement

of human intermediaries in price discovery create a more resilient information infrastructure than the concentrated oracle networks on which DeFi relies (Werner et al., 2022).

The 24/7 global operation of digital finance markets eliminates a feature that traditional financial systems rely upon implicitly: time. When equity markets close overnight and on weekends, stressed institutions gain hours or days to raise capital, negotiate credit lines, or coordinate resolution strategies. Central banks and regulators use these pauses to organize emergency interventions, as the Federal Reserve’s weekend negotiations during the Bear Stearns and Lehman Brothers crises illustrate. Digital finance offers no such respite. The Terra/Luna collapse unfolded over a single weekend, and the FTX bank run played out in approximately 48 hours (Jalan and Matkovskyy, 2023), with customer assets frozen before any regulatory authority could intervene. Liquidity spiral models calibrated to traditional market hours systematically underestimate the speed at which feedback loops operate in continuous markets, because the models assume that margin calls trigger portfolio rebalancing over days rather than triggering automated liquidation within seconds (Klages-Mundt and Minca, 2022).

The measurement infrastructure assumed by quantitative systemic risk tools does not exist in digital finance. CoVaR, MES, and SRISK require time-series data on institution-level returns or balance-sheet quantities reported at regular intervals (Adrian and Brunnermeier, 2016; Acharya et al., 2017). Most CeFi entities do not produce audited financial statements, and many have no regulatory obligation to report positions or exposures (Auer et al., 2023). DeFi protocols do not have “returns” in the conventional sense; their risk is embedded in smart contract logic, parameter configurations, and governance token distributions (Gudgeon et al., 2020). Applying traditional systemic risk measures to digital finance requires either mapping digital-native concepts onto traditional categories a process that inevitably loses critical information about the mechanisms that matter most or developing new measurement approaches tailored to the structural properties of blockchain-based systems.

These gaps are not merely quantitative they do not reflect a shortage of data that better measurement could resolve but structural. Five qualitative differences distinguish digital finance from the environment for which canonical models were designed: the speed of propagation (hours rather than weeks), the automation of critical functions (smart contracts rather than human judgment), composability (permissionless dependency chains rather than negotiated bilateral relationships), the transparency-opacity hybrid (observable on-chain data combined with opaque off-chain positions), and cross-jurisdictional operation (global and borderless rather than nationally regulated) (Financial Stability Board, 2023b; Werner et al., 2022). Each difference represents a qualitative departure from the assumptions of existing models, not merely a parameter shift that recalibration could address.

The taxonomy presented in this paper provides the conceptual foundation for addressing these gaps. By identifying the specific transmission mechanism of each channel, the assumptions from traditional theory that each channel violates, and the digital-native features that each channel introduces, we clarify what precisely needs to be measured and modeled (Xu and Vadgama, 2023). The result is not a call to abandon traditional systemic risk theory but a systematic map of where that theory must be extended, adapted, or replaced to account for the structural realities of a financial system built on permissionless, algorithmically governed, continuously operating blockchain infrastructure.

3 Methodology

The taxonomy presented in this paper 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 of the academic and policy literature to identify candidate channels and assess their coverage, followed by a scoring and selection process that narrows the initial set to the eight channels meeting inclusion thresholds. The hybrid design reflects the nature of the field: digital finance systemic risk is too young for purely data-driven approaches (the major crises span only a decade), too complex for purely theoretical derivation (the mechanisms involve technological features that theory has not yet formalized), and too consequential for purely narrative synthesis. We describe each stage in detail sufficient for replication.

3.1 Systematic Literature Search

The literature search was conducted using OpenAlex, a comprehensive open-access index of scholarly publications covering more than 250 million works across all disciplines, 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. We selected OpenAlex as the primary database because it provides broader coverage of working papers, preprints, and policy documents than Scopus or Web of Science alone. This breadth is critical for a rapidly evolving field in which much of the relevant research appears in working paper series, policy discussion notes, and conference proceedings outside traditional journal outlets. Supplementing OpenAlex with targeted policy repository searches ensured coverage of institutional reports that are often not indexed in academic databases despite their influence on the field. While the search and screening process follows the spirit of PRISMA guidelines for systematic reviews, this study departs from strict PRISMA compliance in two respects: the hybrid nature of the methodology (combining systematic search with theoretical derivation) and the inclusion of policy reports and working papers that fall outside conventional PRISMA-eligible databases. A PRISMA-style flow diagram documenting article counts at each screening stage is available from the author upon request.

The search strategy 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, 2022a](#); [International Monetary Fund, 2021](#)), and the authors' analysis of twenty-five major crisis episodes documented in the crisis chronology database. Each source contributed a distinct perspective: the traditional finance literature identified theoretical mechanisms with potential digital analogs, the policy reports identified practitioner-observed risk categories, and the crisis chronology identified empirically activated channels.

The fourteen candidate channels were: (1) network contagion, (2) liquidity spirals, (3) stablecoin runs, (4) composability risk, (5) liquidation cascades, (6) counterparty concentration, (7) information asymmetry, (8) gateway risk, (9) oracle manipulation, (10) regulatory contagion, (11) governance failure, (12) real-world asset transmission, (13) cross-chain bridge

vulnerability, and (14) validator concentration. For each candidate, we constructed three to four search queries using Boolean combinations of channel-specific terms and general systemic risk vocabulary. Representative queries include: (“systemic risk” OR “financial contagion” OR “financial stability”) AND (“DeFi” OR “decentralized finance” OR “cryptocurrency”) AND (“network” OR “interconnectedness” OR “contagion”) for the network contagion channel, and (“stablecoin” OR “algorithmic stablecoin” OR “USDT” OR “USDC” OR “Terra”) AND (“bank run” OR “de-peg” OR “run dynamics” OR “redemption”) for the stablecoin runs channel.

The search was restricted to works published between 2009 and 2026, capturing the period from the Mt. Gox collapse the first major systemic event in cryptocurrency markets through the most recent available literature. No language restriction was imposed, though the vast majority of retrieved works were in English. The initial search retrieved a combined total of approximately 2,400 unique works across all fourteen channels after deduplication. Each work was screened for relevance based on title and abstract, applying inclusion criteria that required the work to (a) address systemic risk, financial stability, or contagion in the context of digital assets, DeFi, CeFi, or stablecoins, and (b) present original analysis theoretical, empirical, or regulatory rather than solely descriptive or journalistic content.

Works that addressed individual protocol security vulnerabilities (such as reentrancy bugs or flash loan exploits) without connecting them to system-level risk were excluded, as were works focused exclusively on monetary policy implications of CBDCs without addressing financial stability channels. After screening, approximately 620 works met the inclusion criteria and were retained for full-text analysis and channel mapping. This retention rate of approximately 25% reflects the stringency of the inclusion criteria and the breadth of the initial search, which deliberately cast a wide net to avoid omitting relevant work published in non-traditional outlets.

Each retained work was coded against the fourteen candidate channels, with a work assigned to a channel if it analyzed the channel’s transmission mechanism, provided empirical evidence of the channel’s activation, or developed theoretical models relevant to the channel. Many works were assigned to multiple channels, reflecting the interconnected nature of systemic risk analysis in digital finance (Aramonte et al., 2022). The coding yielded a coverage matrix showing the number of works addressing each channel, the distribution of works by publication year, and the citation impact measured as the mean citation count of the ten most-cited works within each channel cluster. The coverage matrix showed compressed variation across channels due to the per-channel retrieval limit of 200 works: ten of fourteen channels returned the maximum count, yielding identical literature volume sub-scores. Among the screened works retained for full-text analysis, liquidity spirals, gateway risk, and stablecoin runs attracted the largest coded clusters, while oracle manipulation, regulatory contagion, and validator concentration had the smallest.

3.2 Taxonomy Construction

The transition from fourteen candidate channels to eight selected channels was governed by a two-stage process. In the first stage, a composite scoring framework was applied to the full set of approximately 2,400 works retrieved from OpenAlex prior to the manual relevance screening described above to provide a preliminary quantitative ranking of candidate chan-

nels. Because the scoring serves as a structured screening heuristic rather than a definitive ranking instrument, breadth of coverage was prioritized over precision in the input corpus. Each candidate channel was scored on three dimensions, with the following weights:

- (i) **Literature volume** (weight: 0.35): The number of works in the channel cluster, normalized to a 0–1 scale. Because the per-channel retrieval limit of 200 works was binding for ten of fourteen candidates, the literature volume sub-score provides limited discrimination: ten channels received identical normalized scores. The four channels whose retrieval exceeded the cap gateway risk (348), liquidity spirals (263), stablecoin runs (218), and oracle manipulation (208) are the only candidates differentiated by this dimension.
- (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. Citation impact serves as a proxy for the degree to which the channel has attracted sustained scholarly attention and influenced subsequent research, as opposed to generating a large volume of descriptive commentary.
- (iii) **Crisis evidence** (weight: 0.30): The number of distinct crisis episodes in the chronology database that activated the channel, weighted by the logarithm of estimated losses associated with each episode. The logarithmic transformation compresses the loss range across events, reducing the dominance of extreme-loss episodes such as the Terra/Luna collapse. Channels with fewer crisis activations receive proportionally lower scores through normalization.

The citation impact dimension is subject to a known limitation arising from the search pipeline. Because OpenAlex queries use Boolean keyword matching and results are sorted by citation count, some retrieved works are topically unrelated to the target channel high-citation papers from adjacent fields inflate the citation impact score for channels whose query terms have broad semantic overlap with other disciplines. A relevance-filtered robustness check, reported in Table 1, recomputes citation impact using only works whose titles contain digital-finance-specific terminology. The filtered analysis confirms that all qualitative retention decisions remain stable: the same eight channels survive selection, and the same six candidates are merged or deferred.

The weights (0.35, 0.35, 0.30) reflect an analytical judgment that scholarly coverage and empirical grounding should carry roughly equal influence, with a slight premium for the two literature-based dimensions because crisis evidence is constrained by the small number of major episodes in the field’s brief history. We do not claim these weights are uniquely optimal; they provide one defensible starting point among many.

We evaluated the sensitivity of channel selection to alternative weight configurations. Under equal weighting (0.33/0.33/0.33), crisis-evidence-dominant weighting (0.25/0.25/0.50), and literature-dominant weighting (0.50/0.25/0.25), all qualitative retention decisions remain invariant: the same eight channels are retained and the same six candidates are merged or deferred. Under relevance-filtered scoring (Table 1), individual rank positions shift substantially bridge vulnerability drops from rank 1 to rank 4, and validator concentration from rank 6 to rank 12 but the qualitative selection is unchanged because it was determined by

mechanism distinctiveness and theoretical grounding, not by rank order. The robustness of the selection to reasonable weight perturbations supports the stability of the taxonomy, though a formal grid-search sensitivity analysis varying weights systematically would further strengthen this claim.

The composite score for each candidate channel was computed as the weighted sum of the three normalized dimension scores. The score served as a structured screening input rather than an algorithmic cutoff. Four lower-scoring candidates fell below the inclusion threshold: oracle manipulation, regulatory contagion, governance failure, and real-world asset transmission. Two additional candidates cross-chain bridge vulnerability (rank 1, composite 0.71) and validator concentration (rank 6, composite 0.56) scored well above the threshold but were nevertheless 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, and validator concentration is a subcase of counterparty concentration ([Bank for International Settlements, 2024](#)).

Oracle manipulation was merged into composability risk on the grounds that oracle dependencies represent a specific instance of the broader phenomenon of permissionless protocol interdependence: when a lending protocol depends on an external price feed, the systemic risk arises from the composed interaction between the protocol and the oracle, not from oracle failure in isolation. Cross-chain bridge vulnerability was similarly merged into composability risk as a specific form of cross-protocol dependency that inherits the general framework. Governance failure was split between composability risk (for protocol-level governance embedded in smart contract parameters) and counterparty concentration (for governance centralization in nominally decentralized systems). Regulatory contagion was distributed between gateway risk (regulatory closure of gateway institutions) and counterparty concentration (regulatory shutdown of concentrated entities), on the grounds that regulatory action is an exogenous trigger activating existing channels rather than a distinct transmission mechanism. Validator concentration was merged into counterparty concentration as infrastructure-layer concentration. Real-world asset transmission was deferred to Section 7 because the tokenized RWA market lacked sufficient scale and crisis evidence as of the analysis cutoff ([Bank for International Settlements, 2024](#)).

We acknowledge partial conceptual overlap between network contagion and counterparty concentration, as both involve transmission through counterparty relationships. The channels are distinguished by their transmission mechanism: network contagion describes the propagation pathway (how shocks travel through the network), while counterparty concentration describes a structural vulnerability (the degree to which critical functions or exposures are concentrated in a small number of entities). A network can exhibit contagion without concentration, and concentration can exist without active contagion propagation. An analogous distinction holds for liquidity spirals and liquidation cascades: the former describes the macroeconomic feedback between funding and market liquidity, while the latter describes the mechanical protocol-level process of collateral seizure and forced selling.

The eight surviving channels were further classified along two dimensions. First, each channel was designated as *novel* (no meaningful traditional-finance analog), *extension* (a direct application of an established traditional-finance theory to digital settings), or *hybrid* (rooted in traditional theory but with digital-native features that substantially transform the mechanism). Two channels composability risk and gateway risk were classified as novel be-

Table 1: Robustness of composite scores to relevance-filtered corpus. The filtered corpus retains only papers whose titles contain digital-finance-specific terminology, removing 40% of the raw OpenAlex retrieval. The retention decision column indicates whether the qualitative selection changes. Boldface indicates channels whose rank shifts by three or more positions.

Channel	Unfiltered Composite	Filtered Composite	Unfilt. Rank	Filt. Rank	Decision Stable?
Bridge Vuln.	0.71	0.56	1	4	Yes (merged)
Liquidity Spirals	0.67	0.73	2	2	Yes
Gateway Risk	0.67	0.75	3	1	Yes
Composability Risk	0.63	0.47	4	6	Yes
Counterparty Conc.	0.58	0.67	5	3	Yes
Validator Conc.	0.56	0.30	6	12	Yes (merged)
Info. Asymmetry	0.45	0.52	7	5	Yes
Liquidation Cascades	0.38	0.28	8	13	Yes (theory)
Network Contagion	0.38	0.36	9	8	Yes (theory)
Governance Failure	0.35	0.38	10	7	Yes (split)
Stablecoin Runs	0.30	0.35	11	9	Yes (theory)
Oracle Manipulation	0.29	0.32	12	10	Yes (merged)
Regulatory Contagion	0.29	0.32	13	11	Yes (dist.)
RWA Transmission	0.28	0.24	14	14	Yes (deferred)

cause neither has a meaningful precedent in traditional financial theory: composability risk arises from a technological feature (permissionless smart contract interaction) absent from traditional systems, and gateway risk exists only at the boundary between two financial systems. Two channels network contagion and counterparty concentration were classified as extensions because they apply established theories (Allen-Gale contagion, too-big-to-fail concentration) to new institutional settings without requiring fundamental theoretical revision. Four channels liquidity spirals, stablecoin runs, liquidation cascades, and information asymmetry were classified as hybrids because each is rooted in canonical theory but incorporates digital-native features (automated execution, reflexive token dynamics, smart contract opacity) that substantially transform the underlying mechanism (Brunnermeier and Pedersen, 2009; Diamond and Dybvig, 1983; Akerlof, 1970). Second, each channel was mapped against the four domains (DeFi, CeFi, Stablecoins, Tokenized TradFi) to identify its cross-domain footprint.

The taxonomy’s organizing principle classification by transmission mechanism was selected over three alternatives after systematic evaluation. A domain-based taxonomy (organizing by DeFi, CeFi, Stablecoins, and Tokenized TradFi) was rejected because most channels operate across multiple domains, which would require extensive cross-referencing and repetition (Financial Stability Board, 2022a). An asset-type taxonomy (organizing by tokens, stablecoins, NFTs, and derivatives) was rejected because the transmission mechanism, not the asset, determines the systemic dynamics: a liquidity spiral operates similarly whether the underlying asset is ETH, a governance token, or a stablecoin LP token. A chronological taxonomy (organizing by the order of channel emergence) was rejected because the channels are not sequential; most coexist and interact simultaneously. The mechanism-based ap-

proach ensures that each channel entry stands alone as a self-contained analytical unit while enabling the cross-channel interaction analysis that is central to this paper’s contribution.

3.3 Conceptual Propositions

The taxonomy’s analytical framework generates four conceptual propositions with testable empirical implications. These propositions formalize the paper’s principal theoretical claims and provide a structured basis for future empirical validation. They are framed as conceptual propositions appropriate to a taxonomy paper rather than as mathematical theorems, reflecting the paper’s contribution as a systematic identification and classification of transmission channels rather than a formal equilibrium model.

Proposition 1 (Cross-Channel Amplification). When two or more systemic risk channels activate simultaneously in the digital finance ecosystem, the aggregate systemic impact exceeds the sum of individual channel effects. Specifically, channels connected by amplifying feedback loops as identified in the cross-channel interaction analysis of Section 5 produce superlinear loss accumulation.

Testable implication: Crisis episodes activating three or more channels simultaneously should exhibit losses disproportionately larger than single-channel events of comparable initial shock magnitude. The 2022 cascade sequence, in which all eight channels activated in temporal succession, provides a candidate test case: total losses exceeded \$2 trillion in market capitalization, vastly exceeding the initial Terra/UST de-peg shock of approximately \$18 billion.

Proposition 2 (Digital-Native Acceleration). The structural features of digital finance automated execution, 24/7 operation, permissionless participation, and on-chain transparency 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.

Testable implication: Measuring time-to-contagion from initial shock to second-order counterparty losses should show DeFi-mediated propagation operating in hours to days versus weeks to months for comparable traditional finance episodes. The contrast between the Bear Stearns rescue (organized over a weekend in March 2008) and the Celsius withdrawal freeze (executed within hours in June 2022) provides an illustrative benchmark.

Proposition 3 (Dual-Layer Opacity). The coexistence of transparent on-chain and opaque off-chain network layers in digital finance 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.

Testable implication: In crisis episodes involving both DeFi and CeFi entities, informed on-chain actors should withdraw from exposed protocols before the broader market reacts, measurable through on-chain withdrawal timing relative to public news. The Terra/Luna and FTX episodes offer candidate data: large wallet withdrawals preceded public announcements by hours in both cases (Briola et al., 2023).

Proposition 4 (Gateway Fragility). The systemic fragility introduced by fiat-crypto gateway concentration increases nonlinearly with gateway bank market share, such that the failure of a gateway bank serving above a critical market-share threshold triggers cascading effects across both digital and traditional financial systems.

Testable implication: The Silvergate and Signature Bank closures of early 2023 should have produced measurably larger stablecoin de-pegging and crypto market disruption than would be predicted by the banks’ balance-sheet size alone, with effects proportional to their crypto-deposit market share rather than their total asset size (Galati and Capalbo, 2023).

These four propositions are testable through future empirical work employing on-chain data analysis, event-study methodology, and cross-crisis comparison. The propositions derive directly from the taxonomy’s analytical structure and provide the conceptual architecture within which formal quantitative models can be developed. We identify the empirical testing of these propositions as a primary avenue for future research in Section 9.

3.4 Limitations and Methodological Considerations

Several limitations of the methodology should be acknowledged. First, channel classification was executed primarily by the first author, with review and validation by the co-authors, introducing the possibility of systematic coding bias in the assignment of works to channels and in the classification of channels as novel, hybrid, or extension. We mitigate this limitation through three mechanisms: the use of a transparent, documented scoring framework with explicit weights and thresholds that enables replication; the grounding of channel classification in crisis case evidence that provides external validation independent of individual researchers’ judgment; and the explicit mapping of each channel to established theoretical categories in the traditional finance literature, which constrains classification to recognized conceptual boundaries. Nonetheless, future work should include independent validation through inter-rater reliability testing, in which fully independent coders classify a sample of works and channels to assess the degree of coding agreement.

Second, the composite scoring framework assigns weights of 0.35 to literature volume, 0.35 to citation impact, and 0.30 to crisis evidence. These weights are analytical judgments, not theoretically derived quantities. The sensitivity analysis reported above demonstrates that the qualitative channel selection is invariant to the four alternative weighting schemes tested, but we acknowledge that no objective criterion exists for calibrating weights in a multi-criteria screening instrument of this type. A formal Analytic Hierarchy Process or expert-elicitation protocol would strengthen the methodological foundation.

Third, the taxonomy prioritizes comprehensive coverage of transmission channels and their interactions over deep formal analysis of any single channel. This breadth-versus-depth design choice reflects the paper’s primary contribution mapping the complete channel structure of digital finance systemic risk, an integration that no prior work has attempted and the observation that the literature’s fragmentation across channel-specific subfields is itself a barrier to understanding systemic dynamics. The tradeoff is that individual channels receive less analytical depth than a dedicated single-channel study would provide. We view the taxonomy as a foundation for subsequent focused studies that can build on the integrated framework to develop formal models, empirical tests, and regulatory applications for specific channels.

Fourth, the taxonomy is grounded in crisis evidence through early 2025. The rapid evolution of digital finance including the growth of tokenized real-world assets, the deployment of AI-driven trading strategies, and the shifting regulatory landscape means that new channels may emerge and existing channel dynamics may shift as the ecosystem matures. The forward-looking assessment in Section 7 identifies prospective channels, but the taxonomy’s empirical grounding is necessarily retrospective.

A further data-collection limitation arises from the search pipeline’s structural features. The per-channel retrieval limit of 200 works, combined with sorting by citation count, means that the literature volume dimension reflects the most-cited subset of the literature for each channel rather than the complete corpus. For the ten channels where the retrieval limit was binding, the literature volume sub-score is identical, substantially reducing the discriminating power of this 35%-weighted dimension. The citation-count sorting also introduces a temporal bias favoring established work over recent publications with fewer accumulated citations. 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. This structural disadvantage reflects the genuine absence of documented crises in this domain as of the analysis cutoff rather than a methodological shortcoming, but it constrains the scoring framework’s ability to distinguish between low-risk channels and those where risk has not yet materialized.

A related limitation concerns the relevance of the retrieved literature corpus. The Boolean keyword queries used to search OpenAlex necessarily retrieve some works that are topically unrelated to digital finance systemic risk papers from other disciplines that happen to match general terms in the query. Because the scoring pipeline operates on the full retrieval rather than the manually screened subset, the citation impact dimension is influenced by off-topic high-citation papers, particularly for channels whose query terms have broad semantic overlap with other fields (such as ‘bridge vulnerability,’ where the term ‘bridge’ matches civil engineering, bioinformatics, and networking literature). Table 1 reports a relevance-filtered robustness check that recomputes scores after removing papers lacking digital-finance-specific terminology from their titles. The filtered analysis confirms that the qualitative selection is stable, though individual rank positions shift substantially for the most affected channels.

The reproducibility of the literature corpus is conditioned on the OpenAlex API state at the time of retrieval. The API does not guarantee stable result ordering for queries with tied citation counts, and the underlying database is updated continuously. While the search parameters, queries, and scoring code are fully documented and the scoring pipeline is deterministic given fixed inputs, the retrieved corpus is archived as a static dataset to ensure reproducibility of downstream analyses.

The crisis evidence dimension is influenced by event multiplicity: the Terra/Luna collapse, the single largest value destruction event in the chronology, activates five channels simultaneously and contributes disproportionate weight. The logarithmic loss weighting compresses the range across events but does not eliminate cross-channel correlation in crisis activation, which is an inherent feature of interconnected financial crises. The literature volume similarly counts each paper in every channel to which it was assigned during the coding process, meaning that broadly relevant works contribute to multiple channels’ volume scores. This multi-channel assignment reflects the cross-cutting nature of systemic risk scholarship and is methodologically intentional, but it means that individual channel volume scores are

not statistically independent.

Threats to Validity. Four categories of validity threat warrant explicit discussion. *Selection bias*: the literature search is limited to English-language works indexed in OpenAlex and supplementary policy repositories maintained by major international institutions. Works published in non-English languages or in venues not indexed by OpenAlex may be underrepresented, though the global nature of digital finance scholarship and the dominance of English-language publication in the field mitigate this concern. *Publication bias*: the published literature may overrepresent channels with abundant empirical evidence and established theoretical frameworks, and underrepresent theoretically important but empirically underdocumented channels. The inclusion of crisis evidence as a scoring dimension partially offsets this bias by rewarding channels with real-world manifestation regardless of publication volume. *Database coverage bias*: OpenAlex provides broader coverage than Scopus or Web of Science for working papers and preprints but may underrepresent certain regional sources and non-English policy documents. The supplementary repository searches mitigate but do not eliminate this limitation. *Temporal bias*: the literature search window of 2009–2026 captures the major crisis episodes but excludes earlier theoretical work on financial networks that predates cryptocurrency. We address this gap through targeted inclusion of seminal traditional finance works (Allen and Gale, 2000; Diamond and Dybvig, 1983; Brunnermeier and Pedersen, 2009; Akerlof, 1970) identified through backward citation tracing, but systematic coverage of pre-2014 network science and financial stability literature is beyond the scope of the search protocol.

4 A Taxonomy of Systemic Risk Channels in Digital Finance

The taxonomy identifies eight channels through which systemic risk transmits in digital finance. Each channel is analyzed through a consistent framework: a formal definition, a detailed description of the transmission mechanism, an assessment of theoretical foundations and their traditional-finance analogs, an identification of digital-native features that distinguish the channel from its traditional counterpart, case evidence from specific crisis episodes, an analysis of cross-domain manifestation, and an assessment of the existing literature and remaining gaps. The channels are ordered by their relationship to the traditional finance literature, beginning with those closest to established theory and proceeding to those that are most digitally native.

4.1 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.

4.1.1 Mechanism

Distress at one node in the digital finance network whether an exchange, a lending protocol, a hedge fund, or a stablecoin issuer transmits losses to connected counterparties through three primary pathways (Allen and Gale, 2000; Acemoglu et al., 2015; Elliott et al., 2014). Direct exposure contagion occurs when a failing entity defaults on obligations owed to creditors, imposing losses on counterparties that hold the failing entity’s debt or have locked collateral with it. Indirect exposure contagion operates through shared asset holdings: when two entities hold significant positions in the same token, a fire sale by one depresses the token price and impairs the other’s balance sheet regardless of any bilateral relationship. Informational contagion arises when one entity’s failure triggers panic-driven withdrawals at connected or seemingly similar platforms, even absent direct financial exposure a pattern observed when the FTX collapse precipitated withdrawal surges at exchanges with no direct FTX exposure.

In the DeFi layer, the network topology is partially observable through on-chain data: token flows between protocols, shared liquidity pool participation, and the dependency graph of composed smart contracts can all be reconstructed from blockchain records (Schär, 2021). Protocols that deposit assets into other protocols a common pattern in yield optimization inherit the credit risk of the downstream protocol, creating contagion edges determined by smart contract logic rather than by bilateral negotiation. Smart contract vulnerabilities compound this mechanical linkage, with attacks across the DeFi ecosystem causing an estimated \$6.45 billion in financial losses (Xu and Vadgama, 2023).

In the CeFi layer, the network topology remains largely opaque. Bilateral OTC lending, prime brokerage relationships, and custodial arrangements lack public disclosure, and the absence of centralized clearing means no single entity holds a complete view of the exposure network (Jalan and Matkovskyy, 2023). The coexistence of a transparent on-chain layer and an opaque off-chain layer creates a dual-network structure with no close analog in traditional finance, where bilateral exposures are at least partially captured by regulatory reporting.

