

Shock Propagation in Decentralized Lending Networks: Evidence from the Compound Protocol*

Natkamon Tovanich

TU Wien, Austria, natkamon.tovanich@tuwien.ac.at

Stefania Marcassa

THEMA (UMR CNRS 8184), CY Cergy Paris Université, France, stefania.marcassa@cyu.fr

Stefan Kitzler

Complexity Science Hub, Austria and AIT Austrian Institute of Technology, Austria, kitzler@csh.ac.at

Christos A. Makridis

Arizona State University, AZ, USA; University of Nicosia, Cyprus, cmakridi@asu.edu

Julien Prat

CREST, CNRS, École Polytechnique, Institut Polytechnique de Paris, France, julien.prat@ensae.fr

We study the propagation of financial shocks in decentralized finance (DeFi) using a novel dataset from Compound, a prominent DeFi lending application on Ethereum. Unlike traditional interbank networks, Compound exhibits a bipartite structure in which users lend and borrow via liquidity pools implemented as smart contracts. We construct daily balance sheets for users and pools from January 2020 to June 2024, model the liability network, and apply a propagation algorithm that captures the impact of liquidations following token-specific price shocks. We assess our model’s ability to predict financial stress and show that network-based measures provide substantial predictive power beyond standard financial indicators. These results underscore the need for topology-aware risk monitoring in algorithmic credit systems.

Key words: decentralized finance; systemic risk; financial networks; decentralized lending. *JEL:* D85, G01, G23.

1. Introduction

Financial fragility is determined not only by the leverage and the size of individual institutions, but also by the structure of the network in which they are embedded. A local shock may remain contained when exposures are diffused and well capitalized, but the same shock can propagate

* *This version: 2026-05-16.* We are grateful to Co-Pierre Georg, Wenqian Huang and Riho Marten Pallum for their valuable comments and suggestions. We also thank conference participants at the Futures of Payment (Bocconi University), CEPR FinTech and Digital Currency, Future of Money II, Grenoble DeFi (INP-UGA), ICCF 2024, Frontiers in DeFi (Complexity Science Hub), Annual Conference on DeFi & Digital Currencies (WBS Gillmore Centre), the Economics of Decentralisation (Waseda) for helpful discussions. Julien Prat gratefully acknowledges financial support from the ANR project “BLOCKFI: Blockchain and Decentralized Finance”- 2024 - CES38. The authors acknowledge the use of LLM assistants to support the writing and coding process.

into systemic stress when it reaches highly connected counterparties. This insight is integral to the literature on financial networks and contagion, which has shown that the severity of shock propagation depends critically on the topology of financial linkages.¹

This paper studies how these insights can be applied to supervise and predict systemic risk in the expanding ecosystem of Decentralized Finance (DeFi). Unlike traditional financial institutions, DeFi platforms operate through algorithmic governance and smart contracts (Capponi et al. 2023), enabling a wide range of financial activities to be conducted without centralized oversight. One of the central promises of DeFi is the mitigation of moral hazard through the substitution of trusted intermediaries with automated execution. However, reliance on algorithms does not eliminate other sources of systemic risk.

The novelty of DeFi’s economic design, together with the very features that foster innovation—namely open access, transparency, and composability—also expand the attack surface of these protocols. This increased exposure renders them susceptible to exploits that can precipitate abrupt liquidity crises.² Crucially, these risks are no longer confined to the DeFi ecosystem. As regulated financial institutions deepen their integration with decentralized infrastructures through tokenized assets and central bank digital currencies, understanding how systemic vulnerabilities emerge and propagate within digital financial networks becomes increasingly important.

The dynamics of risk propagation in DeFi differ from those in traditional finance: whereas inter-bank contagion is dampened by balance sheet opacity and legal frictions (Gennaioli et al. 2013), risk transmission in DeFi occurs through automated mechanisms that operate almost instantaneously. Such algorithmic propagation reduces the buffering effect of discretionary intervention and may therefore accelerate the spread of financial distress. The specificity and novelty of these amplification channels motivate our analysis of shock propagation within DeFi.

Owing to their market size and inherent risk, lending protocols stand out as the most natural candidates for the study of financial contagion in DeFi. With total value locked (TVL) exceeding \$30 billion at the time of writing,³ lending and borrowing are among the most popular activities for DeFi users. These activities are also among the riskiest, as borrowers are exposed to liquidation. Among the various lending protocols, we select Compound for our empirical analysis because it has served as a template for numerous protocols, making its architecture relevant beyond our specific

¹ Seminal models of financial contagion include Allen and Gale (2000) and Elliott et al. (2014), while Glasserman and Young (2016) provide a comprehensive survey of the early literature on the topic. See also Acemoglu et al. (2015) and Gabaix (2011) for theoretical results on the propagation of systemic and idiosyncratic risk in networks, and Di Giovanni et al. (2014) or Carvalho et al. (2021) for empirical results from micro-data.

² Among the risks specific to DeFi, the most extensively studied include smart contract exploits (Sun et al. 2021), flash loan attacks (Gudgeon et al. 2020), and governance challenges (Dotan et al. 2023).

³ See <https://coinmarketcap.com/currencies/compound/> and <https://www.coingecko.com/en/coins/compound> for historical data on the market capitalization of Compound.

context. Furthermore, as one of the earliest decentralized lending platforms, Compound also offers one of the longest available data series (Leshner and Hayes 2019).

In Compound’s financial network, lending pools serve as intermediaries between lenders and borrowers, creating a bipartite structure that differs from conventional interbank networks. Instead of directly lending to one another, users deposit tokens into asset-specific pools and receive tokenized claims, while borrowers draw from those pools after posting collateral. This structure means that distress does not propagate through bilateral pool-to-pool exposures, but through links created by the lending and borrowing activity of users. The mechanism differs from standard interbank contagion. In Compound, distress moves through collateral constraints, automated liquidations, reserve withdrawals, and residual bad debt rather than through direct bilateral default among financial institutions.

Capturing this propagation mechanism requires reconstructing the full network of liabilities. Our first contribution is therefore methodological: we detail a procedure for rebuilding pools’ and users’ balance sheets and show how these can be used to map the network through which financial stress propagates. This representation allows us to identify and characterize the two main sources of risk resulting from a drop in the market value of tokens used as collateral. First, price shocks induce *liquidity risk* as liquidators deplete the reserve of the shocked pool by seizing collateral tokens. Second, price shocks may generate *solvency risk* for the protocol by creating bad debt when borrowers’ positions become undercollateralized. We explain how the design of Compound is largely motivated by the need to contain these two sources of risk. Having clarified the economic logic of the protocol, we propose a propagation algorithm that captures the impact of price shocks, providing a tractable framework for stress testing.

We then apply our methodology to an empirical analysis of Compound. Using the protocol’s indexed subgraph provided by The Graph,⁴ we reconstruct daily snapshots of Compound’s financial network from January 1, 2020, to June 15, 2024. The data indicate that the network evolved in ways that altered its vulnerability to shocks. Early periods (2020–2021) are characterized by dense user participation and elevated leverage, whereas later phases (from 2022 onward) exhibit reduced activity. Despite these changes, five pools—ETH, WBTC, USDC, DAI, and USDT—retain a preponderant role due to their substantially larger size relative to other pools. We also document marked differences in the use of cryptoassets versus stablecoins, with the former serving primarily as collateral and the latter accounting for the bulk of borrowing activity.

⁴ The Graph is a decentralized indexing protocol that organizes and stores blockchain data through open APIs known as subgraphs. It enables developers to query on-chain data from protocols such as Compound in real time, without having to process raw blockchain transactions directly. See <https://thegraph.com/> for more information.

The specific topology of the financial network, together with the automated nature of liquidations, gives rise to novel risk dynamics that require adapting existing stress-testing frameworks. To this end, we apply our propagation algorithm to historical realizations of Compound’s network in order to generate time series measuring the effects of price shocks on the main tokens. We use these series to construct measures of liquidity and solvency risk, which in turn enable us to evaluate an operational framework aimed at *predicting* these two risk measures. To account for the network’s topology, we estimate a series of high-frequency panel regressions based on the Spatial Lag of X (SLX) model framework. This approach allows us to capture localized spillover effects across lending pools. Leveraging daily snapshots of the lending network, we examine how the financial and topological features of pools and users contribute to the liquidity and solvency risk associated with each token, demonstrating through a series of tests the out-of-sample predictive power of our model.

For crypto assets, we apply a 50 percent shock to ETH and WBTC. For stablecoins, we apply a full depegging shock to DAI and USDC, since the relevant stress event is the loss of parity rather than a sequence of price movements. The results show that standard pool indicators capture only part of the risk. In the crypto liquidity exercise, a baseline model using only the standard pool-level indicators reported by Compound explains 23.6 percent of out-of-sample variation in reserve depletion. Adding the exposure-weighted features of collateral-provider raise the test R^2 to 0.869. For liquidity risk, the main information is therefore not simply how levered collateral providers are, but how their positions connect a pool to the rest of the protocol.

Solvency risk follows a different logic. Standard pool indicators explain 39.3 percent of out-of-sample variation in bad debt, while the preferred parsimonious model raises the test R^2 to 0.834. The gain reflects the fact that bad debt depends not only on liquidation-driven withdrawals, but also on whether affected users retain enough collateral value to cover their obligations after the shock. The stablecoin results reinforce this point. Because DAI and USDC are primarily borrowed rather than pledged as collateral, depegging shocks propagate through borrowers as much as through collateral providers. The preferred stablecoin solvency model achieves a test R^2 of 0.929, showing that the relevant side of the network depends on the economic role of the shocked asset.

Across all specifications, network-based measures of connectivity and concentration add substantial predictive power beyond standard balance sheet indicators. Our results therefore demonstrate that counterparty characteristics play a significant role in predicting cross-sectional variations in risk. Pool size, utilization, collateral factors, and aggregate liquidity are useful, but they largely describe pools in isolation. They do not capture how users connect collateral pools to borrowing pools, nor how those connections determine which positions are liquidated after a price shock. The relevant exposure is shaped not only by the characteristics of the pools but also by the composition

and connectivity of their user base. Pools with highly leveraged or highly connected counterparties are more prone to transmitting stress across the protocol. These dynamics explain why risk monitoring in DeFi cannot rely solely on pool size, utilization, or collateral factors, as is commonly done in practice, but must also identify critical nodes and counterparties.

Related Work. There is a large and growing literature on the role of networks in economics and finance, showing that network dependencies help explain how localized shocks—whether idiosyncratic or sectoral—can propagate and amplify into macroeconomic fluctuations. Theoretically, such propagation has been formalized in models of input-output linkages and financial networks (Acmoglu et al. 2015, Gabaix 2011), where the topology of connections and the concentration of exposures shape systemic outcomes. Empirically, production and financial network data have been used to demonstrate how shocks transmit through supplier–buyer chains (Carvalho et al. 2021) and through firm-specific demand shifts (Di Giovanni et al. 2014). A parallel strand of research in finance models interbank markets as networks, highlighting how the structure of claims and exposures influences both contagion and stability (Allen and Gale 2000, Elliott et al. 2014). As explained in Glasserman and Young (2016), these frameworks have given rise to a variety of contagion metrics, such as DebtRank (Battiston et al. 2012, Bardoscia et al. 2015), which quantifies how financial distress propagates through layered interdependencies before defaults occur. Systemic risk in financial networks has also been extensively studied using complex network theory (Caldarelli 2007, Boss et al. 2004). Studies have demonstrated that while increased interconnectivity can mitigate idiosyncratic risks, it also heightens systemic fragility by creating contagion pathways if exposures are clustered among certain groups of users, as in Caccioli et al. (2009).

The relevance of network models has extended beyond traditional finance, finding applications in DeFi, where lending protocols such as Compound and Aave facilitate algorithmic credit intermediation. Unlike banks, DeFi protocols operate via smart contracts, enforcing collateralized borrowing without centralized oversight (Bartoletti et al. 2021). Recent research has examined the unique arbitrage mechanisms within DeFi (Daian et al. 2020, Heimbach et al. 2024), including miner extractable value (MEV) and the role of block proposers and arbitrageurs in shaping market outcomes (Milionis et al. 2023, Azar et al. 2024). While traditional finance relies on established intermediation chains (Green et al. 2007, Di Maggio et al. 2017, Hollifield et al. 2017), DeFi introduces a transparent but fragmented market structure with novel sources of rent extraction.

One of the central mechanisms in DeFi lending networks is the role of collateralization in shaping risk transmission. Most DeFi protocols require overcollateralization, and borrower solvency is continuously assessed through a health ratio—a protocol-specific metric comparing the collateral value, adjusted by a risk factor, to the value of outstanding debt. When this ratio falls below a

liquidation threshold, the borrower’s position is subject to automatic liquidation. While this mechanism enforces discipline, it can also destabilize the protocol during sharp price movements by triggering cascades of liquidations (Sun et al. 2022), with broader spillover effects on asset prices and interlinked positions (Qin et al. 2021, Lehar and Parlour 2022, Gudgeon et al. 2020).

Although prior research in traditional finance has examined the role of collateral requirements and network structure in driving liquidation dynamics (Elsinger et al. 2006, Nier et al. 2007), less is known about how these dynamics play out in automated settings when asset values decline rapidly. While prior work has examined systemic risk in interbank networks using eigenvalue-based stability measures (Caccioli et al. 2018, Glasserman and Young 2016), their applicability to DeFi lending has been limited. Our study contributes to filling this gap by combining a tractable propagation algorithm with real-time health factor metrics to simulate and trace the spread of financial distress in Compound. This approach complements insights from recent models of blockchain-based financial intermediation, which highlight the roles of information asymmetry, arbitrage behavior, and endogenous feedback effects in shaping outcomes in DeFi markets (Cong et al. 2023, Lehar and Parlour 2023).

Finally, our paper builds on the spatial econometrics literature (Anselin 1988, LeSage and Pace 2009). We adopt a simulation-based approach to bypass the well-known challenges of causal inference in networks (Manski 1993, Gibbons and Overman 2012). Rather than analyzing actual contagion episodes, we simulate price shocks of varying magnitudes, propagate them through the network, and regress the resulting stress measures on financial and topological features. Since the outcomes are generated by the simulation itself, standard endogeneity concerns do not apply. Our objective is predictive rather than causal as we aim to identify which features most influence the spread of financial distress under the protocol’s liquidation rules.

Structure of the paper. Section 2 presents the design of Compound and explains how users’ and pools’ balance sheets can be constructed in order to map the protocol’s financial network. Section 3 describes the propagation of price shocks through Compound’s network and introduces a tractable algorithm that captures this transmission mechanism. Section 4 presents the data, while Section 5 outlines our econometric model and the strategy used to assess its predictive performance. Section 6 presents the empirical results, beginning with liquidity risk in subsection 6.1 and then turning to solvency risk in subsection 6.2. Finally, Section 7 concludes, while supplementary material and robustness checks are provided in the Appendix.

2. The Compound Protocol

2.1. Compound’s Financial Network

Compound V2 is a decentralized protocol that operates through smart contracts. It enables users to lend and borrow tokens in a permissionless and automated way.

On the lending side, users supply tokens to the protocol, which aggregates these deposits into *liquidity pools* that contain only one type of tokens. In return for their deposits, lenders receive interest-bearing tokens called **cTokens** (such as cETH or cUSDC). These tokens represent both the user’s principal and the interest accrued over time. Because **cTokens** are themselves transferable and redeemable, they function as a tokenized claim on the underlying assets and should, in principle, be convertible back into the original tokens at any time.

On the borrowing side, users who wish to borrow must first deposit assets into the protocol, receiving **cTokens** in return, and then use these **cTokens** as collateral to borrow other assets. Compound enforces an overcollateralization mechanism to manage risk. Each asset has an associated collateral factor, which determines how much can be borrowed relative to the value of the deposited collateral. If the value of the collateral falls below a required threshold, the position becomes eligible for liquidation, meaning that other participants can repay part of the debt in exchange for a portion of the collateral at a discount.

Table 1 summarizes the balance sheet components of users and lending pools. Users hold inter-node assets in the form of **cTokens** and incur inter-node liabilities through their borrowings. Lending pools, in contrast, carry a combination of tokens and loans — originated through user borrowing — on the asset side, and record user deposits as inter-node liabilities. This representation abstracts from off-protocol positions, such as users’ external holdings of fiat or crypto assets not supplied to Compound, allowing us to isolate the financial linkages within the protocol. Incorporating exogenous assets and their role in risk absorption is a promising direction for future research.

