



HAL
open science

Contagion in Decentralized Lending Protocols: A Case Study of Compound

Natkamon Tovanich, Myriam Kassoul, Simon Weidenholzer, Julien Prat

► **To cite this version:**

Natkamon Tovanich, Myriam Kassoul, Simon Weidenholzer, Julien Prat. Contagion in Decentralized Lending Protocols: A Case Study of Compound. Proceedings of the 2023 Workshop on Decentralized Finance and Security, ACM, Nov 2023, Copenhagen, Denmark. 10.1145/3605768.3623544 . hal-04221228

HAL Id: hal-04221228

<https://hal.science/hal-04221228>

Submitted on 28 Sep 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Copyright

Contagion in Decentralized Lending Protocols: A Case Study of Compound

Natkamon Tovanich

Myriam Kassoul

CREST, CNRS, École Polytechnique
Institut Polytechnique de Paris
Palaiseau, France

Simon Weidenholzer

Department of Economics
University of Essex
Colchester, United Kingdom

Julien Prat

CREST, CNRS, École Polytechnique
Institut Polytechnique de Paris
Palaiseau, France

ABSTRACT

We study financial contagion in Compound V2, a decentralized lending protocol deployed on the Ethereum blockchain. We explain how to construct the balance sheets of Compound’s liquidity pools and use our methodology to characterize the financial network. Our analysis reveals that most users either borrow stablecoins or engage in liquidity mining. We then study the robustness of Compound through a series of stress tests, identifying the pools that are most likely to set off a cascade of defaults.

CCS CONCEPTS

• **Applied computing** → **Economics**; • **Security and privacy** → *Economics of security and privacy*; • **Computing methodologies** → *Simulation tools*.

KEYWORDS

systemic risk, decentralized finance, financial contagion, financial network, stress test

ACM Reference Format:

Natkamon Tovanich, Myriam Kassoul, Simon Weidenholzer, and Julien Prat. 2023. Contagion in Decentralized Lending Protocols: A Case Study of Compound. In *Proceedings of the 2023 Workshop on Decentralized Finance and Security (DeFi '23)*, November 30, 2023, Copenhagen, Denmark. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3605768.3623544>

1 INTRODUCTION

Smart contracts have enabled the rise of Decentralized Finance (DeFi) protocols that offer financial services without relying on an intermediary such as a bank or brokerage house. While it is widely acknowledged that traditional financial systems are vulnerable to contagion through various channels, including bank runs [7] and default cascades [8], little is known about the contagion risks potentially present in DeFi protocols.

To study this question, we focus on Compound V2 but note that alternative protocols, such as AAVE and MakerDAO, share a comparable architecture, potentially exposing them to similar forces. Compound is a decentralized lending protocol built on the

Ethereum blockchain [22]. The protocol manages multiple liquidity pools, each dedicated to a specific token. Lenders can add liquidity to any pool, while borrowers can withdraw liquidity by providing collateral in the form of deposits in other pools. These operations connect the various liquidity pools through a network of financial liabilities.

Our first contribution lies in proposing a methodology for the description of Compound’s financial network. We do so by characterizing the balance sheet of its liquidity pools and identifying how they are connected by the borrowing and collateral obligations of users. Leveraging the public availability of Ethereum’s transaction history, we reconstruct the balance sheet of each pool at any given point in time and without measurement errors. The resulting financial network sheds light on the key functionalities of Compound. Specifically, it indicates that users predominantly utilize Compound for two types of financial operations: borrowing stablecoins and participating in liquidity mining of Compound’s governance token.

Then, we assess the robustness of the protocol. Inspired by the recent literature on financial contagion (e.g., [9, 19]), we investigate how shocks propagate through the financial network. Our first set of stress tests simulates the aftermath of a pool’s default, identifying the pools that pose the highest level of systemic risk. In a second set of simulations, we characterize which liquidity pools default in response to a drop in the price of Bitcoin and Ether. We find that cascading failure is a distinct possibility, albeit requiring fairly sizeable price shocks. The pools of stablecoins are the most likely to default, whereas the pools of Bitcoins and Ethers are the most likely to set off a domino effect.

Related Literature. A growing body of research investigates DeFi protocols in order to assess their robustness and vulnerabilities. Formal analyses of lending pools can be found in [3, 4, 14]. Other studies simulate crash scenarios to explore how lending protocols respond to market price fluctuations [20, 24]. Additionally, researchers have analyzed the resilience of lending protocols to significant market events, such as the Ethereum Merge [18] and governance attacks [13].