The speed of contagion propagation in digital finance significantly exceeds that of traditional financial networks. On-chain contagion operates at blockchain speed: when a lending protocol’s collateral is impaired by an upstream failure, liquidations execute within the same block or the next, propagating losses downstream in seconds rather than the days or weeks typical of traditional interbank settlement. CeFi contagion, while involving human decision-making, is accelerated by 24/7 market operation, social media information dissemination, and the absence of circuit breakers or trading halts (Younis et al., 2024; Cookson et al., 2023; Hanif et al., 2023). The Terra/Luna collapse demonstrated that cross-entity contagion could propagate from an initial stablecoin de-peg to the insolvency of a major hedge fund within approximately three weeks a timeline that compressed months of traditional financial contagion into a fraction of the period.

4.1.2 Theoretical Foundations

Network contagion is the most extensively theorized systemic risk channel in traditional finance. Allen and Gale (2000) established the foundational framework by modeling contagion as the propagation of liquidity shocks through interbank claims networks, demonstrating that systemic vulnerability depends on network structure rather than shock severity alone.

Gai and Kapadia (2010) extended this analysis to show that contagion propagation depends on network topology, with system-wide crises emerging from localized shocks when network connectivity exceeds critical thresholds. Acemoglu et al. (2015) generalized this result by proving a phase-transition property: below a critical connectivity threshold, additional interbank connections diversify risk and enhance stability; above it, the same connections amplify instability. Elliott et al. (2014) extended the framework to cross-ownership, showing that diversification through common asset holdings creates hidden contagion channels that bypass direct bilateral exposures. Building on these models, Battiston et al. (2012) introduced the DebtRank framework, which quantifies the systemic importance of individual nodes based on their network position and the propagation of distress through interbank obligations. The analogy between financial networks and ecological systems, as developed by Haldane and May (2011), suggests that the same network features promoting efficiency in stable conditions amplify fragility during stress an insight with particular resonance for the highly interconnected digital finance ecosystem.

These models apply to digital finance with important modifications. The nodes are not banks but protocols, exchanges, custodians, and hedge funds. The edges include smart contract dependencies, shared liquidity pool participation, and common token holdings alongside bilateral lending (Schär, 2021). The key theoretical extension required is incorporation of the dual-layer network: on-chain edges are observable and deterministic, while off-chain edges are opaque and subject to information asymmetry a structure that amplifies the mechanisms Brunnermeier (2009) identified in the 2007–2008 crisis, where opacity in structured products prevented counterparties from assessing aggregate exposure.

The Acemoglu et al. (2015) phase-transition result carries particular implications for digital finance: the crypto ecosystem’s high and growing interconnectedness documented empirically through volatility spillover analyses (Fang et al., 2022; Hanif et al., 2023; Sakariyahu et al., 2024) suggests that the network may already operate above the critical threshold where additional connections amplify rather than absorb shocks. We identify a significant gap in the theoretical literature: to our knowledge, no existing model captures both the on-chain and off-chain layers in a unified contagion framework. On-chain contagion is well-suited to existing network models, but off-chain CeFi contagion requires exposure information unavailable to researchers and regulators. A complete model must account for both layers simultaneously, including the strategic behavior of informed agents who observe the on-chain layer and attempt to infer the off-chain structure. The absence of such a model limits both academic understanding and regulatory capacity to assess systemic fragility in the digital finance ecosystem.

4.1.3 Digital-Native Features

Three features distinguish digital finance network contagion from its traditional analog. First, on-chain observability creates an asymmetric information environment: sophisticated actors with blockchain analytics can observe the on-chain network in real time, enabling them to front-run contagion by withdrawing from exposed protocols before less informed participants react. Daian et al. (2020) document that arbitrage bots in decentralized exchanges exploit transaction-ordering dependencies and engage in priority gas auctions, creating systemic risks at the consensus layer (Fang et al., 2022).

Second, pseudonymous participation obscures the mapping between on-chain addresses and real-world entities. A single entity may control thousands of addresses across multiple blockchains, and multiple entities may share a multisignature wallet. Constructing the true exposure network from observed on-chain data requires imperfect de-anonymization techniques that are computationally intensive and inherently incomplete (Sakariyahu et al., 2024; Antonakakis et al., 2020). The Three Arrows Capital episode demonstrated this opacity: the fund’s aggregate leverage and counterparty exposure remained invisible until default revealed the network.

Third, composability creates mechanical contagion edges that have no traditional analog. When protocol A holds tokens issued by protocol B, a failure at B automatically impairs A without any decision-making or information processing the contagion is encoded in smart contract logic itself (Werner et al., 2022). These dependency chains can be arbitrarily deep and are assembled permissionlessly, meaning no single entity has visibility into the full dependency graph. Paradoxically, the transparency of DeFi may accelerate contagion by enabling informed actors to act preemptively, transforming potential losses into realized losses before the distressed entity has any opportunity to restructure.

4.1.4 Case Evidence

The Three Arrows Capital (3AC) insolvency of June 2022 provides the clearest illustration of network contagion in digital finance. The hedge fund accumulated concentrated positions in Terra/Luna and staked Ether through bilateral OTC lending with more than twenty counterparties, including Genesis Trading, BlockFi, Voyager Digital, and Celsius Network (Jalan and Matkovskyy, 2023). When Terra’s collapse inflicted portfolio losses, 3AC could not meet margin calls. Because these relationships were bilateral and opaque, counterparties discovered their exposure only as 3AC failed to meet obligations (Jalan and Matkovskyy, 2023). The aggregate counterparty losses from 3AC’s default exceeded \$3 billion across more than twenty creditors, with exposure concentrated among a small number of CeFi lenders. Among the largest individual exposures, Voyager Digital disclosed a \$650 million unsecured loan to 3AC and filed for bankruptcy in July 2022; Genesis Trading, a separate counterparty, absorbed over \$1 billion in losses from its own distinct exposure to 3AC. These figures represent different creditors’ individual losses, not conflicting estimates of a single exposure. The cascade demonstrated that the CeFi lending network exhibits the same vulnerability to concentrated counterparty failure that Allen and Gale (2000) identified in the banking context.

The FTX collapse of November 2022 further illustrated contagion operating across both CeFi and DeFi layers simultaneously. Distress propagated through three distinct channels: direct counterparty exposures (firms with assets custodied on FTX or outstanding loans to Alameda Research), informational channels (withdrawal runs on exchanges with no direct FTX exposure), and price channels (the collapse of the FTT token and associated cross-market selling pressure) (Younis et al., 2024). Social media accelerated the informational channel, compressing the timeline from initial revelation to full-scale withdrawal run into approximately 72 hours far faster than any comparable episode in traditional finance (Cookson et al., 2023). The contagion ultimately reached the traditional banking layer when Silvergate Bank, which served as a primary banking partner for multiple crypto firms, experienced de-

posit outflows that contributed to its closure in March 2023, illustrating how digital finance network contagion can breach the boundary into the regulated financial system.

4.1.5 Cross-Domain Manifestation

Network contagion operates across all four domains of digital finance but manifests differently in each. In DeFi, contagion propagates mechanically through smart contract dependencies and shared liquidity pools at blockchain speed, with the ecosystem’s rapid growth amplifying the attack surface for cascading failures (Fang et al., 2022). In CeFi, it propagates through bilateral lending and custodial arrangements, with speed determined by counterparty risk management processes and information dissemination. In the stablecoin domain, contagion operates through reserve custody relationships and collateral dependencies: a stablecoin issuer’s distress transmits to all protocols and platforms that use the stablecoin as a unit of account or collateral, as the USDC de-peg during the Silicon Valley Bank failure demonstrated. In tokenized traditional finance, contagion channels are emerging as institutional investors take positions that create exposure links between traditional portfolios and digital asset markets (Andryushin, 2024; Haddad and Hornuf, 2023).

4.1.6 Literature Assessment

The theoretical foundations of network contagion are well-established in traditional finance, and an emerging empirical literature applies these frameworks to digital asset markets. Studies measuring dynamic connectedness across crypto and traditional markets document significant volatility spillovers that intensify during crisis periods, with the direction of transmission shifting from traditional-to-crypto toward bidirectional linkages as the ecosystem matures (Antonakakis et al., 2020; Hanif et al., 2023; Younis et al., 2024). These analyses confirm that digital asset markets are not isolated but are increasingly integrated into global financial networks, with the strength of integration varying across time horizons and market conditions.

Intra-ecosystem contagion where distress at one token or protocol spreads to related assets within the same blockchain ecosystem represents a distinctively digital contagion pattern. Sakariyahu et al. (2024) demonstrate significant contagion among tokens within the same ecosystem when adverse news occurs, with effects amplified by investor sentiment dynamics on social platforms. The 2022 crisis sequence confirmed these dynamics empirically, as failures cascaded from Terra/Luna through 3AC to multiple CeFi lenders within weeks (Jalan and Matkovskyy, 2023; Brunnermeier, 2009).

The scale-free topology documented in financial networks (Barabási and Albert, 1999) and the small-world properties (Watts and Strogatz, 1998) that accelerate contagion propagation carry particular implications for digital finance, where protocol interconnections form emergent network structures without central design. Glasserman and Young (2016) provide the authoritative survey of contagion mechanisms in financial networks, while Cont et al. (2013) analyze how network structure interacts with loss-given-default assumptions to determine systemic risk in banking systems. These foundational network science results inform the digital finance context but have not yet been formally adapted to the dual-layer on-chain/off-chain topology that characterizes the crypto ecosystem.

The primary literature gap concerns the integration of on-chain and off-chain network layers into a unified contagion model (Antonakakis et al., 2020). Researchers can reconstruct on-chain contagion through transparent smart contract interactions, but off-chain CeFi contagion propagating through opaque bilateral relationships requires exposure data that remains unavailable to both researchers and regulators. To our knowledge, no comprehensive methodology currently exists for mapping bilateral CeFi counterparty exposures across regulated and unregulated entities spanning multiple jurisdictions. A complete model of digital finance network contagion must account for both layers simultaneously, including the strategic behavior of informed agents who observe the on-chain layer and attempt to infer or exploit the hidden structure of the off-chain layer. Developing such a model one that integrates transparent DeFi linkages with opaque CeFi exposures under heterogeneous information represents perhaps the most pressing theoretical challenge in digital finance systemic risk research.

4.2 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 that is amplified in digital finance by automated market makers, hard-coded liquidation thresholds, and the absence of institutional market makers willing to absorb selling pressure during stress events.

4.2.1 Mechanism

The liquidity spiral mechanism operates through three mutually reinforcing dynamics. The loss spiral begins when declining asset prices reduce the net worth of leveraged participants, forcing position reductions that depress prices further. The margin spiral compounds this pressure: declining collateral values trigger higher margin requirements, forcing additional deleveraging even when positions remain above current liquidation thresholds (Brunnermeier and Pedersen, 2009). Fire-sale externalities amplify the damage, as forced sellers impose negative price externalities on all holders of the same asset (Shleifer and Vishny, 2011). The 2007–2008 crisis demonstrated how these mechanisms interact, as mortgage losses cascaded through the financial system via liquidity dry-ups and forced deleveraging (Brunnermeier, 2009). Arbitrageurs who might stabilize prices face binding funding constraints during crises, eliminating the natural stabilizing force (Gromb and Vayanos, 2010).

In decentralized finance, these spiral dynamics operate with mechanical precision encoded in smart contracts. DeFi lending protocols define each borrowing position by a loan-to-value ratio typically between 75% and 85% embedded directly in on-chain code (Schär, 2021). When a collateral asset’s price breaches the liquidation threshold, any external party can seize the collateral at a protocol-specified discount and sell it on automated market makers, whose constant-product pricing functions mechanically push the price lower in proportion to the sale size relative to pool reserves (Kirste et al., 2024). Each liquidation depresses prices further, pushing additional positions past their thresholds in a self-feeding cascade. The interconnected DeFi ecosystem where identical collateral assets appear across multiple lending protocols ensures that stress propagates rapidly across venues (Arora et al., 2024).

In centralized crypto finance, the spiral mechanism shares its structural logic with DeFi but differs in execution speed and transparency. Centralized exchanges issue margin calls requiring additional collateral within hours rather than seconds (Jalan and Matkovskyy, 2023). CeFi platforms retain discretion over whether to liquidate undercollateralized positions, and some have historically extended forbearance to large clients. The opacity of CeFi margin practices means effective system-wide leverage resists external estimation, and the true scale of forced selling becomes observable only after the fact (Bongini et al., 2025). The FTX collapse of November 2022 revealed that centralized platforms can harbor hidden leverage far exceeding what counterparties anticipated, echoing the opacity problems of the 2007–2008 crisis (Brunnermeier, 2009).

4.2.2 Theoretical Foundations

Brunnermeier and Pedersen (2009) provide the canonical theoretical framework by modeling the interaction between market liquidity the ease of trading without price impact and funding liquidity the ease of financing a position. When funding liquidity tightens, leveraged intermediaries reduce market-making and liquidate positions, depressing market liquidity. Lower market liquidity increases collateral haircuts, tightening funding further and generating multiple equilibria: a stable state with ample liquidity and a crisis state where the feedback loop drives liquidity toward zero. Shleifer and Vishny (2011) complement this by demonstrating that fire-sale externalities impose welfare costs extending beyond direct participants, as depressed prices trigger cascading forced sales. Gromb and Vayanos (2010) show that arbitrageurs face binding funding constraints during crises, and Hautsch et al. (2024) demonstrate that blockchain settlement latency creates additional limits to arbitrage by exposing cross-exchange traders to price risk during the consensus period.

The leverage cycle documented by Adrian and Shin (2010) in which financial intermediaries’ balance sheet constraints amplify asset price movements operates with particular intensity in DeFi, where leverage ratios adjust automatically through smart contract parameters rather than through discretionary risk management decisions. The amplification mechanisms analyzed by Krishnamurthy (2010) where initial losses reduce intermediary capital, tighten funding conditions, and force further asset sales manifest in DeFi through the mechanical interaction of collateral ratios and liquidation thresholds, compressing what were multi-week adjustment processes in traditional finance into hours or minutes.

Digital finance provides a natural laboratory for testing and extending these models because on-chain DeFi markets allow researchers to observe liquidation events, AMM trading volumes, and liquidity withdrawals at transaction-level granularity (Klages-Mundt and Minca, 2022). The key theoretical extension required for DeFi is the incorporation of algorithmic execution: unlike traditional human intermediaries, DeFi liquidation bots execute instantaneously and without discretion, compressing spiral dynamics from days to minutes. The AMM bonding curve introduces a deterministic relationship between trade size and price impact that replaces stochastic liquidity assumptions (Schär, 2021). The game-theoretic dynamics of competing liquidation bots who engage in priority gas auctions to secure execution priority add a layer of strategic interaction absent from traditional fire-sale models (Daian et al., 2020).

4.2.3 Digital-Native Features

Four features distinguish digital finance liquidity spirals from their traditional analogs. First, on-chain liquidation eliminates the time buffer that traditional margin-call processes provide. A borrower whose collateral value declines past the liquidation threshold faces immediate seizure with no notification period, grace window, or opportunity to post additional collateral (Schär, 2021; Bakare et al., 2024). This compression transforms what was historically a multi-day process into an event measured in blockchain confirmation times typically twelve seconds on Ethereum. The mechanical and permissionless nature of DeFi liquidation means no human judgment intervenes between the price decline and the collateral seizure (Oben and Özdamlı, 2024).

Second, AMMs provide continuous but procyclical liquidity. The constant-product pricing function ensures that every sale moves the price, and larger sales produce proportionally greater impact, creating a deterministic amplification mechanism absent from traditional order-book markets where standing limit orders can absorb selling pressure (Kirste et al., 2024). Empirical analysis of decentralized exchange microstructure confirms that liquidity on these platforms exhibits distinct characteristics relative to centralized venues, with lower depth amplifying the price impact of large trades (Wątorrek et al., 2024; Nimalendran et al., 2024). MEV-extracting bots further destabilize prices during stress by front-running liquidation transactions and reordering blocks to maximize their own profit at the expense of orderly price discovery (Choi and Kim, 2024).

Third, liquidity providers can withdraw from AMM pools in real time by removing deposited tokens, unlike traditional market makers who face obligations to maintain quotes during stress. Withdrawal during volatility is rational for individual providers seeking to minimize impermanent loss, but it reduces market depth and amplifies the price impact of liquidation sales, creating a collective-action problem (Oben and Özdamlı, 2024; Bakare et al., 2024). This procyclical withdrawal intensifies when stablecoins serve as collateral, as providers simultaneously face impermanent loss risk and de-peg risk (Häfner et al., 2024). The DeFi ecosystem lacks institutional safeguards such as designated market maker obligations or trading halts that mitigate similar dynamics in traditional finance (Bongini et al., 2025).

Fourth, on-chain transparency allows all market participants to observe the distribution of collateral ratios across lending protocols, enabling liquidation bots to anticipate and prepare for cascades. This visibility, paradoxically, may accelerate spirals as bots compete to execute liquidations as rapidly as possible, engaging in priority gas auctions that bid up transaction costs and crowd out ordinary users (Daian et al., 2020). The scale of this extractive activity is substantial: blockchain extractable value from liquidation-related transactions represents a significant and growing fraction of on-chain economic activity (Qin et al., 2022). The dependence of DeFi lending protocols on external price oracles introduces an additional vulnerability, as oracle manipulation can artificially trigger liquidations even absent genuine market stress (Arora et al., 2024). These digital-native features collectively create a spiral environment that is faster, more mechanistic, and more difficult to interrupt than its traditional counterpart (Szrajber et al., 2025).

4.2.4 Case Evidence

Black Thursday, March 12, 2020, provides the canonical case study of liquidity spirals in DeFi. The COVID-19 panic triggered a roughly 50% single-day decline in Bitcoin and a corresponding crash across all crypto-asset classes. On MakerDAO, the price decline pushed thousands of collateralized debt positions below their 150% liquidation threshold, unleashing automated liquidations. Extreme Ethereum network congestion caused gas prices to spike above levels that liquidation bots were configured to pay; with bots unable to submit transactions, collateral auctions received zero bids for some lots (Klages-Mundt and Minca, 2022). The protocol incurred approximately \$8.3 million in bad debt, requiring new MKR token issuance to recapitalize the system. The episode demonstrated that DeFi liquidity spirals inherit amplification dynamics familiar from traditional crises (Brunnermeier, 2009) while adding novel failure modes tied to blockchain infrastructure constraints (Qin et al., 2022).

The Terra/Luna collapse of May 2022 exhibited liquidity spiral dynamics at vastly larger scale. As UST began to de-peg, holders exchanged UST for LUNA through the mint-and-burn mechanism, creating selling pressure on LUNA. The falling LUNA price reduced the effective backing of UST, triggering further redemptions in a reflexive spiral that drove both tokens toward zero (Santiago et al., 2024). The algorithmic collateral mechanism lacked any circuit breaker to arrest the feedback loop once initiated (Häfner et al., 2024). AMM pools containing UST experienced massive imbalances as one-sided selling drained the dollar-denominated side, reducing exit liquidity. The episode destroyed approximately \$45 billion in market capitalization within five days and transmitted volatility to conventional asset markets through shared portfolios and cross-collateralized margin calls (Bongini et al., 2025).

The Curve Finance crisis of July 2023 illustrated a near-miss liquidity spiral with systemic implications. A Vyper compiler vulnerability allowed attackers to drain several Curve pools, causing the CRV governance token to decline sharply. Curve’s founder had deposited approximately \$168 million of CRV as collateral across Aave, Fraxlend, and Abracadabra. The CRV price decline pushed these positions toward their liquidation thresholds, threatening a cascade in which liquidation sales would further depress the price across protocols (Santiago et al., 2024). The cascade was averted through over-the-counter sales negotiated directly by the founder a resolution relying on personal relationships rather than protocol design. This episode highlighted a system where a single participant’s concentrated positions can create system-wide liquidation risk, and the only effective circuit breaker is ad hoc human intervention (Oben and Özdamlı, 2024).

4.2.5 Cross-Domain Manifestation

Liquidity spirals manifest primarily in DeFi and CeFi but with structural differences that shape their systemic impact. In DeFi, the spiral mechanism is transparent, deterministic, and extremely fast, operating through publicly observable smart-contract interactions where each step can be traced on the blockchain (Szrajber et al., 2025). In CeFi, the mechanism involves human decision-making, discretionary margin calls, and opaque position information, making it slower but potentially larger in absolute scale given the higher trading volumes on centralized platforms (Nimalendran et al., 2024). The stablecoin domain experiences liquidity spirals indirectly, as the contagion channel through which liquidity stress in one

domain transmits to others via shared settlement assets and cross-listed tokens (Bakare et al., 2024). Empirical evidence from the COVID-19 crash and subsequent crises confirms that crypto-market liquidity spirals transmit to conventional asset markets through shared investor portfolios and margin-driven co-movements, blurring the boundary between digital and traditional finance stress (Cui et al., 2025).

4.2.6 Literature Assessment

The liquidity spiral mechanism is well-theorized in traditional finance, and several empirical studies have documented its manifestation in DeFi. Klages-Mundt and Minca (2022) provide a stochastic model of stablecoin depegging and liquidation spirals, while Qin et al. (2022) quantify the scale and frequency of extractable value from on-chain liquidations across major DeFi protocols. Daian et al. (2020) formalize the priority gas auction mechanism through which liquidation bots compete for execution priority, documenting a form of strategic interaction absent from traditional markets. The interaction between market liquidity and funding liquidity described by Brunnermeier and Pedersen (2009) operates with heightened intensity in DeFi, where the deterministic pricing of AMMs and the instantaneous execution of liquidation bots compress the temporal dimension of the spiral from days to minutes.

The primary literature gap concerns formal modeling of AMM-mediated liquidity spirals under stress. Traditional models assume human-mediated market-making with discretionary participation; DeFi liquidity is provided by algorithmic AMMs with deterministic pricing and by non-obligated liquidity providers who withdraw during stress (Kirste et al., 2024). A formal model incorporating the constant-product AMM curve, endogenous liquidity-provider withdrawal, and the game-theoretic dynamics of competing liquidation bots would provide the theoretical foundation for stress-testing DeFi lending protocols (Szrajber et al., 2025). Under-explored dimensions include oracle price feed interactions with liquidation triggers where manipulation can artificially induce cascades (Arora et al., 2024) and blockchain-specific limits to arbitrage that prevent stabilizing capital from entering during stress (Hautsch et al., 2024). We identify developing such integrated models as a first-order priority for future research.

4.3 Stablecoin De-Pegging and Run Dynamics

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

4.3.1 Mechanism

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). Holders possess claims on a reserve pool that, under normal conditions, suffices to honor all redemptions at par. A shock to confidence driven by doubts about reserve adequacy, concerns about

reserve asset quality, or fear that other holders will redeem first triggers self-fulfilling redemption waves (Gorton and Zhang, 2023). When reserves include illiquid or depreciated assets, forced liquidation to meet redemptions realizes losses that impair remaining claims. Ma et al. (2025) demonstrate that concentrated arbitrage structures Tether permits only six agents per month to redeem for cash create a fundamental tradeoff between secondary-market price stability and run vulnerability. The empirical literature confirms that reserve composition and redemption architecture rank among the primary determinants of stablecoin fragility (Ante et al., 2023).

For algorithmic stablecoins, the run mechanism is reflexive and fundamentally more unstable. In the Terra/Luna architecture, UST holders could redeem each UST for \$1 worth of newly minted LUNA tokens. When UST traded below its peg, redemptions created selling pressure on LUNA, which depressed LUNA’s market price, which reduced the implicit backing per UST, which deepened the de-peg and incentivized further redemptions (Uhlig, 2022; Liu et al., 2023). The reflexive loop possessed no natural floor: expanding LUNA supply continuously depressed its price, generating a “death spiral” that drove both tokens to near-zero within five days (Santiago et al., 2024). Unlike reserve-backed runs, where the floor is determined by the liquidation value of reserve assets, algorithmic runs face a theoretically unbounded collapse path because the collateral token’s supply expands endogenously. This reflexive dynamic where the stabilization mechanism itself becomes the transmission channel for collapse introduces a class of run dynamics with no precise traditional analog (Saengchote and Samphantharak, 2024).

The systemic significance of stablecoin runs arises from the settlement-layer function that stablecoins perform across digital finance. Stablecoins denominate the majority of DeFi lending, AMM liquidity, and derivatives positions, serving as the ecosystem’s primary unit of account and settlement medium (Makarov and Schoar, 2022). A de-peg propagates immediately to all protocols and positions collateralized by the affected stablecoin. AMM pools containing the de-pegging token experience imbalances as arbitrageurs drain the non-stablecoin asset, while lending protocols face collateral-value declines that trigger cascading liquidations (Kitzler et al., 2022). Nested protocol compositions amplify these effects: a stablecoin serving as collateral whose receipt tokens are deposited in a yield aggregator creates multi-layer exposure to a single de-peg event. The systemic footprint of a stablecoin run is proportional to its adoption as a settlement and collateral asset, not merely to its market capitalization (Bongini et al., 2025).

4.3.2 Theoretical Foundations

The Diamond and Dybvig (1983) model of bank runs provides the primary theoretical foundation for reserve-backed stablecoin runs. A bank investing in illiquid long-term assets and financing itself with liquid demand deposits faces a coordination problem: if each depositor believes others will withdraw, rational preemptive withdrawal creates a self-fulfilling run. Deposit insurance eliminates this equilibrium, but stablecoin issuers operate without such backstops (Gorton and Zhang, 2023). Gorton and Metrick (2012) extended this framework to shadow banking, demonstrating that maturity-transformation fragility afflicts any institution financing illiquid assets with money-like claims a characterization applying directly to stablecoin issuers holding Treasury bills and commercial paper. Ma et al. (2025) formalize

a stablecoin-specific extension, showing that concentrated arbitrage creates a fundamental tradeoff: efficient redemption improves peg stability in secondary markets but amplifies run risk by reducing investors’ price impact from selling.

Algorithmic stablecoins require a different theoretical treatment because the reflexive mint-and-burn mechanism introduces feedback dynamics that Diamond-Dybvig models do not capture. Uhlig (2022) formally analyze the Terra/Luna death spiral, demonstrating that the system possesses multiple equilibria: a stable equilibrium in which the peg holds and an unstable equilibrium in which the reflexive loop drives both tokens to zero. The transition between equilibria can be triggered by shocks that are small relative to total capitalization but sufficient to overwhelm the stabilization mechanism’s absorption capacity. Increasing system capitalization does not eliminate the unstable equilibrium it merely raises the shock threshold. Blockchain settlement latency compounds instability by limiting the speed of cross-exchange arbitrage that might dampen de-peg dynamics during a run’s critical early phase (Hautsch et al., 2024). Code-based confidence mechanisms prove insufficient under stress because they lack the external enforcement that traditional monetary systems provide (Bodó and Filippi, 2024; Saengchote and Samphantharak, 2024).