	User	Pool
Assets		
External assets	$\emptyset$	Tokens
Inter-node assets	Deposits (cTokens)	Loans
Liabilities		
External liabilities	$\emptyset$	$\emptyset$
Inter-node liabilities	Borrows	Deposits
Equity	Deposits - Borrows	Tokens + Loans - Deposits

Table 1 Balance sheet structure of lending pools and users.

Using this asset-liability decomposition, we can represent Compound’s financial network as a directed bipartite graph, where the nodes consist of lending pools $\mathcal{P} = 1, \dots, p$, and users $\mathcal{U} = p + 1, \dots, p + u$. An edge of weight L_{ij} denotes a liability of node i to node j —that is, node i owes L_{ij} to node j . Due to the bipartite nature of the network, the liability matrix L satisfies $L_{ij} = 0$

for all $i, j \in \mathcal{U}$ or $i, j \in \mathcal{P}$. Hence, liabilities only exist between users and pools, and not within the same group.

Figure 1 provides an illustrative example of the bipartite financial network generated by the Compound protocol. Directed edges represent financial exposures, with arrows pointing from debtors to creditors. Each node is associated with a stylized balance sheet decomposing assets and liabilities. For clarity, node sizes are not drawn to scale. The figure highlights the dual role users play in the system: some users act as pure lenders (e.g., the second user from the left), supplying assets to lending pools without initiating any borrowings, while others both supply and borrow. This intermediation is facilitated by the pools, which aggregate deposits and extend credit to users against collateral. In the pools' balance sheets, deposits appear as liabilities, while outstanding loans appear as inter-node assets. For users, the mirror image holds: cTokens deposits are assets, and borrowings are inter-node liabilities. By representing users and pools as distinct, but financially interlinked agents, this bipartite structure captures the exposure patterns that arise endogenously and serves as the foundation for the analysis of shock propagation.




Figure 1 Bipartite network of Compound.
 Note: Arrows point to creditors. The second user from the left is a pure lender. Node sizes are not proportional to balance sheet sizes.

2.2. Interest Rate Model and Liquidations

The soundness of Compound hinges on its ability to uphold two core commitments: (i) providing on-demand liquidity to lenders who wish to withdraw their tokens; and (ii) ensuring the solvency of its loans by maintaining sufficient collateral backing. The former is underpinned by the protocol’s interest rate model; the latter by its liquidation mechanism, which serves as a safeguard against undercollateralization.

Interest Rate Model. The interest rates at which users can borrow or lend are determined algorithmically by the Interest Rate Model (IRM hereafter) based on supply and demand within each pool. The borrowing rate for reserve i , denoted r_i^B , is a function of the utilization ratio U_i , which measures the share of reserves in pool i that is currently borrowed. The protocol must ensure that the utilization ratio does not approach one, as full utilization would prevent lenders from withdrawing their funds, thereby violating the promise of on-demand liquidity.

The IRM defines, for each reserve, a mapping from the utilization ratio $U_i \in [0, 1]$ to the corresponding borrowing rate $r_i^B \in \mathbb{R}_+$. Interest payments are distributed to lenders proportionally to their deposits, with a fraction $\tau_i \in (0, 1)$, known as the reserve factor, retained by the protocol. Accordingly, the lending rate r_i^L is related to the borrowing rate r_i^B through the following equation

$$r_i^L = (1 - \tau_i) \cdot U_i \cdot r_i^B(U_i).$$

In practice, the mapping $r_i^B(U_i)$ is piecewise linear, with a kink at an optimal utilization level $U_i^{\text{opt}} \in (0, 1)$ determined by governance.⁵ This functional form is designed to approximate a convex, increasing relationship while remaining computationally efficient in order to minimize gas costs.⁶ The kink determines the utilization threshold at which the IRM transitions from a moderate slope to a significantly steeper one. Below this kink, borrowing rates increase gradually with utilization. Beyond this kink, interest rates rise sharply, creating a strong economic incentive for liquidity provision.

⁵ Formally, the borrowing rate is given by

$$r_i^B = \begin{cases} U_i \cdot \frac{r_i^{\text{opt}}}{U_i^{\text{opt}}}, & \text{if } U_i \leq U_i^{\text{opt}}, \\ r_i^{\text{opt}} + (U_i - U_i^{\text{opt}}) \cdot \frac{r_i^{\text{full}} - r_i^{\text{opt}}}{1 - U_i^{\text{opt}}}, & \text{if } U_i > U_i^{\text{opt}}, \end{cases}$$

where r_i^{opt} is the borrowing rate at the optimal utilization ratio, and r_i^{full} is the borrowing rate at full utilization $U_i = 1$.

⁶ The gas cost refers to the fee incurred when executing a transaction on the Ethereum blockchain. See [Bertucci et al. \(2026\)](#) for a thorough analysis of the optimal design of the IRM.

Liquidations. Since Compound’s users are pseudonymous, the only way to guarantee that they have an incentive to repay their loans is by ensuring that their positions remain overcollateralized at all times. The solvency of each user $i \in \mathcal{U}$ is summarized by their *health ratio* H_i , defined as

$$H_i = \frac{\sum_{j \in \mathcal{P}} c_j \cdot L_{ji}}{\sum_{j \in \mathcal{P}} L_{ij}}, \quad (1)$$

where $c_j \in (0, 1]$ is the *collateral factor* of pool j , a parameter that determines what fraction of deposited assets may be borrowed against. The numerator, which weights each collateral position L_{ji} of user i by its collateral factor c_j , measures her borrowing capacity. The denominator represents the user’s total outstanding debt, summing her borrowed amounts L_{ij} across all pools.

Because the protocol accepts volatile assets as collateral, solvency must be actively maintained through a liquidation mechanism that targets positions at risk of becoming undercollateralized. Specifically, when the user’s health ratio defined in (1) falls below one, the position becomes eligible for liquidation: third-party liquidators may repay a portion of the outstanding debt and, in exchange, seize a corresponding amount of the borrower’s collateral at a discount, hereafter referred to as the liquidation bonus LB . The discounted price incentivizes timely intervention and helps restore the position toward solvency, hopefully before the protocol incurs a loss.

To bound the magnitude of any individual intervention, the protocol enforces a *close factor*, which limits the fraction of a borrower’s outstanding debt that may be repaid within a single liquidation event (typically up to 50%). Following the seizure of the maximum admissible collateral under this constraint, the borrower’s position is recomputed. If the resulting health ratio exceeds unity, the liquidation process terminates. Otherwise, subsequent liquidation rounds may be executed on the remaining collateral. This iterative procedure continues until either the position is restored to solvency or the collateral is fully depleted, in which case the protocol incurs residual bad debt.

3. Shock Propagation

3.1. Liquidity and Solvency Risks

Unlike traditional financial markets, which may trigger circuit breakers to suspend operations during periods of extreme stress, smart contracts remain operational at all times. Lending protocols therefore need to be resilient to large price fluctuations, ensuring that sudden market movements do not compromise their solvency or their capacity to provide liquidity. To understand why price fluctuations threaten these two properties, it is instructive to trace how they propagate through the financial network. The left-hand panel of [Figure 2](#) presents a stylized setting in which a user deposits ETH as collateral to borrow stablecoins (USDCs). The horizontal red line crossing the user’s balance sheet represents the liquidation threshold, i.e., the point at which her health ratio falls to one and the position becomes eligible for liquidation.

The right-hand panel of Figure 2 shows how this network is reconfigured after a sudden 20% decline in the price of ETH. The mechanical impact of this price shock is to shrink the assets and liabilities of the ETH pool proportionally, as well as the assets of the borrower who posted ETH as her sole collateral. Since the user’s nominal debt is unchanged, this adjustment lowers her health ratio. In our example, the drop is sufficient to push the position below the liquidation threshold. To enhance visual clarity, we illustrate the case of a full liquidation.




Figure 2 Balance sheet impact of a liquidation event triggered by a drop in the price of ETH. Panel (a) shows the pre-liquidation state; panel (b) shows the post-liquidation state.

As shown in panel (b), the liquidator takes two simultaneous actions. First, it repays the borrower’s outstanding USDC loan, which is recorded on the liability side of the borrower’s balance sheet and on the asset side of the USDC pool. This repayment converts the user’s debt into reserves, thereby lowering the utilization ratio of the USDC pool. In exchange, the liquidator seizes the borrower’s collateral by redeeming her cETH tokens, converting them back into the underlying ETH before withdrawing them from the pool. This withdrawal reduces the deposits of the ETH pool, increasing its utilization ratio.

The overall effect of the liquidation is thus asymmetric across pools: it alleviates liquidity pressure in the borrowing pool while tightening it in the collateral pool. This stylized example highlights that Compound is exposed to two primary categories of financial risk, each arising from distinct but interrelated mechanisms within its liquidation framework:

1. **Liquidity risk:** This risk materializes when liquidators seize and convert tokens used as collateral. In doing so, they withdraw liquidity from the corresponding collateral pools. If liquidation activity becomes sufficiently large, it can deplete the available reserves. In such a scenario, depositors may face delays or, in extreme cases, an inability to redeem their cTokens for the underlying asset, thereby compromising the protocol’s promise of on-demand liquidity.

2. **Solvency risk:** When the value of a borrower’s collateral is insufficient to cover the outstanding debt plus the liquidation bonus, the liquidation mechanism cannot fully recover the borrowed amount, and the resulting shortfall is recorded as *bad debt*.⁷ If left unmitigated, persistent bad debt can erode confidence in the protocol by creating a mismatch between users’ deposits and recoverable assets.

Compound’s design is predicated on mitigating these two risks. To this end, the protocol’s governance can adjust the two key parameters introduced above: (i) the kink in the interest rate curve and (ii) the collateral factor. By calibrating both the position of the kink and the steepness of the post-kink slope, governance can influence the likelihood that pools reach critically high utilization levels. This mechanism addresses liquidity risk by helping ensure that sufficient reserves remain available. The collateral factor, by contrast, is central to the management of both risks. A higher collateral factor improves capital efficiency for borrowers but reduces the buffer available to absorb adverse price movements. By fine-tuning this parameter across asset classes, governance can adjust the probability that liquidations deplete reserves and generate bad debt.

These parameters form a complementary risk management framework: the IRM primarily regulates short-term liquidity dynamics, while the collateral factor governs the structural resilience of the protocol against both illiquidity and insolvency.

3.2. Propagation Algorithm

Having established the central role of price shock containment in the management of protocol’s risk, we now introduce a simulation algorithm that goes beyond qualitative insights to quantify both liquidity and solvency risks. To ensure tractability of the propagation mechanism, we impose two simplifying assumptions:

- A1 Liquidation sizes:** Liquidated amounts are set such that, whenever feasible, the borrower’s health factor is restored exactly to one.
- A2 Multi-asset positions:** For users with collateral or liabilities across multiple pools, liquidations are assumed to be executed simultaneously across all pools, with liquidation shares proportional to the relative size of liabilities in each pool.

Assumption [A1](#) leads to an underestimation of actual liquidation volumes. In Compound, liquidators may seize up to 50% of the collateral even when a smaller amount would suffice to restore a position to a healthy level. This rule introduces a discontinuity at the liquidation threshold, whereas our framework generates smooth liquidation profiles.⁸ As a result, our simulations provide

⁷ See [subsection 3.2](#) and [Tovanich et al. \(2023\)](#).

⁸ The inefficiency of discontinuous liquidation rules is well documented among practitioners and has been addressed by Aave, a leading lending protocol closely related to Compound, in its recent V4 update of the liquidation engine (see <https://aave.com/blog/aave-v4-liquidations>). This update implements a design conceptually similar to that underlying our propagation algorithm.

conservative estimates that should be interpreted as a lower bound on total liquidation volumes. Importantly, this approximation primarily affects positions with health factors close to one prior to liquidation and therefore does not materially bias estimates of bad debt generated by price shocks.

Assumption A2 abstracts from the fact that liquidators may selectively target a subset of pools to minimize gas costs and manage risk exposure. However, because pool selection depends on idiosyncratic liquidator strategies, explicitly modeling such behavior would require a detailed characterization of these decision rules. We leave this promising avenue for future research and instead adopt an agnostic specification that treats all pools symmetrically.

Relaxing either assumption would require agent-based methods (e.g., Tovanich et al. (2023)), which introduce stochastic variation into simulations and complicate the identification of robust predictive features. Our approach instead strikes a pragmatic balance between realism and tractability by abstracting from the behavioral heterogeneity of liquidators, whose explicit modeling would require a dedicated structural framework.

To identify positions eligible for liquidation, we first compute total collateral and total debt for each user $i \in \mathcal{U}$:

$$C_i = \sum_{j=1}^p L_{ji}, \quad D_i = \sum_{j=1}^p L_{ij}, \quad (2)$$

where $p = |\mathcal{P}|$ is the number of pools. We then define *pledgeable collateral* as the collateral-factor-weighted sum of collateral positions

$$CL_i = \sum_{j=1}^p c_j \cdot L_{ji}. \quad (3)$$

Since the health factor is given by $H_i = CL_i/D_i$, user i is eligible for liquidation whenever $CL_i < D_i$.

Let LB denote the liquidation bonus for liquidators. We assume throughout that

$$\max_j \{c_j\} \cdot (1 + LB) < 1, \quad (4)$$

since, as discussed below, it is a necessary condition for the solvency of the protocol.

The propagation rule depends on whether the user's position generates bad debt:

1. **No bad debt:** The user's collateral is sufficient to cover her debt, i.e.,

$$C_i \geq (1 + LB) \cdot D_i. \quad (5)$$

Condition (4) implies that no bad debt arises when the health factor is exactly equal to one.⁹ This condition is therefore necessary for protocol solvency: without it, positions with a

⁹ For a user with $H_i = 1$, we have $CL_i = D_i$. Using (4) and the definition of CL_i , it follows that

$$C_i \geq \frac{CL_i}{\max_j \{c_j\}} = \frac{D_i}{\max_j \{c_j\}} > (1 + LB) \cdot D_i.$$

health factor at or even above one would generate bad debt upon liquidation. In practice, the restriction is largely satisfied by Compound’s parameters.

Under the proportionality rule [A2](#), the liquidation ratio Δ_i rescales user i ’s positions as

$$\begin{cases} L'_{ij} = (1 - \Delta_i) \cdot L_{ij}, \\ L'_{ji} = \left[1 - (1 + LB) \cdot \Delta_i \cdot \frac{D_i}{C_i} \right] \cdot L_{ji}, \end{cases}$$

where we denote post-liquidation values with a prime. Assumption [A1](#) further requires that Δ_i restore the health factor to one, i.e., $CL'_i = D'_i$. This yields the closed-form expression

$$\Delta_i = \frac{D_i - CL_i}{D_i \cdot \left[1 - (1 + LB) \cdot \frac{CL_i}{C_i} \right]} \in [0, 1].$$

The liquidation ratio Δ_i starts at 0 for users that are just eligible for liquidation and increases linearly to 1 at the boundary of the bad debt region where all collateral is seized by liquidators.

2. **Bad debt:** When $C_i < (1 + LB) \cdot D_i$, the user’s collateral is insufficient to cover her debt. All positions are then fully unwound:

$$\begin{cases} L'_{ij} = 0, & \forall j \in \llbracket 1, p \rrbracket, \\ L'_{ji} = 0, & \forall j \in \llbracket 1, p \rrbracket, \end{cases}$$

and the resulting bad debt allocated by user i to pool j is given by

$$B_{ij} = \left[1 - \frac{C_i}{(1 + LB) \cdot D_i} \right] \cdot L_{ij}.$$

These updating rules allow us to simulate the impact of price shocks by identifying which positions are liquidated and, if so, to what extent. Aggregating post-liquidation liabilities across users then yields the resulting pool balance sheets, including their updated reserves and accumulated bad debt. For brevity, we relegate the aggregation procedure to [Appendix A](#) and turn to the empirical implementation of the propagation algorithm.

4. Data

4.1. Data Description

Our analysis relies on daily snapshots of the Compound V2 protocol collected between January 1, 2020, and June 15, 2024. The data is sourced from The Graph, which provides an indexed subgraph of the Compound V2 smart contracts. After June 2024, this hosted subgraph was deprecated following The Graph’s migration to a decentralized network.¹⁰ However, by that time, Compound V2 was effectively deprecated as well, with almost negligible user activity remaining.