The examination of participants’ behavior highlights severe liquidation risks due to their leverage [17] and risk appetite [23]. [6, 25] show that using debt-financed collateral fosters interconnectivity, whereas [5] explains why rigid haircut rules are likely to cause price-liquidity feedback loops. Empirical studies of liquidations reveal vulnerabilities leading to fire sales [24] and liquidation spirals [29] that could potentially endanger the stability of the DeFi ecosystem [21].

Instances of illiquidity have been documented, especially in newly established platforms [13, 15, 27]. Our paper investigates

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

DeFi '23, November 30, 2023, Copenhagen, Denmark.

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0261-7/23/11...\$15.00

<https://doi.org/10.1145/3605768.3623544>

whether liquidity pools are likely to become illiquid and explores how such an occurrence might propagate across Compound's network. To model these scenarios, we draw upon the extensive research that studies the mechanisms governing the transmission of shocks and distress across financial markets [1, 9–11, 16]. Given the extensive scope of the literature on financial contagion, we direct interested readers to two comprehensive surveys [12, 19]. Our approach, like these studies, uses the balance sheets of financial institutions to capture their connections and derive the corresponding network structure. In this context, [2, 26] are closely related to our research, as they use network analysis to assess the decentralization of DeFi and the interconnectedness of various protocols. By contrast, our paper pioneers the investigation of contagion risks and network effects *within* lending protocols.

2 DESCRIPTION OF COMPOUND

We start by describing how users interact with Compound since their actions determine the structure of the financial network.

2.1 Lending

When a user adds liquidity to a pool by depositing tokens, she receives an equivalent amount of cTokens in return. Essentially, cTokens are tokenized proofs of the deposit that can be redeemed at any time. Users are incentivized to provide liquidity because cTokens are deflationary and tend to increase in value relative to the underlying asset, offering a rate of return on deposits.

For ease of notation, we will omit time indexes for all variables. We use $d_i^u \geq 0$ to denote the amount of tokens held by user u as deposits in pool i and gather all the deposits in the matrix d .¹ The row vector $d^u = (d_1^u, d_2^u, \dots, d_k^u)$, where k is the number of pools managed by Compound, represents user u 's deposits across the different pools. The column vector $d_i = (d_i^1, d_i^2, \dots, d_i^n)$, where n is the number of active users, lists all the deposits in pool i . Consequently, the sum of deposits in pool i can be calculated as $\bar{d}_i = \sum_u d_i^u$.

To compare deposit values across pools, we convert them into a common unit of account. Let p_i represent the price of token i in US Dollars.² By stacking all prices into the vector $p = (p_1, \dots, p_k)$, we can express the value of user u 's deposits as $v(d^u, p) = d^u p^T = \sum_{i=1}^k d_i^u p_i$. Additionally, the vector of deposit values across all users is given by $v(d, p) = (v(d^1, p), \dots, v(d^n, p)) = dp^T$.

2.2 Borrowing

Users have two ways of withdrawing tokens from liquidity pools. As mentioned before, they can redeem their cTokens, which are then burned by the protocol, effectively releasing their deposited tokens. Alternatively, they can borrow tokens by using a portion of their deposits as collateral. When borrows are backed by cTokens from different pools, the loans create a web of liabilities interconnecting the various liquidity pools.

Using a notation similar to that used for deposits, we denote the amount of asset i borrowed by user u as b_i^u and collect these

¹For simplicity, we present deposits in terms of the underlying token. In practice, when you deposit funds into Compound V2, the protocol internally converts the underlying tokens into an equivalent amount of cTokens based on the current exchange rate.

²Compound primarily relies on Chainlink's *Open Price Feed* as its price oracle. The protocol preforms sanity checks by comparing Chainlink's price feeds to the prices quoted by Uniswap V2.

borrow amounts in the matrix b . Before borrowing an asset, users must select which cTokens they wish to use as collateral from the various tokens they have supplied. When a user enters a market, all cTokens they hold in that specific asset class are considered collateral. Let e be the matrix of dimension $n \times k$, where the element e_i^u represents the "enterMarket" option chosen by user u for token i . It takes a value of 1 if the user intends to use this asset class as collateral and 0 otherwise. Consequently, the collateral matrix c is given by $c = e \odot d$ where \odot denotes the Hadamard product.