4.3.3 Digital-Native Features

Stablecoin runs differ from traditional bank runs in several respects that amplify fragility. First, on-chain redemption is instantaneous and permissionless: any holder can convert stablecoins to reserves or to newly minted volatile tokens without queuing, notification, or regulatory intervention (Ma et al., 2025). The sequential-service constraint that creates a first-mover advantage in bank runs is replaced by a technological speed advantage holders with faster smart-contract interactions and lower-latency network connections exit before others, creating a different but equally pernicious form of inequitable loss allocation (Hautsch et al., 2024). Second, the global and pseudonymous holder base renders coordination to prevent a run impractical; no single regulator holds jurisdiction over the complete holder population (Makarov and Schoar, 2022). Third, algorithmic stabilization mechanisms introduce reflexive dynamics with no traditional analog: the mint-and-burn process that maintains the peg under normal conditions becomes the death-spiral transmission channel under stress (Ante et al., 2023).

Fourth, composability means a de-peg event transmits instantaneously to every protocol using the stablecoin as collateral, unit of account, or liquidity-pool component creating systemic consequences that vastly exceed the issuer’s direct liabilities (Kitzler et al., 2022). DeFi protocols compose permissionlessly, and the absence of circuit breakers or trading halts means no institutional mechanism can pause cascading liquidations during a run (Galati and Capalbo, 2023). Fifth, stablecoin holders operate without deposit insurance, orderly resolution frameworks, or lender-of-last-resort facilities. Unlike bank depositors in jurisdictions with credible deposit guarantees, stablecoin holders have no institutional backstop to prevent a run from proceeding to completion (Bakare et al., 2024). These five features collectively compress run dynamics from the days or weeks observed in traditional banking crises to hours or minutes in digital finance.

4.3.4 Case Evidence

The Iron Finance/TITAN collapse of June 2021 served as an early warning of algorithmic stablecoin fragility. Iron Finance’s partially algorithmic stablecoin IRON was backed by a combination of USDC and the protocol’s native TITAN token. When large IRON redemptions triggered selling pressure on TITAN, the TITAN price collapsed from \$65 to near zero within hours (Briola et al., 2023). The reflexive spiral followed the death-spiral pattern precisely: IRON redemptions minted TITAN, TITAN selling depressed its price, and the reduced collateral value triggered further IRON redemptions. The speed of the collapse measured in hours rather than days illustrated how on-chain execution compresses run dynamics relative to traditional bank runs. The event destroyed approximately \$2 billion in value but served as a proof of concept for the death-spiral mechanism that would engulf Terra/Luna at hundred-fold larger scale one year later (de Carvalho et al., 2025).

The Terra/UST collapse of May 2022 destroyed approximately \$50 billion in combined market capitalization within five days, making it the largest stablecoin run in history (Liu et al., 2023). Large withdrawals from Anchor Protocol which held approximately \$14 billion in UST deposits at an unsustainable 20% yield initiated de-peg pressure. The Luna Foundation Guard’s deployment of Bitcoin reserves proved insufficient, and the reflexive feedback between UST redemptions and LUNA minting drove both tokens to near-zero. Santiago et al. (2024) document increased volatility correlations among digital assets during the collapse. The systemic propagation extended far beyond Terra: Three Arrows Capital became insolvent within weeks, Celsius and Voyager halted withdrawals within months, and the resulting confidence shock contributed to cascading failures culminating in the FTX collapse (Saengchote and Samphantharak, 2024).

We draw three lessons from the Terra episode that generalize to all algorithmic stablecoin architectures. First, wealthier and more sophisticated investors exited first, while smaller holders bore disproportionate losses replicating the inequitable loss allocation observed in traditional bank runs but with the additional feature that blockchain transparency allowed sophisticated actors to monitor each other’s withdrawal behavior in real time (Briola et al., 2023). Second, arbitrage bots that ordinarily stabilize prices through cross-venue trading amplified the death spiral by front-running redemption transactions and extracting value during the collapse (Daian et al., 2020). Third, the governance concentration documented across DeFi protocols meant that the system’s nominal decentralization provided no effective check on the buildup of risk within the Anchor Protocol (Bank for International Settlements, 2021b). The broader consequence was a lasting erosion of confidence in algorithmic stabilization mechanisms, contributing to the market-wide contraction that followed (Freitas et al., 2025).

The SVB/USDC de-peg of March 2023 demonstrated that reserve-backed stablecoins face run risk originating in the traditional banking system. When Silicon Valley Bank failed, Circle disclosed that \$3.3 billion of USDC’s reserves were held at the bank. USDC de-pegged to \$0.87 on decentralized exchanges within hours as holders rushed to convert to other stablecoins or fiat (Galati and Capalbo, 2023; de Carvalho et al., 2025). The de-peg cascaded through DeFi as protocols using USDC as collateral experienced automatic liquidations and pool imbalances, and the instability spread to other stablecoin pairs (Gregory et al., 2024). The run was arrested only when the FDIC announced that all SVB depositors including Cir-

cle would be made whole, an extraordinary government intervention that stablecoin holders cannot assume in future episodes (Bongini et al., 2025). This event revealed the bidirectional exposure between stablecoin issuers and traditional banks that defines the gateway risk channel discussed in Section 4.8.

4.3.5 Cross-Domain Manifestation

Stablecoin runs originate in the stablecoin domain but propagate rapidly across digital finance through three channels. In DeFi, propagation occurs through collateral impairment and AMM pool imbalances, as stablecoins underpin lending, exchange, and yield protocols as fundamental building blocks (Kitzler et al., 2022). In CeFi, settlement disruption and withdrawal pressure cascade to exchanges and lending platforms holding stablecoin balances, compounding confidence effects (Makarov and Schoar, 2022). At the gateway boundary, stablecoin reserve custody arrangements create bidirectional exposure between issuers and traditional banks, transmitting shocks in both directions as the SVB episode demonstrated (Mejia, 2024). Cross-ecosystem contagion amplifies these domain-level effects: a run on one stablecoin can trigger informational runs on others as holders reassess the safety of all stablecoin designs simultaneously (Sakariyahu et al., 2024). Because stablecoins intermediate the vast majority of on-chain value transfer, a de-peg disrupts not merely the issuer’s liabilities but the entire digital finance payment and settlement infrastructure.

4.3.6 Literature Assessment

The theoretical treatment of stablecoin runs has advanced rapidly but remains incomplete. Gorton and Zhang (2023) connect stablecoin fragility to the broader history of private money issuance, arguing that unregulated stablecoin markets replicate the instability of nineteenth-century wildcat banking. Lyons and Viswanath-Natraj (2023) analyze the arbitrage mechanisms that maintain stablecoin pegs and identify the conditions under which these mechanisms fail, providing empirical evidence on the fragility of peg maintenance during periods of market stress. The narrow banking literature, particularly Pennacchi (2012), provides a theoretical lens for evaluating reserve-backed stablecoins, which approximate the narrow banking model of full-reserve deposits but without the regulatory infrastructure that such proposals traditionally assume. Catalini and de Gortari (2019) examine the economic design choices underlying stablecoin architectures and their implications for stability, documenting the tradeoffs between capital efficiency and run vulnerability that different designs embody. Empirical studies of the Terra collapse have documented run mechanics at transaction-level granularity, establishing that the run was a complex multi-chain phenomenon precipitated by growing doubts about the system’s sustainability (Liu et al., 2023). Ante et al. (2023) identify algorithmic stablecoins as a major gap in the empirical literature, with only nascent work addressing their reflexive dynamics and equilibrium properties. The literature on reserve-backed stablecoins has developed more fully, but even there the formal treatment of redemption architecture and its implications for run vulnerability remains at an early stage.

The primary remaining gaps concern three areas. First, the formal modeling of simultaneous (as opposed to sequential) withdrawal dynamics in stablecoin runs, where Diamond-

Dybvig’s sequential-service assumption does not apply. Second, the equilibrium analysis of optimal reserve composition under the threat of runs, which requires modeling the interaction between reserve liquidity, redemption speed, and holder beliefs. Third, cross-stablecoin contagion whether a run on one token triggers runs on others through informational or liquidity channels represents a critical open question that the SVB/USDC episode brought into sharp focus (Gregory et al., 2024; de Carvalho et al., 2025). The interaction between stablecoin run dynamics and traditional bank runs, illustrated by the SVB episode, remains particularly undertheorized because it requires a model that endogenizes both run types simultaneously (da Cunha et al., 2021).

4.4 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.

4.4.1 Mechanism

The composability of DeFi protocols allows any smart contract to invoke the functions of any other smart contract without authorization, creating the capacity to build arbitrarily complex financial products by stacking simple building blocks a property widely termed “money legos” (Schär, 2021; Kim et al., 2022; Kitzler et al., 2022; Salami, 2021). A canonical example illustrates the mechanism: a user deposits ETH into Lido, receiving stETH; deposits the stETH into Aave as collateral to borrow USDC; deposits the USDC into a Curve liquidity pool, receiving LP tokens; and deposits those LP tokens into Convex Finance for yield optimization. The resulting position spans four protocol layers, and a failure at any layer propagates mechanically to all dependent layers.

We characterize this transmission as deterministic and near-instantaneous. When protocol B suffers an exploit and its token value drops to zero, protocol A’s balance sheet which holds tokens issued by protocol B degrades automatically, without assessment, negotiation, or forbearance (Gudgeon et al., 2020; Wen et al., 2024). The impairment triggers automatic liquidation of positions in protocol A that used protocol B’s tokens as collateral. The speed of propagation is limited only by blockchain block times (approximately twelve seconds on Ethereum) and the gas required to execute the relevant transactions. No human can intervene in the seconds between an exploit and the resulting cascade.

The dependency graph created by composability assembles permissionlessly, meaning no single entity has visibility into the complete set of dependencies (Kitzler et al., 2022; Collibus et al., 2022; Salami, 2021; Bank for International Settlements, 2021b). Protocol A’s developers may be unaware that protocol C which they have never interacted with holds tokens issued by protocol A as collateral and could trigger selling pressure on protocol A’s token upon a liquidation event. The resulting hidden correlations parallel the opaque securitization chains that contributed to the 2007–2008 financial crisis, but with a critical difference: in DeFi, the dependencies are encoded in publicly auditable smart contracts and could, in principle, be mapped by anyone with sufficient blockchain analytics capability. In

practice, the combinatorial complexity of the dependency graph and the rapid pace at which new protocols launch make comprehensive mapping a formidable computational challenge.

4.4.2 Theoretical Foundations

Composability risk constitutes a genuinely novel systemic risk channel with no close analog in traditional finance. Traditional financial products assemble through bilateral contracts, legal agreements, and regulatory approvals that impose natural limits on the depth and complexity of dependency chains (Schär, 2021, 2020; Bank for International Settlements, 2021b; Kim et al., 2022). A securitization chain in traditional finance originator to SPV to trustee to tranche investor typically involves three to five layers, each governed by documented legal agreements and subject to rating-agency scrutiny. In DeFi, an anonymous participant can assemble the equivalent dependency chain in minutes, extend it to arbitrary depth, and govern it solely by the logic encoded in smart contracts whose correctness relies on informal code auditing. The absence of gatekeepers at any layer of the stack represents the fundamental architectural difference that makes composability risk irreducible to existing financial contagion frameworks.

To our knowledge, no formal theoretical model of cascading failure in permissionlessly composed systems currently exists. The closest analogs in the academic literature are network contagion models that assume bilateral exposure networks with stochastic shock propagation (Dave et al., 2021), but composability failures propagate deterministically through mechanically linked smart contracts (Gudgeon et al., 2020; Daian et al., 2020). Reliability engineering models of cascading system failure offer a complementary perspective, but these typically assume centralized failure detection and response mechanisms absent from permissionless DeFi. The theoretical gap is significant: a formal model would need to capture the combinatorial structure of the dependency graph, deterministic propagation through smart contract logic, the game-theoretic behavior of liquidation bots and MEV searchers, and the endogenous formation of dependency structures through permissionless protocol deployment.

4.4.3 Digital-Native Features

We identify composability risk as entirely digital-native, exhibiting features with no counterpart in traditional finance. First, the depth and complexity of dependency chains are unbounded (Schär, 2020; Kim et al., 2022). Any protocol can compose with any other protocol, and the resulting chain can extend to arbitrary depth without approval, oversight, or even awareness by the protocols involved. No regulator grants permission; no compliance officer reviews the dependency; no risk manager assesses the systemic implications. Second, the propagation of failure is mechanical and deterministic: smart contracts execute exactly as coded, with no possibility of discretionary intervention, forbearance, or negotiated restructuring. A protocol whose collateral degrades from an upstream failure executes liquidations automatically, regardless of whether doing so amplifies a cascade or destabilizes the broader ecosystem.

Third, the dependency graph is dynamic: yield aggregators and strategy vaults continuously reallocate capital across protocols based on return optimization algorithms, meaning the network of dependencies changes constantly and cannot be captured in a static snapshot

(Xu and Vadgama, 2023; Collibus et al., 2022; Dave et al., 2021). Price oracle dependencies introduce an additional dimension, as many DeFi protocols rely on external price feeds from other protocols a dependency that creates failure pathways even without direct asset holdings. Fourth, cross-chain composability enabled by bridge protocols that transfer assets and messages between blockchains extends the dependency graph across blockchain boundaries, creating failure pathways that span multiple consensus mechanisms and security models. These four features distinguish composability risk from all traditional financial contagion channels and establish it as a uniquely digital phenomenon.

4.4.4 Case Evidence

The DAO Hack of June 2016, while predating the modern DeFi ecosystem, established the foundational case for composability risk. An attacker exploited a reentrancy vulnerability in The DAO’s smart contract code a flaw arising from the interaction between the withdrawal function and Ethereum’s call mechanism draining approximately 3.6 million ETH (\$60 million at the time) (Dotan et al., 2023; Wen et al., 2024; Gudgeon et al., 2020). The exploit demonstrated that smart contract vulnerabilities serve as systemic attack vectors when the affected contract holds assets on behalf of thousands of participants. The subsequent contentious hard fork that split Ethereum into Ethereum and Ethereum Classic illustrated that resolving composability failures may itself introduce systemic governance risk.

The Euler Finance exploit of March 2023 demonstrated composability risk propagation in a mature DeFi ecosystem. An attacker exploited a vulnerability in Euler’s donation and liquidation logic to steal \$197 million from the lending protocol (Werner et al., 2022; Xu and Vadgama, 2023). The exploit drained multiple token markets simultaneously and disrupted downstream protocols that had integrated Euler as a yield source or collateral manager. Protocols such as Balancer, Angle, and Idle Finance experienced secondary losses because their smart contracts held eTokens Euler’s interest-bearing deposit tokens that became worthless upon the exploit. The incident demonstrated that a single-protocol vulnerability can propagate losses across the dependency graph to protocols whose own code is entirely secure.

The Curve Finance pool exploit of July 2023 illustrated how composability risk extends beyond direct protocol dependencies to the governance token layer. A Vyper compiler vulnerability a failure not in any protocol’s code but in the programming language used to write the code allowed attackers to drain several Curve pools (Daian et al., 2020; Dotan et al., 2023; Gan et al., 2022). The resulting CRV price decline threatened cascading liquidations across Aave, Fraxlend, and Abracadabra, where Curve’s founder had deposited CRV as collateral. The episode revealed a novel form of composability risk in which a single governance token serves as shared collateral across multiple independent lending protocols, creating a hidden common exposure that no individual protocol’s risk assessment captures.

Cross-chain bridge exploits constitute a subclass of composability risk with particularly large loss magnitudes. The Ronin Bridge hack (\$625 million, March 2022), the Wormhole hack (\$326 million, February 2022), the Nomad hack (\$190 million, August 2022), and the Poly Network hack (\$611 million, August 2021) all involved failures in the composed interaction between two blockchain environments (Collibus et al., 2022; Schär, 2020; Gan et al., 2022). Bridges aggregate large pools of locked assets behind smart contract code that must correctly implement consensus verification logic across heterogeneous blockchains

a technically demanding requirement that has proven exceptionally difficult to secure in practice. Aggregate bridge losses exceeded \$1.7 billion in 2021–2022 alone, establishing cross-chain composability as the highest-magnitude attack surface in the DeFi ecosystem.

4.4.5 Cross-Domain Manifestation

Composability risk concentrates in the DeFi domain, where permissionless protocol interaction is a core architectural feature. CeFi platforms are not directly composable in the DeFi sense, but they interact with DeFi protocols through yield strategies, token custody, and market-making activities that create indirect composability exposures (Dotan et al., 2023; Salami, 2021; Kitzler et al., 2022). Stablecoins serve as a composability medium: because stablecoins are accepted across virtually all DeFi protocols, a failure in any protocol holding large stablecoin reserves can transmit to the stablecoin itself through selling pressure or reserve impairment. The FTX collapse of November 2022 illustrated this indirect channel, as CeFi yield strategies that had deployed customer funds into DeFi protocols transmitted losses back through the dependency chain. Tokenized TradFi assets that become embedded in composed protocol stacks will inherit composability risk as DeFi integration deepens, potentially creating a transmission pathway from DeFi composability failures into traditional financial markets.

4.4.6 Literature Assessment

The academic literature on composability risk is growing rapidly but remains primarily empirical and taxonomic. Existing surveys systematize DeFi security threats, attack vectors, and protocol primitives (Werner et al., 2022; Xu and Vadgama, 2023). Empirical work on composability itself has begun to decompose protocol interactions into nested building blocks, demonstrating how stablecoin failures propagate through composition trees and identifying swaps as the most frequently nested primitive across DeFi protocols. The primary literature gap is the absence of a formal theoretical model of cascading failure in permissionlessly composed systems a model that would need to capture the combinatorial structure of the dependency graph, deterministic propagation through smart contract logic, and the endogenous formation of the dependency structure. This gap distinguishes composability risk from all other channels we examine: every other systemic risk channel possesses at least a partial theoretical foundation, whereas composability risk rests entirely on empirical observation and taxonomic classification.

Formal verification approaches to composable protocols (Tolmach et al., 2021) represent an emerging direction for identifying dependency vulnerabilities before they manifest as cascading failures. These efforts complement the growing empirical literature on composability exploits and provide a potential pathway toward ex ante risk assessment.

Formal verification efforts represent a nascent but promising approach to mitigating composability risk at the protocol level. Dave et al. (2021) present verified AMM contracts with proven bounds on manipulation costs, while Wen et al. (2024) develop attack synthesis tools targeting deep logical vulnerabilities that arise from cross-contract interactions (Schär, 2021). These contributions address individual protocol correctness but do not yet extend to system-level properties of composed protocol stacks. We identify the development of a formal

cascading-failure model for permissionlessly composed systems as the most pressing theoretical gap in the composability risk literature. Such a model would be analogous to network contagion models in traditional finance but must account for the mechanical, deterministic, and unbounded nature of smart contract dependency chains.

4.5 Leverage and Liquidation Cascades

Liquidation cascades occur when the automated liquidation of undercollateralized positions in on-chain lending protocols triggers a chain reaction of further liquidations, driven by permissionless liquidation bots, amplified by MEV extraction, and compressed into timelines measured in minutes rather than the days or weeks of traditional margin-call cascades.

4.5.1 Mechanism

On-chain lending protocols such as Aave, Compound, and MakerDAO permit users to deposit collateral assets and borrow other assets against them up to a protocol-defined loan-to-value ratio, typically between 75% and 85% (Alamsyah et al., 2024). Each borrowing position carries a collateralization ratio the value of deposited collateral divided by the value of outstanding debt and a liquidation threshold below which the position becomes eligible for seizure by any external actor. When the market price of the collateral asset declines sufficiently to push the collateralization ratio below this threshold, any participant can partially repay the debt, seize a proportional share of the collateral at a protocol-specified discount of typically 5–10%, and profit from the difference between the collateral’s market value and the repayment cost (Bodó and Filippi, 2024). The cascade dynamic emerges when liquidation of one position generates selling pressure that drives other positions past their thresholds. A liquidator seizes collateral from an undercollateralized position and immediately sells it on an automated market maker or centralized exchange to realize the liquidation bonus in a stable asset. The sale depresses the collateral price, reducing collateralization ratios for all borrowers holding the same asset. Positions that were previously above the threshold cross below it, becoming eligible for liquidation in turn (Klages-Mundt and Minca, 2022). The cycle continues until the price stabilizes at a level where surviving positions remain sufficiently collateralized or all vulnerable positions have been liquidated. Correlated collateral holdings across protocols amplify this feedback: when many borrowers pledge the same asset, a single price decline simultaneously threatens positions across multiple venues, creating the potential for cross-protocol synchronized cascades (Li et al., 2023a; Kopytov, 2023).

The game-theoretic structure of permissionless liquidation adds a distinctive layer to the cascade mechanism. Liquidation bots automated programs that continuously monitor the blockchain for positions approaching thresholds compete to seize collateral the instant a threshold is breached (Daian et al., 2020). Multiple bots target the same liquidation opportunity, and the winner is determined by the gas price offered, network latency, and the ability to negotiate preferential transaction ordering with block builders through private channels such as Flashbots. This competition generates Maximal Extractable Value (MEV), a form of profit extraction in which validators or builders reorder, insert, or censor transactions to capture value (Qin et al., 2022; Li et al., 2023b). During cascades, MEV activity amplifies price instability through front-running of liquidation sales, sandwich attacks on large trades,

and congestion of block space that crowds out ordinary user transactions (Xue et al., 2023).

4.5.2 Theoretical Foundations

Liquidation cascades are a hybrid channel rooted in the traditional margin-call literature but with substantial digital-native extensions. Brunnermeier and Pedersen (2009) model the margin spiral within the broader liquidity spiral framework: declining asset prices trigger higher margin requirements, forcing borrowers to post additional collateral or liquidate positions, which depresses prices further and tightens funding conditions in a self-reinforcing feedback loop. Traditional margin calls involve human judgment at each stage the lender decides whether to issue a call, the borrower chooses how to respond, and market participants decide whether to buy the liquidated assets (Brunnermeier, 2009). Each decision point introduces time delays and the possibility of forbearance that can moderate cascade severity, as lenders may extend grace periods to borrowers they wish to retain (Kopytov, 2023).

DeFi liquidations eliminate all of these decision points. The liquidation threshold is hard-coded in the smart contract and cannot be modified in response to market conditions without a governance vote that typically requires days of deliberation and community consensus (Weingärtner et al., 2023; Adamyk et al., 2025). Autonomous bots execute liquidations with no capacity for discretion or forbearance. The selling of seized collateral is automated through AMM interactions that mechanically depress prices according to their constant-product bonding curves. The entire chain from threshold breach to market impact operates within a single Ethereum block approximately twelve seconds or across a small number of consecutive blocks (Bank for International Settlements, 2021b). This compression from a multi-day human process involving multiple decision points to an algorithmic event measured in seconds represents a qualitative transformation of the margin cascade mechanism, not merely a quantitative acceleration.

4.5.3 Digital-Native Features

Five features distinguish on-chain liquidation cascades from traditional margin-call cascades. First, the full transparency of on-chain positions allows all market participants to observe the distribution of collateralization ratios across a lending protocol in real time, down to the individual position level. Liquidation bots identify positions approaching thresholds and prepare transactions in advance, creating a “liquidation wall” effect in which the market collectively anticipates the cascade and may front-run it by selling the collateral asset before liquidations begin (Daian et al., 2020). This anticipatory dynamic, absent in traditional margin systems where position information is private and held bilaterally, accelerates the onset and severity of the cascade (Kopytov, 2023).

Second, the permissionless nature of liquidation removes any designated liquidator with a relationship to the borrower. Any actor can liquidate any position, eliminating the possibility of lender forbearance that often moderates traditional margin cascades (Adamyk et al., 2025). Third, the liquidation bonus the discount at which the liquidator acquires collateral, typically 5–10% of the seized amount incentivizes aggressive and rapid execution: delay risks losing the opportunity to a competing bot. These design choices embed procyclicality directly into the protocol architecture, rewarding speed over stability and creating a race condition that

accelerates the cascade (Kirişci, 2025).

Fourth, MEV extraction during liquidation events amplifies price volatility through adversarial transaction ordering. Front-running bots detect pending liquidation transactions in the mempool and insert trades ahead of them that profit from the anticipated price movement, worsening the net impact on other market participants (Li et al., 2023b; Xue et al., 2023). Sandwich attacks where a bot places buy and sell orders around a large liquidation sale extract additional value from the price dislocation at the liquidated borrower’s expense. The scale of this extractive activity is substantial and growing: MEV from liquidation-related transactions represents a meaningful and increasing share of on-chain economic activity (Qin et al., 2022). Fifth, the concentration of liquidation infrastructure among a small number of sophisticated bot operators introduces counterparty concentration risk within the liquidation mechanism itself. If a dominant operator experiences technical failure during a cascade, liquidation efficiency collapses and the protocol accumulates bad debt, as occurred during MakerDAO’s Black Thursday (Kirişci, 2025).

4.5.4 Case Evidence

Black Thursday, March 12, 2020, remains the canonical case study of DeFi liquidation cascades. The COVID-19 market crash triggered a roughly 50% single-day decline in ETH, pushing thousands of MakerDAO collateralized debt positions below their 150% liquidation threshold and unleashing cascading liquidations that overwhelmed the protocol’s auction mechanism. Extreme Ethereum network congestion caused gas prices to spike above levels that most liquidation bots were configured to pay (Schueffel, 2025). With primary bidding infrastructure incapacitated, a small number of actors submitted zero-bid transactions and won auctions unopposed, acquiring collateral for free (Klages-Mundt and Minca, 2022). MakerDAO incurred approximately \$8.3 million in protocol losses, requiring new MKR token issuance to recapitalize the system. The episode revealed that the liquidation mechanism, while sound under normal conditions, fails catastrophically when infrastructure constraints gas prices, network throughput interact with cascade dynamics (Kaur et al., 2023).

The Terra/Luna collapse of May 2022 triggered cascading liquidations on the Anchor Protocol that amplified the death spiral between UST and LUNA. As UST’s peg deteriorated, borrowers on Anchor whose positions were collateralized by bonded LUNA (bLUNA) faced liquidation as the LUNA price collapsed in parallel with the stablecoin (Kokorin, 2023). The liquidation of bLUNA positions generated additional LUNA selling pressure on secondary markets, which further reduced UST’s algorithmic backing, deepened the de-peg, and triggered yet more Anchor liquidations in a reflexive feedback loop. The interaction between the stablecoin run channel (Section 4.3) and the liquidation cascade channel was the core amplification mechanism of the collapse, destroying approximately \$45 billion in market capitalization within five days (Li et al., 2023a; Brunnermeier and Pedersen, 2009).

The Curve Finance near-miss of July 2023 illustrated how concentrated collateral positions create cross-protocol liquidation cascade risk. Curve’s founder held approximately \$168 million of CRV tokens as collateral distributed across Aave (approximately \$60 million), Fraxlend (approximately \$55 million), and Abracadabra (approximately \$50 million). As the Vyper compiler exploit drove the CRV price downward, these positions approached liquidation thresholds on all three protocols simultaneously (Weingärtner et al., 2023; Klages-

Mundt and Minca, 2022). A cascade on any one protocol would have generated CRV selling pressure sufficient to push positions on the remaining two past their thresholds, creating a synchronized cross-protocol cascade with potential to destabilize the broader DeFi lending ecosystem. The crisis was ultimately averted through OTC sales negotiated directly between the founder and individual buyers, demonstrating that DeFi’s automated liquidation mechanism can be circumvented by actors with sufficient social capital and market access (Kokorin, 2023).