Each daily snapshot is taken at the final block of the day in Coordinated Universal Time (UTC) and contains comprehensive information on both lending pool characteristics and user positions:

¹⁰ See: <https://thegraph.com/blog/sunrise-here-the-graph-network/> and <https://thegraph.com/docs/en/archived/sunrise/>

- **Lending pool characteristics:** For each lending pool, we collect snapshot data on total deposits, total borrows, available liquidity (cash), reserves, and collateral factor. The pool’s external assets are defined as the sum of cash and reserves, as shown in Table 1.
- **Price oracles:** Compound uses the token price data from Chainlink’s decentralized price oracles, which serve as the protocol’s source of truth for calculating asset valuations.¹¹ We collect the token price in U.S. dollars as reported at the final block of each day.
- **User positions:** For each user, we record their deposit and borrow positions across supported assets, i.e., the amount each user lent and borrowed from each lending pool. These values are normalized in U.S. dollars using daily token prices to construct the liability matrix L_{ij} as defined in subsection 2.1.

4.2. Descriptive Statistics

We examine the evolution of Compound V2 using the aggregate lending and borrowing positions of users, as shown in Figure 3. The protocol experienced its peak activity in 2021, followed by a gradual decline in usage over subsequent years. The figures in the left and middle panels reveal persistent user preferences: ETH and USDC consistently serve as the dominant supplied assets, while borrowing activity remains primarily concentrated in USDC and DAI. On the right panel, the total equity of the protocol increased over time, reflecting the accumulation of profits retained in the pool reserves.




Figure 3 Evolution of deposit and borrow amounts of the top lending pools in Compound V2.

Figure 4 documents additional trends in the number of active users, health ratios, self-borrowing prevalence,¹² and utilization ratios over time. To focus on economically significant users, we consider

¹¹ Compound v2 Price Feed: <https://docs.compound.finance/v2/prices/>, and Chainlink Data Feeds: <https://data.chain.link/feeds>

¹² Self-borrowing refers to the practice of depositing an asset as collateral and borrowing the same asset from the same pool. Users engage in this strategy primarily to extract rewards from protocol incentives.

only lenders and borrowers with positions exceeding 100 USD. The number of active participants grew steadily from early 2020, reaching a peak of approximately 35,000 lenders and nearly 5,000 borrowers by the end of 2021. Both figures declined sharply in mid-2022 and stabilized after 2023, albeit at much lower levels of total value locked.




Figure 4 Evolution of Compound V2: lenders and borrowers, health ratio, self-borrowing and utilization ratios. Note: Left panel: number of active lenders (blue) and borrowers (red) with positions exceeding \$100. Middle panel: median (solid) and exposure-weighted average (dashed) health ratio of borrowers with positions exceeding \$100, with the shaded band showing the interquartile range (25th–75th percentile). Right panel: share of self-borrowing (blue) and utilization ratio (red).

The sharp increase in borrowing activity coincides with the launch of the COMP incentive program in June 2020, which distributed rewards to users engaged in borrowing (Park and Stinner 2023). This led to a sudden surge in both self-borrowing and the utilization ratio. However, as the level of incentives declined, self-borrowing dropped to nearly zero, and the utilization ratio stabilized between 20% and 30% from 2022 onward.

The middle panel of the figure displays the median and weighted average of the health ratio of borrowers, along with interquartile bands. Following the introduction of the incentive program, the median health ratio hovered around 1.5, suggesting that users maintained their positions near the liquidation threshold. The weighted average health ratio was even lower than the median during 2021 and 2022, implying that large borrowers were more aggressively optimizing their leverage positions. As self-borrowing diminished after 2022 (right panel), the weighted average health ratio increased, reflecting the exit of large, highly leveraged participants from the protocol.

The top 5 pools. Lending and borrowing activity is distributed across a set of pools authorized by the protocol’s governance. The number of active pools has increased over time, ranging from 6 pools at the protocol’s inception in May 2019 to a total of 19 pools.¹³ However, deposits remain

¹³ See Leshner, R. (2019), “Compound v2 is Live,” <https://medium.com/compound-finance/compound-v2-is-live-157db0b7cfc8>

highly concentrated within a small subset of pools. Figure 5 shows that the three stablecoin pools (DAI, USDC, and USDT) and the two primary crypto pools (BTC and ETH) consistently account for more than 90% of the total value locked in Compound. Among these pools, USDT cannot be used as collateral since its collateral factor is set to zero. We therefore focus on the other four pools since they are the only ones propagating losses induced by liquidations.




Figure 5 Share of Total Value Locked in Compound for the Top 5 Pools.

Crypto vs stablecoin pools. In addition to being concentrated within a small subset of pools, lending and borrowing activity is also segmented between crypto-assets and stablecoins. More precisely, the vast majority of loans are originated to borrow stablecoins against crypto-assets used as collateral. In other words, most borrowers enter into repurchase-like agreements in which the cash leg is replaced by a transfer of stablecoins.

The dominance of this trading pattern is illustrated in the left panel of Figure 6, which shows that, at any given point in time, between 80% and 90% of all outstanding loans in Compound are used to borrow one of the three major stablecoins.¹⁴ The right panel of Figure 6 corroborates this finding by showing that the utilization ratio—defined as the share of loaned-out tokens relative to the total reserves of each pool—is systematically higher for stablecoin pools. Particularly striking is the very low utilization ratio of ETH and BTC pools, which hovers around 10%, highlighting the limited demand for borrowing these assets. As we will show in subsequent sections, these marked differences have important implications for the nature of the risks associated with stablecoins versus crypto-assets, thereby motivating a separate analysis for the two asset classes.

¹⁴ See Tovanich et al. (2023) for detailed evidence on the widespread use of crypto-assets as collateral for stablecoin borrowing.




Figure 6 Share of total borrowing and utilization ratios by asset class.

5. Empirical Analysis

Overview. The empirical strategy proceeds in three steps. First, we simulate sudden token-price shocks of size s applied separately to each of the top four pools and propagate them through the liability network using the algorithm of subsection 3.2. Second, for every pool j and date t , we extract two outcomes from the simulation output: liquidity risk, $\ell_{t,j}$, defined as the share of pool assets withdrawn by liquidators, and solvency risk, $b_{t,j}$, defined as the share of pool assets left as bad debt once positions are unwound. Third, we treat $\ell_{t,j}$ and $b_{t,j}$ as separate target variables and regress each on a panel of pool-level and depositor-level characteristics through a Spatial Lag of X (SLX) model presented in subsection 5.1. The features that enter the regressors are defined in subsection 5.2, and the LASSO-based estimation and out-of-sample evaluation procedure is described in subsection 5.3.

5.1. Regression Specification

Simulated outcomes. We simulate sudden price shocks, denoted s , of varying magnitudes (ranging from 10% to 100%) applied daily to each of the top five pools. For every scenario, we construct two outcome measures that capture exposure to liquidity and solvency risk, respectively.

For a given shock size, liquidity risk is defined as:

$$\ell_{t,j} = \frac{T_{t,j} - T'_{t,j}}{Assets_{t,j}}, \quad (6)$$

where $T_{t,j}$ denotes the value of pool j 's reserves immediately after the price shock but prior to any liquidations, and $T'_{t,j}$ denotes the reserve value after both the shock and the ensuing liquidations. The numerator captures the net destruction of reserves attributable solely to the liquidation process. Normalizing by the total $Assets_{t,j}$ of pool j , defined as the sum of reserves and loans in the

pool,¹⁵ enables meaningful comparisons across pools of heterogeneous size. By construction, $\ell_{t,j}$ belongs to $[0, 1]$, with larger values indicating greater relative liquidity losses.

Our second outcome of interest is solvency risk, defined as:

$$b_{t,j} = \frac{BD_{t,j}}{Assets_{t,j}}, \quad (7)$$

where $BD_{t,j}$ denotes the total bad debt generated in the financial network following the shock to pool j .¹⁶ This measure isolates losses that materialize when the liquidation mechanism fails to fully recover the value of undercollateralized positions. Higher values of $b_{t,j}$ therefore reflect more severe solvency stress, as a larger share of loans extended by the protocol becomes unrecoverable. Both outcome variables are stationary: Augmented Dickey-Fuller tests reject the null of a unit root for both ETH and WBTC at $s = 50\%$.¹⁷

Relationship between simulated and realized stress. We rely on simulated outcomes because estimating a predictive model of realized liquidation events is not feasible given the rarity of large price crashes. Stable variable selection and reliable out-of-sample performance require samples of several hundred observations. In practice, however, large crypto price crashes occur at most a few times per year, making it impossible to assemble a test set of sufficient size to evaluate predictive accuracy with any statistical confidence. The simulation-based approach circumvents this limitation by generating a rich panel of stress scenarios from the observed daily network.

A natural question is whether the simulated outcomes are informative about real-world financial stress. Our static simulation design abstracts from two countervailing channels that operate in real time. On the one hand, price declines are often gradual rather than instantaneous, giving proactive users time to top up their collateral or reduce their borrowing, thereby avoiding liquidation. On the other hand, intra-day price volatility can trigger successive waves of liquidations within the same trading day, amplifying total losses beyond what a single end-of-day snapshot would suggest. Because these two channels operate in opposite directions, our simulations may over- or underestimate realized liquidation volumes depending on which effect dominates in a given episode.

We assess the quantitative distance between simulated and realized liquidation volumes in [Appendix D](#), using Precision-Recall analysis and episode-level comparisons across the most extreme liquidation days in our sample. Since our end-of-day simulations abstract from intra-day price dynamics, we apply a scalar multiplier $m \geq 1$ to the observed price drop to account for intra-day

¹⁵ See [Appendix A](#) for an explicit decomposition of the pools' balance sheets.

¹⁶ See equation (10) in [Appendix A](#) for a description of how $BD_{t,j}$ is computed.

¹⁷ ADF statistics for $s = 50\%$: liquidity risk, -6.34 for ETH ($p < 0.001$) and -2.87 for WBTC ($p = 0.049$); solvency risk, -3.17 for ETH ($p = 0.022$) and -3.40 for WBTC ($p = 0.011$). Lag length selected by AIC.

amplification. Treating the simulated liquidation-to-total-assets ratio as a continuous risk score, the benchmark multiplier $m = 1.50$ achieves a PR-AUC of 0.309, representing a $31\times$ improvement over a random classifier and confirming that the simulation has meaningful signal for identifying extreme liquidation days. Among the largest crashes occurring in the data, the model matches the COVID crash (March 2020), the Bitcoin flash crash (May 2021), the January 2022 market correction, the Terra/Luna collapse (May–June 2022), and the FTX crisis (November 2022), with the peak correctly identified on or near the day of each crash.

These results support the interpretation of the simulation as isolating the structural sensitivity of the network: by holding user behavior fixed and applying a clean price shock, the simulation isolates how much of the realized stress is attributable to the topology and leverage of the network itself, independently of behavioral responses. This structural sensitivity is informative for risk monitoring, where the primary concern is identifying vulnerable configurations rather than replicating the exact dynamics of a specific episode.

Econometric model. Having defined the targets, we now specify the regression model that links them to pool- and user-level characteristics. The propagation mechanism described in Section 3 suggests a natural structure. When a token price falls, it is the positions of depositors who have pledged that token as collateral that become undercollateralized and can trigger liquidation cascades. The scale of the resulting liquidity losses and bad debt therefore depends not only on the pool’s own balance sheet, but on the financial condition and network connectivity of the users linked to it: how leveraged they are, how close to the liquidation threshold, and how broadly their positions extend across the rest of the protocol.

We capture this counterparty channel through a *Spatial Lag of X (SLX) model* (Elhorst 2014), which augments pool-level covariates with exposure-weighted averages of user characteristics:

$$Y_t = \mathbf{X}_t\beta + \mathbf{W}_t^D\mathbf{Z}_t\theta^D + \mathbf{W}_t^B\mathbf{Z}_t\theta^B + \mu + u_t, \quad (8)$$

where $Y_t \in \mathbb{R}^N$ stacks the chosen risk outcome—either $\ell_{t,j}$ from (6) or $b_{t,j}$ from (7)—across the $N = |\mathcal{P}|$ lending pools. Equation (8) is estimated separately for each outcome. The matrix $\mathbf{X}_t \in \mathbb{R}^{N \times K}$ collects K pool-level financial and topological characteristics (groups 1.1 and 1.2 in Section 5.2). $\mathbf{Z}_t \in \mathbb{R}^{M \times K}$ is the matrix of time-varying user characteristics (groups 2.1 and 2.2) at time t , where $M = |\mathcal{U}|$ denotes the number of users.¹⁸ The products $\mathbf{W}_t^D\mathbf{Z}_t$ and $\mathbf{W}_t^B\mathbf{Z}_t$ aggregate depositors’ and

¹⁸ Note that the number of active users varies over time. For notational consistency, we define $M = |\mathcal{U}|$ as the maximum number of users observed at any time t . Accordingly, the matrices $\mathbf{Z}_t \in \mathbb{R}^{M \times K}$, $\mathbf{W}_t^D \in \mathbb{R}^{N \times M}$, and $\mathbf{W}_t^B \in \mathbb{R}^{N \times M}$ are constructed with M rows or columns, with entries corresponding to inactive users set to zero as needed.

borrowers' characteristics by weighting each user's contribution by their relative exposure to the pool. The weight matrices are row-normalized, with entries

$$w_{ij,t}^D = \frac{L_{ij,t}}{\sum_k L_{ik,t}} \quad \text{and} \quad w_{ij,t}^B = \frac{L_{ji,t}}{\sum_k L_{ki,t}}, \quad (9)$$

reflecting the share of pool i 's liabilities owed to depositor j and the share of pool i 's assets lent to borrower j at time t . By construction, $\mathbf{W}_t^B$ captures the relative importance of each borrower to pool i , while $\mathbf{W}_t^D$ reflects the relative importance of each depositor. The coefficient vector β in (8) measures the marginal effect of the pool's characteristics on Y_t , while θ^D and θ^B capture the marginal effects of the exposure-weighted depositor and borrower covariates, respectively. The vector $\mu \in \mathbb{R}^N$ collects pool fixed effects that absorb time-invariant heterogeneity across pools, such as differences in collateral factor design or governance history, and Newey–West HAC standard errors correct for residual serial dependence in the disturbance vector $\mathbf{u}_t \in \mathbb{R}^N$.

The SLX framework is well suited to our setting for two reasons. First, it avoids the reflection problem emphasized by Manski (1993): because propagation losses are computed independently for each pool via simulation, there is no simultaneity in Y_t across pools, so the identifying assumption of the SLX model holds by construction.¹⁹ Second, SLX produces directly interpretable coefficients: θ measures the marginal effect of exposure-weighted depositor characteristics on pool-level risk without requiring the inversion of a spatial multiplier matrix, an operation that would be both computationally demanding and difficult to interpret in a network with daily-varying topology and thousands of active users (Elhorst 2014). Hence, the liability network plays the role that geographic distance occupies in traditional spatial models: it defines who influences whom, with exposure-weighted financial links replacing physical proximity as the channel of interaction (Abdallah et al. 2022, Paelinck and Klaassen 1979).

5.2. Features Extraction

We now describe the variables that populate $\mathbf{X}_t$ and $\mathbf{Z}_t$ in equation (8). A token-price shock affects systemic risk through two linked channels. The first is the pool-level balance sheet channel: pools differ in utilization, equity buffers, collateral parameters, and the concentration of claims. The second is the counterparty channel: the users connected to a pool differ in leverage, health ratios, collateral usage, and network connectivity. These user characteristics determine whether a shock triggers liquidations, how large those liquidations are, and how losses are transmitted across pools. We organize the predictors into four groups, which map directly into the regression: groups 1.1 and 1.2 enter as $\mathbf{X}_t$, while groups 2.1 and 2.2 enter as $\mathbf{Z}_t$ before being weighted by $\mathbf{W}_t^D$ and $\mathbf{W}_t^B$.

¹⁹ Because the liability network is jointly shaped by user behavior, protocol incentives, and market conditions, the exposure-weighted regressors ($\mathbf{W}\mathbf{Z}_t$) are interpreted descriptively rather than causally. Our objective is prediction rather than identification of structural peer effects.