Each cToken has its own collateral factor, indicating the proportion of the underlying asset value that can be borrowed. These collateral factors are determined and set by the governance of the protocol. In general, tokens with a small market capitalization tend to have a low collateral factor. Formally, let $\kappa = (\kappa_1, \dots, \kappa_k)$ where $\kappa_i \in [0, 1)$ for all i , be a vector representing the collateral factors associated with each pool.³ The maximal collateral value available to user u for borrowing is given by $v(c^u, \kappa p) = c^u (\kappa \odot p)^T = \sum_{i=1}^k c_i^u \kappa_i p_i$. This value represents the borrowing capacity of user u since she faces the credit constraint $v(b^u, p) \leq v(c^u, \kappa p)$.

The financial health of each user is quantified by the ratio $h^u(p) = v(c^u, \kappa p) / v(b^u, p)$. Whenever a user's health ratio falls below 1, which can occur due to various reasons, such as an increase in the price of the borrowed tokens, her collateral assets become eligible for liquidation.

2.3 Borrowing and lending rates

Borrowers pay interest on their borrowed tokens, while lenders receive interest for providing their assets. The interest rates for borrowing and lending are determined algorithmically based on the utilization rates of each pool, i.e., the proportion of the pool's tokens that have been borrowed [22]. The protocol maintains a positive interest rate difference between borrowing and lending to reward stakeholders and create liquidity reserves. Furthermore, Compound introduced a liquidity mining program on June 16, 2020, incentivizing user engagement through the distribution of its governance token (COMP). The details of this distribution mechanism are subject to governance control and may vary over time.⁴

3 COMPOUND'S FINANCIAL NETWORK

The liquidity pools within Compound's financial network are interconnected through loans, as users lock tokens in one pool to borrow from another. This mechanism is similar to repurchase agreements (repos) between banks. Liquidity pools replace financial institutions, with each loan representing a claim from the pool of the borrowed asset towards the pool(s) providing collateral.

Compound's financial network comprises a set $\mathcal{K} = \{1, \dots, k\}$ of pools or nodes. The balance sheets of the pools are constructed as follows:

- On the liabilities side, we allocate the certificates of deposit (cTokens) across three categories. Firstly, we collect the market value of all deposits (D_i) that have not entered any market and are, therefore, external liabilities. Secondly, the interpool liabilities of pool i are the sum of all liabilities towards other pools ($\sum_{j=1}^K L_{ij}$). Specifically, when a user decides to lock

³See: <https://docs.compound.finance/v2/comptroller/#collateral-factor>

⁴See: <https://compound.finance/governance/comp>

their cTokens as collateral, an amount corresponding to the debt becomes a liability of the pool towards the pool from which the tokens have been borrowed. Thirdly, the remaining portion of cTokens that have entered the market but are in excess of the value of the debt is allocated to the pool's buffer (B_i). We separate these cTokens from deposits because they are not external liabilities. Instead, the buffer can be mobilized to secure the debt of users with a deteriorating health ratio.

- On the assets side, we aggregate the market value of all the available tokens in each pool (T_i),⁵ including the reserves set aside by the protocol (R_i) and the value of the interpool assets ($\sum_{j=1}^K L_{ji}$), i.e., the value of all tokens from other pools locked as collateral for loans originating from pool i .

We use the following procedure to identify undercollateralized claims. First, we define $\alpha^u = v(b^u, p)/v(c^u, p)$ as user u 's borrowing-to-collateral ratio.⁶ Assuming proportional allocation of collateral, if the user has non-negative net worth ($\alpha^u \leq 1$), then the effective liability of pool i towards pool j corresponds in value to the nominal debt, implying $l_{ij}^u = \alpha^u \beta_j^u c_i$ where $\beta_j^u \equiv p_j b_j^u / v(b^u, p)$ is the share of u 's borrowing in pool j . However, if the user has a negative net worth ($\alpha^u > 1$), the pool can only recover $l_{ij}^u = \beta_j^u c_i^u$ units of asset i before depleting all the funds set aside in u 's buffer. Hence, the liabilities and buffer of user u read

$$l_{ij}^u = \min\{1, \alpha^u\} \beta_j^u c_i^u \text{ and } B_{ij}^u = (1 - \min\{1, \alpha^u\}) \beta_j^u c_i^u. \quad (1)$$

The matrices \bar{L} and L summarize the nominal and effective interpool liabilities. Nominal values represent the *promised payments* associated with each claim, while effective values take into account users' solvency by adjusting for all undercollateralized claims. Therefore, $\bar{L}_{ij} \in \mathbb{R}^+$ ($L_{ij} \in \mathbb{R}^+$) represents the nominal (effective) liabilities of pool i towards pool j . These values can be calculated by aggregating the nominal and effective liabilities of all users as follows: $\bar{L}_{ij} = \sum_{u=1}^n \alpha^u \beta_j^u c_i^u p_i$ and $L_{ij} = \sum_{u=1}^n l_{ij}^u p_i$. Finally, the buffer of each pool in \mathcal{K} is given by $B_i = \sum_j \sum_u B_{ij}^u p_i$.