4.5.5 Cross-Domain Manifestation

Liquidation cascades originate primarily in DeFi, where on-chain lending protocols provide the institutional setting for automated, permissionless liquidation. CeFi platforms also execute margin-call liquidations, but the process involves human discretion, opaque execution, and potential forbearance making it structurally distinct from DeFi cascades (Kokorin, 2023; Weingärtner et al., 2023). The cross-domain transmission operates primarily through price impact: forced selling on DeFi protocols depresses asset prices that are arbitrated onto centralized exchanges within seconds, transmitting the cascade’s effects to CeFi traders and spot markets (Li et al., 2023a). The stablecoin domain absorbs liquidation cascade effects when stablecoin-collateralized positions are liquidated or when cascade-driven volatility triggers stablecoin de-pegging, as the Terra/Anchor episode demonstrated (Section 4.3). This cross-domain propagation ensures that DeFi liquidation cascades can destabilize asset classes and venues far beyond the protocol where the cascade originates.

4.5.6 Literature Assessment

Perez et al. (2021) provide the first systematic empirical analysis of on-chain liquidation events, documenting that liquidation cascades in DeFi lending protocols exhibit self-reinforcing dynamics where each liquidation depresses collateral prices, triggering further liquidations. Early quantitative analysis of market risk in lending protocols by Kao et al. (2020) established that parameter configurations collateral ratios, liquidation penalties interact with market conditions to produce nonlinear cascade risk, a finding that informed subsequent protocol design but has not eliminated the fundamental vulnerability.

The empirical literature on DeFi liquidation cascades is well developed relative to other digital finance risk channels, owing in part to the full observability of on-chain transactions that enables precise event reconstruction. Stochastic models of liquidation spirals in stablecoin systems provide the most advanced formal treatment of the cascade mechanism, while empirical studies quantify blockchain extractable value during stress episodes at transaction-level granularity (Qin et al., 2022). The MEV literature documents the competitive dynamics of liquidation bots and priority gas auctions, with systematic catalogs of extraction strategies in Flashbots bundles revealing seventeen distinct MEV activity types (Li et al., 2023b). Traditional margin-cascade theory, led by Brunnermeier and Pedersen (2009), provides the theoretical scaffolding but assumes human decision-making with information-processing delays that DeFi eliminates (Kaur et al., 2023).

The primary theoretical gap concerns formal modeling of cascade dynamics in an environment with algorithmic execution, transparent position information, and MEV extraction.

To our knowledge, traditional margin-cascade models do not account for the game-theoretic competition among permissionless liquidation bots, the interaction between on-chain gas auctions and liquidation execution speed, or the role of MEV in amplifying price dislocations during stress (Xue et al., 2023; Adamyk et al., 2025). A model that jointly captures the liquidation mechanism, the AMM price-impact function, bot strategic behavior, and MEV dynamics would provide the basis for optimal protocol parameter design including liquidation thresholds, bonus percentages, and auction mechanisms calibrated to minimize cascade risk while preserving protocol solvency (Kirişci, 2025). The systemic role of concentrated liquidation bot operators as de facto critical infrastructure actors also remains unexplored, as does the welfare analysis of alternative liquidation designs such as gradual liquidation, Dutch auctions, and protocol-owned liquidation reserves.

4.6 Counterparty and Concentration Risk

Counterparty concentration risk in digital finance arises when a small number of entities centralized exchanges, lending platforms, custodians, market makers, oracle providers, or bridge validators 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.

4.6.1 Mechanism

Despite the decentralization ethos that motivates blockchain technology, the digital finance ecosystem exhibits extreme concentration at multiple layers of its institutional stack. A small number of centralized exchanges historically dominated by Binance, followed by Coinbase, OKX, and (before its collapse) FTX account for the majority of global cryptocurrency trading volume (Ante and Saggi, 2024). At the market-making layer, a handful of firms provide the bulk of liquidity across both centralized and decentralized venues. At the custody and infrastructure layers, dominant oracle providers, stablecoin issuers, and bridge operators concentrate critical system functions in single entities or small groups, creating what Dionysopoulos et al. (2023) term the “decentralisation illusion” nominally distributed systems with operationally centralized control (Bank for International Settlements, 2021b).

The failure of a concentrated counterparty propagates losses through three mechanisms. Direct exposure losses affect counterparties with assets deposited at, lent to, or owed by the failing entity. When FTX collapsed, counterparties with funds on the platform or outstanding loans to Alameda Research suffered immediate losses that triggered cascading insolvencies across the ecosystem (Liu et al., 2023). Liquidity withdrawal effects arise as market participants reduce exposure to entities perceived as connected to the failed entity, draining liquidity from the broader market (Brunnermeier, 2009). Confidence effects operate through the informational channel: the failure of a major entity erodes trust in the broader ecosystem, reducing participation, trading volume, and capital inflows (Haddad and Hornuf, 2023).

The too-big-to-fail problem that concentrated counterparties create in digital finance differs from its traditional-finance counterpart in a critical respect: no institutional mechanism for resolution exists. In traditional finance, the failure of a systemically important institution triggers a resolution process managed by a designated authority that can transfer critical

functions, impose orderly losses on creditors, and provide temporary liquidity support. No equivalent framework governs centralized crypto platforms, most of which operate across multiple regulatory perimeters and commingle customer assets with proprietary positions (Mikhaylov, 2023). International regulators have acknowledged this gap, noting that national authorities must develop frameworks comparable to those in traditional finance, yet implementation remains fragmented and incomplete (Adisa et al., 2024; Maple et al., 2023).

4.6.2 Theoretical Foundations

The traditional-finance literature on too-big-to-fail institutions and counterparty concentration provides the theoretical foundation for this channel. The endogenous network formation models of Babus (2016) predict that financial networks evolve toward concentrated structures when counterparties seek to minimize redundant connections a prediction consistent with the concentration patterns observed in CeFi lending, where a small number of bilateral relationships channeled the majority of credit. Upper (2011) surveys the simulation-based literature on interbank contagion, demonstrating that the failure of a large, concentrated counterparty in the interbank network can trigger cascading defaults across the system. The simulations reveal that contagion risk depends critically on the structure of the bilateral exposure network: concentrated hub-and-spoke structures in which many institutions hold exposures to a single dominant counterparty generate far larger cascades than diversified networks (Brunnermeier, 2009). Duffie (2010) formalizes the concept of slow-moving capital, showing that when intermediary capital is destroyed, asset prices deviate from fundamental values for extended periods because replacement capital enters slowly, creating persistent dislocations that impose welfare losses on the broader economy. The slow-moving capital framework applies with particular force to digital finance, where the failure of a concentrated intermediary removes not only capital but also the operational infrastructure order books, matching engines, API connections, and custodial relationships that other market participants depend on for execution.

The application to digital finance is direct: entities that concentrate critical functions exchange, custody, market-making, oracle provision impose systemic externalities on the broader ecosystem when they fail, because their failure removes infrastructure that other participants depend on and cannot easily replace. The extension required for digital finance concerns the multi-layer nature of concentration. In traditional finance, concentration risk is typically analyzed at the institutional level: banks that are too big, too interconnected, or too complex to fail. In digital finance, concentration operates simultaneously at multiple layers a single entity may serve as the dominant exchange, the largest market maker, the primary OTC lender, and the custodian of assets for dozens of other firms, as FTX/Alameda was (Fang et al., 2022). The failure of such a multi-function entity removes multiple infrastructure services simultaneously, creating a disruption whose scope exceeds what single-layer concentration analysis would predict (Gramlich et al., 2023).

We identify a theoretical tension that the existing literature has not fully resolved. The economic logic of intermediation creates strong forces toward concentration: centralized intermediaries reduce search frictions, provide liquidity, and offer custodial convenience that decentralized alternatives cannot yet match (Momtaz, 2024; Maple et al., 2023). The equilibrium outcome in digital finance may therefore be persistent concentration even in a nominally

decentralized architecture, as the empirical evidence suggests a prediction consistent with the “decentralisation illusion” documented empirically (Schär, 2020).

4.6.3 Digital-Native Features

Three features distinguish counterparty concentration in digital finance from its traditional analog. First, the absence of regulatory licensing and capital requirements in most jurisdictions allows concentrated entities to grow without prudential constraints. An exchange can achieve dominant market share without meeting capital adequacy requirements, maintaining reserve ratios, or submitting to supervisory examination (Bakare et al., 2024). The inadequacy of existing oversight mechanisms enables unchecked growth of systemic importance without corresponding safeguards (Mikhaylov, 2023).

Second, the commingling of functions exchange, custodian, market maker, lender, and proprietary trader within a single entity is prohibited or heavily regulated in traditional finance through rules requiring separation of client assets and licensing for different activities. In digital finance, these separations largely do not exist, enabling a single entity to accumulate concentration risk across multiple functions simultaneously (Bakare et al., 2024). The resulting co-location of functions means that a single governance failure, security breach, or liquidity event can disable all services simultaneously, as the FTX case demonstrated with losses exceeding six billion dollars in a single week (Kaur et al., 2023).

4.6.4 Case Evidence

The Mt. Gox collapse of February 2014 established the archetype of concentrated exchange failure in digital finance. At the time of its failure, Mt. Gox handled approximately 70% of global Bitcoin trading volume a level of concentration inconceivable in any regulated securities market. The exchange revealed losses on the order of \$460 million in customer Bitcoin holdings due to long-running security failures that went undetected (Fang et al., 2022; Ante and Saggiu, 2024). Bitcoin’s price fell more than 50% in the weeks following the bankruptcy, and the loss of the dominant trading venue severely disrupted price discovery and market liquidity across the nascent cryptocurrency market (Zhu et al., 2024).

The FTX collapse of November 2022 demonstrated that a decade of market evolution had not resolved the concentration problem. FTX was the third-largest cryptocurrency exchange by trading volume and, through Alameda Research, was simultaneously one of the largest market makers, OTC lenders, and venture investors in the ecosystem. When a CoinDesk report revealed that Alameda’s balance sheet was dominated by FTT FTX’s proprietary token Binance announced it would liquidate its FTT holdings, triggering a bank run on FTX. Within 48 hours, FTX halted withdrawals and filed for bankruptcy, revealing that customer funds had been commingled with Alameda’s trading positions (Jalan and Matkovskyy, 2023). The collapse propagated through counterparty exposures to Genesis Trading, BlockFi, and dozens of portfolio companies, eliminating a significant fraction of market-making liquidity across centralized and decentralized venues (Liu et al., 2023).

The Three Arrows Capital (3AC) collapse of June 2022 revealed concentration risk in the bilateral OTC lending network that connected CeFi intermediaries. The Singapore-based hedge fund had accumulated leveraged positions exceeding \$10 billion across multiple coun-

terparties, with concentrated exposure to the Terra/Luna ecosystem. When Terra collapsed in May 2022, 3AC’s losses cascaded through unsecured bilateral loans to Voyager Digital, Genesis Trading, and BlockFi entities that had extended credit without adequate collateral requirements or exposure limits (Haddad and Hornuf, 2023; Liu et al., 2023). The resulting chain of insolvencies demonstrated that counterparty concentration in digital finance operates not only at the platform level but also through opaque bilateral lending networks whose structure becomes visible only after a major failure (Gramlich et al., 2023; Mikhaylov, 2023).

The Multichain bridge collapse of July 2023 illustrated counterparty concentration at the infrastructure layer. The cross-chain bridge ceased operations after its CEO was reportedly detained by Chinese authorities, leaving approximately \$126 million in bridged assets inaccessible (Siam et al., 2025). The incident revealed that a nominally decentralized cross-chain protocol was in practice controlled by a single individual who held the administrative keys required for bridge operation a concentration of control that eliminated the redundancy supposedly provided by decentralized governance (Zhou et al., 2020).

4.6.5 Cross-Domain Manifestation

Counterparty concentration manifests most acutely in the CeFi domain, where exchanges, custodians, and market makers concentrate critical functions in corporate entities. In DeFi, concentration operates at the protocol and infrastructure layers: dominant AMMs such as Uniswap and Curve concentrate trading liquidity, dominant lending protocols such as Aave concentrate lending activity, and dominant oracle providers such as Chainlink concentrate price-feed services (Kaur et al., 2023). In the stablecoin domain, the duopoly of Tether and Circle concentrates issuance, creating dependency on two entities whose reserve management practices and regulatory relationships determine the stability of the primary settlement asset (Ante et al., 2023; Dionysopoulos et al., 2023). In tokenized TradFi, the early-stage nature of the market means that a small number of issuance platforms and custody providers concentrate the infrastructure for tokenized asset markets, creating nascent concentration risks whose systemic implications will grow with market scale (Andryushin, 2024). The cross-domain pattern reveals that concentration migrates rather than disappears: efforts to decentralize one layer of the stack frequently result in re-concentration at an adjacent layer, as participants seek the efficiency gains that centralized intermediation provides.

4.6.6 Literature Assessment

The too-big-to-fail literature in traditional finance is extensive, and several studies have applied its insights to digital finance. Upper (2011) provides the simulation methodology for assessing interbank contagion, while Duffie (2010) establishes the theoretical framework for systemic externalities from concentrated intermediary failure. Fang et al. (2022) survey cryptocurrency trading with attention to the concentrated platform structure, and Adisa et al. (2024) document concentration patterns across exchanges and market makers.

The primary literature gaps concern three areas. First, to our knowledge, no formal model captures multi-layer concentration where a single entity simultaneously concentrates functions across exchange, custody, market-making, and lending and the amplified cascade dynamics this creates. Second, monitoring metrics for counterparty concentration in digital

finance remain undeveloped, because most entities do not report exposures and counterparty relationships span regulated and unregulated entities across multiple jurisdictions (Gramlich et al., 2023). Third, the absence of a resolution framework for failed crypto intermediaries remains a critical institutional gap that the academic literature has identified but not resolved: no study has proposed a feasible resolution mechanism that accounts for the cross-jurisdictional, commingled, and partially on-chain nature of crypto intermediary balance sheets (Haddad and Hornuf, 2023; Maple et al., 2023). Future work should address the interaction between on-chain transparency and off-chain opacity in counterparty networks, developing monitoring tools that can track concentration in real time across both observable and opaque layers of the ecosystem.

4.7 Information Asymmetry and Opacity

Information asymmetry in digital finance manifests as the gap between what insiders know about the solvency, risk exposure, and operational integrity of centralized platforms and what depositors, counterparties, and regulators can observe a gap that enables fraud, misrepresentation, and hidden risk-taking, and that amplifies the severity of crises when concealed information is eventually revealed.

4.7.1 Mechanism

The information asymmetry channel operates through three interacting components: the production of private information by platform insiders, the failure of disclosure mechanisms to convey that information to counterparties, and the sudden revelation of concealed information that triggers confidence shocks and run dynamics. Unlike traditional finance, where regulatory mandates compel periodic disclosure of audited financial statements, the digital finance ecosystem lacks a coherent disclosure regime. Platform insiders executives, controlling shareholders, or key personnel at centralized exchanges, lending platforms, and custodians possess private information about the platform’s true financial condition. The relevant private information includes the extent of rehypothecation of customer assets, the quality and composition of loan collateral, the existence of undisclosed liabilities, the true reserves backing stablecoins, and the degree to which customer funds are commingled with proprietary trading positions (Akerlof, 1970; Choi and Kim, 2024; Bodó and Filippi, 2024; Tan and Saraniemi, 2022).

Depositors and counterparties face a classic adverse-selection problem because the disclosure mechanisms available in digital finance are structurally inadequate. Voluntary proof-of-reserves attestations show assets but not liabilities and can be manipulated by temporarily borrowing assets to inflate the snapshot. Third-party audits are rare, and several high-profile engagements such as Mazars’ proof-of-reserves work for Binance were terminated under scrutiny, eroding their credibility. Regulatory reporting obligations exist in some jurisdictions but remain fragmented, inconsistent, and enforceable only within their own jurisdictions (Schuler et al., 2024; Adamyk et al., 2025; Nabben and Filippi, 2024). The result is an information environment in which insiders can engage in fraud or excessive risk-taking with low probability of detection until the platform fails.

When concealed information is eventually revealed typically through investigative journalism, whistleblower disclosures, or the cascading failure of connected counterparties the resulting confidence shock triggers run dynamics amplified by the information void. Depositors who discover that a platform’s true condition is worse than represented cannot assess the magnitude of losses without complete disclosure, which is rarely provided voluntarily during a crisis. The uncertainty triggers worst-case assumptions and mass withdrawal, pushing platforms that are impaired but solvent into insolvency (Stiglitz and Weiss, 1981; Oben and Özdamli, 2024; Bakare et al., 2024). The absence of deposit insurance eliminates the safety net that prevents self-fulfilling bank runs in traditional banking, leaving depositors to bear the full cost of the information failure. This dynamic transforms what might be a localized solvency problem into a system-wide confidence crisis, as the revelation of opacity at one platform leads depositors at other platforms to reassess their own information deficits.

4.7.2 Theoretical Foundations

Akerlof (1970) established the foundational framework for understanding how private information about quality drives adverse selection and market failure. Applied to digital finance, the framework predicts that platforms with strong risk management and full reserves cannot compete with platforms offering higher yields financed through hidden risk-taking, because depositors cannot distinguish between the two. This 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, and why equilibrium pricing in crypto markets reflects extrinsic volatility unrelated to fundamentals (Fang et al., 2022).

Stiglitz and Weiss (1981) extended the analysis to credit markets, demonstrating that information asymmetry between borrowers and lenders produces credit rationing and adverse selection in loan portfolios. These dynamics apply directly to CeFi lending platforms that extended unsecured loans to counterparties such as Three Arrows Capital without adequate assessment of creditworthiness or aggregate leverage. The lending platforms could not price the true risk of their loan portfolios because borrowers possessed private information about their total leverage across multiple platforms. The resulting credit misallocation concentrated losses in precisely the institutions least equipped to absorb them (Allen et al., 2022).

The shadow-banking literature provides a complementary theoretical lens for understanding the CeFi failures of 2022. The CeFi lending platforms that failed Celsius, Voyager, BlockFi replicated shadow-banking architecture: they accepted demand-withdrawable deposits promising high yields, deployed them in illiquid and opaque strategies (DeFi yield farming, bilateral lending, directional trading), and failed when depositors discovered the gap between promised and actual risk profiles (Bank for International Settlements, 2021b; Adisa et al., 2024; Mirdala, 2024). This pattern mirrors the 2007–2008 financial crisis, in which sudden uncertainty about the value of collateral backing short-term funding triggered runs on shadow banking institutions and propagated losses across the financial system. The key distinction is that CeFi platforms operated outside any prudential supervisory perimeter, with no capital requirements, no liquidity buffers, and no resolution frameworks to manage their orderly unwinding.

4.7.3 Digital-Native Features

The information environment of digital finance creates a paradoxical hybrid of transparency and opacity that has no close analog in traditional finance. DeFi protocols are transparent by design: their smart contract code is publicly auditable, on-chain transactions are permanently recorded, and reserve compositions can be verified in real time by anyone with blockchain analytics capability. CeFi platforms, by contrast, operate with opacity that often exceeds traditional financial institutions, because most are unregulated, not subject to auditing requirements, and not obligated to disclose balance sheets or risk exposures (Schär, 2021). This coexistence generates an information environment where the degree of observability varies sharply by domain, platform type, and participant sophistication. The paradox is structural: the very existence of on-chain transparency in DeFi may have provided false comfort about the overall ecosystem’s information quality, masking the deeper opacity of the CeFi layer where the largest losses ultimately materialized.

The coexistence of transparent and opaque layers creates a distinctive information stratification in which sophisticated participants extract actionable intelligence from on-chain data that retail participants cannot access or interpret. Easley et al. (2019) document how transaction fee markets in blockchain systems create information asymmetries between sophisticated miners/validators and ordinary users, extending the classical market microstructure framework of informed versus uninformed traders to the digital asset setting. This information advantage enables MEV extraction and front-running that impose direct costs on less-informed participants: arbitrage bots competitively bid up transaction fees to obtain priority execution, extracting value from ordinary users’ trades (Choi and Kim, 2024; Hautsch et al., 2024). The result is a two-tier information environment in which on-chain data are nominally public but practically accessible only to participants with the technical infrastructure to interpret and act upon them in real time. This form of information asymmetry is particularly insidious because it operates under the appearance of transparency, giving retail participants a false sense of informational parity with sophisticated actors.

Proof-of-reserves the industry’s primary voluntary disclosure mechanism epitomizes the inadequacy of existing information infrastructure in CeFi. A standard attestation demonstrates that an exchange controls on-chain addresses holding assets matching reported customer liabilities at a specific point in time. The attestation does not capture off-chain liabilities, does not prevent temporary inflation of balances through borrowing, and does not address commingling of customer assets with proprietary positions or pledging as collateral for undisclosed loans (Hautsch et al., 2024; Fang et al., 2022). The FTX case demonstrated that even platforms appearing well-capitalized in aggregate could be deeply insolvent once off-balance-sheet liabilities and commingled positions were accounted for. As of early 2025, the industry has not yet developed attestation standards that adequately address these gaps.

4.7.4 Case Evidence

The QuadrigaCX collapse of February 2019 established the canonical case for information asymmetry in crypto custody. The Canadian exchange claimed that its founder’s unexpected death left \$190 million in customer funds locked in cold wallets to which only he held the private keys. The Ontario Securities Commission’s investigation revealed the exchange had

operated as a Ponzi scheme: the founder traded customer deposits on other exchanges, suffered losses, and covered the shortfall with new deposits (Ontario Securities Commission, 2020; Carpentier-Desjardins et al., 2025; Bodó and Filippi, 2024). No voluntary disclosure mechanism or third-party audit had detected the fraud during the exchange’s years of operation. The case established that single-person control over custody keys, combined with zero external verification, creates conditions for total information asymmetry between operator and depositor.

The Celsius Network and Voyager Digital failures of mid-2022 demonstrated information asymmetry in the CeFi lending model at systemic scale. Both platforms marketed themselves as safe alternatives to bank deposits, offering yields of 8–18% on cryptocurrency deposits while deploying funds in opaque, high-risk strategies including unsecured bilateral lending, illiquid DeFi positions, and directional trading. Celsius revealed approximately \$1.2 billion in undisclosed losses when it halted withdrawals in June 2022; Voyager revealed a \$650 million unsecured loan to Three Arrows Capital that it had not disclosed to depositors (Allen et al., 2022). The gap between depositors’ information and reality constituted classic adverse selection: opacity-enabled risk-taking attracted deposits that would not have been made under full information. Neither platform had disclosed its aggregate risk exposure to any regulator or auditor prior to failure.

The FTX collapse of November 2022 combined information asymmetry with outright fraud at the largest scale yet observed in digital finance. Customer funds were transferred to Alameda Research for proprietary trading, losses were concealed through fabricated accounting, and the proprietary token FTT was used as undisclosed collateral for loans. The information asymmetry was total: insiders knew the platform was deeply insolvent, while depositors and even sophisticated institutional investors believed it to be among the most well-capitalized exchanges in the industry (Schuler et al., 2024; Adisa et al., 2024). The FTX case demonstrated that reputational capital and high-profile venture capital investment provided no substitute for verifiable disclosure, as the platform’s perceived legitimacy rested entirely on unaudited representations.

4.7.5 Cross-Domain Manifestation

Information asymmetry concentrates in the CeFi domain, where centralized entities possess private information that depositors and counterparties cannot verify, but it manifests in distinct forms across all domains of digital finance. In DeFi, smart contract transparency mitigates traditional information asymmetry but creates a different form: the complexity of composed protocol stacks means that even sophisticated participants may not fully understand the risk profile of their positions, and pseudonymous participation prevents aggregate counterparty exposure assessment (Nabben and Filippi, 2024; Adamyk et al., 2025). In the stablecoin domain, information asymmetry concerns reserve composition and custody arrangements, as the persistent controversy over Tether’s reserve disclosures illustrates. CeFi platforms can observe DeFi positions and adjust strategies accordingly, while DeFi participants cannot observe the CeFi exposures that may ultimately determine their counterparty risk a directional information advantage that compounds systemic fragility.

4.7.6 Literature Assessment

The adverse-selection framework of [Akerlof \(1970\)](#) and the credit-rationing analysis of [Stiglitz and Weiss \(1981\)](#) provide strong theoretical foundations for understanding information asymmetry in CeFi, and the empirical literature has documented multiple cases of disclosure failure. Detailed post-mortems of QuadrigaCX ([Ontario Securities Commission, 2020](#)) and FTX ([Jalan and Matkovskyy, 2023](#)) have enriched the case literature. The primary literature gap concerns the hybrid information environment created by the coexistence of on-chain transparency and off-chain opacity. Existing adverse-selection models do not account for partial observability, in which some actions are visible on-chain while others remain hidden off-chain, and in which the ability to extract information varies across participants based on analytical sophistication. A model that endogenizes information production in this hybrid environment capturing the strategic choices of platforms regarding what to place on-chain versus off-chain would contribute significantly to understanding the equilibrium level of opacity in digital finance. Such a model would need to incorporate the heterogeneous analytical capabilities of participants, the endogenous choice of information architecture by platforms, and the welfare implications of the resulting information stratification for regulatory design.

4.8 Fiat-Crypto Gateway and Banking Channel Risk

Gateway risk refers to the systemic fragility introduced by the small number of institutions that mediate between the traditional financial system and the digital asset ecosystem, creating bidirectional contagion channels through which traditional bank failures can destabilize crypto markets and crypto firm failures can impose losses on traditional banks.

4.8.1 Mechanism

The digital finance ecosystem does not operate in isolation from the traditional financial system. Fiat currency enters and exits crypto markets through a narrow set of gateway institutions: banks that maintain accounts for cryptocurrency exchanges, stablecoin issuers whose reserves reside in traditional banks, regulated broker-dealers providing institutional access, and payment processors enabling retail fiat-to-crypto conversion ([Gorton and Zhang, 2023](#)). We argue that the concentration of these gateway functions in a small number of entities creates a structural bottleneck through which bidirectional contagion can flow. When gateway institutions fail, both the traditional banking system and the digital asset ecosystem absorb losses simultaneously, because the gateway serves as a shared node in two otherwise distinct financial networks ([Brunnermeier, 2009](#)).

In the crypto-to-traditional direction, the failure of a major crypto firm an exchange collapse, a stablecoin run, or the insolvency of a large trading firm imposes direct losses on the traditional bank providing gateway services. Deposit withdrawals by the crypto firm, asset freezes during bankruptcy proceedings, and reputational damage from servicing a failed crypto client can trigger a run on the gateway bank itself ([Gorton and Zhang, 2023](#)). We observe that contagion amplifies when multiple crypto firms concentrate at the same gateway bank, because the failure of one client triggers preemptive withdrawals by other crypto clients, creating a coordination problem among the bank’s depositor base ([Fang et al.,](#)

2022). The concentration of crypto deposits at a single bank creates a correlated withdrawal risk that traditional deposit insurance frameworks were not designed to address.

In the traditional-to-crypto direction, the failure of a gateway bank freezes stablecoin reserves, halts fiat settlement for exchanges, and severs fiat on-ramp infrastructure for all dependent crypto firms. [Diop et al. \(2024\)](#) demonstrate this mechanism through the SVB/USDC episode of March 2023: when Silicon Valley Bank failed, the \$3.3 billion of USDC reserves held at the bank became temporarily inaccessible, causing USDC to de-peg by 13% and triggering DeFi liquidations, AMM pool imbalances, and contagion to other stablecoins using USDC as a reserve asset ([Diop et al., 2024](#)). The episode revealed that traditional banking stress transmits to digital finance at blockchain speed hours rather than weeks because stablecoin de-pegging cascades through smart contracts without human intervention ([Bodó and Filippi, 2024](#)).