1. **Pool characteristics** (populate $\mathbf{X}_t$):

- 1.1. **Financial features** summarize the pool’s balance sheet and protocol-defined risk parameters. These include the collateral factor, utilization ratio, equity share, pool size, the share of supplied assets used as collateral, and the self-borrowing rate. These variables capture the pool’s capacity to absorb withdrawals, liquidations, and bad debt following a price shock.
- 1.2. **Topological features** describe the pool’s position in the liability network. We include measures of connectivity, such as in-degree and out-degree, concentration measures based on Herfindahl-Hirschman indices of incoming and outgoing exposures, and centrality measures such as PageRank and hub scores. These variables capture whether a pool is broadly connected or dependent on a concentrated set of users.

2. **User characteristics** (populate $\mathbf{Z}_t$, then aggregated to the pool level via $\mathbf{W}_t^D$ and $\mathbf{W}_t^B$):

- 2.1. **Financial features** summarize the condition and behavior of users connected to each pool. These variables are constructed as exposure-weighted averages at the pool level and include users’ health ratios, leverage, collateral usage, and related balance sheet measures. They capture how close connected users are to liquidation and how strongly their positions depend on the shocked asset.
- 2.2. **Topological features** describe the connectivity and exposure concentration of users connected to each pool. These variables are also aggregated to the pool level using exposure weights and include user in-degree, out-degree, flow weights, and the corresponding Herfindahl-Hirschman concentration indices. They capture whether a pool is connected to users whose positions link many parts of the protocol or to users whose exposures are more isolated.

In short, these four groups constitute the initial set of $P = 40$ candidate variables submitted to the estimation procedure. Full variable definitions are provided in [Table 3](#) in [Appendix B](#).

5.3. Predictive Modelling Strategy

Each daily network snapshot provides a cross-sectional observation of Compound’s liability network at the end of a trading day. Lending pools are observed repeatedly over time, so the data have a panel structure. Our empirical objective is not to estimate dynamic causal effects. We use the repeated daily snapshots to examine how contemporaneous pools’ and exposure-weighted users’ characteristics predict simulated liquidity and solvency risk. The estimates should therefore be interpreted as predictive associations that identify which features are most informative about systemic vulnerability under the propagation algorithm.

To evaluate predictive performance, we use a strict chronological train/test split, LASSO variable selection, and expanding-window (rolling-origin) cross-validation for penalty tuning. This design

reflects the forecasting problem faced by risk managers: whether information available in the current network can predict simulated liquidity losses and bad debt under large token-price declines. The chronological split ensures that models are trained on past observations and evaluated on future observations, while LASSO penalizes model complexity and reduces overfitting. Coefficient estimates should be read as measures of predictive association, not as causal effects. [Appendix H](#) confirms that the linear specification is not restrictive: the preferred OLS model achieves comparable accuracy than random forest benchmarks, while preserving the interpretability of the estimated coefficients.

Sample splitting. We construct a panel dataset integrating the features described above for the major crypto assets listed on the Compound protocol, observed daily from 1 January 2021 to 1 October 2023.²⁰ The primary outcome variables are the liquidity risk $\ell_{t,j}$ and the solvency risk $b_{t,j}$, constructed under an exogenous price decline applied separately to each asset.

Although the simulated outcomes are conditionally independent given the features, the features themselves may exhibit autocorrelation. A misspecified model will fail to fully absorb this structure, inducing autocorrelation in the residuals and violating the independence assumption underlying standard cross-validation. We therefore adopt time series validation techniques, which respect temporal ordering and provide a more reliable evaluation in this setting ([Bergmeir and Benítez 2012](#)). The sample is divided chronologically in an 80/20 ratio: the first 80 % of observations constitute the training set ($N_{\text{train}} = 1,607$ observations, covering 1 January 2021 to 15 March 2023) and the remaining 20 % the held-out test set ($N_{\text{test}} = 401$ observations, covering 15 March 2023 to 1 October 2023).²¹ This ensures that the model is always estimated on past data and evaluated on future observations, as in a genuine forecasting exercise.²²

We adopt an expanding-window (rolling-origin) design rather than a fixed-width rolling window for two reasons. First, a rolling window would reduce the effective training sample and potentially destabilise LASSO variable selection, which already operates on a relatively limited sample. Second, since the simulated outcomes are conditionally independent across time, horizon-specific forecast accuracy is not a concern, removing one of the principal arguments in favour of a fixed-width rolling

²⁰ Pre-2021 observations are excluded because the protocol operated in a structurally different regime: the collateral factor averaged 0.49 in 2020 versus 0.70–0.76 from 2021 onward following a governance change, the borrower count was approximately five times smaller, and pool sizes were an order of magnitude lower. Including pre-2021 observations materially reduces out-of-sample R^2 , confirming the structural break.

²¹ The test period is not a low-variance regime: both mean and variance of $\ell_{t,j}$ and $b_{t,j}$ are substantially higher in the test period than in the training period (Levene $p < 0.001$), confirming that the high test R^2 reflects genuine predictive accuracy.

²² All predictors are standardized to zero mean and unit variance using moments computed on the training set only, so that coefficient magnitudes are directly comparable across predictors and no information from the test period contaminates the scaling step.

window (Tashman 2000). The remaining limitation of expanding-window evaluation is sensitivity to occurrences unique to a particular origin. We address this concern in Appendix E.2 through a subsampling robustness check, randomly dropping 10 % of training observations across 50 draws. Both the selected variables and the out-of-sample R^2 remain largely unchanged, providing direct evidence against parameter instability over the training period.

Variable selection and estimation. We apply LASSO to the full set of $P = 40$ standardized predictors,²³ to identify a sparse and robust subset that generalises well to the test period. The penalty parameter λ is selected by expanding-window (rolling-origin) time-series cross-validation with $K = 5$ folds, which respects the temporal ordering of the data by always training on past observations and validating on the subsequent block. A grid of 200 values of λ is evaluated, and the value minimising mean cross-validation MSE across folds ($\lambda_{\min}$) is retained for variable selection. LASSO is used exclusively as a selection device; OLS is then refit on the selected variables with Newey–West HAC standard errors ($L = 4$ lags), removing the shrinkage bias inherent in LASSO coefficients. A final trimming step retains only predictors whose post-selection OLS coefficient is individually significant at $\alpha = 0.001$, ensuring that the final model contains only variables with robust statistical support.

Performance is evaluated on the held-out test set using three metrics: the coefficient of determination R^2 , the root mean squared error RMSE, and the mean absolute error MAE. A substantial gap between training and test R^2 is interpreted as evidence of overfitting.

Shock size. The analysis is conducted separately for each shock size $s \in \{10\%, 20\%, \dots, 100\%\}$. For crypto assets (ETH and WBTC), we focus our main results on $s = 50\%$. This choice is informed by Figure 7, which reports test and training R^2 of the trimmed model for each shock size. Two patterns emerge. First, model fit is low for $s \leq 20\%$. This is intuitive since price shocks of 20% or less are smaller than the overcollateralization ratios required by the protocol. Hence, few positions become undercollateralized following such shocks and the outcome variable has limited cross-sectional variation, leaving little signal for the predictors to exploit. Second, test R^2 rises steeply from $s = 30\%$ onward, peaks at $s = 50\%$, and then declines at larger shock sizes despite training R^2 continuing to rise monotonically. The reason is that as the shock size increases beyond 50%, an ever-larger share of positions is fully liquidated and liquidity losses converge across pools: the standard deviation of $\ell_{t,j}$ peaks at $s = 50\%$ ($\sigma = 0.047$) and falls sharply thereafter, reaching $\sigma = 0.013$ at $s = 90\%$. When outcomes are nearly identical across observations, the genuine heterogeneity in network and

²³ Zero-variance predictors are dropped prior to estimation; $P = 40$ holds across all shock sizes in our sample. For shock sizes where no predictor survives at the significance level $\alpha = 0.001$, the threshold is relaxed to $p < 0.05$; this occurs only at $s \in \{0.10, 0.20\}$ where the signal is weakest.

financial characteristics that the model is designed to capture disappears. The model instead fits residual noise in the training data, producing high in-sample fit that does not generalize to the test period. A shock size of $s = 50\%$ therefore maximizes both outcome heterogeneity and out-of-sample predictive accuracy, while corresponding to a severe but plausible stress scenario for crypto assets, consistent with historical drawdowns observed during the Terra/Luna collapse (May 2022) and the FTX crisis (November 2022). This choice is further validated in [Appendix D](#): the observed intra-day price declines during major crypto crashes, once scaled by a multiplier of 1.50–1.75 to account for intra-day amplification, are consistent with the $s = 50\%$ benchmark shock used in the simulations. For completeness, results across the full range of shock sizes are reported in [Appendix G](#).




Figure 7 Predictive performance by shock size — trimmed model (no asset fixed effect), liquidity risk $\ell_{t,j}$ (ETH and WBTC). Note: Coral dots show test R^2 and blue squares show training R^2 on the held-out test set. The shaded red region marks the low-signal regime ($s \leq 20\%$) where few positions are liquidated and the outcome has limited cross-sectional variation.

For stablecoin pools (DAI and USDC), the analysis focuses on a full de-pegging scenario ($s = 100\%$) rather than on shocks of varying magnitudes. This reflects the fundamentally different nature of risk for stablecoins relative to crypto assets. In crypto pools, the shock impairs collateral values and triggers liquidation cascades of increasing severity; varying s therefore traces out a meaningful risk frontier. In stablecoin pools, by contrast, the relevant stress event is a de-pegging in which the stablecoin loses its parity with the US dollar. Partial shocks to stablecoin prices (e.g., $s = 20\%$) do not correspond to a plausible economic scenario, since stablecoins are designed to maintain a fixed parity. In practice, either the peg holds or it breaks, and the economically relevant question is what

determines the protocol’s vulnerability *conditional* on a full de-pegging, not how risk scales with partial price declines. We therefore set $s = 100\%$ for the stablecoin analysis and report results for both DAI and USDC.

6. Results

We present results for liquidity and solvency risk, reporting the variables selected by LASSO and retained after trimming, the predictive performance of the preferred model on the test set, and an interpretation of the key coefficient estimates in terms of the economic channels identified in [subsection 3.1](#).

6.1. Liquidity Risk

We begin by analyzing the factors that drive liquidity risk arising from a decline in the price of one of the main tokens. As discussed in [section 3](#), this risk materializes when liquidators seize and convert tokens used as collateral. To preserve lenders’ ability to redeem their deposits, the system must prevent liquidators from fully depleting the available reserves of a given pool.

Cryptos vs stablecoins. Token withdrawals triggered by liquidations are substantially larger in crypto pools than in stablecoin pools, reflecting the much higher share of crypto deposits pledged as collateral. [Figure 8](#) illustrates this dichotomy by reporting changes in pool available reserves following a 50 % price drop. For each pool, we report two daily statistics: (i) the change when the price drop affects the pool’s own token, and (ii) the average change when the price drop affects other top-5 tokens.

A shock to a crypto pool’s own token generates substantial reserve losses, typically around 20% of total assets, as a large fraction of deposits serves as collateral. When other tokens are shocked instead, the impact on crypto pool reserves is comparatively modest, as the two opposing channels—*inflows from repaid loans and outflows from liquidated cross-collateral positions*—largely cancel out. For stablecoin pools, the dynamics are markedly different: shocks to other tokens replenish stablecoin reserves as loans are repaid, while own-token shocks sometimes have negligible impact, most notably for DAI, because stablecoins are less often pledged as collateral. Two main lessons follow: liquidity risk materializes primarily through own-token shocks, and crypto pools face greater exposure to liquidation-driven reserve depletion. Since the price shocks of primary interest for stablecoins are full de-pegging events that deplete available reserves regardless of propagation, we restrict the liquidity risk analysis to the two main crypto assets, WBTC and ETH.

Variable selection and predictive performance. We first assess whether borrowers’ characteristics improve the predictive power of the model. The robustness checks relegated to [Appendix E.1](#) show that adding them does not improve predictive performance. We therefore



Figure 8 Changes in pool available reserves following a price drop of 50 % (ETH, WBTC, DAI, and USDC). Note: For each pool, the blue solid line shows the daily change in available reserves relative to total assets when the price drop affects the pool's own token; the red dashed line shows the average change when the price drop affects one of the other top-5 tokens. Upper panels correspond to crypto pools, whereas lower panels correspond to stablecoin pools.

restrict the set of user covariates to the depositor side and exclude borrower characteristics ($P = 30$). This finding is consistent with the logic of the propagation algorithm presented in Section 3.2: price shocks impair the collateral of depositors, so it is their connectivity and leverage that primarily determine whether an initial devaluation escalates into a liquidation cascade that generates liquidity losses. We then apply LASSO, which selects 20 variables at $\lambda_{\min}$. After trimming at $\alpha = 0.001$, 14 predictors remain. The asset fixed effect for WBTC is excluded from the preferred specification, as it is not statistically significant once pool and collateral characteristics are controlled for ($\hat{\gamma} = +0.0039$, $p = 0.337$), indicating that the two assets do not differ systematically in liquidity losses beyond what is explained by the observable predictors.

The preferred model achieves a training $R^2 = 0.835$ and a test $R^2 = 0.868$, with $\text{RMSE} = 0.0205$ and $\text{MAE} = 0.0164$ on the held-out test set. The test R^2 exceeding the training R^2 reflects the lower variance of $\ell_{t,j}$ in the test period and confirms that the model generalises well without overfitting. A comparison of all specifications is relegated to Appendix F.1. The full OLS benchmark on all 30 predictors achieves a test $R^2 = 0.784$, substantially below the trimmed model despite using twice as many variables, confirming that LASSO-based selection improves out-of-sample performance by reducing overfitting.

Figure 9 reports the HAC coefficient estimates for the preferred model, organized by predictor group, along with the four diagnostic panels.²⁴ Three channels stand out.




Figure 9 Diagnostic plots for the preferred model — liquidity risk $\ell_{t,j}$ (ETH and WBTC, $s = 50\%$, 14 variables, Train $R^2 = 0.835$, Test $R^2 = 0.868$).
Note: Upper-left: actual against fitted values on the training set (blue). Upper-right: actual against fitted values on the held-out test set (coral). Lower-left: HAC-estimated coefficients with 95 % confidence intervals, ordered by magnitude. Lower-right: residuals over time for the training (blue) and test (coral) periods, separated by the dashed vertical line at the train/test boundary (March 2023).

The most quantitatively important channel operates through the characteristics of *collateral providers*. The value-weighted collateral utilization rate²⁵ carries the largest positive coefficient (+0.0303), indicating that pools whose collateral base is more heavily pledged sustain larger liquidity losses. The value-weighted in-flow weight concentration also carries a large positive coefficient (+0.0257), indicating that pools whose collateral base is more concentrated face amplified liquidity losses. The value-weighted health ratio carries a positive coefficient (+0.0153): pools whose

²⁴ See Table 8 in Appendix F.1 for a detailed summary of the coefficients and their predictive power.

²⁵ This feature corresponds to the share of deposits used as collateral across lenders, averaged with weights proportional to the size of their deposits. See Appendix B for more details on its construction and that of other features.

collateral providers have higher health ratios—meaning they hold more collateral relative to their debt—generate greater reserve destruction upon liquidation, as larger collateral buffers translate into larger liquidation volumes when prices fall sufficiently.

The second channel operates through *own-pool financial characteristics*. Collateral utilization (+0.0131) and the utilization ratio (+0.0057) both carry positive coefficients, confirming the finding that pools with more actively deployed reserves are more exposed. By contrast, the collateral factor (−0.0122) carries a negative coefficient: pools enabling higher collateralization ratios sustain smaller losses because the larger buffers required of borrowers reduce the probability that liquidations generate residual shortfalls.

The third channel operates through the *cross-pool network structure*. Greater heterogeneity in pool connectivity, measured by in-degree dispersion, is associated with larger losses (+0.0106), while the value-weighted average out-degree across pools carries a negative coefficient (−0.0083), consistent with better diversification of liquidation risk in more outwardly connected protocols.

The diagnostic panels in Figure 9 confirm the model’s stability. Residuals in the training period show transient spikes during the Terra/Luna collapse (May 2022) and the FTX crisis (November 2022), both of which were associated with unusually large liquidation waves. In the test period, residuals remain centred around zero with substantially lower volatility, and the transition at the train/test boundary is smooth, confirming structural stability across the two sub-periods.

Forecast error dynamics. Figure 10 examines whether prediction accuracy improves as the model accumulates more data, using two complementary approaches. The left and centre panels plot absolute residuals over time from the preferred model, for the training and test periods respectively.