With this at hand, we can express the net worth of pool i as

$$V_i = \sum_j L_{ji} + R_i + T_i - \sum_j L_{ij} - B_i - D_i. \quad (2)$$

Figure 1 contains a schematic financial network with only two liquidity pools. The arrows connecting the pools represent the direction in which payments flow.

4 DATA AND DESCRIPTIVE STATISTICS

4.1 Data

The public availability of Ethereum's transaction history allows us to reconstruct the pools' balance sheets at any desired time point. By collecting the daily snapshots of users' positions from Compound V2's subgraph,⁷ we obtain a comprehensive list of all users,

⁵More precisely, we calculate the total market value by multiplying the number of tokens T_i available in pool i by their market price p_i .

⁶We define α as the borrowing-to-collateral ratio, rather than the collateral-to-borrowing ratio, in order to prevent infinite values for pure lenders.

⁷<https://the-graph.com/hosted-service/subgraph/graphprotocol/compound-v2>




Figure 1: Pools' balance sheets and interpool linkages.

denoted as \mathcal{U} , along with their respective lending (d_i^u) and borrowing (b_i^u) balances for each asset, including accrued interests from both lending and borrowing. In order to convert all balances to US Dollars, we rely on the market prices (p) of the tokens provided by Compound's oracle. This results in a sizable dataset containing the daily positions of 422,459 users across 19 pools, spanning from January 1, 2020, to June 30, 2023. Using this information, we construct the liability matrices and balance sheets for each daily snapshot, following the procedure outlined in Section 3.

4.2 Financial Network

Our procedure generates daily snapshots of the balance sheet of each liquidity pool. Figure 2 presents a cross-section of the top 10 pools on September 7, 2021.⁸ Two observations stand out. Firstly, the majority of all deposits are concentrated within five main pools, which can be categorized into two groups: (i) stablecoin-pools (cUSDC, cUSDT, cDAI), and (ii) crypto-pools (cETH, cWBTC2).⁹ Secondly, we notice consistent disparities between the balance sheets of crypto and stablecoin pools. Specifically, stablecoin-pools hold significant interpool-assets, whereas crypto-pools have minimal interpool-assets. This observation suggests that users deposit cryptoassets in Compound primarily to borrow stablecoins.

This intuition is confirmed by Figure 3, which presents the financial network on our reference day (September 7, 2021). The size of the circles in the graph is proportional to the values in US Dollars of the pools' interpool assets. The arrows connecting the pools indicate the direction of payment flows, while the size of the arrows is proportional to the value of the claims. Upon analyzing Figure 3, it becomes evident that the majority of interpool links originate from the crypto-pools and connect to the stablecoin-pools. This observation supports the notion that most borrowers utilize Compound as a protocol to engage in repurchase agreements (repos) to trade their cryptoassets for stablecoins.

⁸We selected this snapshot because it is the day on which Compound reached its highest Total Value Locked (TVL). We show below that the network structure exhibits persistent features, allowing us to extrapolate general insights from this specific day.

⁹WBTCs are wrapped bitcoins, i.e., ERC-20 tokens on the Ethereum blockchain pegged to Bitcoin. The cWBTC2 pool replaced cWBTC in April 2021, following an upgrade of the cToken contract implementation (See: <https://compound.finance/governance/proposals/41>).




Figure 2: Balance sheets of the top 10 pools on Sept. 7, 2021.




Figure 3: Interpool liability network on Sept. 7, 2021. The percentage of self-borrow represents the proportion of interpool assets that consist of tokens from the same pool $(L_{ii} / \sum_{j=1}^K L_{ji})$.

Figure 3 further illustrates an intriguing pattern: users often engage in borrowing from the same pool they have previously lent to. This self-borrowing behavior is captured by the color of the nodes, with darker shades indicating a higher proportion of cTokens being used for self-borrowing. Such a strategy can be financially advantageous, despite the gap separating the borrowing from the lending interest rates, because Compound encourages liquidity provision by distributing its governance token (COMP). Additionally, the benefits of liquidity mining motivate certain users to utilize one stablecoin to borrow another, with cUSDC and cDAI being particularly notable in this regard.