The bidirectional nature of gateway risk creates feedback potential: a crypto-to-traditional shock that weakens a gateway bank reduces the bank’s capacity to serve as a gateway, which disrupts stablecoin reserves and fiat access for crypto firms, which amplifies the original crypto stress. [Khoury et al. \(2023\)](#) document analogous spillover dynamics between fintech and traditional financial indices, showing that contagion flows in both directions during periods of market stress. The feedback loop we identify is constrained only by the number and diversity of gateway banks and the evidence from 2023 demonstrates that this number is dangerously small ([Gromb and Vayanos, 2010](#)).

4.8.2 Theoretical Foundations

Gateway risk constitutes a novel systemic risk channel with no close analog in traditional finance, because the channel exists only at the boundary between two financial systems that were, until recently, largely separate. The closest theoretical analog is the correspondent banking literature, which analyzes the concentration of cross-border payment and settlement functions in a small number of global banks and the systemic risks this concentration creates. Just as the failure of a major correspondent bank can sever payment access for hundreds of downstream institutions, the failure of a crypto-friendly gateway bank severs fiat access for the entire digital asset ecosystem that depends on it. [Schuler et al. \(2024\)](#) demonstrate that even nominally decentralized financial infrastructure exhibits centralization vectors including reliance on traditional banking for fiat settlement that create regulatory-relevant chokepoints. The key insight from this literature, that concentration of gateway functions creates fragility because the failure of a gateway removes access for all dependent institutions, applies directly to the fiat-crypto boundary ([Bodó and Filippi, 2024](#); [Abdullah, 2024](#); [Guo et al., 2024](#)).

The political economy dimension of gateway risk deserves particular attention. Regulatory actions that deliberately narrow the set of banks willing to serve crypto clients described in industry commentary as “Operation Choke Point 2.0” reduce the number of gateway institutions and thereby increase gateway concentration risk ([Ozili and Alonso, 2024](#)). We identify an equilibrium tension between regulatory incentives to limit banks’ crypto exposure and financial stability considerations: a smaller number of gateways concentrates risk in fewer institutions and increases the systemic impact of any individual gateway failure. [Chu and Rathbun \(2025\)](#) analyze an analogous tension in the CBDC context, where competing models for cross-border payment platforms reveal trade-offs between monetary sovereignty

and interoperability. No formal model captures these dynamics for the fiat-crypto gateway, though the theoretical ingredients regulatory capital requirements, network externalities, and concentration risk are individually well understood (Guo et al., 2024).

4.8.3 Digital-Native Features

Gateway risk is entirely digital-native: the channel exists only because digital finance and traditional finance are distinct systems connected by institutional bridges. Four features distinguish gateway risk from other forms of interconnection risk. First, gateway concentration arises not from economies of scale (as in traditional too-big-to-fail dynamics) but from regulatory attrition: the compliance burden and reputational risk of serving crypto clients have driven many banks from the market, leaving a small number of crypto-friendly banks as the only available gateways (Ozili and Alonso, 2024). Second, gateway institutions serve simultaneously as chokepoints for fiat settlement, stablecoin reserves, and institutional access, meaning that the failure of a single gateway can disrupt all three functions at once. No comparably concentrated multi-function dependency exists in traditional correspondent banking (Financial Stability Board, 2022a).

Third, the bidirectional nature of gateway risk means that traditional banking stress can be imported into digital finance (SVB to USDC) and digital finance stress can be exported to traditional banking (FTX to Silvergate), creating a contagion pathway that regulators on neither side of the boundary fully monitor (Brunnermeier and Pedersen, 2009). Fourth, the speed of contagion through the gateway channel is asymmetric: crypto-to-traditional contagion operates at banking speed (days to weeks, mediated by deposit insurance and resolution processes), while traditional-to-crypto contagion operates at blockchain speed (hours, driven by stablecoin de-pegging and automated DeFi liquidations) (Khoury et al., 2023).

4.8.4 Case Evidence

The Silvergate Bank failure of March 2023 illustrates the crypto-to-traditional direction of gateway risk. Silvergate, a California-chartered bank, had transformed itself into the primary banking partner for the cryptocurrency industry, operating the Silvergate Exchange Network (SEN) a real-time payment platform enabling crypto firms to transfer US dollars between accounts 24/7. Following the FTX collapse in November 2022, Silvergate experienced \$8.1 billion in deposit outflows as crypto firms withdrew funds, forcing the bank to sell securities at a loss (Galati and Capalbo, 2023). The losses impaired Silvergate’s capital position and the bank announced voluntary liquidation in March 2023, removing the SEN as critical fiat settlement infrastructure for the crypto industry (Chu and Rathbun, 2025).

The SVB/USDC de-peg of March 2023 demonstrates the traditional-to-crypto direction with particular clarity. Silicon Valley Bank’s failure, triggered by a run driven by concerns about unrealized losses on its held-to-maturity securities portfolio, revealed that Circle held \$3.3 billion of USDC reserves at the bank. USDC de-pegged to \$0.87 on decentralized exchanges within hours of the FDIC’s receivership announcement (Diop et al., 2024). The de-peg cascaded through DeFi: AMM pools containing USDC became imbalanced as holders dumped USDC for other stablecoins; lending protocols using USDC as collateral triggered liquidations; and DAI, which uses USDC as a significant fraction of its collateral, experienced

secondary de-peg pressure (Abdullah, 2024). The USDC peg was restored only after the Federal Deposit Insurance Corporation announced that all SVB depositors would be made whole an extraordinary government intervention that, as a stablecoin stability mechanism, is neither reliable nor scalable.

The Signature Bank closure of March 2023 completed the triptych of gateway failures. New York state regulators closed Signature Bank on March 12, 2023, two days after SVB's failure, citing concerns about the bank's crypto-related deposits and the viability of its Signet real-time payments platform a competitor to Silvergate's SEN that provided real-time fiat settlement for crypto firms (Cookson et al., 2023). The closure further narrowed the set of banks willing to serve crypto clients and eliminated the second of two major real-time fiat settlement platforms available to the crypto industry. The combined loss of Silvergate, SVB (as a stablecoin reserve custodian), and Signature within a single month demonstrated the extreme concentration of gateway functions and the speed at which gateway infrastructure can vanish. Whether the regulatory decision to close Signature reflected genuine solvency concerns or a deliberate strategy to sever banking access for crypto firms remains contested, but the systemic consequence a dramatic reduction in gateway capacity was identical under either interpretation.

4.8.5 Cross-Domain Manifestation

Gateway risk operates primarily at the boundary between CeFi and traditional finance, but its effects propagate across all domains of digital finance. The stablecoin domain faces direct exposure through reserve custody: the stability of reserve-backed stablecoins depends on the solvency and accessibility of the banks holding reserves (Abdullah, 2024; Guo et al., 2024). The DeFi domain faces indirect exposure through the settlement-layer function of stablecoins: a gateway-induced stablecoin de-peg transmits to all DeFi protocols using the affected stablecoin as collateral or liquidity (Werner et al., 2022). The tokenized TradFi domain will face increasing gateway risk as tokenized real-world assets require fiat settlement, custody, and regulatory interface functions that gateway institutions at the traditional-digital boundary must provide. As institutional adoption deepens through Bitcoin ETFs, tokenized treasuries, and regulated custody, the volume of value flowing through gateway institutions will grow, amplifying the systemic consequences of any future gateway failure.

4.8.6 Literature Assessment

Gateway risk is the least theorized of the systemic risk channels we examine, reflecting its recent emergence as a distinct risk category. Gorton and Zhang (2023) discuss the relationship between stablecoin reserves and the banking system in proposing stablecoin regulation. Auer et al. (2023) document the concentration of banking relationships among crypto firms, while Schuler et al. (2024) provide a framework for assessing the factual decentralization of blockchain-based financial infrastructure, revealing the centralization vectors including banking dependence that gateway risk exploits. Policy reports from the BIS acknowledge the interconnection between crypto markets and the banking system but do not formalize gateway risk as a distinct transmission channel (Schuler et al., 2024).

The primary literature gaps we identify are substantial. No formal model captures the

bidirectional contagion dynamics between crypto markets and gateway banks (Bodó and Filippi, 2024; Ozili and Alonso, 2024). The equilibrium analysis of gateway provision how many gateway banks should exist, what capital requirements they should face, and how regulatory policy affects gateway concentration remains unaddressed (Chu and Rathbun, 2025). The political economy of gateway provision, including the role of regulatory actions in narrowing or widening gateway access, is also undertheorized (Khoury et al., 2023). As institutional connections between digital and traditional finance deepen through tokenization, ETFs, and institutional custody, the gateway risk channel will become increasingly systemically significant, and the theoretical gap will become increasingly consequential. The March 2023 triptych of gateway failures provides an empirical foundation for future theoretical work, but the formal apparatus has yet to be constructed.

5 Cross-Channel Interactions and Amplification Mechanisms

No major crisis in digital finance has ever activated a single risk channel in isolation. The eight channels analyzed in Sections 4.1–4.8 interact, compound, and amplify losses far beyond what any one channel would generate independently the system-level risk exceeds the sum of its parts (Huan and Renn, 2025).

We organize the cross-channel analysis around five themes: a taxonomy of interaction types, the five strongest pairwise feedback loops, three detailed crisis case studies demonstrating multi-channel activation, structural features of digital finance that accelerate cascade propagation, and the cross-domain contagion pathways that connect the four domains of the ecosystem.

5.1 Interaction Taxonomy: Amplifying, Dampening, and Sequencing

We distinguish three interaction types between channels. *Amplifying interactions* occur when activation of one channel increases the severity or speed of another: a liquidity spiral deepens a liquidation cascade, which further deepens the liquidity spiral. *Dampening interactions* occur when one channel’s activation reduces the intensity of another: a stablecoin run that triggers rapid deleveraging may reduce the leverage available to fuel subsequent liquidation cascades, providing a partial natural brake. *Sequential interactions* occur when one channel’s resolution triggers a different channel with a time lag: the FTX collapse in November 2022 activated counterparty concentration immediately, but the resulting gateway risk at Silvergate materialized weeks later as deposit outflows accumulated (Acemoglu et al., 2015; Shleifer and Vishny, 2011).

Amplifying interactions dominate the empirical record. Of the $8 \times 8 = 64$ directed pairwise relationships between channels, we identify 18 amplifying pairs, 4 dampening pairs, and 12 sequential pairs; the remaining 30 pairs exhibit weak or negligible interaction. Appendix B catalogs each pair with supporting crisis evidence and theoretical basis.

Because amplifying interactions outnumber dampening ones by more than four to one,

the system-level risk of digital finance is superadditive: the aggregate risk exceeds the sum of individual channel risks, a property that single-channel risk assessments systematically understate (Allen and Gale, 2000; Elliott et al., 2014; Financial Stability Board, 2022a).

Why does the distinction between simultaneous and sequential activation matter for crisis management? Simultaneous activation as observed in the Terra/Luna collapse, where stablecoin runs, liquidity spirals, and liquidation cascades fired within hours of each other overwhelms intervention capacity because multiple channels must be addressed at once. Sequential activation as in the 2022 cascade from Terra to Three Arrows Capital to FTX over six months provides windows for intervention that the digital finance ecosystem currently lacks the institutional capacity to exploit. Traditional finance crisis management relies implicitly on these intervention windows; their absence in the simultaneous-activation case explains the severity of multi-channel events (Gudgeon et al., 2020; International Monetary Fund, 2023).

5.2 The Five Strongest Feedback Loops

The tightest feedback loop operates between liquidity spirals and liquidation cascades (Sections 4.2 and 4.5). A price decline triggers automated liquidations in on-chain lending protocols such as Aave and Compound. Liquidation bots seize collateral and sell it immediately on AMMs, depressing prices further along the bonding curve. The lower prices push additional borrowing positions past their liquidation thresholds, triggering the next round of forced sales, and the cycle repeats. The interaction between these two channels is multiplicative rather than additive: the speed of automated liquidation feeds the intensity of the liquidity spiral, and the spiral’s depth determines the breadth of the next liquidation wave (Klages-Mundt and Minca, 2022; Qin et al., 2022).

Both Black Thursday 2020 and the Terra/Luna collapse exhibited this feedback loop as their core amplification mechanism. On-chain transparency enables market participants to observe positions approaching liquidation thresholds in real time, creating incentives to front-run the cascade through preemptive selling or MEV extraction paradoxically accelerating the very dynamics that participants seek to exploit or avoid (Daian et al., 2020; Qin et al., 2022; Mohan, 2022). Speed is the critical differentiator: where the traditional-finance analog of margin calls triggering fire sales operates over days with human discretion at each step, the DeFi version compresses the entire cycle to minutes with no human involvement.

The second feedback loop operates between counterparty concentration and network contagion (Sections 4.6 and 4.1). When a failing entity occupies a central position in the network by virtue of its exchange volume, market-making activity, lending relationships, or venture investments the contagion radius scales dramatically with the entity’s centrality. FTX’s simultaneous roles as a top-three exchange, the operator of Alameda Research (one of the largest market makers), and a venture investor in dozens of crypto firms meant that its failure activated network contagion across exchange, OTC lending, market-making, and venture capital channels simultaneously, rather than through a single transmission pathway (Upper, 2011).

As Acemoglu et al. (2015) demonstrate, highly connected network nodes can flip a financial system from shock-absorbing to shock-amplifying a result that applies with particular force when the highly connected node is also a concentrated counterparty. In traditional

finance, regulatory designation of systemically important financial institutions (SIFIs) partially addresses this dynamic through enhanced capital requirements and resolution planning. Digital finance lacks equivalent designation mechanisms, leaving the concentration-contagion feedback loop unmitigated by institutional safeguards ([Financial Stability Board, 2022a](#); [International Monetary Fund, 2023](#)).

The third feedback loop connects stablecoin runs and gateway risk (Sections 4.3 and 4.8). This loop is structurally distinctive because the contagion operates at the boundary between digital and traditional finance. When a gateway bank holding stablecoin reserves fails, the stablecoin de-pegs, and the de-peg propagates through DeFi and CeFi via stablecoins' settlement-layer function. Conversely, when a stablecoin run forces rapid liquidation of reserves held at traditional banks, the resulting selling pressure can impair the bank's balance sheet and trigger a traditional bank run. The SVB/USDC episode of March 2023 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](#); [Bank of England, 2023](#); [Financial Stability Board, 2023a](#)).

Opacity enables hidden concentration this is the fourth feedback loop, connecting information asymmetry and counterparty concentration (Sections 4.7 and 4.6). Counterparties that accumulate excessive leverage, commingle customer funds, or take undisclosed proprietary risks are precisely those with the strongest incentive to prevent disclosure. The information asymmetry channel shields concentrated counterparties from market discipline until cascading failures or investigative journalism forces revelation, at which point the combined shock exceeds what either channel would produce independently: the gap between depositors' beliefs and reality determines the run's severity ([Akerlof, 1970](#); [Stiglitz and Weiss, 1981](#)).

The fifth feedback loop connects composability risk and liquidation cascades (Sections 4.4 and 4.5). When an exploit drains a base-layer protocol, all protocols that hold tokens issued by the exploited protocol experience collateral impairment. The impairment triggers automated liquidations in lending protocols that accepted the impaired tokens as collateral, and the resulting forced sales depress asset prices, which can trigger further liquidations in other protocols holding the same assets. The Curve/Vyper episode of July 2023 demonstrated this compounding: a compiler vulnerability enabled exploitation of Curve pools, which depressed the CRV token, which threatened cascading liquidations across Aave, Fraxlend, and Abracadabra where CRV served as collateral for over \$100 million in loans ([Chaliasos et al., 2024](#)).

5.3 Crisis Case Study: Terra/Luna 2022

The Terra/Luna collapse of May 2022 activated at least five channels simultaneously, making it the most complex single-event cascade in digital finance history before the FTX episode. The crisis began when large UST redemptions on the Anchor Protocol triggered the algorithmic stablecoin's reflexive death spiral: UST de-pegging caused LUNA to be minted for redemptions, selling pressure on LUNA reduced UST's effective backing ratio, and further de-pegging accelerated the cycle in a self-reinforcing loop. This activated the stablecoin run channel with a mechanism reflexive algorithmic collapse that has no traditional analog and that existing Diamond-Dybvig models do not capture ([Liu et al., 2023](#); [Briola et al., 2023](#);

Uhlig, 2022).

When the Luna Foundation Guard sold approximately \$3.5 billion in Bitcoin reserves in a failed attempt to defend the peg, the stablecoin run activated the liquidity spiral channel directly. Forced selling depressed Bitcoin prices by over 25% in a single week, transmitting the shock from the Terra ecosystem to the broader cryptocurrency market and activating the liquidity spiral across all major trading pairs. Simultaneously, the liquidation cascade channel activated on the Anchor Protocol as collateral values collapsed, triggering tens of thousands of automated liquidations that further depressed LUNA prices and amplified the spiral (Shleifer and Vishny, 2011; Gudgeon et al., 2020; Klages-Mundt and Minca, 2022). Composability risk compounded the damage: protocols holding UST or LUNA as collateral or liquidity pool components experienced mechanical balance-sheet impairment across Ethereum, Avalanche, and Cosmos, regardless of their own code quality. Network contagion then propagated the shock from DeFi to CeFi Three Arrows Capital, which held concentrated positions in LUNA, stETH, and Grayscale Bitcoin Trust, discovered that its portfolio correlation was far higher during stress than its models assumed, a finding consistent with the correlation-in-crisis literature in traditional finance (Aramonte et al., 2022).

The destruction of approximately \$45 billion in Terra/Luna market capitalization within five days demonstrated that five-channel activation produces consequences qualitatively different from single-channel events. The sequential propagation to 3AC, Celsius, Voyager, and ultimately FTX over the following six months showed that multi-channel activation in one event creates the preconditions depleted capital buffers, heightened counterparty suspicion, reduced liquidity for multi-channel activation in subsequent events (International Monetary Fund, 2021; Liu et al., 2023).

5.4 Crisis Case Study: FTX 2022

The FTX collapse of November 2022 activated at least four channels simultaneously and demonstrated the counterparty-concentration-meets-information-asymmetry dynamic at its most destructive scale. FTX processed approximately 10% of global cryptocurrency spot volume and operated Alameda Research, one of the largest quantitative trading firms and market makers in the ecosystem. The CoinDesk report revealing Alameda’s concentrated holdings of FTT (FTX’s exchange token) triggered a bank run that drained \$6 billion in customer withdrawals within approximately 72 hours, far exceeding FTX’s available liquid assets (Duffie, 2010; Jalan and Matkovskyy, 2023; International Monetary Fund, 2022).

The counterparty concentration channel activated because FTX and Alameda collectively occupied a uniquely central position in the ecosystem’s trading, lending, custody, and venture networks. The information asymmetry channel activated because the commingling of approximately \$8 billion in customer funds, the fabrication of balance sheets, and the accumulation of undisclosed proprietary positions had been concealed behind corporate opacity and inadequate governance. The combined shock massive hidden risk-taking by the single most systemically central entity produced a cascade proportional to the gap between the market’s beliefs about FTX’s solvency and the reality of its balance sheet (Allen and Gale, 2000; Jalan and Matkovskyy, 2023; Upper, 2011).

Network contagion propagated through FTX’s lending and investment relationships: Genesis Global Trading halted withdrawals in November 2022 and filed for bankruptcy in

January 2023; BlockFi filed for bankruptcy within weeks of FTX; and dozens of portfolio companies lost both their primary market maker and a significant investor. The liquidity spiral channel activated as the collapse of FTT from approximately \$25 to under \$1 and the withdrawal of Alameda’s market-making liquidity produced a market-wide selling event that depressed total cryptocurrency capitalization by over \$200 billion. Gateway risk activated as Silvergate Bank, already stressed from deposit outflows following earlier 2022 failures, experienced further deterioration from FTX-related counterparty losses (Gorton and Zhang, 2023).

5.5 Crisis Case Study: Black Thursday 2020

Black Thursday, March 12, 2020, was the first systemic event to activate DeFi-specific channels and demonstrated that DeFi had inherited traditional liquidity-spiral dynamics while adding novel failure modes. The COVID-19 market panic triggered a 50% single-day Bitcoin decline that cascaded into DeFi through three channels simultaneously. The liquidation cascade channel activated as MakerDAO positions fell below collateralization thresholds, but Ethereum network congestion prevented normal liquidation auction participation, producing the zero-bid liquidation episode in which a small number of liquidators acquired \$8.3 million in ETH collateral for zero DAI. The liquidity spiral channel activated as forced selling on thinly-capitalized AMMs amplified price declines (Gudgeon et al., 2020).

The composability risk channel activated as protocols holding DAI or other MakerDAO-issued tokens experienced downstream disruption from the liquidation dysfunction. The event was smaller in absolute terms than the 2022 crises total DeFi TVL was approximately \$1 billion at the time but it provided the first empirical demonstration that DeFi had inherited the liquidity-spiral dynamics theorized by Brunnermeier and Pedersen (2009) while adding failure modes that traditional models do not contemplate: gas-price auctions that priced out legitimate participants, network congestion that prevented timely liquidation, and oracle-lag that caused stale prices to persist during rapid moves (Klages-Mundt and Minca, 2022).

Black Thursday activated three channels simultaneously. The 2022 crises subsequently demonstrated that four-channel and five-channel activation produces qualitatively more severe outcomes, confirming the superadditive relationship between channel count and crisis severity that the interaction taxonomy predicts. The escalating severity across the three case studies Black Thursday (three channels, \$ 8 million in direct protocol losses), Terra/Luna (five channels, \$ 45 billion in market capitalization destroyed), FTX (four channels, \$ 8 billion in customer funds lost) provides empirical support for the cascade threshold hypothesis (Ozili, 2022).

5.6 Market Structure and Cascade Acceleration

Three structural features of digital finance markets accelerate cross-channel cascades beyond the pace at which they would operate in traditional finance.

First, 24/7 global operation eliminates the cooling-off periods that overnight market closures provide in traditional equity and derivative markets. Consider the timeline: the Terra collapse unfolded substantially over a weekend when traditional markets were closed, and

the FTX bank run played out in approximately 48 hours without pause. Overnight halts in traditional finance provide time for stressed institutions to raise capital, negotiate credit lines, and coordinate with regulators; the absence of these pauses in digital finance removes a crisis-management tool upon which the traditional system relies implicitly ([Financial Stability Board, 2023a](#); [Chu and Rathbun, 2025](#)).

Second, on-chain transparency cuts both ways. During normal conditions, the ability of all market participants to observe liquidation thresholds, stablecoin reserve compositions, protocol dependency structures, and large-position concentrations in real time promotes market efficiency and price discovery. During stress events, however, that same observability creates incentives for preemptive selling and MEV extraction that accelerates the cascade transparency becomes procyclical, amplifying the dynamics participants seek to exploit or protect against. This procyclical transparency effect has no close analog in traditional finance, where position-level data is typically available only to the holder and its regulators ([Mohan, 2022](#)).

Third, pseudonymous participation prevents the kinds of targeted intervention moral suasion directed at specific institutions, coordinated creditor rollovers, selective capital injection, temporary forbearance agreements that central banks and financial regulators routinely employ to arrest cascading failures in traditional finance. When the identities of distressed entities are unknown or shielded behind wallet addresses, targeted intervention is structurally impossible even if the regulatory will and resources exist. The combination of continuous operation, real-time position observability, and pseudonymous participation creates an environment in which cascades propagate faster, attract more exploitative behavior, and terminate later than comparable cascades in traditional markets ([International Monetary Fund and Financial Stability Board, 2023](#); [Siam et al., 2025](#)).

5.7 Cross-Domain Contagion Pathways

How do cross-channel interactions translate into cross-domain contagion? Three pathways are currently active and well-documented; a fourth is emerging. The DeFi-to-CeFi pathway operates when DeFi protocol failures composability exploits, liquidation cascades generate price declines that impair CeFi firms' balance sheets and trigger CeFi counterparty failures, as the 2022 cascade from Anchor liquidations through Three Arrows Capital to Genesis, BlockFi, Celsius, and Voyager demonstrated ([Aramonte et al., 2022](#); [Financial Stability Board, 2022a](#)). Running in the opposite direction, the CeFi-to-DeFi pathway operates when CeFi firm failures remove market-making liquidity, trigger selling pressure on tokens held by the failed firm, and depress collateral values in DeFi lending protocols, thereby triggering on-chain liquidation cascades. A third pathway stablecoin-gateway contagion connects the stablecoin domain to the traditional financial system through reserve custody relationships. The SVB/USDC episode of March 2023 activated both the CeFi-to-DeFi and stablecoin-gateway pathways simultaneously, demonstrating their potential for co-activation and mutual reinforcement ([Gorton and Zhang, 2023](#); [Bank of England, 2021](#)).

As the tokenized TradFi domain grows, a fourth pathway connecting tokenized real-world assets to both DeFi composability risk and traditional-finance balance-sheet risk will emerge. A failure in a tokenization layer could transmit to traditional asset markets through forced selling of underlying assets, while a disruption in traditional markets could impair

tokenized collateral in DeFi protocols. This pathway has not yet been tested by a major crisis episode, but the structural elements required for its activation tokenized Treasuries used as DeFi collateral, institutional holders of tokenized products are accumulating (Bank for International Settlements, 2024; Cisar et al., 2025; International Monetary Fund, 2022).

5.8 Interaction Matrix and Channel Coupling Intensity

We quantify channel coupling intensity using a three-level scale based on crisis evidence. *Strong coupling* denotes bidirectional amplification documented in multiple independent crisis episodes. *Moderate coupling* denotes unidirectional amplification or bidirectional amplification documented in a single crisis. *Weak coupling* denotes theoretical interaction without clear crisis evidence. Of the 28 unique undirected channel pairs, 5 exhibit strong coupling: liquidity spirals–liquidation cascades, counterparty concentration–network contagion, stablecoin runs–gateway risk, information asymmetry–counterparty concentration, and composability risk–liquidation cascades. Eight pairs exhibit moderate coupling, and the remaining 15 pairs exhibit weak or negligible interaction (Huan and Renn, 2025; Jalan and Matkovskyy, 2023).

The five strongly coupled pairs share a common structural property: the output of one channel mechanically serves as the input to the other. Liquidation cascades produce forced selling that feeds liquidity spirals; network contagion transmits the failure of concentrated counterparties to their creditors; stablecoin runs deplete reserves held at gateway banks. This mechanical coupling where the transmission is automatic and requires no behavioral intermediation distinguishes strongly coupled pairs from moderately coupled pairs, where the interaction is mediated by behavioral responses such as panic selling or flight-to-quality that introduce variability and potential dampening effects (Acemoglu et al., 2015; Elliott et al., 2014; Aramonte et al., 2022).

Crisis evidence suggests a nonlinear relationship between the number of simultaneously active channels and the severity of the outcome. Single-channel events such as a protocol exploit activating only composability risk typically remain contained to the directly affected protocols. Two-channel events such as Black Thursday 2020 activating liquidation cascades and liquidity spirals produce significant but ultimately recoverable losses.