Figure 10 Forecast error dynamics — liquidity risk $\ell_{t,j}$ (ETH and WBTC, $s = 50\%$).
 Note: Left: absolute residuals from the fixed preferred model over the training period (blue dots), with a 30-day rolling mean and an OLS trend line. Centre: absolute residuals over the test period (coral dots), with a 15-day rolling mean and an OLS trend line. Right: monthly mean absolute error from a recursive expanding-window exercise in which the full LASSO → OLS pipeline is refit at the start of each month using all data available up to that point; the dashed line is the OLS trend.

The right panel reports monthly mean absolute errors from a recursive expanding-window exercise in which the full LASSO → OLS pipeline is refit at the start of each month using all available

data up to that point. Three findings emerge. First, training residuals show a statistically significant downward trend (slope = -3×10^{-6} , $p < 0.001$), indicating that the fixed model’s in-sample fit improves gradually over the training period. Second, this improvement does not translate into better out-of-sample performance: the recursive expanding-window exercise in the right panel shows that monthly MAE does not decline significantly as the training set grows (slope = -1.38×10^{-4} , $p = 0.457$), indicating that the data-generating process is stable and early observations remain informative. The in-sample improvement therefore reflects the model fitting the training data more tightly over time rather than genuinely learning a changing relationship. Third, the test period residuals show an upward trend (slope = 5.0×10^{-5} , $p < 0.001$), but this reflects a short 7-month window and elevated errors concentrated in a specific episode toward the end of the sample rather than a systematic deterioration of predictive accuracy. Taken together, these results support the expanding-window design: since the recursive exercise shows no meaningful improvement from accumulating more data, a fixed-width rolling window that discards early observations would offer no advantage.

6.2. Solvency Risk

We now turn to the factors that drive solvency risk. As discussed in [section 3](#), this risk arises when a decline in collateral values is so severe that liquidators are unable to fully recover borrowers’ nominal debt, thereby leaving residual bad debt in the system.

6.2.1. Crypto Assets

Variable selection and predictive performance. LASSO at $\lambda_{\min}$ selects 18 variables. After trimming at $\alpha = 0.001$ and including the asset fixed effect for WBTC, the preferred specification retains 10 predictors. Unlike the liquidity risk model, the WBTC fixed effect is retained here: WBTC pools generate systematically higher bad debt than ETH pools beyond what is explained by the structural covariates ($\hat{\gamma} = +0.0086$, $p < 0.001$), likely reflecting the lower collateral factor of WBTC (0.687 vs 0.804) and the greater concentration of its user base.

The preferred model achieves a training $R^2 = 0.838$ and a test $R^2 = 0.834$, with RMSE = 0.0072 and MAE = 0.0058. The near-equality of training and test R^2 confirms that the determinants of bad debt are stable across the two sub-periods. A comparison of all specifications is reported in [Table 9](#), which we relegate to [Appendix F.2](#). The full OLS benchmark, using all 30 predictors, achieves a test $R^2 = 0.835$, essentially identical to that of the trimmed model despite employing three times as many variables. This confirms that LASSO successfully compresses the predictive information in the full set of candidates into a parsimonious specification without sacrificing out-of-sample accuracy.

Coefficient estimates. Figure 11 reports the coefficient estimates, organized by predictor group.²⁶ The results confirm three channels, broadly mirroring those identified for liquidity risk but with distinct leading predictors.




Figure 11 Diagnostic plots for the preferred model — solvency risk $b_{t,j}$ (ETH and WBTC, $s = 50\%$, 10 variables, Train $R^2 = 0.838$, Test $R^2 = 0.834$).

Note: Upper-left: actual against fitted values on the training set (blue). Upper-right: actual against fitted values on the held-out test set (coral). Lower-left: HAC-estimated coefficients with 95 % confidence intervals, ordered by magnitude. The WBTC fixed effect is the largest single predictor. Lower-right: residuals over time for the training (blue) and test (coral) periods, separated by the dashed vertical line at the train/test boundary (March 2023).

The most important channel operates through the characteristics of *collateral providers*. The value-weighted collateral utilization rate carries the largest coefficient in absolute value (+0.0071) as pools with more heavily pledged collateral bases generate more bad debt. The value-weighted in-flow weight concentration (+0.0011) reinforces this finding: concentrated collateral structures amplify solvency exposure because the failure of a single large position generates more unrecoverable shortfalls.

²⁶ See Table 10 in Appendix F.2 for a complete description of the coefficient estimates and their predictive power.

The second channel operates through *own-pool network position*. PageRank centrality (+0.0069) is the most important own-pool predictor, reflecting the tendency of central pools to attract larger and more leveraged positions. Its coefficient nearly doubles from +0.0033 (no FE) to +0.0069 (with FE), as the fixed effect disentangles the asset-level bad debt differential from the network-position effect, sharpening identification of the structural channel.

The WBTC fixed effect ($\hat{\gamma} = +0.0086$, $p < 0.001$) absorbs a residual asset-level difference in solvency exposure not captured by the structural covariates. The lower collateral factor of WBTC ($c = 0.687$ versus $c = 0.804$ for ETH) is a plausible contributor to this gap alongside differences in user base composition and liquidity depth. Given that the collateral factor is mostly time-invariant within each asset, its cross-asset effect is largely absorbed into the fixed effect rather than identified separately. Coefficient signs are stable across specifications with and without the fixed effect, confirming that its inclusion sharpens rather than distorts the estimated predictive associations (see Table 10 in Appendix F.2).

6.2.2. Stablecoin Pools

Variable selection and predictive performance. LASSO at $\lambda_{\min}$ selects variables from the full candidate set of $P = 40$ predictors, which for stablecoin pools includes both collateral provider and borrower characteristics. After trimming at $\alpha = 0.001$, 11 predictors are retained. No asset fixed effect is included, as the dummy constructed for DAI versus USDC carries a zero coefficient and an undefined p -value, confirming that the two stablecoin pools do not differ systematically in bad debt exposure beyond what is captured by the covariates.

The preferred model achieves a training $R^2 = 0.958$ and a test $R^2 = 0.929$, with RMSE = 0.0237, on the held-out test set. Both the high level and the near-equality of training and test R^2 confirm that the model generalizes well to the out-of-sample period without overfitting, and that the determinants of stablecoin bad debt under a de-pegging scenario are highly predictable from the observable network and financial characteristics of the pool and its users.

Coefficient estimates. Figure 12 displays the diagnostic plots for the preferred model. The 11 retained predictors span all four feature groups, with three channels of particular interest.

The most prominent channel operates through *own-pool financial characteristics*. Collateral utilization (+) and the collateral factor (+) carry the largest own-pool coefficients, indicating that pools accepting more collateral at higher loan-to-value ratios generate more bad debt when the stablecoin de-pegs. The self-borrowing rate also enters positively, consistent with internally recycled positions being more exposed to a sudden loss of parity.




Figure 12 Diagnostic plots for the selected model — solvency risk $b_{t,j}$ (DAI and USDC, $s = 100\%$, 11 variables, Train $R^2 = 0.958$, Test $R^2 = 0.929$).
Note: Upper-left: actual against fitted values on the training set (blue). Upper-right: actual against fitted values on the held-out test set (coral). Lower-left: HAC-estimated coefficients with 95 % confidence intervals, ordered by magnitude. Lower-right: residuals over time for the training period (blue) and the test period (coral), separated by the dashed vertical line at the train/test boundary (March 2023).

The second channel operates through *collateral network features*. The value-weighted out-weight HHI of collateral providers enters significantly, indicating that more concentrated collateral structures amplify bad debt under a de-pegging event, mirroring the finding from the crypto solvency analysis.

The third and distinctive channel, absent in the crypto models, operates through *borrower characteristics*. Three borrower variables are retained by LASSO: the value-weighted in-weight HHI of borrowers, their value-weighted out-degree, and their value-weighted collateral utilization rate. This confirms the theoretical argument that, in stablecoin pools, solvency exposure is shaped not only by the network position and leverage of collateral providers, but also by those of *borrow-*

ers. When the stablecoin de-pegs, borrowers with more concentrated and leveraged positions are less able to adjust their obligations, generating larger unrecoverable shortfalls. The inclusion of borrower features is thus not only theoretically justified but empirically validated by LASSO’s selection across a candidate set of 40 predictors.

6.3. Incremental Value of Liability-Side Information

A key question motivating our empirical design is whether reconstructing the liability network from on-chain data adds predictive content beyond what can be inferred from the asset side of the pools’ balance sheets as already published by Compound itself. To answer this directly, we estimate a sequence of nested models defined in Table 2, expanding the candidate feature set progressively from a single asset-side predictor to the full liability-side specification. Figure 13 displays the corresponding R^2 profiles.

Table 2 Incremental predictive value of liability-side information.

Model	Liquidity risk ($\ell_{t,j}$)			Solvency risk ($s_{t,j}$)		
	Test R^2	Δ Test R^2	Vars	Test R^2	Δ Test R^2	Vars
M1: Asset side	0.236	—	1	0.393	—	3
M2: + Pool features	0.523	+0.287	11	0.731	+0.338	8
M3: + Collateral network	0.869	+0.346	10	0.749	+0.018	11
M4: + User features (full model)	0.868	≈ 0	14	0.834	+0.085	10

Note: M1: collateral factor and utilization ratio. M2: + pool financial characteristics, centrality, and cross-pool linkages. M3: + exposure-weighted connectivity of collateral providers. M4: + financial condition of collateral providers (health ratio, collateral utilization, self-borrowing rate, deposit share). Δ Test R^2 : gain relative to previous row. LASSO penalty selected by expanding-window cross-validation ($K = 5$ folds); OLS refit on selected variables with Newey–West HAC errors ($L = 4$ lags); trimming retains predictors significant at $\alpha = 0.001$. Predictors standardized using training-set moments only.




Figure 13 Incremental predictive value of liability-side information (ETH and WBTC, $s = 50\%$).

Note: M1 includes the collateral factor and utilization ratio as asset-side predictors. M2 adds pool financial characteristics, centrality measures, and cross-pool linkages. M3 adds exposure-weighted connectivity measures of collateral providers, constructed from the bipartite liability network. M4 adds the financial condition of collateral providers (health ratio, collateral utilization, self-borrowing rate, deposit share), completing the full model. Solid circles show test R^2 ; dashed squares show train R^2 . Left panel: liquidity risk $\ell_{t,j}$. Right panel: solvency risk $s_{t,j}$.

Liquidity risk. The asset-side baseline (M1), which includes only the collateral factor and utilization ratio, explains 23.6% of out-of-sample variation in liquidity risk ($\ell_{t,j}$). Adding the full set of pool features — financial characteristics, centrality measures, and cross-pool linkages — at M2 raises test R^2 to 0.523 (+28.7 percentage points). The dominant gain in the sequence comes at M3, where adding collateral network features raises test R^2 to 0.869 (+34.6 percentage points). Adding the financial condition of depositors at M4 yields essentially no further improvement ($R^2 = 0.868$), indicating that once depositor connectivity is known, their leverage ratios carry no additional predictive content. This pattern confirms that, for reserve depletion, *connectivity is the primary channel*: knowing how broadly depositors are linked across the protocol is sufficient to predict liquidation-driven liquidity losses, and knowing how leveraged they are adds little beyond that.

Solvency risk. The pattern is less stark for solvency risk ($b_{t,j}$). The asset-side baseline (M1) explains 39.3% of out-of-sample variation. Adding pool features at M2 raises test R^2 to 0.731 (+33.8 percentage points) — a larger gain than in the liquidity case, reflecting that pool-level network position (particularly PageRank centrality) is an important predictor of bad debt. Adding the collateral network at M3 delivers only a modest further gain to 0.749 (+1.8 percentage points), and adding the financial condition of collateral providers at M4 yields a final improvement to 0.834 (+8.5 percentage points). The financial condition of collateral providers thus contributes more at the final step than their network connectivity, the reverse of the pattern for liquidity risk. This reflects the different transmission mechanism: bad debt arises when collateral *value* fails to cover outstanding debt, a condition that depends jointly on the network position of depositors — which determines the scale of liquidation cascades — and on their financial condition, which determines whether individual positions generate unrecoverable shortfalls.

The asset-side gap. Taken together, these results establish that the information published directly by Compound, namely the collateral factor and utilization ratio that constitute our asset-side baseline, provides only a partial basis for risk monitoring. A surveillance system relying solely on these observable pool-level statistics would explain less than a quarter of the variation in liquidity risk ($R^2 = 0.236$) and less than 40% of the variation in solvency risk ($R^2 = 0.393$). Even augmenting the asset side with the full set of pool financial characteristics, centrality measures, and cross-pool linkages (M2) leaves substantial predictive gaps of 34.5 percentage points for liquidity risk and 10.3 percentage points for solvency risk relative to the full model. These findings underscore that topology-aware monitoring is not a mere refinement of standard risk dashboards but a necessary input for assessing systemic vulnerability in decentralized lending protocols.

7. Conclusion

This paper provides a high-frequency, network-based empirical analysis of systemic risk in decentralized lending protocols using a full reconstruction of the Compound protocol’s liability network from 2020 to 2024. After computing several topological and financial measures, we quantify the effects of interdependencies and find that network measures are as informative as conventional balance sheet indicators in predicting the transmission of systemic stress. These findings suggest that systemic risk in DeFi is not a function of size or collateral coverage alone, but of how users and assets are embedded in the network of obligations.

Our findings speak to a broader opportunity created by DeFi’s architecture. Unlike traditional financial systems, where network reconstruction depends on confidential exposures and reporting lags, DeFi protocols record every transaction on a public blockchain in real time, allowing balance sheets and liability links to be mapped continuously. Furthermore, smart contracts encode the system’s responses to stress through liquidation rules and interest rate adjustments, making it possible to simulate the propagation of shocks with a precision that is unattainable in traditional finance. Together, these features transform the feasibility of systemic risk assessment: rather than relying on periodic snapshots of aggregated balance sheet indicators, regulators and protocol governance bodies could monitor the full network of obligations and simulate system responses in real time. Our methodology is designed to exploit this opportunity, with results suggesting that doing so captures dimensions of systemic vulnerability that aggregate indicators cannot reveal. As policymakers and central banks consider how new infrastructures—such as CBDCs, Regulated Liability Networks, and tokenized capital markets—might interact with DeFi ecosystems, our findings indicate that understanding network externalities will be critical to evaluating financial resilience.

Several directions remain open for future research. Our simulations could be enriched by incorporating more complex price dynamics, including gradual price declines and intra-day volatility, as well as user-level behavioral responses such as collateral refilling and endogenous deleveraging. Extending the analysis to block-level data would allow the model to be calibrated against intra-day liquidation dynamics, opening the door to very short-term predictions relevant to traders and arbitrageurs rather than the regulatory and governance audience targeted here. Finally, our framework focuses on liquidation-driven propagation but does not model endogenous liquidity runs, in which depositors withdraw en masse, pushing utilization ratios toward full depletion and trapping lenders. The recent liquidity crunch in Aave triggered by the KelpDAO exploit of April 2026 shows that such risk is a practical concern.²⁷ Extending our balance sheet reconstruction and propagation

²⁷ A bridge vulnerability in KelpDAO’s infrastructure enabled the minting of unbacked rsETH tokens, which were used as collateral on Aave to borrow ETH. Once the hack became public, it triggered cascading withdrawals across Aave’s ETH, USDT and USDC pools, driving utilization to 100% and generating substantial bad debt, with contagion spreading from directly exposed ETH lenders to stablecoin lenders (Bristow and Macchiavelli 2026). See also [Aave’s incident report](https://governance.aave.com/t/rsETH-incident-report-april-20-2026/24580) (<https://governance.aave.com/t/rsETH-incident-report-april-20-2026/24580>) for further details.

framework to incorporate depositor run dynamics would therefore be a natural and policy-relevant next step.