4.3 Centrality

The node-link diagram presented in Figure 3 may not readily apply to generalization across multiple snapshots. To achieve this, we build a scalar measure for each pool that effectively encapsulates

their centrality within the financial network. Centrality measures can be computed from the liabilities matrix \bar{L} .¹⁰ Specifically, there are two measures of eigenvector centrality of node j which can be obtained in the following way [12]:

$$\lambda v_j^L = \sum_{i=1}^k v_i^L \bar{L}_{ij} \text{ and } \lambda v_j^R = \sum_{i=1}^k \bar{L}_{ij} v_i^R, \quad (3)$$

where λ represents the dominant eigenvalue of \bar{L} . The left eigenvector v^L measures *funding centrality*, attributing more centrality to nodes that hold claims on nodes with higher centrality. Conversely, the right eigenvector v^R measures *borrowing centrality*, assigning more centrality to nodes that hold obligations towards nodes with greater centrality.

Figure 4 reports the evolution of the two centrality measures for the dominant pools over time. Although we observe some variations in the relative importance of each pool, an underlying pattern consistent with the earlier snapshot becomes apparent: Borrowing centrality attributes most of the weights to crypto-pools (ETH and either WBTC1 or WBTC2) while funding centrality is concentrated within the stablecoin-pools (DAI, USDC, and USDT).




Figure 4: Centrality of the six main pools over time.

We identified only two exceptions to this rule, occurring in July 2020 and September 2022, where stablecoins gained a higher borrowing centrality than cryptoassets. In June 2020, following the launch of the liquidity mining program, cDAI and cUSDC experienced increased usage to borrow other stablecoins. Most of these loans were repaid within a span of one to two weeks. In September 2022, a significant ETH price drop triggered a flight to stable assets, temporarily making cUSDC the main collateral pool. Towards the end of September, as ETH stabilized, cUSDC's borrowing centrality returned to its usual level.

To summarize, a descriptive analysis of Compound's network reveals that it serves two primary purposes. Firstly, it provides a decentralized protocol enabling users to issue repurchase agreements of cryptoassets against stablecoins. Consequently, the main risk to the stability of the protocol is a decline in the value of cryptoassets, particularly ETH and BTC, as it would undermine the value of the collateral supporting the majority of loans. Secondly, Compound allows its investors to engage in liquidity mining, either by borrowing from the same pool they have lent to or by utilizing one

¹⁰Since we focus on interpool links, we exclude self-borrowing by setting the diagonal entries of the liability matrix to zero, i.e., $L_{ii} = 0$ for all $i \in \mathcal{K}$

stablecoin to borrow another. The latter strategy carries little risk, except in cases where the stablecoin used as collateral undergoes a depegging episode. We now investigate whether these two sources of risks are likely to propagate across the network.

5 CONTAGION

Financial networks are prone to contagion episodes, wherein the default of a financial intermediary triggers a cascade of failures [12, 19, 28]. We focus on the following contagion mechanism, which proceeds in two steps. Initially, a wave of liquidations is set off whenever a pool defaults on its obligations by suspending the convertibility of its cTokens. If the shock is large enough, the liquidation process fails to restore the value of all interpool assets, thereby burdening the balance sheets of connected pools with bad loans. This may lead to other pools becoming insolvent, further amplifying the initial shock and causing a domino effect.

5.1 Liquidations

Compound's liquidation process safeguards lenders by maintaining an adequate level of collateralization. When a borrower's health ratio falls below 1, the liquidation of her position is automatically triggered. Liquidators repay a portion of the debt, known as the "close factor" (denoted by γ), and receive collateral at the current price plus a liquidation discount factor $\lambda \in (0, 1)$. Multiple rounds of liquidation may occur until the borrower's health ratio is restored above 1.

Let's consider a scenario where user u becomes eligible for liquidation after a change in price from p to p' . For tractability, we assume that liquidators follow a proportional rule, seizing all borrowed assets based on their share of the user's total debt.¹¹ Under this assumption, the user's borrowing balance is reduced by $\gamma v(b^u, p')$ and the liquidator acquires collateral worth $(1 + \lambda)\gamma v(b^u, p')$ in return. As a result, all asset holdings of user u after liquidation are diminished by $\psi^u(p')c_j^u$, where $\psi^u(p') \equiv (1 + \lambda)\gamma v(b^u, p')/v(c^u, p')$. The liquidation process is explained in more detail in Appendix A where we describe the algorithm used for its simulation.