Beyond that threshold, outcomes change qualitatively. Events that activate three or more channels Terra/Luna and FTX produce systemic consequences with lasting structural impact on the ecosystem. We characterize this pattern as a *cascade threshold*: below three simultaneous channels, the system absorbs the shock; above three, the system amplifies it through cross-channel feedback (Liu et al., 2023; Huan and Renn, 2025).

5.9 Implications for Systemic Risk Assessment

What does the cross-channel analysis imply for systemic risk assessment and regulatory design? Three consequences stand out. First, channel-level monitoring, while necessary, is insufficient for capturing system-level risk. A monitoring dashboard that tracks each channel’s stress indicators in isolation but ignores pairwise coupling intensities will systematically underestimate risk precisely during the periods when accurate assessment matters most the

lead-up to multi-channel cascade events (Financial Stability Board, 2023a; International Monetary Fund, 2023; International Organization of Securities Commissions, 2023).

Second, and perhaps most actionable, the cascade threshold finding implies that regulatory interventions targeting the most strongly coupled channel pairs may yield disproportionate system-level benefits. Breaking or dampening even one link in a strongly coupled pair for example, requiring minimum collateralization ratios that reduce the severity of the liquidation-cascade-to-liquidity-spiral coupling could prevent the transition from two-channel to three-channel activation that transforms containable events into systemic crises (Financial Stability Board, 2023b).

Third, the cross-domain contagion pathways require that monitoring and regulation span the domain boundaries that currently fragment the regulatory landscape. The interaction between stablecoin runs and gateway risk, for instance, requires coordinated monitoring across crypto-specific regulators and banking supervisors a coordination that current institutional arrangements in most jurisdictions do not provide. The taxonomy’s mechanism-based organization, which categorizes channels by how distress transmits rather than by which domain it originates in, offers a template for structuring such cross-domain regulatory coordination (Schueffel, 2025).

6 Evolutionary Dynamics of Systemic Risk

Since Bitcoin’s launch, the systemic risk profile of digital finance has evolved substantially driven by changes in market structure, institutional complexity, and the scale of interconnections both within the ecosystem and between digital and traditional finance (Ozili, 2022). This section examines how the eight taxonomy channels have emerged, intensified, and interacted across four distinct eras, providing context for assessing which channels are growing in systemic significance, which are declining, and which represent permanent structural features of the architecture.

6.1 Early Era: Exchange Failures and Bilateral Risk (2011–2017)

Centralized exchange failures and bilateral counterparty exposures dominated the first era. The ecosystem was small in absolute terms Bitcoin’s total market capitalization did not exceed \$20 billion until 2017 but highly concentrated, with a small number of exchanges intermediating the vast majority of global trading volume. Only three of the eight taxonomy channels were materially active: counterparty concentration at the exchange layer, information asymmetry regarding exchange reserves and security practices, and gateway risk arising from the fragile banking relationships that connected early exchanges to the fiat financial system (Carpentier-Desjardins et al., 2025).

Consider how concentrated the early market was: Mt. Gox handled approximately 70% of global Bitcoin trading volume at the time of its February 2014 failure. Its bankruptcy constituted a systemic event for the entire Bitcoin market despite the market’s small absolute size. The bankruptcy revealed catastrophic failures in custodial security and reserve verification, activating the counterparty concentration and information asymmetry channels simultaneously, while the severing of Mt. Gox’s banking relationships with alternative fiat

gateways scarce provided an early manifestation of gateway risk (Jalan and Matkovskyy, 2023).

In June 2016, the DAO Hack introduced the composability risk channel to the digital finance vocabulary though the DeFi ecosystem in which composability risk would become systemically significant did not yet exist. The reentrancy exploit that drained \$60 million from The DAO’s smart contract demonstrated that vulnerabilities in composed smart contract interactions could serve as attack vectors with cascading consequences. The subsequent hard fork, splitting Ethereum into Ethereum and Ethereum Classic, established that governance mechanisms for handling composability failures were themselves sources of systemic uncertainty, a dynamic that would recur in more complex forms as DeFi matured (Chaliasos et al., 2024,?).

Regulatory enforcement could also trigger systemic channels, as the BTC-e seizure of July 2017 showed. The exchange’s abrupt shutdown by US law enforcement stranded user funds and disrupted trading, particularly in Eastern European markets that relied on BTC-e for liquidity, foreshadowing the regulatory dimension of gateway risk that would become prominent in 2023 with the closures of Silvergate, SVB, and Signature Bank (Foley et al., 2019; Bank for International Settlements, 2021b). Three of the eight taxonomy channels composability risk, liquidity spirals, and liquidation cascades remained largely dormant during this era, as the DeFi protocols through which these channels operate had not yet been developed at meaningful scale.

6.2 DeFi Summer and Emerging Complexity (2020–2021)

From approximately \$1 billion in June 2020 to over \$100 billion by November 2021 two orders of magnitude in eighteen months. The rapid growth of DeFi fundamentally altered the systemic risk landscape by introducing composability, on-chain leverage, and automated liquidation as structural features of the digital finance ecosystem, creating for the first time a system with sufficient scale and interconnection to produce genuinely systemic events. The pace of this growth outstripped the development of monitoring tools, risk frameworks, and regulatory attention (Bank for International Settlements, 2021b; International Monetary Fund, 2021; Alamsyah et al., 2024).

Black Thursday, March 12, 2020, was the first systemic event to activate DeFi-specific channels. The COVID-19 market panic triggered a 50% single-day decline in Bitcoin that cascaded into DeFi through the liquidation cascade and liquidity spiral channels. MakerDAO’s liquidation auctions malfunctioned under Ethereum network congestion, producing the zero-bid liquidation episode that demonstrated the fragility of on-chain liquidation mechanisms under stress. The event revealed that DeFi had inherited the liquidity-spiral dynamics theorized by Brunnermeier and Pedersen (2009) while adding novel failure modes gas-price auctions, network congestion, and oracle lag that the traditional theoretical models do not contemplate (Klages-Mundt and Minca, 2022; Gudgeon et al., 2020).

A proof of concept for the algorithmic stablecoin death spiral arrived with the Iron Finance/TITAN collapse of June 2021 one year before Terra/Luna. IRON’s partially algorithmic stablecoin lost its peg when large redemptions triggered the reflexive mint-and-sell dynamic in its TITAN collateral token, which fell from \$65 to near zero within hours, activating the stablecoin run, liquidity spiral, and composability risk channels simultaneously

(Briola et al., 2023; Mohan, 2022). Two months later, the Poly Network hack extracted \$611 million through a cross-chain bridge vulnerability, demonstrating that composability risk extended across blockchain boundaries and that bridge protocols concentrated systemic risk at the cross-chain layer.

By the end of 2021, all eight channels in the taxonomy were active. The interconnections between them and particularly the linkages between DeFi and CeFi through shared counterparty exposures and correlated token holdings had not yet been tested by a severe stress event, but the ecosystem had grown to a scale where a multi-channel cascade was not only possible but, given the accumulation of leverage, opacity, and concentrated counterparty exposures in the CeFi layer, increasingly probable (Chaliasos et al., 2024; Siam et al., 2025).

6.3 The Great Unraveling (2022)

Then came 2022. All eight channels in the taxonomy activated simultaneously over the course of the year, producing the most consequential cascade in digital finance history (documented in detail in Section 5). The sequence Terra/Luna in May, Three Arrows Capital in June, Celsius and Voyager in June–July, and FTX in November demonstrated that the channels do not merely coexist but compound in ways that transform individual failures into system-wide crises (Ozili, 2022). What made the 2022 crisis sequence exceptional was not only its scale total losses exceeding \$60 billion in direct destruction of market capitalization but the diversity of channels activated and the speed of cross-channel propagation: Terra/Luna combined stablecoin run dynamics with liquidity spirals and liquidation cascades; Three Arrows Capital demonstrated network contagion through opaque bilateral OTC lending; Celsius and Voyager revealed information asymmetry within the shadow-banking architecture of CeFi lending platforms; and FTX activated counterparty concentration, information asymmetry, network contagion, and gateway risk simultaneously (Aramonte et al., 2022).

Perhaps most revealing was the mechanism by which the 2022 cascade finally stopped not institutional intervention but exhaustion. Traditional financial systems have developed lender-of-last-resort facilities, deposit insurance, orderly resolution frameworks, and coordinated intervention protocols that arrest cascading failures before they become systemic. Digital finance possesses none of these mechanisms. By the end of 2022, the entities most vulnerable to contagion had already failed, and surviving entities had reduced exposures through forced deleveraging and withdrawal (Financial Stability Board, 2023a; Brunnermeier, 2009).

6.4 Maturation and Institutional Entry (2023–Present)

Since 2023, three developments have reshaped the systemic risk landscape: regulatory response, institutional entry, and market-structure evolution. The FSB finalized its high-level recommendations for crypto-asset regulation in July 2023, the EU’s Markets in Crypto-Assets (MiCA) regulation entered into force in 2024, and several jurisdictions introduced or strengthened licensing requirements for crypto exchanges and stablecoin issuers. These regulatory developments address the information asymmetry and counterparty concentration channels, though their effectiveness in preventing the next crisis remains to be tested in practice (Financial Stability Board, 2023b; International Organization of Securities Commissions, 2023; Schueffel, 2025).

Yet March 2023 demonstrated that gateway risk had not been addressed by the 2022 experience or the regulatory response that followed. Three crypto-friendly banks Silvergate, SVB, and Signature failed within a single month, eliminating critical fiat-to-crypto infrastructure and triggering the largest reserve-backed stablecoin de-peg to date, with USDC falling to \$0.87 before recovery (Galati and Capalbo, 2023). Because the gateway risk channel operates at the institutional boundary between traditional and digital finance, crypto-only regulation cannot mitigate it.

Institutional connections have continued to deepen. The approval of Bitcoin and Ethereum spot exchange-traded funds in the United States in 2024 and the growth of tokenized real-world assets create additional conduits through which shocks can transmit between digital and traditional finance, introducing traditional-finance participants pension funds, asset managers, insurance companies, banks as counterparties whose exposure to digital finance risk channels is mediated by regulated products but nevertheless real. The systemic significance of these connections will depend on the scale of institutional adoption and the adequacy of the regulatory frameworks governing the interface products (Bank for International Settlements, 2024; International Monetary Fund, 2024; Cisar et al., 2025).

6.5 Channel Evolution Trajectories

Across the four eras, the channels have followed distinct evolutionary trajectories. Counterparty concentration and information asymmetry, the dominant channels of the early era, remain significant but are now partially addressed by regulatory initiatives requiring licensing, disclosure, and the separation of customer assets. Network contagion has intensified as the ecosystem has grown and interconnections have multiplied, though the composition of the network has shifted from purely CeFi bilateral relationships to a hybrid of on-chain DeFi interactions and off-chain CeFi exposures (Financial Stability Board, 2023b).

Liquidity spirals and liquidation cascades, which emerged with DeFi in 2020, have become structurally embedded and will persist as long as on-chain lending with automated liquidation exists. Composability risk continues to grow in significance as the DeFi stack deepens, cross-chain bridges multiply, and yield-optimization strategies increase the complexity of protocol dependency graphs (Gudgeon et al., 2020; Mohan, 2022). Meanwhile, the stablecoin run channel has been partially addressed by regulatory attention to reserve composition and redemption mechanisms, but the fundamental fragility private money without deposit insurance or lender-of-last-resort support remains unresolved.

Where is systemic risk heading? Gateway risk is arguably increasing in systemic significance as institutional adoption creates broader exposure to the fiat-crypto boundary while the number of gateway institutions remains constrained by regulatory caution and reputational risk. The broader evolutionary trajectory suggests that the center of gravity is shifting from CeFi opacity and fraud (the dominant risk sources of the early and 2022 eras) toward structural risks embedded in DeFi architecture (composability, liquidation cascades) and boundary risks connecting digital and traditional finance (gateway risk). The taxonomy accommodates this evolution because channels are defined by transmission mechanism rather than by the institutional features of any particular era (Bank for International Settlements, 2024).

7 Emerging Channels and Forward-Looking Assessment

Grounded in observed crisis episodes and established theory, the eight taxonomy channels capture the current systemic risk landscape but digital finance continues to evolve in ways that may give rise to new channels or substantially alter existing dynamics (Kirişci, 2025). Five emerging risk areas have not yet produced a major systemic event but exhibit structural characteristics warranting prospective analysis. The section concludes by assessing the taxonomy’s durability in accommodating these developments.

7.1 Tokenized Real-World Assets and TradFi Bridges

By late 2024, tokenized US Treasuries alone exceeded \$1 billion in value, and major financial institutions including BlackRock, Franklin Templeton, and JPMorgan had launched tokenized asset products. The tokenization of real-world assets (RWAs) government bonds, money market fund shares, corporate credit, real estate, and equities represents the most significant structural development in the digital finance landscape since the emergence of DeFi, creating novel systemic risk pathways that cut across the existing taxonomy (Bank for International Settlements, 2024; Cisar et al., 2025,?).

Where does the primary risk lie? When a tokenized Treasury bond serves as collateral in a DeFi lending protocol, the position inherits both the credit risk of the underlying asset and the smart contract risk, composability risk, and liquidation cascade risk of the DeFi layer. A failure in the DeFi layer an exploit, a liquidation cascade, or a composability failure could force the sale of tokenized Treasuries, transmitting selling pressure from DeFi to the traditional Treasury market. Conversely, a disruption in the traditional asset market a sudden repricing of credit risk, a settlement failure, or a custodial disruption would impair the collateral value of tokenized positions in DeFi, potentially triggering on-chain liquidation cascades (International Monetary Fund, 2022; Alamsyah et al., 2024).

Redemption introduces an additional fragility with no precedent in native crypto markets. Unlike native tokens traded continuously on-chain, redemption of a tokenized Treasury for the underlying asset requires coordination between the on-chain tokenization layer and the off-chain custodial and settlement infrastructure. Delays, failures, or capacity constraints in this process could cause the tokenized asset to trade at a discount to its underlying value a de-peg analogous to stablecoin de-pegging triggering the run dynamics analyzed in Section 4.3 but applied to a broader and potentially more systemically significant class of assets (Bank for International Settlements, 2024).

7.2 CBDC Interoperability Risks

More than 130 jurisdictions are actively developing or pilot-testing central bank digital currencies (CBDCs), with the digital yuan, the digital euro, and the Bank of England’s digital pound at various stages of design and consultation. Although CBDCs are designed to reduce systemic risk by providing a risk-free digital settlement asset, the interoperability arrangements required to connect them with existing payment systems, commercial bank money, and potentially DeFi protocols introduce novel risk channels that extend the taxonomy’s gateway

risk and stablecoin run categories (Bank for International Settlements, 2021a; International Monetary Fund, 2024; Bordo and Levin, 2017).

Cross-border CBDC interoperability enabling atomic or near-atomic settlement of foreign exchange transactions using CBDCs from different jurisdictions requires technical bridges that resemble the cross-chain bridges in DeFi, inheriting similar composability and concentration risks. A failure in the interoperability layer could disrupt cross-border payment flows, with systemic consequences that scale with the volume of transactions processed through the CBDC infrastructure (Bank for International Settlements, 2023; Sethaput and Innet, 2023; Chu and Rathbun, 2025).

A second source of risk emerges from the interaction between CBDCs and private stablecoins. If CBDCs achieve widespread adoption, they may partially or fully displace private stablecoins as the settlement layer of digital finance. During the transition period when CBDCs and private stablecoins coexist and compete holders might rapidly shift from private stablecoins to CBDCs, triggering stablecoin runs driven not by concerns about reserve adequacy but by rational substitution toward a risk-free alternative. Design choices in CBDC holding limits, remuneration rates, and conversion mechanisms will critically shape the stability of this coexistence equilibrium (Bank of England, 2021; Board of Governors of the Federal Reserve System, 2022b; Allen et al., 2022).

7.3 AI-Driven Trading and Algorithmic Herding

What happens when dozens of trading firms deploy AI models trained on similar data sets, using similar neural network architectures, and optimizing similar objective functions? The resulting strategies may exhibit correlated behavior that amplifies market movements in ways that no individual model was designed to produce a new form of systemic risk best described as algorithmic herding. The risk arises not from any single AI system malfunctioning but from the collective behavior of many AI systems producing emergent instability through correlated action (Olanrewaju, 2025; Auer et al., 2023).

Two pathways drive this mechanism. First, AI-driven trading strategies that rely on similar signals on-chain transaction data, social media sentiment, order flow patterns may generate correlated buy or sell decisions that amplify price movements beyond what fundamental conditions warrant, accelerating liquidity spirals by removing liquidity precisely when human participants are also withdrawing. Second, AI systems designed to optimize execution speed may interact with MEV extraction infrastructure in ways that amplify the game-theoretic instability of on-chain markets, increasing the frequency and severity of adverse transaction ordering during liquidation cascades (Daian et al., 2020; Mohan, 2022).

Compounding these dynamics, AI opacity deepens the information asymmetry channel identified in Section 4.7. When a trading firm’s strategy is determined by a neural network whose internal logic is not interpretable even by its operators, counterparties face a form of information asymmetry that extends beyond traditional adverse selection: they cannot assess the nature, magnitude, or correlation structure of the risk being taken because the risk-taking process itself is opaque (Akerlof, 1970). Regulatory frameworks designed for human-interpretable trading strategies may prove inadequate for supervising AI-driven market participation in digital finance.

7.4 Cross-Chain Composability and Layer-2 Fragmentation

Optimism, Arbitrum, Base, zkSync, Solana, Avalanche, Sui the proliferation of Layer-2 scaling solutions and alternative Layer-1 blockchains has created a fragmented landscape in which liquidity, DeFi protocols, and user activity are distributed across dozens of semi-independent execution environments. Cross-chain bridges connect these environments, but as the bridge exploits documented in Section 4.4 demonstrate, bridges are among the most vulnerable components of the digital finance infrastructure (Jourenko et al., 2025; Chaliasos et al., 2024; Siam et al., 2025).

This fragmentation introduces a novel form of liquidity risk that extends the taxonomy’s liquidity spiral channel. Assets that appear liquid in aggregate across all chains may be illiquid on any individual chain, and the bridges that enable cross-chain liquidity movement may experience congestion, delays, or outright failures during stress events meaning a liquidation cascade on one Layer-2 could be amplified by the inability to source liquidity from other chains in time to absorb the selling pressure.

A further concentration risk lurks in sequencer infrastructure. Most Layer-2 networks rely on a single sequencer operator to order transactions, creating a single point of failure for all DeFi activity on the chain (Chaliasos et al., 2024,?).

Cross-chain composability the ability for a smart contract on one chain to invoke functions on another chain through messaging protocols extends the composability risk channel across blockchain boundaries. A composed position that spans multiple chains inherits the security properties of each chain, each bridge, and each messaging protocol in the dependency chain. The weakest link determines the security of the entire position, and the combinatorial complexity of multi-chain dependency graphs exceeds that of single-chain composability by orders of magnitude, creating a risk surface that no current monitoring tool can fully map (Alamsyah et al., 2024).

7.5 Framework Durability

Can the taxonomy survive the market evolution described above? Its organizing principle transmission mechanism rather than institutional features or specific technologies is designed for exactly this purpose. Network contagion, liquidity spirals, stablecoin runs, composability risk, liquidation cascades, counterparty concentration, information asymmetry, and gateway risk describe fundamental pathways through which distress propagates in interconnected financial systems. While the specific institutional manifestation of each channel will evolve as the market structure changes, the underlying transmission mechanisms are structural properties of any financial system that combines leverage, interconnection, maturity transformation, and information asymmetry (Diamond and Dybvig, 1983).

The emerging channels analyzed in this section do not require additions to the taxonomy so much as extensions of existing channels to new institutional settings. Tokenized RWA risk extends composability risk and gateway risk to a new asset class. CBDC interoperability risk extends gateway risk and stablecoin run dynamics to a new institutional arrangement. AI-driven herding extends liquidity spirals and information asymmetry to a new market-making technology. Cross-chain fragmentation extends composability risk and liquidation cascades to a new infrastructure architecture. The taxonomy’s mechanism-based organizing

principle ensures that these extensions can be accommodated within the existing framework without requiring new channels, provided that the emerging risks do not introduce genuinely novel transmission mechanisms a condition that, as of the analysis cutoff date of March 2026, appears to hold (Bank for International Settlements, 2024; Financial Stability Board, 2023a).

8 Policy Implications

Beyond its analytical contributions, the taxonomy carries implications for the design of regulatory frameworks, monitoring architectures, and institutional arrangements for digital finance (van der Linden and Shirazi, 2023). We draw three categories of policy implications: regulatory gaps revealed by the channel structure, monitoring challenges arising from the heterogeneity of channel observability, and institutional design considerations for the evolving landscape. The discussion is deliberately measured this is a research paper, not a policy paper, and the implications are presented as analytical findings rather than prescriptive recommendations.

8.1 Regulatory Gaps Highlighted by the Taxonomy

Not all channels receive equal regulatory attention. The FSB’s high-level recommendations for crypto-asset regulation, finalized in July 2023, focus primarily on the counterparty concentration and information asymmetry channels, requiring licensing of crypto service providers, segregation of customer assets, disclosure of reserve compositions, and governance standards for stablecoin issuers. The EU’s Markets in Crypto-Assets regulation addresses similar channels, with specific provisions for stablecoin issuers (reserve requirements, redemption rights) and crypto-asset service providers (authorization, conduct rules, prudential requirements) (Financial Stability Board, 2023b,a).

Stablecoin runs receive more targeted treatment. The Bank of England’s framework for systemic stablecoins directly targets this channel through requirements for full reserve backing with liquid assets, liquidity management standards, and orderly wind-down planning, while IOSCO’s policy recommendations complement these efforts with standards for market intermediaries and trading platforms. Taken together, these regulatory initiatives address the channels that dominated the 2022 crisis sequence counterparty concentration, information asymmetry, and stablecoin runs but leave significant gaps in channels structurally embedded in DeFi architecture (Bank of England, 2023; Financial Stability Board, 2023c; Allen et al., 2022).

Where are the gaps? Composability risk the most digitally native channel, with no traditional analog is largely unaddressed by existing regulatory frameworks. The permissionless nature of DeFi protocol composition means that dependency chains form without authorization or oversight, and no current regulatory proposal requires protocols to map, disclose, or limit their dependency exposures. Liquidation cascades present a similar blind spot: DeFi lending protocols set their own liquidation thresholds, collateralization ratios, and liquidation bonuses without regulatory input, and no framework exists for stress-testing these parameters against historical tail-risk scenarios such as Black Thursday or the Terra collapse

(International Organization of Securities Commissions, 2022; European Central Bank, 2023; Chaliasos et al., 2024).

Gateway risk occupies perhaps the most consequential blind spot, because it operates at the boundary between two regulatory perimeters. Crypto-specific regulators (where they exist) focus on exchange licensing, stablecoin reserves, and custody standards; banking regulators focus on capital requirements, deposit insurance, and liquidity management. Neither side fully captures the bidirectional contagion dynamics that the gateway channel creates, and no framework coordinates the regulatory response across both sides of the boundary.

The SVB/USDC episode made this gap concrete: a banking-perimeter failure transmitted to the crypto perimeter through stablecoin reserves. The Silvergate episode demonstrated the reverse transmission (Board of Governors of the Federal Reserve System, 2022b; International Monetary Fund and Financial Stability Board, 2023).

A structural pattern emerges from this gap analysis. Channels most effectively addressed by regulation are those with traditional-finance analogs for which regulatory templates already exist (counterparty concentration, information asymmetry, stablecoin runs). Channels least effectively addressed are those that are digitally native (composability risk), that combine on-chain and off-chain dynamics in novel ways (liquidation cascades), or that span regulatory boundaries (gateway risk). Regulatory effectiveness in digital finance, it appears, is constrained not primarily by political will but by the availability of conceptual templates from traditional financial regulation (Werner et al., 2022; Alamsyah et al., 2024; Kirişci, 2025).

8.2 Macroprudential Tools for Digital Finance

Because the systemic risk of digital finance exceeds the sum of its individual channels (Section 5), channel-specific regulation while necessary is insufficient. The regulatory architecture must also address the interactions between channels that produce multi-channel cascades. Macroprudential oversight, which takes a system-wide view rather than an entity-specific or instrument-specific view, is the natural institutional framework for this task (International Monetary Fund, 2023; Huan and Renn, 2025).

Countercyclical leverage limits represent the most direct tool. DeFi lending protocols currently allow collateralization ratios as low as 110% for certain assets, creating positions vulnerable to liquidation from modest price declines. Requiring minimum collateralization ratios calibrated to asset volatility and increasing those ratios during periods of rapid TVL growth or declining market liquidity would reduce the density of positions near liquidation thresholds and dampen the liquidation-cascade-to-liquidity-spiral feedback loop identified in Section 5.2 (Bank for International Settlements, 2021b; Alamsyah et al., 2024).

Concentration thresholds offer a second tool. Since the 2008 crisis, traditional finance has designated systemically important financial institutions (SIFIs) and subjected them to enhanced capital requirements; digital finance has no equivalent. A threshold-based framework triggered by market share in trading volume, lending activity, custodial assets, or stablecoin issuance would enable preemptive monitoring and enhanced prudential requirements for entities whose failure would activate multiple channels simultaneously (Auer et al., 2023; Jalan and Matkovskyy, 2023).

Third, stress-testing requirements adapted to digital finance would force both CeFi intermediaries and DeFi protocols to demonstrate resilience against historical tail scenarios. 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.

Circuit breakers raise distinct implementation challenges. Protocol-level or network-level circuit breakers are widely used in traditional equity and derivative markets but absent in 24/7 digital finance. Who triggers the breaker in a permissionless protocol? How does a single-chain pause interact with cross-chain positions? These questions warrant dedicated research ([Financial Stability Board, 2023b](#); [Mohan, 2022](#)).

8.3 Microprudential Considerations

Turning from system-level to entity-level regulation: for stablecoin issuers, the taxonomy analysis supports requirements for full reserve backing with high-quality liquid assets, real-time reserve attestation (moving beyond periodic audits), redemption-right guarantees with defined settlement timeframes, and orderly wind-down planning. Because algorithmic stablecoins face reflexive collapse dynamics that reserve requirements cannot address (Section 4.3), the policy implication is that algorithmic designs warrant either prohibition or substantially higher regulatory scrutiny than reserve-backed designs ([Financial Stability Board, 2023c](#); [Schueffel, 2025](#)).

As the FTX case study (Section 5.4) demonstrated, centralized exchanges require mandatory customer asset segregation verified by independent custodians, regular independent audits, restrictions on proprietary trading by exchange-affiliated entities, and disclosure of material affiliations between exchanges and market makers. MiCA and comparable frameworks in other jurisdictions are beginning to implement these measures ([International Organization of Securities Commissions, 2023](#); [Carpentier-Desjardins et al., 2025](#)).

For DeFi protocols, microprudential regulation faces the fundamental challenge that many protocols operate without legal entities, identifiable operators, or jurisdictional nexus. Protocol-level governance standards including mandatory security audits by independent firms, formal verification of smart contract logic for critical functions, tiered collateralization requirements calibrated to asset volatility, and emergency pause mechanisms with defined activation criteria represent a middle ground between full entity-level regulatory coverage and the current unregulated state. The composability risk analysis in Section 4.4 further supports requirements for dependency-graph disclosure and downstream-risk notification to protocols that hold tokens issued by other protocols ([Chaliasos et al., 2024](#); [Siam et al., 2025](#)).