References

- ABDALLAH, N. B., H. DABBOU, AND G. M. IMEN (2022): “The Contribution of Spatial Econometrics in the Field of Empirical Finance,” *Review of Economics and Finance*, 20, 17–31, iSSN: 1923-7529.
- ACEMOGLU, D., A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2015): “Systemic Risk and Stability in Financial Networks,” *American Economic Review*, 105, 564–608.
- ALLEN, F. AND D. GALE (2000): “Financial Contagion,” *Journal of Political Economy*, 108, 1–33.
- ANGRIMAN, E., A. VAN DER GRINTEN, M. HAMANN, H. MEYERHENKE, AND M. PENSCHUCK (2023): “Algorithms for large-scale network analysis and the NetworKit toolkit,” in *Algorithms for Big Data: DFG Priority Program 1736*, Springer Nature Switzerland Cham, 3–20.
- ANSELIN, L. (1988): *Spatial Econometrics: Methods and Models*, vol. 4 of *Studies in Operational Regional Science*, Springer, 1st ed.
- AZAR, P. D., A. CASILLAS, AND M. FARBOODI (2024): “Information and Market Power in DeFi Intermediation,” Staff Report 1102, FRB of New York, revised September 2024.
- BARDOSCIA, M., S. BATTISTON, F. CACCIOLI, AND G. CALDARELLI (2015): “DebtRank: A microscopic foundation for shock propagation,” *PloS one*, 10, e0130406.
- BARTOLETTI, M., J. H.-Y. CHIANG, AND A. L. LAFUENTE (2021): “SoK: Lending Pools in Decentralized Finance,” in *Financial Cryptography and Data Security. FC 2021 International Workshops: CoDecFin, DeFi, VOTING, and WTSC, Virtual Event, March 5, 2021, Revised Selected Papers*, Berlin, Heidelberg: Springer-Verlag, 553–578.
- BATTISTON, S., M. PULIGA, R. KAUSHIK, P. TASCA, AND G. CALDARELLI (2012): “Debtrank: Too central to fail? financial networks, the fed and systemic risk,” *Scientific reports*, 2, 541.
- BERGMEIR, C. AND J. M. BENÍTEZ (2012): “On the use of cross-validation for time series predictor evaluation,” *Information Sciences*, 191, 192–213, data Mining for Software Trustworthiness.
- BERTUCCI, C., L. BERTUCCI, M. G. DELAUNAY, O. GUÉANT, AND M. LESBRE (2026): “Agents’ Behavior and Interest Rate Model Optimization in DeFi Lending,” *Mathematical Finance*, 36, 374–396.
- BEUTEL, J., S. LIST, AND G. VON SCHWEINITZ (2019): “Does machine learning help us predict banking crises?” *Journal of Financial Stability*, 45, 100693.
- BOSS, M., H. ELSINGER, M. SUMMER, AND S. THURNER (2004): “Network Topology of the Interbank Market,” *Quantitative Finance*, 4, 677–684.
- BRISTOW, L. AND M. MACCHIAVELLI (2026): “Crypto Hacks and DeFi Runs,” <https://bpi.com/crypto-hacks-and-defi-runs/>, accessed May 2026.

- CACCIOLI, F., P. BARUCCA, AND T. KOBAYASHI (2018): “Network models of financial systemic risk: a review,” *Journal of Computational Social Science*, 1, 81–114.
- CACCIOLI, F., M. MARSILI, AND P. VIVO (2009): “Eroding Market Stability by Proliferation of Financial Instruments,” *The European Physical Journal B: Condensed Matter and Complex Systems*, 71, 467–479.
- CALDARELLI, G. (2007): *Scale-Free Networks: Complex Webs in Nature and Technology*, Oxford University Press.
- CAPPONI, A., G. IYENGAR, AND J. SETHURAMAN (2023): “Decentralized Finance: Protocols, Risks, and Governance,” *arXiv preprint*.
- CARVALHO, V. M., M. NIREI, Y. U. SAITO, AND A. TAHBAZ-SALEHI (2021): “Supply Chain Disruptions: Evidence from the Great East Japan Earthquake,” *The Quarterly Journal of Economics*, 136, 1255–1321.
- CONG, L., K. TANG, Y. WANG, AND X. ZHAO (2023): “Inclusion and Democratization Through Web3 and DeFi? Initial Evidence from the Ethereum Ecosystem,” Tech. rep., SSRN.
- DAIAN, P., S. GOLDFEDER, T. KELL, Y. LI, X. ZHAO, I. BENTOV, L. BREIDENBACH, AND A. JUELS (2020): “Flash Boys 2.0: Frontrunning in Decentralized Exchanges, Miner Extractable Value, and Consensus Instability,” in *2020 IEEE Symposium on Security and Privacy (SP)*, IEEE, 910–927.
- DAVIS, J. AND M. GOADRICH (2006): “The relationship between Precision-Recall and ROC curves,” in *Proceedings of the 23rd International Conference on Machine Learning*, New York, NY: ACM, 233–240.
- DI GIOVANNI, J., A. A. LEVCHENKO, AND I. MEJEAN (2014): “Firms, Destinations, and Aggregate Fluctuations,” *Econometrica*, 82, 1303–1340.
- DI MAGGIO, M., A. KERMANI, AND Z. SONG (2017): “The Value of Trading Relations in Turbulent Times,” *Journal of Financial Economics*, 124, 266–284.
- DOTAN, M., A. YAISH, H.-C. YIN, E. TSYTKIN, AND A. ZOHAR (2023): “The Vulnerable Nature of Decentralized Governance in DeFi,” in *Proceedings of the 2023 Workshop on Decentralized Finance and Security*, New York, NY, USA: Association for Computing Machinery, DeFi ’23, 25–31.
- ELHORST, J. P. (2014): *Spatial Econometrics: From Cross-sectional Data to Spatial Panels*, Heidelberg, New York, Dordrecht, London: Springer.
- ELLIOTT, M., B. GOLUB, AND M. O. JACKSON (2014): “Financial Networks and Contagion,” *American Economic Review*, 104, 3115–3153.
- ELSINGER, H., A. LEHAR, AND M. SUMMER (2006): “Risk Assessment for Banking Systems,” *Management Science*, 52, 1301–1314.
- GABAIX, X. (2011): “The Granular Origins of Aggregate Fluctuations,” *Econometrica*, 79, 733–772.
- GENNAIOLI, N., A. SHLEIFER, AND R. W. VISHNY (2013): “A Model of Shadow Banking,” *Journal of Finance*, 68, 1331–1363.

- GIBBONS, S. AND H. G. OVERMAN (2012): “Mostly Pointless Spatial Econometrics?” *Journal of Regional Science*, 52, 172–191.
- GLASSERMAN, P. AND H. P. YOUNG (2016): “Contagion in Financial Networks,” *Journal of Economic Literature*, 54, 779–831.
- GREEN, R. C., B. HOLLIFIELD, AND N. SCHÜRHO (2007): “Dealer Intermediation and Price Behavior in the Aftermarket for New Bond Issues,” *Journal of Financial Economics*, 86, 643–682.
- GUDGEON, L., D. PEREZ, D. HARZ, B. LIVSHITS, AND A. GERVAIS (2020): “The Decentralized Financial Crisis,” in *2020 Crypto Valley Conference on Blockchain Technology (CVCBT)*, 1–15.
- HEIMBACH, L., V. PAHARI, AND E. SCHERTENLEIB (2024): “Non-Atomic Arbitrage in Decentralized Finance,” Working Paper.
- HOLLIFIELD, B., A. NEKLYUDOV, AND C. SPATT (2017): “Bid-Ask Spreads, Trading Networks, and the Pricing of Securitizations,” *The Review of Financial Studies*, 30, 3048–3085.
- KLEINBERG, J. M. (1999): “Authoritative sources in a hyperlinked environment,” *Journal of the ACM (JACM)*, 46, 604–632.
- LEHAR, A. AND C. A. PARLOUR (2022): *Systemic fragility in decentralized markets*, Bank for International Settlements, Monetary and Economic Department.
- (2023): “Battle of the Bots: Flash Loans, Miner Extractable Value and Efficient Settlement,” Tech. rep., SSRN.
- LESAGE, J. P. AND R. K. PACE (2009): *Introduction to Spatial Econometrics*, Statistics: A Series of Textbooks and Monographs, Boca Raton, FL: Chapman & Hall/CRC, illustrated ed.
- LESHNER, R. AND G. HAYES (2019): “Compound: The money market protocol,” Tech. rep.
- MANSKI, C. F. (1993): “Identification of Endogenous Social Effects: The Reflection Problem,” *The Review of Economic Studies*, 60, 531–542.
- MILIONIS, J., C. C. MOALLEMI, T. ROUGHGARDEN, AND A. L. ZHANG (2023): “Automated Market Making and Loss-Versus-Rebalancing,” Working Paper.
- NIER, E., J. YANG, T. YORULMAZER, AND A. ALENTORN (2007): “Network Models and Financial Stability,” *Journal of Economic Dynamics and Control*, 31, 2033–2060.
- PAELINCK, J. H. P. AND L. H. KLAASSEN (1979): *Spatial Econometrics*, vol. 1 of *Studies in Spatial Analysis*, Farnborough, UK: Saxon House, illustrated, reprinted ed., originally from the University of Michigan. Digitized on July 24, 2009.
- PAGE, L., S. BRIN, R. MOTWANI, AND T. WINOGRAD (1999): “The PageRank citation ranking: Bringing order to the web.” Tech. rep., Stanford infolab.
- PARK, A. AND J. STINNER (2023): “Phantom liquidity in decentralized lending,” *SSRN Electronic Journal*.

- QIN, K., L. ZHOU, P. GAMITO, P. JOVANOVIĆ, AND A. GERVAIS (2021): “An empirical study of DeFi liquidations: incentives, risks, and instabilities,” in *Proceedings of the 21st ACM Internet Measurement Conference*, New York, NY, USA: Association for Computing Machinery, IMC ’21, 336–350.
- SAITO, T. AND M. REHMSMEIER (2015): “The Precision-Recall Plot Is More Informative than the ROC Plot When Evaluating Binary Classifiers on Imbalanced Datasets,” *PLOS ONE*, 10, e0118432.
- SUN, X., S. LIN, V. SJÖBERG, AND J. JIE (2021): “How to Exploit a DeFi Project,” in *Financial Cryptography and Data Security. FC 2021 International Workshops*, ed. by M. Bernhard, A. Bracciali, L. Gudgeon, T. Haines, A. Klages-Mundt, S. Matsuo, D. Perez, M. Sala, and S. Werner, Berlin, Heidelberg: Springer, Lecture Notes in Computer Science, 162–167.
- SUN, X., C. STASINAKIS, AND G. SERMPINIS (2022): “Liquidity risks in lending protocols: Evidence from aave protocol,” *arXiv preprint arXiv:2206.11973*.
- TASHMAN, L. J. (2000): “Out-of-sample tests of forecasting accuracy: an analysis and review,” *International Journal of Forecasting*, 16, 437–450, the M3- Competition.
- TOVANICH, N., M. KASSOUL, S. WEIDENHOLZER, AND J. PRAT (2023): “Contagion in Decentralized Lending Protocols: A Case Study of Compound,” in *Proceedings of the 2023 Workshop on Decentralized Finance and Security, DeFi 2023, Copenhagen, Denmark, 30 November 2023*, ed. by K. Qin and F. Zhang, ACM, 55–63.

Appendix A. Aggregation Procedure

We explain in this appendix how to reconstruct the pool's balance sheet after liquidations. We denote by $\mathbf{C}_j$ and $\mathbf{D}_j$ the amounts borrowed and deposited in pool j , so that:

$$\mathbf{C}_j = \sum_{i=1}^M L_{ij}, \quad \mathbf{D}_j = \sum_{i=1}^M L_{ji},$$

where $M = |\mathcal{U}|$ is the number of users. We also use $\mathbf{T}_j$ to denote the available tokens in pool j (i.e., deposited but not borrowed), and $\mathbf{E}_j$ the equity of pool j . Considering that pools have no bad debt initially, the balance sheet of pool j is given by

$$\begin{cases} \text{Assets}_j &= \mathbf{T}_j + \mathbf{C}_j, \\ \text{Liabilities}_j &= \mathbf{D}_j + \mathbf{E}_j. \end{cases}$$

We now characterize the impact of a price shock. We denote by $\mathcal{L}$ and $\mathcal{L}'$ the subsets of liquidated users without and with bad debt after the shock, respectively. As explained in [subsection 3.2](#), the liquidated amounts Δ_i for each user are given by

$$\Delta_i = \begin{cases} \frac{D_i - CL_i}{D_i \cdot \left[1 - (1 + LB) \cdot \frac{CL_i}{C_i} \right]}, & \text{when } i \in \mathcal{L}, \\ \frac{C_i}{(1 + LB) \cdot D_i}, & \text{when } i \in \mathcal{L}', \end{cases}$$

where C_i , D_i , and CL_i denote the collateral, debt, and pledgeable collateral of user i , as defined in (2) and (3) respectively. Aggregating their values across all users yields the following updating rules for the components of the pools' balance sheets:

$$\begin{aligned} \mathbf{T}'_j &= \mathbf{T}_j + \sum_{i \in \mathcal{L} \cup \mathcal{L}'} \Delta_i \cdot L_{ij} - (1 + LB) \cdot \sum_{i \in \mathcal{L} \cup \mathcal{L}'} \Delta_i \cdot \frac{D_i}{C_i} \cdot L_{ji}, \\ \mathbf{C}'_j &= \mathbf{C}_j - \sum_{i \in \mathcal{L}'} L_{ij} - \sum_{i \in \mathcal{L}} \Delta_i \cdot L_{ij}, \\ \mathbf{D}'_j &= \mathbf{D}_j - (1 + LB) \cdot \sum_{i \in \mathcal{L} \cup \mathcal{L}'} \Delta_i \cdot \frac{D_i}{C_i} \cdot L_{ji}, \\ \mathbf{B}'_j &= \sum_{i \in \mathcal{L}'} (1 - \Delta_i) \cdot L_{ij}, \end{aligned}$$

where primes denote post-shock values and $\mathbf{B}_j$ is the bad debt in pool j . By definition, we have $\mathbf{E}'_j = \mathbf{E}_j - \mathbf{B}'_j$. Using k to index the shocked pool, the total bad debt generated by the price shock to k is given by

$$BD_k = \sum_{j=1}^p \mathbf{B}'_j, \tag{10}$$

where $p = |\mathcal{P}|$ is the number of pools. When liquidation volumes are large enough to deplete the reserves of a pool, the equations above must be adjusted to account for the inability of liquidators

to redeem their **cTokens**. Since unredeemed **cTokens** remain in the pools, they are added to the post-liquidation deposits:

$$\begin{aligned} \mathbf{T}'_j &= \max \left\{ 0, \mathbf{T}_j + \sum_{i \in \mathcal{L} \cup \mathcal{L}'} \Delta_i \cdot L_{ij} - (1 + LB) \cdot \sum_{i \in \mathcal{L} \cup \mathcal{L}'} \Delta_i \cdot \frac{D_i}{C_i} \cdot L_{ji} \right\}, \\ \mathbf{D}'_j &= \mathbf{D}_j - (1 + LB) \cdot \sum_{i \in \mathcal{L} \cup \mathcal{L}'} \Delta_i \cdot \frac{D_i}{C_i} \cdot L_{ji} \\ &\quad + \max \left\{ 0, (1 + LB) \cdot \sum_{i \in \mathcal{L} \cup \mathcal{L}'} \Delta_i \cdot \frac{D_i}{C_i} \cdot L_{ji} - \mathbf{T}_j - \sum_{i \in \mathcal{L} \cup \mathcal{L}'} \Delta_i \cdot L_{ij} \right\}. \end{aligned}$$

Appendix B. Feature Extraction

The bipartite network represents liability relationships between user and lending pool nodes, denoted as L_{ij} , where node i has a liability to node j . Based on the daily asset prices, the edge weights correspond to the liability amounts converted to U.S. dollars.

Node features. We extract node-level features as summarized in Table 3. We further compute the share of each lending pool’s equity, external assets, and size to assess the relative scale of each lending pool. These shares are calculated by normalizing each variable by the total across all pool nodes, i.e., $share_j = variable_j / \sum_k variable_k$, where j indexes a specific lending pool and k ranges over all pools.

Table 3 Node features for users and lending pools

Feature	User Node	Lending Pool Node
In-degree	Number of deposited assets	Number of borrowing users
In-weight	Total deposited amount	Total borrowed amount
In-weight HHI	Herfindahl index of deposited assets	Herfindahl index of borrowing users
Out-degree	Number of borrowed assets	Number of depositing users
Out-weight	Total borrowed amount	Total deposited amount
Out-weight HHI	Herfindahl index of borrowed assets	Herfindahl index of depositing users
In/Out Ratio	Deposit-to-borrow ratio	Borrow-to-deposit ratio
Health Ratio	Health ratio of users	N/A
External assets	0 (assuming no external assets)	Available liquidity (cash) + reserves
Equity	External assets + In-weight – Out-weight	
Size	In-weight + Out-weight	
% Self-borrowing	Percentage of self-borrowed amount	

Note: The Herfindahl-Hirschman Index (HHI) is calculated as $HHI = \sum_i s_i^2$, where s_i is the share of value associated with asset i (for users: deposits or borrows; for pools: users depositing or borrowing). A higher HHI indicates greater concentration.