8.4 Monitoring Architecture

Monitoring is only as good as the data it can access and the eight taxonomy channels differ substantially in their observability. Three channels liquidity spirals, liquidation cascades, and composability risk operate primarily on-chain and are, in principle, fully observable through blockchain analytics: collateralization ratio distributions, AMM liquidity pool depth and

composition, and composed protocol dependency structures can all be reconstructed from public blockchain data in real time.

For these channels, the monitoring challenge is not data availability but analytical capacity developing early-warning indicators from the vast volume of on-chain data (Gudgeon et al., 2020; Qin et al., 2022; Adamyk et al., 2025).

Off-chain channels present a fundamentally harder problem. Counterparty concentration, information asymmetry, and network contagion in the CeFi layer involve exposures that blockchain analytics cannot capture and that regulatory reporting requirements only partially address. The bilateral OTC lending relationships, custodial arrangements, and proprietary trading positions constituting the off-chain network are disclosed only voluntarily (if at all), typically surfacing only in bankruptcy proceedings as the 2022 crisis sequence demonstrated. Monitoring these channels requires either mandatory reporting analogous to trade repositories for OTC derivatives in traditional finance, or inferential techniques that estimate off-chain exposures from observable on-chain data and market prices (Auer et al., 2023; Financial Stability Board, 2023b; Brunnermeier, 2009).

Gateway risk demands monitoring that spans both the crypto and banking regulatory perimeters tracking the concentration of crypto firm deposits at individual banks, the composition and custody arrangements of stablecoin reserves, and the volume of fiat settlement flowing through gateway institutions. The data required for effective gateway monitoring exists: banks report deposit concentrations to banking regulators, and stablecoin issuers disclose reserve compositions to their respective overseers. What is missing is consolidation. The information is fragmented across multiple regulatory agencies and jurisdictions and is not currently assembled into a unified framework capable of detecting bidirectional contagion dynamics (International Monetary Fund and Financial Stability Board, 2023; Board of Governors of the Federal Reserve System, 2022a).

Stablecoin runs occupy an intermediate position in the observability landscape. On-chain behavior redemptions, secondary-market trading, DeFi withdrawals, concentration patterns is fully observable through blockchain analytics, but reserve adequacy depends on off-chain information about reserve composition, custodial arrangements, and the liquidity of reserve assets under stress. Bridging these two data sources is the core challenge: current monitoring architectures, designed for either on-chain or off-chain data but not both simultaneously, are not equipped for the data fusion that effective stablecoin-run surveillance requires (Bank for International Settlements, 2021b).

8.5 International Coordination Challenges

Regulatory arbitrage is not hypothetical. Digital finance operates across jurisdictional boundaries with minimal friction, and as the post-2022 migration of exchange volume to offshore venues demonstrated, a framework effective in one jurisdiction may simply redirect activity to less regulated jurisdictions. The FSB, BIS, and IOSCO have issued coordinated recommendations, but implementation timelines, scope, and enforcement mechanisms remain fragmented and uneven (Financial Stability Board, 2023a; International Monetary Fund and Financial Stability Board, 2023; International Organization of Securities Commissions, 2023; Jalan and Matkovskyy, 2023).

For DeFi protocols which may have no identifiable legal entity, no physical presence, and

no clear regulatory nexus the jurisdictional challenge is particularly acute. Activity-based regulation, targeting the function performed (lending, trading, custody) rather than the entity performing it, offers a conceptual solution, but enforcement against permissionless protocols operating without identifiable operators remains a fundamental challenge. The debate is converging toward hybrid approaches that combine entity-level regulation for identifiable CeFi intermediaries with activity-level standards for the broader DeFi ecosystem (Schueffel, 2025; European Central Bank, 2023; International Organization of Securities Commissions, 2022).

Three coordination priorities emerge from the taxonomy analysis. First, harmonized standards for stablecoin regulation to prevent runs driven by cross-jurisdictional reserve fragmentation or regulatory arbitrage in issuer domiciles. Second, coordinated gateway monitoring to track bidirectional contagion across the traditional-digital boundary, requiring information sharing between crypto regulators and banking supervisors across jurisdictions. Third, mutual recognition frameworks for exchange licensing to reduce regulatory arbitrage while maintaining minimum prudential standards for customer protection and systemic risk management (Auer et al., 2023; Chu and Rathbun, 2025).

8.6 Institutional Design Priorities

Three institutional design priorities emerge from the taxonomy analysis. First and most urgent is a consolidated monitoring framework integrating on-chain analytics, off-chain regulatory reporting, and cross-perimeter gateway data into a unified systemic risk surveillance capability. The current fragmentation on-chain data available to anyone with analytical tools, off-chain data siloed within specific regulatory agencies, gateway data split between crypto and banking supervisors prevents the holistic risk assessment that cross-channel interactions demand (International Monetary Fund, 2023; Financial Stability Board, 2023a; Adamyk et al., 2025).

Second, macroprudential tools countercyclical leverage limits, concentration thresholds, stress-testing requirements, and potentially circuit-breaker mechanisms must be extended to the digital finance ecosystem, calibrated to the specific dynamics of the eight taxonomy channels. Adapting these tools is nontrivial: traditional macroprudential instruments were designed for banking systems with known counterparties, periodic reporting cycles, and standing regulatory intervention capacity. Digital finance’s 24/7, pseudonymous, partially permissionless architecture requires substantial methodological and institutional innovation (Auer et al., 2023).

Third, the development of resolution frameworks for systemically important digital finance entities particularly stablecoin issuers, dominant exchanges, and critical infrastructure providers such as bridge operators and oracle networks that provide for orderly failure management without requiring the extraordinary government interventions (such as the FDIC’s guarantee of SVB deposits) that current crises have necessitated. The FTX bankruptcy proceedings, conducted under ordinary corporate bankruptcy law without the specialized tools available for bank resolution, demonstrated the inadequacy of existing frameworks for managing the failure of a systemically important crypto entity (Financial Stability Board, 2023c; Bank of England, 2023).

9 Conclusion

Digital finance has transitioned from a peripheral experiment to a financial system whose stability constitutes a matter of macroprudential concern (Bank for International Settlements, 2022). This paper develops a comprehensive taxonomy of eight systemic risk transmission channels that operate across four domains of the digital finance ecosystem decentralized finance, centralized crypto platforms, stablecoin ecosystems, and the emerging tokenized traditional finance sector and analyzes their interactions, evolutionary dynamics, and policy implications.

9.1 Taxonomy Contributions

Two genuinely novel channels, four hybrid channels, and two extensions of established theory to digital settings emerge from the taxonomy. Composability risk and gateway risk are novel. Composability risk arises from the permissionless composition of smart contract protocols, creating dependency chains that propagate failures mechanically and at blockchain speed a mechanism with no traditional analog. Gateway risk arises from the bidirectional contagion dynamics at the institutional boundary between traditional and digital finance, where a small number of banks and intermediaries serve as chokepoints through which distress transmits in both directions. Both novel channels lack formal theoretical treatment and represent the paper’s primary conceptual contributions to the systemic risk literature (Allen and Gale, 2000; Acemoglu et al., 2015; Werner et al., 2022).

Each of the four hybrid channels liquidity spirals, stablecoin runs, liquidation cascades, and information asymmetry extends a traditional mechanism with significant digital-native features that alter its speed, determinism, and amplification structure. Liquidity spirals operate through AMM bonding curves and automated liquidation bots rather than through human market makers exercising judgment. Stablecoin runs proceed without the deposit insurance, sequential-service constraints, or lender-of-last-resort backstops that moderate traditional bank runs. Liquidation cascades are executed by permissionless bots competing for MEV rather than by margin departments exercising forbearance. Information asymmetry manifests in a hybrid environment where on-chain DeFi activity is fully transparent but off-chain CeFi operations frequently operate with opacity exceeding that of traditional financial institutions (Diamond and Dybvig, 1983; Gudgeon et al., 2020).

Network contagion and counterparty concentration represent direct extensions of established theory to digital settings. The contagion models of Allen and Gale (2000) and Acemoglu et al. (2015) apply to the digital finance network, but the network’s structure is a hybrid of on-chain (observable) and off-chain (opaque) connections that no existing model integrates into a unified framework. Counterparty concentration follows the too-big-to-fail logic analyzed by Upper (2011), but digital finance concentration spans regulated and unregulated entities across multiple jurisdictions, complicating both measurement and mitigation (Elliott et al., 2014; Upper, 2011).

9.2 Key Findings

Three principal findings emerge from the analysis. First, the systemic risk of digital finance is not reducible to any single channel or domain. The 2022 crisis sequence spanning the Terra/Luna collapse, the Three Arrows Capital insolvency, the Celsius and Voyager failures, and the FTX collapse activated all eight channels and demonstrated that compounding interactions between channels produce systemic consequences far exceeding the sum of individual channel effects. The cross-channel interaction matrix documents these dynamics, identifying five strongly coupled pairs and a cascade threshold above which the system transitions from shock-absorbing to shock-amplifying (Liu et al., 2023; Financial Stability Board, 2022a; Auer et al., 2023).

Second, the structural features of digital finance composability, 24/7 operation, automated liquidation, pseudonymous participation, and algorithmic governance fundamentally alter the dynamics of systemic risk transmission. Even channels rooted in established theory (liquidity spirals, bank runs, counterparty cascades) operate differently in digital finance than in their traditional settings, typically with compressed timelines, mechanical rather than discretionary propagation, and amplification mechanisms such as MEV extraction and procyclical transparency that have no traditional analog. The taxonomy’s classification of each channel as novel, extension, or hybrid provides a precise mapping of how digital finance inherits, transforms, and creates systemic risk mechanisms (Klages-Mundt and Minca, 2022; Daian et al., 2020).

Third, the existing regulatory and monitoring infrastructure is unevenly matched to the channel structure of digital finance systemic risk. Channels that dominated the 2022 crisis sequence counterparty concentration and information asymmetry are increasingly addressed by emerging frameworks including the FSB recommendations and MiCA. By contrast, channels embedded in DeFi architecture (composability risk, liquidation cascades) remain largely outside the regulatory perimeter, and the gateway risk channel, which operates at the boundary between digital and traditional finance, falls between the jurisdictions of crypto regulators and banking supervisors.

Observability compounds this mismatch. Some channels are fully captured by on-chain data; others require off-chain disclosure that existing surveillance architectures are not designed to collect or integrate. Closing these gaps demands monitoring infrastructure built around the channel structure itself, rather than adaptations of reporting regimes designed for traditional intermediaries (Financial Stability Board, 2023b; International Monetary Fund, 2023; International Organization of Securities Commissions, 2023).

9.3 Limitations

The taxonomy is subject to several limitations that should be acknowledged. The analysis is bounded by a cutoff date of March 2026, and the digital finance landscape continues to evolve in ways that may alter the relative significance of the eight channels. The literature search, while systematic, is limited to works indexed in OpenAlex and policy repositories maintained by major international bodies; relevant research published in non-indexed venues, working papers not yet in the public domain, or policy documents from jurisdictions not covered by our repository may have been missed. The crisis evidence is necessarily retrospective, and

channels that have not yet been activated by a major crisis episode such as the tokenized RWA channel or the AI-driven herding channel may be underweighted relative to their prospective significance (Bank for International Settlements, 2024; Schueffel, 2025; Alamsyah et al., 2024).

The breadth of the taxonomy spanning eight channels across four domains necessarily limits the analytical depth devoted to any single channel. An alternative design analyzing fewer channels in greater formal depth would offer different advantages. We chose comprehensive coverage because the primary gap in the literature is the absence of an integrated framework identifying all major transmission channels and their interactions; deep formal analysis of individual channels is better addressed in subsequent focused studies building on the taxonomic foundation established here.

The composite scoring methodology used to select the eight channels from the initial set of fourteen candidates involves weighting choices that, while transparent and justified in Section 3, inevitably reflect analytical judgment. Alternative weighting schemes could produce different channel boundaries, though we note that the eight selected channels collectively account for the transmission mechanisms observed in every major crisis episode in our sample. The quantitative thresholds used in the channel selection process are calibrated to the current literature’s coverage density and may require adjustment as the literature matures and the crisis record lengthens (Jalan and Matkovskyy, 2023; Adamyk et al., 2025).

9.4 Future Research Directions

The taxonomy opens five avenues for future research. First, the development of formal theoretical models for the two novel channels composability risk and gateway risk would fill the most significant gaps in the current literature. A model of cascading failure in permissionlessly composed systems, capturing the combinatorial structure of dependency graphs and the deterministic propagation of smart contract failures, would provide the theoretical foundation for stress-testing DeFi protocols and quantifying the system-level risk that composability creates. A model of bidirectional contagion through gateway institutions, capturing the equilibrium dynamics of gateway provision under regulatory constraints and the conditions under which gateway concentration becomes systemically dangerous, 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 the specific dynamics of each channel extending CoVaR, MES, and SRISK to digital finance settings or developing new measures tailored to on-chain data availability would enable the empirical monitoring of channel-level and system-level risk in real time. The on-chain observability of DeFi channels creates opportunities for real-time risk measurement that traditional finance does not offer, but the methodological development required to translate raw blockchain data into actionable risk indicators has not kept pace with data availability (Qin et al., 2022; Auer et al., 2023; Mohan, 2022).

Third, two emerging technological developments will reshape the channel structure in ways that the current literature has not yet systematically analyzed. As the tokenized real-world asset market grows to a scale where crisis testing becomes possible, the interaction between tokenized RWAs and the existing channel structure requires dedicated study. Inte-

grating traditional-finance credit risk with DeFi composability risk and liquidation cascade dynamics creates contagion pathways whose theoretical properties remain uncharacterized and whose empirical behavior is untested including the scenario in which a DeFi-originated cascade transmits to traditional Treasury or corporate bond markets through tokenized-asset forced selling. Separately, the impact of artificial intelligence on trading behavior, information processing, and market stability warrants prospective analysis before AI-driven herding becomes a source of realized systemic risk; the interaction between AI-generated trading correlations and the liquidity-spiral and information-asymmetry channels could amplify known vulnerabilities in ways that current models do not capture (Bank for International Settlements, 2024; Cisar et al., 2025; Allen et al., 2022; Olanrewaju, 2025).

Fourth, empirical validation of the taxonomy could proceed along three pathways. Structured case-study analysis applying the channel framework to future crisis episodes would test whether the eight channels provide sufficient coverage of observed transmission mechanisms. Quantitative analysis using on-chain data could test the propositions regarding cross-channel amplification and digital-native acceleration by measuring contagion speed and loss magnitude across episodes with varying numbers of simultaneously active channels. Expert survey or Delphi-method assessment of channel boundaries and classification would provide inter-rater validation of the taxonomy structure, addressing the single-researcher coding limitation acknowledged in Section 3.4.

Fifth, the design of macroprudential tools and resolution frameworks specifically adapted to the multi-channel, cross-domain risk architecture documented in this taxonomy represents a critical policy research agenda. Traditional macroprudential tools assume known counterparties, periodic reporting, and the institutional capacity for regulatory intervention assumptions that digital finance’s 24/7, pseudonymous, partially permissionless architecture violates. Developing tools that function effectively within these constraints, while preserving the innovation benefits and financial inclusion potential of permissionless systems, is among the most consequential applied research questions in financial regulation today (Financial Stability Board, 2023a; International Monetary Fund, 2023; Financial Stability Board, 2023c).

9.5 Closing Perspective

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. Composability risk and gateway risk have no traditional analog; they are products of the architectural choices permissionless composition, concentrated institutional bridging that define digital finance as a distinct financial system. Liquidity spirals, bank runs, counterparty cascades, and information asymmetry are inherited from traditional finance, but digital finance transforms their speed, mechanism, and amplification structure in ways that existing theory captures only partially and that existing regulation addresses only incompletely (Akerlof, 1970; Schueffel, 2025).

Three policy priorities follow. Regulatory frameworks for the novel channels must be developed from first principles, since no traditional-finance analog supplies even a conceptual template. Macroprudential tools must be adapted to the accelerated dynamics of hybrid channels, particularly the liquidation-cascade and liquidity-spiral feedback loop that serves

as the primary amplification mechanism during crises. Cross-domain monitoring capacity must be built to match the cross-domain contagion pathways documented in the interaction analysis. Each priority requires innovation beyond the adaptation of existing tools ([Financial Stability Board, 2023a](#); [Allen et al., 2022](#)).

Across all of these applications, the analytical foundation offered here is the same: a mechanism-based account of how systemic risk propagates in digital finance, how risk channels evolve as market structure changes, and how regulatory and institutional responses can be calibrated to the risks that each channel poses. As digital finance continues to mature, as its connections to traditional finance deepen through tokenization and institutional adoption, and as novel technologies such as AI reshape market dynamics, the need for a rigorous, mechanism-based framework for systemic risk assessment will only grow. We offer this taxonomy as a contribution toward that framework ([Werner et al., 2022](#); [Auer et al., 2023](#); [International Monetary Fund, 2024](#)).

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. The authors also acknowledge financial support from the Swiss National Science Foundation 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. Their significant contribution has been instrumental in advancing the research and fostering collaboration within the digital finance field across Europe.

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.

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. We 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.

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

- Abdullah, Adam**, “Monetary Reform and Central Bank Digital Currencies: The Impact on Retail Banking,” 2024.
- Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi**, “Systemic Risk and Stability in Financial Networks,” *American Economic Review*, 2015, *105* (2), 564–608.
- Acharya, Viral V., Lasse Heje Pedersen, Thomas Philippon, and Matthew Richardson**, “Measuring Systemic Risk,” *Review of Financial Studies*, 2017, *30* (1), 2–47.
- Adamyk, Bogdan, Vladlena Benson, Bogdan Adamyk, and Oksana Liashenko**, “Risk Management in DeFi: Analyses of the Innovative Tools and Platforms for Tracking DeFi Transactions,” 2025.
- Adisa, Olawale, Bamidele Segun Ilugbusi, Ogugua Chimezie, Kehinde Fernami Awonuga, Odunayo Adewunmi Adelekan, Onyeka Franca Asuzu, and Ndubuisi Leonard Ndubuisi**, “Decentralized Finance (DEFI) in the U. S. economy: A review: Assessing the rise, challenges, and implications of blockchain-driven financial systems,” 2024.
- Adrian, Tobias and Hyun Song Shin**, “Liquidity and Leverage,” *Journal of Financial Intermediation*, 2010, *19* (3), 418–437.
- and **Markus K. Brunnermeier**, “CoVaR,” *American Economic Review*, 2016, *106* (7), 1705–1741.
- Akerlof, George A.**, “The Market for “Lemons”: Quality Uncertainty and the Market Mechanism,” *Quarterly Journal of Economics*, 1970, *84* (3), 488–500.
- Alamsyah, Andry, Gede Natha Wijaya Kusuma, and Dian Puteri Ramadhani**, “A Review on Decentralized Finance Ecosystems,” 2024.
- Allen, Franklin and Douglas Gale**, “Financial Contagion,” *Journal of Political Economy*, 2000, *108* (1), 1–33.
- , **Xian Gu, and Julapa Jagtiani**, “Fintech, Cryptocurrencies, and CBDC: Financial Structural Transformation in China,” 2022.
- Andryushin, S.**, “Tokenization of real assets: classification, platforms, applications, opportunities and challenges of development,” 2024.
- Ante, Lennart and Aman Saggi**, “Time-Varying Bidirectional Causal Relationships between Transaction Fees and Economic Activity of Subsystems Utilizing the Ethereum Blockchain Network,” 2024.
- , **Ingo Fiedler, Jan Marius Willruth, and Fred Steinmetz**, “A Systematic Literature Review of Empirical Research on Stablecoins,” 2023.

- Antonakakis, Nikolaos, Ioannis Chatziantoniou, and David Gabauer**, “Refined Measures of Dynamic Connectedness based on Time-Varying Parameter Vector Autoregressions,” 2020.
- Aramonte, Sirio, Wenqian Huang, and Andreas Schrimpf**, “DeFi Leverage,” Quarterly Review Article BIS Quarterly Review, December 2022, Bank for International Settlements 2022.
- Arora, Sanidhay, Yingjiu Li, Yebo Feng, and Jiahua Xu**, “SecPLF: Secure Protocols for Loanable Funds against Oracle Manipulation Attacks,” 2024.
- Auer, Raphael, Giulio Cornelli, Sebastian Doerr, Jon Frost, Leonardo Gambacorta, Jon Frost, and Leonardo Gambacorta**, “Crypto Trading and Bitcoin Prices: Evidence from a New Database of Retail Adoption,” 2023.
- Babus, Ana**, “The Formation of Financial Networks,” *RAND Journal of Economics*, 2016, 47 (2), 239–272.
- Bakare, Felix Adebayo, J. T. Omojola, and Augustine Chibuzor Iwuh**, “Blockchain and decentralized finance (DEFI): Disrupting traditional banking and financial systems,” 2024.
- Bank for International Settlements**, “CBDCs: An Opportunity for the Monetary System,” Annual Report Chapter Annual Economic Report 2021, Chapter III, Bank for International Settlements 2021.
- , “DeFi Risks and the Decentralisation Illusion,” Quarterly Review Article BIS Quarterly Review, December 2021, Bank for International Settlements 2021. Authors: Sirio Aramonte, Wenqian Huang, Andreas Schrimpf.
- , “Crypto Risk and Macro-Financial Stability: A Literature Review,” Working Paper BIS Working Paper No. 1036, Bank for International Settlements 2022.
- , “Blueprint for the Future Monetary System: Improving the Old, Enabling the New,” Annual Report Chapter Annual Economic Report 2023, Chapter III, Bank for International Settlements 2023.
- , “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 2024.
- Bank of England**, “New Forms of Digital Money,” Discussion Paper, Bank of England June 2021.
- , “Regulatory Regime for Systemic Payment Systems Using Stablecoins and Related Service Providers,” Discussion Paper DP 2023/4, Bank of England November 2023.
- Barabási, Albert-László and Réka Albert**, “Emergence of Scaling in Random Networks,” *Science*, 1999, 286 (5439), 509–512.

- Battiston, Stefano, Michelangelo Puliga, Rahul Kaushik, Paolo Tasca, and Guido Caldarelli**, “DebtRank: Too Central to Fail? Financial Networks, the FED and Systemic Risk,” *Scientific Reports*, 2012, *2*, 541.
- Board of Governors of the Federal Reserve System**, “Financial Stability Report: Cryptoasset and Stablecoin Risks,” Financial Stability Report, Board of Governors of the Federal Reserve System November 2022. Section on crypto-asset risks.
- , “Money and Payments: The U.S. Dollar in the Age of Digital Transformation,” Discussion Paper, Board of Governors of the Federal Reserve System January 2022.
- Bodó, Balázs and Primavera De Filippi**, “Trust in context: The impact of regulation on blockchain and <scp>DeFi</scp>,” 2024.
- Bongini, Paola, Francesca Mattassoglio, Alessia Pedrazzoli, and Silvio Vismara**, “Crypto ecosystem: navigating the past, present, and future of decentralized finance,” 2025.
- Bordo, Michael D. and Andrew Levin**, “Central Bank Digital Currency and the Future of Monetary Policy,” 2017.
- Briola, Antonio, David Vidal-Tomás, Yuanrong Wang, and Tomaso Aste**, “Anatomy of a Stablecoin’s Failure: The Terra-Luna Case,” *Finance Research Letters*, 2023, *51*, 103358.
- Brunnermeier, Markus K.**, “Deciphering the Liquidity and Credit Crunch 2007–2008,” 2009.
- **and Lasse Heje Pedersen**, “Market Liquidity and Funding Liquidity,” *Review of Financial Studies*, 2009, *22* (6), 2201–2238.
- Carpentier-Desjardins, Catherine, Masarah Paquet-Clouston, Stefan Kitzler, and Bernhard Haslhofer**, “Mapping the DeFi crime landscape: an evidence-based picture,” 2025.
- Catalini, Christian and Alonso de Gortari**, “On the Economic Design of Stablecoins,” Working Paper, MIT Sloan School of Management 2019.
- Chaliasos, Stefanos, Marcos Antonios Charalambous, Liyi Zhou, Rafaila Galanopoulou, Arthur Gervais, Dimitris Mitropoulos, and Benjamin Livshits**, “Smart Contract and DeFi Security Tools: Do They Meet the Needs of Practitioners?,” 2024.
- Choi, Nakhoon and Heeyoul Kim**, “Decentralized Exchange Transaction Analysis and Maximal Extractable Value Attack Identification: Focusing on Uniswap USDC3,” 2024.
- Chu, Yu-Ming and Nina Srinivasan Rathbun**, “Monetary Sovereignty and Central Bank Digital Currencies: Competing Models for Future Cross-Border Payment Platforms,” 2025.

- Cisar, David, Benjamin Schellinger, Jens-Christian Stoetzer, Nils Urbach, Florian Weiß, Vincent Gramlich, and Tobias Guggenberger**, “Designing the future of bond markets: Reducing transaction costs through tokenization,” 2025.
- Collibus, Francesco Maria De, Matija Piškorec, Alberto Partida, and Claudio J. Tessone**, “The Structural Role of Smart Contracts and Exchanges in the Centralisation of Ethereum-Based Cryptoassets,” 2022.
- Cont, Rama, Amal Moussa, and Edson Bastos e Santos**, “Network Structure and Systemic Risk in Banking Systems,” in “Handbook on Systemic Risk,” Cambridge University Press, 2013, pp. 327–368.
- Cookson, J. Anthony, Corbin Fox, Javier Gil-Bazo, Juan Felipe Imbet, and Christoph Schiller**, “Social Media as a Bank Run Catalyst,” 2023.
- Cui, Jinxin, Aktham Maghyereh, and Salem Adel Ziadat**, “Crude oil, forex, and stock markets: unveiling the higher-order moment and cross-moment risk spillovers in times of turmoil,” 2025.
- da Cunha, Paulo Rupino, Paulo Melo, and Hélder Sebastião**, “From Bitcoin to Central Bank Digital Currencies: Making Sense of the Digital Money Revolution,” 2021.
- Daian, Philip, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels**, “Flash Boys 2.0: Frontrunning in Decentralized Exchanges, Miner Extractable Value, and Consensus Instability,” in “IEEE Symposium on Security and Privacy” 2020, pp. 910–927.
- Dave, Kinnari, Vilhelm Sjöberg, and Xin-Yuan Sun**, “Towards Verified Price Oracles for Decentralized Exchange Protocols,” 2021.
- de Carvalho, Rubens Moura, Helena Inácio, and Rui Pedro Marques**, “Stablecoin: A Story of (In)Stabilities and Co-Movements Written Through Wavelet,” 2025.
- Diamond, Douglas W. and Philip H. Dybvig**, “Bank Runs, Deposit Insurance, and Liquidity,” *Journal of Political Economy*, 1983, 91 (3), 401–419.
- Dionysopoulos, Lambis, Miriam Marra, and Andrew Urquhart**, “Central bank digital currencies: A critical review,” 2023.
- Diop, Papa Ousseynou, Julien Chevallier, and Bilel Sanhaji**, “Collapse of Silicon Valley Bank and USDC Depegging: A Machine Learning Experiment,” 2024.
- Dotan, Maya, Aviv Yaish, Hsin-Chu Yin, Eytan Tsytkin, and Aviv Zohar**, “The Vulnerable Nature of Decentralized Governance in DeFi,” 2023.
- Duffie, Darrell**, “Presidential Address: Asset Price Dynamics with Slow-Moving Capital,” *Journal of Finance*, 2010, 65 (4), 1237–1267.