In addition to these attributes, we compute several centrality measures for each node to characterize its structural role in the network:

- *Closeness Centrality*, estimated using approximate algorithms provided by the NetworkKit library (Angriman et al. 2023), due to computational limitations of exact methods;
- *Hubs and Authorities*, computed via the HITS algorithm (Kleinberg 1999);
- *PageRank*, which we use as a practical proxy for eigenvector centrality, given its faster convergence and scalability (Page et al. 1999).

Edge features. For each user–pool edge, we compute the share of deposits and borrows to quantify the user’s relative exposure to the pool. These features are defined as:

$$\text{shareDeposit}_{u,i} = \frac{\text{deposit}_{u,i}}{\sum_i \text{deposit}_{u,i}}, \quad \text{shareBorrow}_{u,i} = \frac{\text{borrow}_{u,i}}{\sum_i \text{borrow}_{u,i}}, \quad (11)$$

where $\text{deposit}_{u,i}$ and $\text{borrow}_{u,i}$ denote the amounts deposited to or borrowed from pool i by user u . These ratios capture the proportion of a user’s activity attributed to a specific pool.

Lenders and borrowers features: We compute the weighted average of node attributes for lenders (nodes with incoming edges from the pool) and borrowers (nodes with outgoing edges to the pool), using the out-weight and in-weight, respectively, as weighting factors. This allows us to capture the average characteristics of users interacting with each pool, adjusted by the magnitude of their deposited or borrowed amounts.

Table 4 lists the features included in the regression models. The features include topological metrics and financial attributes for the shocked pool, as well as user-aggregated statistics for lenders and borrowers.

Table 4 Predictor variables used in the regression models for liquidity risk ($\ell_{t,j}$) and solvency risk ($b_{t,j}$). Variables are organized by predictor group. Collateral provider and borrower statistics are value-weighted averages computed at the pool level.

Variable	Pool level		Collateral (wavg.)		Borrowers (wavg.)
	Own-pool	Cross-pool	Network	Financial	Network & Financial
<i>Topological features</i>					
In-degree	✓		✓		✓
Out-degree	✓		✓		✓
In-weight HHI	✓		✓		✓
Out-weight HHI	✓		✓		✓
In-flow weight			✓		✓
In-flow weight HHI			✓		✓
Out-flow weight HHI			✓		✓
PageRank	✓				
Hub score	✓				
In-degree dispersion		✓			
In-degree (wavg.)		✓			
Out-degree (wavg.)		✓			
Deposit concentration (HHI)		✓			
Borrowing concentration (HHI)		✓			
Self-borrowing dispersion		✓			
Self-borrowing share (wavg.)		✓			
<i>Financial features</i>					
Collateral utilization	✓			✓	✓
Utilization ratio	✓				
Self-borrowing rate	✓			✓	✓
Collateral factor	✓				
Health ratio				✓	✓
Deposit share				✓	✓

Notes: HHI: Herfindahl–Hirschman index. Wavg.: value-weighted average over collateral providers or borrowers at the pool level. Cross-pool variables are protocol-level and identical across pools on any given date. “Self-borrowing rate” refers to distinct variables at the pool, collateral provider, and borrower levels. Borrower features are included in the stablecoin specification and in the robustness checks; they are excluded from the baseline crypto specification.

Appendix C. Descriptive Statistics

Table 5 Descriptive statistics (Panel A) — Pool characteristics and collateral provider features, 1 January 2021 – 15 June 2024

	ETH		WBTC		DAI		USDC	
Variable	Mean	Std	Mean	Std	Mean	Std	Mean	Std
<i>Outcome variables</i>								
Liquidity risk (ℓ)	0.1294	0.0413	0.0941	0.0449	—	—	—	—
Bad debt share (b)	0.0236	0.0151	0.0077	0.0062	0.0168	0.0192	0.1036	0.0696
<i>Own-pool characteristics</i>								
Collateral utilization	0.3764	0.0393	0.3472	0.0441	0.3961	0.1841	0.3346	0.1131
Self-borrowing rate	0.0161	0.0155	0.0205	0.0286	0.3744	0.2061	0.2115	0.1778
Utilization ratio	0.0409	0.0222	0.0350	0.0334	0.5991	0.1523	0.6042	0.1610
Collateral factor	0.8043	0.0328	0.6869	0.0272	0.8073	0.0362	0.8196	0.0445
In-degree	1310.81	218.82	449.57	19.35	2786.53	136.81	2578.44	250.72
In-weight HHI	0.1988	0.0972	0.2577	0.1054	0.1782	0.0651	0.1034	0.0343
Out-degree	70102.56	4105.48	3310.23	218.33	18512.81	924.06	217109.73	998.96
Out-weight HHI	0.0508	0.0191	0.0703	0.0316	0.1290	0.0530	0.0682	0.0413
<i>Pool centrality</i>								
PageRank (scaled)	0.4887	0.1033	0.1372	0.0132	0.9956	0.0168	0.9098	0.0577
Hub score (scaled)	0.4776	0.4219	0.3832	0.4244	0.3936	0.4561	0.2264	0.3530
<i>Cross-pool linkages^a</i>								
	All assets							
Deposit concentration (HHI)	0.2188	0.0205						
Borrowing concentration (HHI)	0.3393	0.0296						
In-degree dispersion	901.00	79.34						
Out-degree dispersion	54425.86	4130.41						
Self-borrowing dispersion	0.2726	0.0836						
Out-degree (wavg.)	68577.29	14517.57						
In-degree (wavg.)	2309.51	273.74						
Self-borrowing share (wavg.)	0.3847	0.2249						
<i>Collateral network features</i>								
In-degree (wavg.)	2.0017	0.4203	2.3651	0.3465	1.7121	0.5971	1.7050	0.4578
In-flow weight (wavg.)	160.95M	217.55M	122.86M	128.69M	256.37M	276.89M	193.90M	247.64M
In-flow weight HHI (wavg.)	0.8273	0.0391	0.7649	0.0538	0.9101	0.0964	0.9022	0.0685
Out-degree (wavg.)	1.7533	0.3142	1.9325	0.2936	1.4151	0.3986	1.4778	0.3929
Out-flow weight (wavg.)	66.24M	83.92M	49.40M	51.67M	181.14M	206.45M	136.46M	180.42M
Out-flow weight HHI (wavg.)	0.8311	0.0782	0.8082	0.0814	0.9496	0.0511	0.9397	0.0567
<i>Collateral financial features</i>								
Self-borrowing rate (wavg.)	0.0645	0.0752	0.0455	0.0530	0.8661	0.2235	0.5350	0.3170
Collateral utilization (wavg.)	0.4117	0.0355	0.3800	0.0389	0.6424	0.0840	0.6630	0.0572
Health ratio (wavg.)	2.3947	0.3486	2.2618	0.3540	1.3535	0.2889	1.3780	0.2316
Deposit share (wavg.)	0.8240	0.0442	0.7716	0.0645	0.9005	0.0984	0.9008	0.0722

Table 6 Descriptive statistics (Panel B) — Borrower features, 1 January 2021 – 15 June 2024

Variable	ETH		WBTC		DAI		USDC	
	Mean	Std	Mean	Std	Mean	Std	Mean	Std
<i>Borrower network features</i>								
In-degree (wavg.)	3.6521	1.4271	3.7100	1.5268	2.3169	0.5319	2.2937	0.4527
In-flow weight (wavg.)	70.35M	84.72M	67.76M	89.94M	244.01M	266.08M	209.36M	232.33M
In-flow weight HHI (wavg.)	0.7682	0.1383	0.6725	0.1761	0.8268	0.0918	0.7856	0.0780
Out-degree (wavg.)	3.2595	0.9393	3.9186	1.1931	2.0375	0.3783	2.1007	0.3229
Out-flow weight (wavg.)	38.87M	52.83M	35.83M	56.32M	148.47M	163.70M	118.17M	137.12M
Out-flow weight HHI (wavg.)	0.7044	0.1192	0.6710	0.1493	0.8236	0.0833	0.7883	0.0747
<i>Borrower financial features</i>								
Self-borrowing rate (wavg.)	0.4250	0.2251	0.4649	0.2557	0.4933	0.2169	0.3118	0.2024
Collateral utilization (wavg.)	0.5198	0.0771	0.5040	0.0850	0.5664	0.0684	0.5068	0.0657
Health ratio (wavg.)	1.6371	0.3213	1.6919	0.3209	1.6265	0.3156	1.7489	0.2562
Deposit share (wavg.)	—	—	—	—	—	—	—	—

Notes: Post-break sample, 1 January 2021 – 15 June 2024. Crypto columns use $s = 50\%$; stablecoin columns use $s = 100\%$. Health threshold = 10: collateral provider averages include only users with health ratio below this value; results are insensitive to this choice within the empirically relevant range. $\ell_{t,j}$ and $b_{t,j}$ are defined in Section 5; liquidity risk is not reported for stablecoins as a full de-pegging produces zero reserve losses by construction. Wavg.: value-weighted average over collateral providers or borrowers. HHI: Herfindahl–Hirschman index. Flow weights in millions of dollars.

^a Cross-pool variables are protocol-level and identical across assets on any given date; reported once.

Appendix D. Simulated vs. Realized Liquidations

We compare simulated liquidation volumes from our propagation algorithm to realized daily liquidations on Compound over the full sample period. Realized liquidation volumes are obtained directly from the Compound V2 subgraph. For each day, we simulate liquidations using the observed end-of-day network snapshot under a range of shock sizes.

As discussed in Section 3.2, our algorithm provides a conservative lower bound on realized liquidation volumes due to the close-factor discontinuity (Assumption A1) and the omission of intra-day liquidation cascades. To quantify this gap, we identify the price shock $s^* = s \times m$, where s is the inter-day price shocks observed in the data, and $m \geq 1$ is a scalar multiplier that best matches realized liquidation volumes. To assess whether it correctly identifies the days on which extreme liquidation activity occurs, we use the Precision-Recall AUC (PR-AUC) as our evaluation metric (Beutel et al. 2019, Davis and Goadrich 2006, Saito and Rehmsmeier 2015).

We define a realized extreme day as one in the top 1% of the realized liquidation-to-total-assets distribution, resulting in the selection of 17 days within our sample. For each multiplier $m \in [1.0, 4.0]$, we treat the simulated liquidation-to-total-assets ratio as a continuous risk score and evaluate it against the binary indicator of realized extreme days. The PR curve (Figure 14) traces the tradeoff between two metrics as the flagging threshold varies. Precision measures the share of simulation-flagged days that are genuinely extreme in the data. Recall measures the share of the 17 extreme days correctly flagged. High thresholds produce few but reliable alerts (top-left of the

curve), while low thresholds flag most extreme days at the cost of more false positives (bottom-right). A random classifier has PR-AUC equal to the base rate of $17/1,600 \approx 0.010$; a perfect classifier has $\text{PR-AUC} = 1.0$.

All multipliers substantially outperform the random baseline. The benchmark multiplier $m = 1.50$ achieves the highest PR-AUC of 0.309 — a $31\times$ improvement over random. At low recall thresholds (top-left of the curve), precision exceeds 60–80%: when the simulation flags only its most extreme days, the vast majority correspond to stress episodes in the data. This confirms that the simulation provides a reliable early warning signal for tail liquidation risk, even when calibrated using only end-of-day network snapshots and a single fixed shock size.




Figure 14 Precision-Recall curve — simulation as early warning signal for extreme liquidation days. Note: The x-axis shows recall (share of the 17 realized extreme days correctly flagged); the y-axis shows precision (share of flagged days that are genuinely extreme). Each point corresponds to a specific threshold on the simulated liquidation-to-total-assets ratio. The dashed line is the random classifier baseline ($\text{PR-AUC} = 0.010$); $m = 1.50$ achieves the highest PR-AUC of 0.309 ($31\times$ above baseline).

Figure 15 reports the six best-matched episodes, spanning the COVID crash (March 2020), the Bitcoin flash crash (May 2021), the January 2022 market correction, the Terra/Luna collapse (May–June 2022), and the FTX crisis (November 2022). In each case, the scaled simulated series closely tracks the shape and timing of realized liquidations, with the peak correctly identified on or near the day of the crash. This confirms that our propagation mechanism was economically active during realized stress events and that our simulations, once adjusted for intra-day fluctuations,

provide a reasonable benchmark for the scale of losses. In accordance with this interpretation, the lower multiplier required for November 2022 ($m = 1.50$) relative to May 2021 ($m = 1.75$) is consistent with the more gradual nature of the FTX-induced price decline compared to the sharp Bitcoin flash crash.²⁸




Figure 15 Realized vs. simulated liquidations around six of the largest stress episodes in the sample. Note: Each panel shows daily liquidation volumes as a share of total protocol assets (%) over a ± 7 -day window around the peak stress day (vertical dotted line). Dark blue solid: realized liquidations; orange dashed: simulated liquidations with episode-specific multiplier m (shown in legend). Dollar values indicate peak realized liquidation volumes.

Appendix E. Robustness Checks

Appendix E.1. Robustness to inclusion of borrower features — crypto assets.

The baseline specification for crypto pools excludes borrower characteristics from the candidate set, on the theoretical grounds that risk transmission in crypto pools operates through the collateral channel rather than through borrowers. When a crypto price falls, it is the collateral posted by lenders that becomes impaired, triggering liquidation cascades regardless of borrower behaviour. To assess whether this exclusion materially affects the results, we re-estimate the full pipeline for $s = 50\%$ augmenting the candidate set with the borrower network and financial features used in the stablecoin specification.

Figure 16 reports the comparison. Adding borrower features leaves the collateral-side coefficients largely stable in both sign and magnitude: no previously significant variable changes sign, and

²⁸ Bitcoin's intraday high-low spread was approximately 30% on May 19, 2021 and 15% on November 9, 2022 (source: CoinGecko historical OHLC data, <https://www.coingecko.com>).

the coefficient on the dominant predictor — value-weighted collateral utilization (-0.0303) — is virtually unchanged.




Figure 16 Robustness to inclusion of borrower features — trimmed model (no asset fixed effect), liquidity shock $\ell_{t,j}$ (ETH and WBTC, $s = 50\%$).

Note: Each column shows the same three diagnostics for the two specifications side by side. Row 1: actual against fitted values on the held-out test set. Row 2: HAC-estimated coefficients with 95 % confidence intervals, ordered by magnitude; grey bars indicate variables present in the union of both specifications but not selected in that particular model. Row 3: residuals over time for the training (blue) and test (coral) periods, separated by the dashed vertical line at the train/test boundary (March 2023).

Out-of-sample predictive performance is also unaffected, with test R^2 declining marginally from 0.868 (without borrowers) to 0.827 (with borrowers), consistent with the additional variables introducing noise rather than signal in the crypto setting. LASSO retains the same 14 collateral-side predictors in both specifications, with the borrower variables either not selected or dropped at the trimming stage. This confirms that borrower characteristics carry no independent predictive content for crypto liquidity risk once collateral-side features are controlled for, and validates the parsimonious baseline specification.

Appendix E.2. Robustness to training subsampling.

Figure 17 reports the distribution of out-of-sample R^2 across 50 subsampling draws for each shock size $s \in \{10\%, \dots, 90\%\}$, where each draw randomly drops 10% of training observations while keeping the test set and variable set fixed. The figure therefore simultaneously addresses two robustness questions: whether results are sensitive to the choice of shock size, and whether they are sensitive to which specific training observations are included.




Figure 17 Robustness to random subsampling of the training set — trimmed+FE model, outcome: liquidity loss share $\ell_{t,j}$ (ETH and WBTC).
Note: Each box plot summarises the distribution of R^2 across 50 draws in which 10% of training observations are dropped at random; the red dot marks the baseline result using the full training set. Left: out-of-sample (test) R^2 . Right: in-sample (training) R^2 .

On the first question, the figure confirms an inverted-U pattern consistent with Figure 7: model fit is low and highly dispersed for $s \leq 20\%$, where few positions are liquidated and the outcome has limited cross-sectional variation. From $s = 40\%$ onward, the baseline test R^2 stabilises, with the $s = 50\%$ benchmark corresponding to the peak of predictive accuracy.