- Easley, David, Maureen O’Hara, and Soumya Basu**, “From Mining to Markets: The Evolution of Bitcoin Transaction Fees,” *Journal of Financial Economics*, 2019, 134 (1), 91–109.
- Elliott, Matthew, Benjamin Golub, and Matthew O. Jackson**, “Financial Networks and Contagion,” *American Economic Review*, 2014, 104 (10), 3115–3153.
- European Central Bank**, “Crypto-Assets: Implications for Financial Stability, Monetary Policy, and Payments and Market Infrastructures,” Occasional Paper No. 223, European Central Bank 2019. Crypto-Assets Task Force.
- , “Stablecoins’ Role in Crypto and Beyond: Functions, Risks and Policy,” Macroprudential Bulletin Article 2, July 2022, European Central Bank 2022.
- , “Decrypting DeFi: The Role of Governance and Smart Contracts in the Future of Finance,” Occasional Paper, European Central Bank 2023. Crypto-Assets Task Force.
- Fang, Fan, Carmine Ventre, Michail Basios, Leslie Kanthan, David Martínez-Rego, Fan Wu, and Lingbo Li**, “Cryptocurrency trading: a comprehensive survey,” 2022.
- Financial Stability Board**, “Crypto-Asset Markets: Potential Channels for Future Financial Stability Implications,” Report, Financial Stability Board October 2018.
- , “Assessment of Risks to Financial Stability from Crypto-Assets,” Report, Financial Stability Board February 2022.
- , “Review of the FSB High-Level Recommendations of the Regulation, Supervision and Oversight of “Global Stablecoin” Arrangements,” Consultative Document, Financial Stability Board October 2022.
- , “Global Regulatory Framework for Crypto-Asset Activities,” Report, Financial Stability Board July 2023.
- , “High-Level Recommendations for the Regulation, Supervision and Oversight of Crypto-Asset Activities and Markets,” Final Report, Financial Stability Board July 2023.
- , “High-Level Recommendations for the Regulation, Supervision and Oversight of Global Stablecoin Arrangements: Final Report,” Final Report, Financial Stability Board July 2023.
- Foley, Sean, Jonathan R. Karlsen, and Tālis J. Putniņš**, “Sex, Drugs, and Bitcoin: How Much Illegal Activity Is Financed through Cryptocurrencies?,” *Review of Financial Studies*, 2019, 32 (5), 1798–1853.
- Freitas, Luís Pedro, Jorge Cerdeira, and Diogo Lourenço**, “Hayekian Hurdles: Challenges to Cryptocurrency as a Viable Basis for a New Monetary Order,” 2025.
- Gai, Prasanna and Sujit Kapadia**, “Contagion in Financial Networks,” *Proceedings of the Royal Society A*, 2010, 466 (2120), 2401–2423.

- Galati, Luca and Francesco Capalbo**, “Silicon Valley Bank bankruptcy and Stablecoins stability,” 2023.
- Gan, Rundong, Le Wang, Xiangyu Ruan, and Xiaodong Lin**, “Understanding Flash-Loan-based Wash Trading,” 2022.
- Glasserman, Paul and H. Peyton Young**, “Contagion in Financial Networks,” *Journal of Economic Literature*, 2016, 54 (3), 779–831.
- Gorton, Gary and Andrew Metrick**, “Securitized Banking and the Run on Repo,” *Journal of Financial Economics*, 2012, 104 (3), 425–451.
- Gorton, Gary B. and Jeffery Y. Zhang**, “Taming Wildcat Stablecoins,” *University of Chicago Law Review*, 2023, 90 (3), 909–971.
- Gramlich, Vincent, Tobias Guggenberger, Marc Principato, Benjamin Schellinger, and Nils Urbach**, “A multivocal literature review of decentralized finance: Current knowledge and future research avenues,” 2023.
- Gregory, Gadzinski, Castello Alessio, Liuzzi Vito, and Sargenti Patrice**, “Break a peg! A study of stablecoin co-instability,” 2024.
- Gromb, Denis and Dimitri Vayanos**, “Limits of Arbitrage,” *Annual Review of Financial Economics*, 2010, 2, 251–275.
- Gudgeon, Lewis, Daniel Pérez, Dominik Harz, Benjamin Livshits, and Arthur Gervais**, “The Decentralized Financial Crisis,” 2020.
- Guo, Sky, Joseph Kreitem, and Thomas Moser**, “DLT Options for CBDC¹,” 2024.
- Haddad, Christian and Lars Hornuf**, “How do fintech start-ups affect financial institutions’ performance and default risk?,” 2023.
- Haldane, Andrew G. and Robert M. May**, “Systemic Risk in Banking Ecosystems,” *Nature*, 2011, 469, 351–355.
- Hanif, Waqas, Hee-Un Ko, Linh Pham, and Sang Hoon Kang**, “Dynamic connectedness and network in the high moments of cryptocurrency, stock, and commodity markets,” 2023.
- Hautsch, Nikolaus, Christoph Scheuch, and Stefan Voigt**, “Building trust takes time: limits to arbitrage for blockchain-based assets,” 2024.
- Huan, Liu and Ortwin Renn**, “Polycrisis and Systemic Risk: Assessment, Governance, and Communication,” 2025.
- Häfner, Matthias, Marco Henriques Pereira, Helmut Dietl, and Juan Beccuti**, “The Four Types of Stablecoins: A Comparative Analysis,” 2024.

International Monetary Fund, “The Crypto Ecosystem and Financial Stability Challenges,” GFSR Chapter Global Financial Stability Report, October 2021, Chapter 2, International Monetary Fund 2021.

– , “Cryptoasset Risks: Is a Regulatory Response Warranted?,” GFSR Chapter Global Financial Stability Report, April 2022, Chapter 3, International Monetary Fund 2022.

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

– , “Central Bank Digital Currency: Progress, Plans, and Risks,” Staff Discussion Note, International Monetary Fund 2024.

International Monetary Fund and Financial Stability Board, “IMF-FSB Synthesis Paper: Policies for Crypto-Assets,” Synthesis Paper, International Monetary Fund and Financial Stability Board September 2023.

International Organization of Securities Commissions, “IOSCO Decentralized Finance Report,” Report, IOSCO March 2022.

– , “Policy Recommendations for Crypto and Digital Asset Markets,” Final Report, IOSCO November 2023.

Jalan, Akanksha and Roman Matkovskyy, “Systemic risks in the cryptocurrency market: Evidence from the FTX collapse,” 2023.

Jourenko, Maxim, Kanta Kurazumi, Mario Larangeira, and Keisuke Tanaka, “SoK: A Taxonomy for Layer-2 Scalability Related Protocols for Cryptocurrencies,” 2025.

Kao, Hsien-Tang, Tarun Chitra, Rei Chiang, and John Morrow, “An Analysis of the Market Risk to Participants in the Compound Protocol,” Working Paper, Gauntlet Networks 2020.

Kaur, Sandeepa, Simarjeet Singh, Sanjay Gupta, and Sangeeta Wats, “Risk analysis in decentralized finance (DeFi): a fuzzy-AHP approach,” 2023.

Khoury, Rim El, Nohade Nasrallah, Khaled Hussainey, and Rima Assaf, “Spillover analysis across FinTech, ESG, and renewable energy indices before and during the Russia–Ukraine war: International evidence,” 2023.

Kim, Hyongsung, Hyun-Sik Kim, and Yong-Suk Park, “Perpetual Contract NFT as Collateral for DeFi Composability,” 2022.

Kirişci, Murat, “An integrated decision-making process for risk analysis of decentralized finance,” 2025.

Kirste, Daniel, Alexander Poddey, Niclas Kannengießner, and Ali Sunyaev, “On the influence of conventional and automated market makers on market quality in crypto-economic systems,” 2024.

- Kitzler, Stefan, Friedhelm Victor, Pietro Saggese, and Bernhard Haslhofer**, “Disentangling Decentralized Finance (DeFi) Compositions,” 2022.
- Klages-Mundt, Aariah and Andreea Minca**, “While stability lasts: A stochastic model of noncustodial stablecoins,” *Mathematical Finance*, 2022, 32 (4), 943–981.
- Kokorin, Ilya**, “The anatomy of crypto failures and investor protection under MiCAR,” 2023.
- Kopytov, Alexandr**, “Booms, Busts, and Common Risk Exposures,” 2023.
- Krishnamurthy, Arvind**, “Amplification Mechanisms in Liquidity Crises,” *American Economic Journal: Macroeconomics*, 2010, 2 (3), 1–30.
- Li, Lei, Kun Qin, and Desheng Wu**, “A Hybrid Approach for the Assessment of Risk Spillover to ESG Investment in Financial Networks,” 2023.
- Li, Zihao, Jianfeng Li, Zheyuan He, Xiapu Luo, Ting Wang, Xiaoze Ni, Wenwu Yang, Xi Chen, and Ting Chen**, “Demystifying DeFi MEV Activities in Flashbots Bundle,” 2023.
- Liu, Jia-Geng, Igor Makarov, and Antoinette Schoar**, “Anatomy of a Run: The Terra Luna Crash,” 2023.
- Lyons, Richard K. and Ganesh Viswanath-Natraj**, “What Keeps Stablecoins Stable?,” *Journal of International Money and Finance*, 2023, 131, 102777.
- Ma, Yiming, Yao Zeng, and Anthony Lee Zhang**, “Stablecoin Runs and the Centralization of Arbitrage,” 2025.
- Makarov, Igor and Antoinette Schoar**, “Cryptocurrencies and Decentralized Finance (DeFi),” 2022.
- Maple, Carsten, Łukasz Szpruch, Gregory Epiphaniou, Kalina Staykova, S. Basanta Singh, William Penwarden, Y Wen, Zijian Wang, Jagdish Hariharan, and Pavle Avramović**, “The AI Revolution: Opportunities and Challenges for the Finance Sector,” 2023.
- Mejia, Julian Fernandez**, “Essays in International Finance,” 2024.
- Mikhaylov, Alexey**, “Understanding the risks associated with wallets, depository services, trading, lending, and borrowing in the crypto space,” 2023.
- Mirdala, Rajmund**, “Revolutionizing Finance: Decentralized Finance as a Disruptive Challenge to Traditional Finance,” 2024.
- Mohan, Vijay**, “Automated market makers and decentralized exchanges: a DeFi primer,” 2022.

- Momtaz, Paul P.**, “Decentralized finance (DeFi) markets for startups: search frictions, intermediation, and the efficiency of the ICO market,” 2024.
- Nabben, Kelsie and Primavera De Filippi**, “Accountability protocols? On-chain dynamics in blockchain governance,” 2024.
- Nimalendran, Mahendrarajah, Praveen Pathak, Mariia Petryk, and Liangfei Qiu**, “Informational Efficiency of Cryptocurrency Markets,” 2024.
- Oben, Remy Jonkam and Fezile Özdamlı**, “Decentralized Finance (DeFi): Benefits, Risks, and RiskMitigation Strategies,” 2024.
- Olanrewaju, Ayobami Gabriel**, “Artificial Intelligence in Financial Markets: Optimizing Risk Management, Portfolio Allocation, and Algorithmic Trading,” 2025.
- Ontario Securities Commission**, “QuadrigaCX: A Review by Staff of the Ontario Securities Commission,” Technical Report, Ontario Securities Commission 2020.
- Ozili, Peterson K.**, “Decentralized finance research and developments around the world,” 2022.
- **and Sergio Luis Nández Alonso**, “Central Bank Digital Currency Adoption Challenges, Solutions, and a Sentiment Analysis,” 2024.
- Pennacchi, George**, “Narrow Banking,” *Annual Review of Financial Economics*, 2012, 4, 141–159.
- Perez, Daniel, Sam M. Werner, Jiahua Xu, and Benjamin Livshits**, “Liquidations: DeFi on a Knife-Edge,” in “Financial Cryptography and Data Security (FC 2021)” 2021.
- Qin, Kaihua, Liyi Zhou, Benjamin Livshits, and Arthur Gervais**, “Quantifying Blockchain Extractable Value: How Dark is the Forest?,” *IEEE Symposium on Security and Privacy*, 2022.
- Saengchote, Kanis and Krislert Samphantharak**, “Digital money creation and algorithmic stablecoin run,” 2024.
- Sakariyahu, Rilwan, Rodiat Lawal, Rasheed A. Adigun, Audrey Paterson, and Sofia Johan**, “One crash, too many: Global uncertainty, sentiment factors and cryptocurrency market,” 2024.
- Salami, Iwa**, “Challenges and Approaches to Regulating Decentralized Finance,” 2021.
- Santiago, Venator, Michel Charifzadeh, and Tim Alexander Herberger**, “Risks of decentralized finance and their potential negative effects on capital markets: the Terra-Luna case,” 2024.
- Schueffel, Patrick**, “What colors are the bricks? Unboxing the DeFi model- A literature survey, empirical study, and taxonomy of decentralized finance,” 2025.

- Schuler, Katrin, Ann Sofie Cloots, and Fabian Schär**, “On DeFi and On-Chain CeFi: How (Not) to Regulate Decentralized Finance,” 2024.
- Schär, Fabian**, “Decentralized Finance: On Blockchain- and Smart Contract-based Financial Markets,” 2020.
- , “Decentralized Finance: On Blockchain- and Smart Contract-Based Financial Markets,” 2021.
- Sethaput, Vijak and Supachate Innet**, “Blockchain application for central bank digital currencies (CBDC),” 2023.
- Shleifer, Andrei and Robert Vishny**, “Fire Sales in Finance and Macroeconomics,” *Journal of Economic Perspectives*, 2011, 25 (1), 29–48.
- Siam, K., Bilash Saha, Md Mehedi Hasan, Md Jobair Hossain Faruk, Nafisa Anjum, Sharaban Tahora, Aiasha Siddika, and Hossain Shahriar**, “Securing Decentralized Ecosystems: A Comprehensive Systematic Review of Blockchain Vulnerabilities, Attacks, and Countermeasures and Mitigation Strategies,” 2025.
- Stiglitz, Joseph E. and Andrew Weiss**, “Credit Rationing in Markets with Imperfect Information,” *American Economic Review*, 1981, 71 (3), 393–410.
- Szrajber, Bentzion Edgardo, Ilan Alon, and Shalom Levy**, “Systematic Analysis of Decentralized Finance,” 2025.
- Tan, Teck Ming and Saila Saraniemi**, “Trust in blockchain-enabled exchanges: Future directions in blockchain marketing,” 2022.
- Tolmach, Palina, Yi Li, Shang-Wei Lin, and Yang Liu**, “Formal Analysis of Composable DeFi Protocols,” in “Financial Cryptography Workshops” 2021.
- Uhlig, Harald**, “A Luna-tic Stablecoin Crash,” Working Paper 30256, National Bureau of Economic Research 2022.
- Upper, Christian**, “Simulation Methods to Assess the Danger of Contagion in Interbank Markets,” *Journal of Financial Stability*, 2011, 7 (3), 111–125.
- van der Linden, M. and Tina Shirazi**, “Markets in crypto-assets regulation: Does it provide legal certainty and increase adoption of crypto-assets?,” 2023.
- Watts, Duncan J. and Steven H. Strogatz**, “Collective Dynamics of ‘Small-World’ Networks,” *Nature*, 1998, 393 (6684), 440–442.
- Weingärtner, Tim, Fabian Fasser, Pedro Costa, and Walter Farkas**, “Deciphering DeFi: A Comprehensive Analysis and Visualization of Risks in Decentralized Finance,” 2023.

- Wen, Hongbo, Hanzhi Liu, Jiaxin Song, Yanju Chen, Wenbo Guo, and Yu Feng**, “FORAY: Towards Effective Attack Synthesis against Deep Logical Vulnerabilities in DeFi Protocols,” 2024.
- Werner, Sam M., Daniel Pérez, Lewis Gudgeon, Arian Klages-Mundt, Dominik Harz, and William J. Knottenbelt**, “SoK: Decentralized Finance (DeFi),” 2022.
- Wątopek, Marcin, Marcin Królczyk, Jarosław Kwapien, Tomasz Stanisz, and Stanisław Drożdż**, “Approaching Multifractal Complexity in Decentralized Cryptocurrency Trading,” 2024.
- Xu, Jiahua and Nikhil Vadgama**, “SoK: Decentralized Finance (DeFi) Attacks,” *arXiv preprint*, 2023.
- Xue, Yue, Dunqiu Fan, Shen Su, Jialu Fu, Ning Hu, Wenmao Liu, and Zhihong Tian**, “A Review on the Security of the Ethereum-Based DeFi Ecosystem,” 2023.
- Younis, Ijaz, Himani Gupta, Anna Min Du, Waheed Ullah Shah, and Waqas Hanif**, “Spillover dynamics in DeFi, G7 banks, and equity markets during global crises: A TVP-VAR analysis,” 2024.
- Zhou, Qiheng, Huawei Huang, Zibin Zheng, and Jing Bian**, “Solutions to Scalability of Blockchain: A Survey,” 2020.
- Zhu, Jason, Arijit Khan, and Cüneyt Gürçan Akçora**, “Data depth and core-based trend detection on blockchain transaction networks,” 2024.

A Literature Search Protocol

The systematic literature search was conducted using OpenAlex, an open scholarly metadata catalogue indexing more than 250 million academic works. OpenAlex was chosen for its comprehensive coverage of both peer-reviewed journals and preprint repositories (notably SSRN and arXiv), its freely accessible API, and its structured metadata including citation counts, topic classifications, and author affiliations which allow reproducible, large-scale bibliometric queries without the licensing constraints of proprietary databases such as Scopus or Web of Science.

Channel-Specific Search Queries

Each of the 14 candidate channels was queried with a tailored set of search strings designed to capture the relevant theoretical and empirical literature. Table 2 reports the full query set.

Table 2: Search queries by candidate channel.

Channel	Search Queries
Network Contagion	“systemic risk contagion cryptocurrency network”; “DeFi interconnectedness systemic risk”; “crypto exchange network contagion”; “financial network topology cryptocurrency”
Liquidity Spirals	“liquidity spiral cryptocurrency”; “AMM liquidity crisis DeFi”; “exchange run crypto”; “fire sale digital assets”
Stablecoin Runs	“stablecoin de-peg systemic risk”; “algorithmic stablecoin collapse”; “stablecoin bank run”; “USDT USDC de-pegging contagion”
Oracle Manipulation	“oracle manipulation DeFi”; “price feed contagion blockchain”; “oracle dependency systemic risk”; “Chainlink oracle attack”
Composability Risk	“DeFi composability risk”; “smart contract exploit cascade”; “money legos systemic risk”; “DeFi protocol dependency”
Liquidation Cascades	“on-chain liquidation cascade”; “DeFi leverage systemic risk”; “cascading liquidation cryptocurrency”; “margin call crypto”
Counterparty Concentration	“counterparty risk cryptocurrency exchange”; “FTX contagion systemic”; “centralized exchange failure”; “crypto concentration risk”
Regulatory Contagion	“regulatory shock cryptocurrency”; “crypto regulation financial stability”; “compliance contagion digital finance”
Gateway Risk	“fiat crypto gateway systemic”; “banking channel cryptocurrency”; “on-ramp off-ramp fragility”; “Silvergate SVB crypto”
Governance Failure	“DAO governance attack”; “crypto governance systemic risk”; “hard fork contagion”; “protocol governance failure”
Information Asymmetry	“proof of reserves cryptocurrency”; “CeFi opacity systemic risk”; “information asymmetry crypto exchange”
RWA Transmission	“tokenized assets systemic risk”; “real world asset DeFi bridge”; “tokenization financial stability”
Bridge Vulnerability	“cross-chain bridge exploit”; “bridge hack systemic risk”; “Wormhole Ronin bridge contagion”
Validator Concentration	“mining concentration systemic risk”; “validator centralization blockchain”; “consensus layer risk”

Broad Queries

In addition to the channel-specific queries, seven broad queries were executed to capture cross-cutting contributions: “systemic risk digital finance”; “systemic risk cryptocurrency”; “DeFi financial stability”; “stablecoin systemic risk”; “crypto contagion”; “tokenization systemic risk”; “CBDC financial stability.”

Search Parameters

All queries shared the following parameters: `max_results_per_query = 200`, sorted by `cited_by_count:desc`, restricted to `publication_year ≥ 2009`, and filtered to English-language works (`language = en`).

Search Implementation

The core search logic iterates over the query set for each channel, deduplicates results by OpenAlex work identifier, and accumulates papers up to a per-channel limit. Listing 1 shows the simplified search routine.

```
1 def search_channel(client, channel_id, channel_info,
2                   per_channel_limit):
3     queries = channel_info.get("queries", [])
4     seen_ids = set()
5     channel_papers = []
6     for query in queries:
7         if len(channel_papers) >= per_channel_limit:
8             break
9         response = client.search_works(
10            search=query,
11            filter_params={"publication_year": "2009-2026"},
12            per_page=min(
13                200,
14                per_channel_limit - len(channel_papers)),
15            sort="cited_by_count:desc",
16        )
17        for work in response.get("results", []):
18            paper = extract_paper(work)
19            if paper["id"] not in seen_ids:
20                seen_ids.add(paper["id"])
21                channel_papers.append(paper)
22    return channel_papers
```

Listing 1: Core channel search routine (simplified from `openalex_search.py`).

Channel strength was assessed with a composite score combining three normalised dimensions literature volume, citation impact, and crisis-event evidence weighted 0.35, 0.35, and 0.30, respectively. Listing 2 shows the scoring formula.

```
1 # Channel strength composite scoring
2 W_LIT, W_CIT, W_CRISIS = 0.35, 0.35, 0.30
3
4 lit_volume = {ch: count / max_count
5               for ch, count in paper_counts.items()}
6 cit_impact = {ch: mean_top10_cites / max_mean
7               for ch, mean in channel_means.items()}
8 crisis_ev = {ch: log10_weighted_sum / max_sum
9               for ch, log10_weighted_sum
10                in crisis_sums.items()}
11
12 composite = {ch: lit_volume[ch] * W_LIT
13              + cit_impact[ch] * W_CIT
14              + crisis_ev[ch] * W_CRISIS
15              for ch in all_channels}
```

Listing 2: Composite channel-strength scoring (simplified from `channel_mapper.py`).

Selection Flow

The literature search and channel-selection process proceeded through five stages:

- Stage 1: Identification.** Approximately 2,400 unique works were retrieved across the 14 channel-specific and 7 broad query sets via the OpenAlex API.
- Stage 2: Deduplication.** Cross-channel deduplication by OpenAlex work identifier reduced the corpus to 2,433 unique papers.
- Stage 3: Screening.** Title and abstract screening retained approximately 620 papers (a retention rate of approximately 25%), removing works that referenced digital finance only peripherally or lacked a systemic-risk dimension.
- Stage 4: Channel scoring.** Each of the 14 candidate channels was scored on the composite metric described above, combining literature volume, citation impact, and crisis-event evidence.
- Stage 5: Channel merger and selection.** The six channels falling below the 40th percentile of the composite score distribution oracle manipulation, regulatory contagion, governance failure, bridge vulnerability, validator concentration, and RWA transmission were not discarded but *merged* into thematically related surviving channels whose scope subsumes the core mechanisms of the absorbed channel. For example, oracle manipulation was folded into composability risk, and bridge vulnerability was absorbed into network contagion. This merger process yielded the final set of eight channels analysed in the paper.

B Complete Channel Interaction Matrix

Table 3 presents the full 8×8 undirected interaction matrix for the eight systemic risk channels. Coupling strength is rated on a three-level scale: strong (+++), indicating that the activation of one channel reliably triggers or materially amplifies the other; moderate (++) , indicating a well-documented but conditional linkage; and weak (+), indicating a plausible but indirect or episodic connection. The matrix is symmetric each pair is rated identically in both directions.

Table 3: 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; theoretical link exists but empirical evidence is limited.

Blank: Negligible or no meaningful interaction pathway.

C Crisis Event Database

Table 4 catalogues all 25 systemic risk events identified during the review, sorted chronologically. Channel abbreviations are: NC = Network Contagion, LS = Liquidity Spirals, SR = Stablecoin Runs, CR = Composability Risk, LC = Liquidation Cascades, CC = Counterparty Concentration, IA = Information Asymmetry, GR = Gateway Risk, BV = Bridge Vulnerability, VC = Validator Concentration, OM = Oracle Manipulation, RC = Regulatory Contagion, GF = Governance Failure.

Table 4: Crisis event database, 2009–2026.

Date	Event	Domain	Channels	Est. Losses
2014-02	Mt. Gox Collapse	CeFi	CC, IA, GR	\$460M
2016-06	The DAO Hack	DeFi	CR, GF	\$60M
2016-08	Bitfinex Hack	CeFi	CC, IA	\$72M
2017-07	BTC-e Seizure and Shut-down	CeFi	CC, RC, GR	n/a

Continued on next page

Date	Event	Domain	Channels	Est. Losses
2019-02	QuadrigaCX Collapse	CeFi	CC, IA, GR	\$190M
2020-03	Black Thursday (COVID Crash)	DeFi/CeFi	LC, LS, NC, GR	n/a
2020-09	SushiSwap Vampire Attack	DeFi	CR, LS, GF	n/a
2021-06	Iron Finance / TITAN Collapse	DeFi	SR, LS, CR	\$2B
2021-08	Poly Network Hack	DeFi	BV, CR	\$611M
2022-02	Wormhole Bridge Hack	DeFi	BV, CR, LS	\$326M
2022-03	Ronin Bridge Hack	DeFi/CeFi	BV, VC, CC	\$625M
2022-05	Terra/Luna Collapse	Stbl./DeFi	SR, LS, NC, CR, CC	\$45B
2022-06	Three Arrows Capital Insolvency	CeFi	CC, NC, LS, IA	\$3.5B
2022-06/07	Celsius Network and Voyager Digital Failures	CeFi	CC, LS, IA, GR	\$5.4B
2022-08	Nomad Bridge Hack	DeFi	BV, CR	\$190M
2022-10	Mango Markets Exploit	DeFi	OM, LC, CR	\$114M
2022-11	FTX Collapse	CeFi	CC, IA, NC, GR, LS	\$8.7B
2023-03	Euler Finance Exploit	DeFi	CR, LC, LS	\$197M
2023-03	SVB Failure and USDC De-Peg	Stbl./TradFi	GR, SR, NC, LS	n/a
2023-07	Curve Finance Pool Exploit	DeFi	CR, LS, LC	\$62M
2023-07	Multichain Bridge Collapse	DeFi	BV, CC, GF	\$126M
2024-05	DMM Bitcoin Exchange Hack and Closure	CeFi	CC, IA, GR	\$305M
2024-07	WazirX Multi-Sig Exploit	CeFi	CC, BV, GR	\$235M
2025-02	Bybit \$1.5B Hack	CeFi	CC, BV, IA, LS	\$1.5B
2025-03	Hyperliquid Whale Manipulation Events	DeFi	LC, OM, LS	\$15M