On the second question, the box plots are narrow across all shock sizes from $s = 40\%$ onward, with the baseline red dots sitting on or near the median in every case. This confirms that the baseline results are not driven by specific training observations and that there is no indication of temporal

parameter instability that would motivate a fixed-width rolling window over the expanding-window design. The exception is $s = 10\%$, where the wider distribution simply reflects the low-signal environment noted above.

Note that the figure is generated using the trimmed+FE specification for comparability across shock sizes. Results are virtually identical for the preferred no-FE specification, given the non-significance of the fixed effect ($\hat{\gamma} = -0.0039$, $p = 0.337$) and the negligible difference in test R^2 (0.868 vs 0.871).

Appendix F. Model Comparison

Appendix F.1. Liquidity Risk

Table 7 summarizes the predictive performance of all specifications considered for the liquidity loss share $\ell_{t,j}$. The full OLS benchmark on all 30 predictors achieves the lowest training R^2 relative to its complexity, with a test R^2 of 0.784 substantially below the LASSO-based specifications, confirming that unrestricted OLS overfits the training data. The LASSO model at $\lambda_{\min}$ selects 20 variables and achieves the highest test R^2 (0.878), while the trimmed model retains only the 14 variables significant at $\alpha = 0.001$ with a test R^2 of 0.868 and only a marginal increase in RMSE. The asset fixed effect for WBTC is not significant ($\hat{\gamma} = -0.0039$, $p = 0.337$) and its inclusion in the trimmed model improves test R^2 by only 0.003, confirming that the preferred specification is the trimmed model without fixed effect.

Table 7 Model comparison — Liquidity Risk ($\ell_{t,j}$), ETH and WBTC,
 $s = 50\%$

Model	Train R^2	Test R^2	RMSE	Vars
Full OLS (all vars)	0.855	0.784	0.0263	30
LASSO $\lambda_{\min}$, no FE	0.839	0.878	0.0198	20
Trimmed, no FE (preferred)	0.835	0.868	0.0205	14
Trimmed + FE	0.836	0.871	0.0203	15

Notes: OLS on training set; Newey–West HAC errors ($L = 4$ lags); test metrics on held-out test set. WBTC fixed effect not significant ($\hat{\gamma} = -0.0039$, $p = 0.337$); excluded from preferred specification. Trimming removes predictors with $p \geq 0.001$.

Table 8 and Figure 18 report the coefficient estimates for the trimmed model without and with the asset fixed effect. Coefficients are remarkably stable across the two specifications: in no case does the sign change, and the magnitude changes are economically negligible, confirming that the exclusion of the fixed effect does not distort the estimated predictive associations.

Table 8 Coefficient comparison — Trimmed model without and with asset fixed effect, liquidity risk ($\ell_{t,j}$)

Variable	No FE		With FE	
	Coef.	<i>p</i> -value	Coef.	<i>p</i> -value
<i>Own-pool characteristics</i>				
Collateral utilization	+0.0131	< 0.001	+0.0130	< 0.001
Self-borrowing rate	−0.0105	< 0.001	−0.0107	< 0.001
Utilization ratio	+0.0057	< 0.001	+0.0060	< 0.001
Collateral factor	−0.0122	< 0.001	−0.0105	< 0.001
<i>Cross-pool linkages</i>				
In-degree dispersion (cross-pool)	+0.0106	< 0.001	+0.0108	< 0.001
Out-degree (wavg.)	−0.0083	< 0.001	−0.0083	< 0.001
<i>Collateral network features</i>				
In-degree (wavg., collateral)	+0.0074	< 0.001	+0.0073	< 0.001
In-flow weight (wavg., collateral)	−0.0169	< 0.001	−0.0167	< 0.001
In-flow weight HHI (wavg., collateral)	+0.0257	< 0.001	+0.0262	< 0.001
Out-degree (wavg., collateral)	+0.0092	< 0.001	+0.0090	< 0.001
Out-flow weight HHI (wavg., collateral)	+0.0094	< 0.001	+0.0093	< 0.001
<i>Collateral financial features</i>				
Self-borrowing rate (wavg., collateral)	−0.0043	< 0.001	−0.0040	< 0.001
Collateral utilization (wavg., collateral)	+0.0303	< 0.001	+0.0300	< 0.001
Health ratio (wavg., collateral)	+0.0153	< 0.001	+0.0151	< 0.001
<i>Asset fixed effect</i>				
WBTC dummy (d_a)	—		+0.0039	0.337
Constant	+0.1037	< 0.001	+0.1018	< 0.001
Train R^2	0.835		0.836	
Test R^2	0.868		0.871	
RMSE	0.0205		0.0203	

Notes: OLS on training set; Newey–West HAC errors ($L = 4$ lags); predictors standardized using training-set moments. Preferred specification: trimmed model without WBTC fixed effect ($\hat{\gamma} = +0.0039$, $p = 0.337$).

Appendix F.2. Solvency Risk

Table 9 summarizes the predictive performance of all specifications considered for the bad debt share $s_{t,j}$. The full OLS benchmark on all 30 predictors achieves a test R^2 of 0.835 but at the cost of using three times as many variables as the preferred specification, confirming that unrestricted OLS retains a large number of redundant predictors without improving out-of-sample accuracy. The LASSO model at $\lambda_{\min}$ selects 18 variables and achieves a test R^2 of 0.822, while the trimmed model without the fixed effect retains 13 variables with a test R^2 of 0.819. The addition of the asset fixed effect for WBTC improves test R^2 substantially to 0.849 and reduces RMSE from 0.0076 to 0.0069, confirming that WBTC pools generate systematically higher bad debt than ETH pools beyond what is explained by the structural covariates. After post-HAC re-trimming at $\alpha = 0.001$, the preferred specification retains 10 variables and is the trimmed model with the asset fixed effect.




Figure 18 Model comparison — Trimmed model without and with asset fixed effect, liquidity risk $\ell_{t,j}$ (ETH and WBTC, $s = 50\%$).

Note: Each row of panels shows the same four diagnostics for the two specifications side by side. Columns 1 and 3: actual against fitted values on the training set (blue) and the held-out test set (coral). Columns 2 and 4: HAC-estimated coefficients with 95 % confidence intervals, ordered by magnitude, and residuals over time for the training (blue) and test (coral) periods separated by the dashed vertical line at the train/test boundary (October 2023).

Table 9 Model comparison — Solvency risk ($b_{t,j}$), ETH and WBTC, $s = 50\%$

Model	Train R^2	Test R^2	RMSE	Vars
Full OLS (all vars)	0.867	0.835	0.0072	30
LASSO $\lambda_{\min}$, no FE	0.836	0.822	0.0075	18
Trimmed, no FE	0.834	0.819	0.0076	13
Trimmed + FE (preferred)	0.838	0.834	0.0072	10

Notes: OLS on training set; Newey–West HAC errors ($L = 4$ lags); test metrics on held-out test set. WBTC fixed effect significant ($\hat{\gamma} = +0.0086$, $p < 0.001$); retained in preferred specification. Re-trimming after fixed effect inclusion reduces active set from 13 to 10 predictors.

Table 10 and Figure 19 report the coefficient estimates for the trimmed model without and with the asset fixed effect. The inclusion of the WBTC fixed effect produces meaningful changes in several coefficients: the constant falls from 0.0115 to 0.0071 as the asset-level intercept shift is absorbed by the dummy, and the PageRank centrality coefficient nearly doubles (from +0.0033 to +0.0069), reflecting that the fixed effect disentangles the asset-level bad debt differential from the network-position effect. Signs are stable across all shared variables, confirming that the fixed effect does not distort the estimated predictive associations but sharpens their identification.

Table 10 Coefficient comparison — Trimmed model without and with asset fixed effect, solvency risk ($b_{t,j}$)

Variable	No FE		With FE (preferred)	
	Coef.	p -value	Coef.	p -value
<i>Own-pool characteristics</i>				
PageRank (scaled)	+0.0033	< 0.001	+0.0069	< 0.001
Utilization ratio	+0.0011	< 0.001	—	
In-weight HHI	+0.0005	< 0.001	—	
<i>Cross-pool linkages</i>				
Deposit concentration (HHI)	+0.0013	< 0.001	+0.0015	< 0.001
Self-borrowing share (wavg.)	−0.0013	< 0.001	−0.0015	< 0.001
Self-borrowing dispersion	−0.0021	< 0.001	—	
Out-degree (wavg.)	−0.0006	0.003	−0.0010	< 0.001
<i>Collateral network features</i>				
In-flow weight HHI (wavg., collateral)	+0.0015	< 0.001	+0.0011	< 0.001
Out-flow weight (wavg., collateral)	−0.0017	< 0.001	−0.0023	< 0.001
<i>Collateral financial features</i>				
Self-borrowing rate (wavg., collateral)	−0.0018	< 0.001	−0.0017	< 0.001
Collateral utilization (wavg., collateral)	+0.0067	< 0.001	+0.0071	< 0.001
Health ratio (wavg., collateral)	+0.0026	< 0.001	+0.0028	< 0.001
Deposit share (wavg., collateral)	−0.0014	< 0.001	—	
<i>Asset fixed effect</i>				
WBTC dummy (d_a)	—		+0.0086	< 0.001
Constant	+0.0115	< 0.001	+0.0071	< 0.001
Train R^2	0.834		0.838	
Test R^2	0.819		0.834	
RMSE	0.0076		0.0072	

Notes: OLS on training set; Newey–West HAC errors ($L = 4$ lags); predictors standardized using training-set moments. Preferred: trimmed model with WBTC fixed effect ($\hat{\gamma} = +0.0086$, $p < 0.001$). Variables marked — excluded after post-HAC re-trimming at $\alpha = 0.001$. Out-degree (wavg.) retained at $p < 0.01$ in no-FE specification under relaxed threshold.

Appendix G. Predictive Performance Across Shock Sizes

Table 11 reports the predictive performance of the trimmed model (no asset fixed effect) for liquidity risk $\ell_{t,j}$ across the full range of shock sizes $s \in \{10\%, 20\%, \dots, 90\%\}$. The table confirms the inverted-U pattern visible in Figure 7: test R^2 rises from near zero at $s = 10\%$ to a peak of 0.868 at $s = 50\%$, then declines as larger shocks push an increasing share of positions into full liquidation, reducing the explanatory power of the pre-shock network and financial characteristics. The number of retained predictors also varies non-monotonically, dropping to as few as 1 at $s = 10\%$ — where the outcome has near-zero variance — and reaching 15 at $s = 90\%$, where a small number of variables captures the residual variation among the minority of positions not yet fully liquidated.




Figure 19 Model comparison — Trimmed model without and with asset fixed effect, solvency risk $b_{t,j}$ (ETH and WBTC, $s = 50\%$).

Each pair of columns shows the same four diagnostics for one specification. Columns 1–2: Trimmed model without fixed effect. Columns 3–4: Trimmed model with WBTC fixed effect. Within each pair: upper-left: actual against fitted values on the training set (blue); upper-right: actual against fitted values on the held-out test set (coral); lower-left: HAC-estimated coefficients with 95 % confidence intervals, ordered by magnitude; lower-right: residuals over time for the training (blue) and test (coral) periods, separated by the dashed vertical line at the train/test boundary (March 2023).

Table 11 Predictive performance across shock sizes
— trimmed model (no FE), liquidity risk ($\ell_{t,j}$), ETH and WBTC

s	N	Train R^2	Test R^2	RMSE	Vars
0.10	2,008	0.079	0.116	0.0003	1
0.20	2,008	0.460	0.463	0.0151	7
0.30	2,008	0.666	0.661	0.0261	11
0.40	2,008	0.816	0.799	0.0228	13
0.50	2,008	0.835	0.868	0.0205	14
0.60	2,008	0.835	0.834	0.0215	14
0.70	2,008	0.868	0.672	0.0247	7
0.80	2,008	0.887	0.722	0.0107	9
0.90	2,008	0.920	0.472	0.0055	15

Notes: Each row: separate model on training set ($N_{\text{train}} = 1,607$), evaluated on held-out test set ($N_{\text{test}} = 401$). LASSO at $\lambda_{\min}$ with HAC trimming at $\alpha = 0.001$ ($L = 4$ lags); no WBTC fixed effect. Preferred specification ($s = 50\%$) in bold. $s = 100\%$ omitted: no predictors survive trimming. Low RMSE at $s \geq 0.80$ reflects near-zero outcome variance, not improved accuracy; test R^2 confirms deterioration.

Appendix H. Robustness to nonlinear specification

We estimate two random forest models for each outcome using the same chronological train/test split as the preferred OLS specification: one on the full candidate set of 30 variables and one restricted to the LASSO-selected variables. The comparison serves two purposes. First, it provides a model-free benchmark for predictive performance: if the linear OLS model outperforms the random forest out of sample, the linear specification is not missing important nonlinearities. Second, permutation importance scores from the full random forest provide an independent ranking of variable importance that does not depend on the linear specification.

Appendix H.1. Liquidity Risk

Figure 20 reports the results for liquidity risk. The trimmed OLS model achieves a test R^2 of 0.868, outperforming both the full random forest (0.773) and the restricted random forest (0.820). The superior out-of-sample performance of the linear model suggests that the relationship between the selected predictors and liquidity risk is well approximated by a linear specification, and that the random forest overfits the training data despite regularization. Permutation importance scores from the full random forest independently confirm the variable ranking implied by the OLS coefficients: collateral utilization of depositors and in-flow concentration emerge as the two dominant predictors, consistent with Figure 9, and variables excluded by LASSO rank uniformly lower. The agreement between the two approaches strengthens confidence that the LASSO-selected linear model identifies the true predictive structure of the data.




Figure 20 OLS vs Random Forest benchmark — liquidity risk $\ell_{t,j}$ (ETH and WBTC, $s = 50\%$, same chronological train/test split).

Note: Left panel: permutation importance scores on the held-out test set from a Random Forest estimated on the full candidate set of 30 variables (500 trees, $\text{max_features} = \sqrt{p}$, $\text{min_samples_leaf} = 5$). Blue bars indicate variables selected by LASSO and retained in the preferred OLS specification; grey bars indicate variables excluded by LASSO. Error bars show one standard deviation across 20 permutation repeats. Right panel: test R^2 on the held-out test set for three specifications — trimmed OLS (preferred), Random Forest on all 30 candidate variables, and Random Forest restricted to the LASSO-selected variables.

Appendix H.2. Solvency Risk

Figure 21 reports the results for solvency risk. The trimmed OLS model again outperforms both random forests, achieving a test R^2 of 0.834 against 0.727 for the full random forest and 0.687 for the restricted random forest. The larger gap between OLS and the restricted random forest — compared to the liquidity case — suggests that for solvency risk the predictive content of the LASSO-selected variables resides in their linear combination rather than in their individual marginal contributions, which the random forest exploits less efficiently.

The permutation importance ranking reveals an important nuance. The top-ranked variable by RF importance — exposure-weighted collateral utilization (importance = 0.117) — is also retained by LASSO, confirming it as the dominant predictor of solvency risk. However, several variables excluded by LASSO — notably own-pool collateral utilization, In-degree, PageRank (raw), Out-degree, and the collateral factor — also rank highly, unlike in the liquidity case where LASSO-excluded variables ranked uniformly low. The likely explanation is multicollinearity: LASSO dropped these variables not because they are uninformative, but because their predictive content overlaps with that of the retained variables, particularly the exposure-weighted collateral utilization and PageRank (scaled). The random forest, which selects splitting variables one at a time, recovers their marginal importance. The OLS still outperforms because the linear model already captures this shared information efficiently through the retained regressors. Notably, the WBTC fixed effect ranks tenth in permutation importance (importance = 0.014), confirming that it captures a genuine asset-level difference in solvency exposure.




Figure 21 OLS vs Random Forest benchmark — solvency risk $b_{t,j}$ (ETH and WBTC, $s = 50\%$, same chronological train/test split).

Note: Left panel: permutation importance scores on the held-out test set from a Random Forest estimated on the full candidate set of 30 variables (500 trees, $\text{max_features} = \sqrt{p}$, $\text{min_samples_leaf} = 5$). Blue bars indicate variables selected by LASSO and retained in the preferred OLS specification; grey bars indicate variables excluded by LASSO. Error bars show one standard deviation across 20 permutation repeats. Right panel: test R^2 on the held-out test set for three specifications — trimmed OLS (preferred), Random Forest on all 30 candidate variables, and Random Forest restricted to the LASSO-selected variables.