The health ratio of user u after t rounds of liquidation, which we denote by $h_t^u(p')$, obeys the following law-of-motion

$$h_{t+1}^u(p') = \frac{1 - \psi^u(p')}{1 - \gamma} h_t^u(p'). \quad (4)$$

This expression yields an intuitive threshold condition:

- (1) $v(c^u, p') > (1 + \lambda)v(b^u, p')$: In this case, the collateral value exceeds the borrowed value multiplied by one plus the liquidation discount. Here, liquidation improves the health factor as the reduction in borrowing exceeds the reduction in collateral ($\gamma > \psi^u$). After potentially multiple rounds of liquidation, the health of the account will eventually be restored ($h^u(p') > 1$).
- (2) $v(c^u, p') < (1 + \lambda)v(b^u, p')$: Here, the collateral value is lower than the borrowed value multiplied by one plus the

liquidation discount. Then liquidation *worsens* the health factor, triggering multiple rounds of liquidation until all of user u 's collateral is liquidated [29]. Consequently, there will be $\left(1 - \frac{v(c^u, p')}{v(b^u, p')(1 + \lambda)}\right) b^u$ bad loans left in the various pools where u chose to invest.

5.2 Cascades

The liquidation process has the following impact on the pools' balance sheets. On the liabilities side, collateral assets are transferred from borrowers to liquidators, effectively converting the interpool liabilities of the pools where the borrowers held collateral into deposits. On the assets side, all pools in which users borrow are replenished in proportion to the amounts borrowed, thereby converting the interpool assets into tokens. However, when the shock is so severe that condition 2 mentioned above holds, a share of the interpool assets is not fully repaid, which induces a fall in the net worth of the lending pool.

The shortfall may result in a negative net worth, indicating that the pool owes more than it owns. We follow the literature on financial contagion in assuming that this signal triggers a run wherein depositors try to withdraw all the available funds [7]. If borrowers can mobilize external cash to repay their loans, they are able to withdraw all the liabilities of the pool. But, since the pool has a negative net worth, it cannot honor its commitment and must suspend the convertibility of its cTokens.¹²

To avoid overestimating the likelihood of cascades, we consider a more conservative scenario where users do not have access to external funds.¹³ In this scenario, users can only redeem a portion of the pool's liabilities, specifically the cTokens held as deposit and excess buffer.¹⁴ Then, the run on cTokens leads to default solely when the combined value of the deposits and excess buffers redeemed by users surpasses the value of the tokens and reserves held by the pool.

We emphasize that our simulations should be seen as offering a conservative estimate of the impact of liquidations. This is because we do not account for the potential further price declines of the liquidated asset. Empirical studies [5, 21, 30] indicate that this is an optimistic scenario, as markets do not always have sufficient liquidity to absorb selling pressure effectively.

5.3 Simulations

We combine the mechanisms discussed in the preceding subsections to design an algorithm, outlined in Appendix A, that simulates the spread of a pool default throughout the entire network. Running this algorithm for each pool identifies those that pose the highest level of systemic risk. We will defer the examination of the factors that might initiate a particular pool's default until the conclusion of this subsection.

Focusing on our reference day (September 7, 2022), we present the predictions of the contagion algorithm in Figure 5. Each row

¹¹In practice, liquidators have the flexibility to choose the token for repayment, the amount to repay within the factor limit, and the collaterals to seize. We impose proportional liquidations because it enables us to derive analytical results. We assess the impact of this assumption in Appendix B.

¹²The suspension of convertibility is explicitly handled on lines 518-521 of the smart contract CToken.sol for pool management.

¹³This assumption also has the advantage of being consistent with the premise that users do not prevent liquidations by recapitalizing their undercollateralized positions.

¹⁴The excess buffer refers to the cTokens that have entered the market but can be redeemed without triggering additional liquidations. In other words, the excess buffer equals the buffer minus the haircut associated with the collateral factor of the loans.




Figure 5: Default cascades on Sept. 7, 2021.

Initial defaulting pool indicated by the row, affected pools indicated by the columns. The pools with a negative net worth are encoded as a dot, while the defaulting pools are encoded as an outer circle. The color indicates the round in which a pool is affected.

tracks the contagion resulting from the default of a specific pool. As expected, contagion risks are concentrated within the main pools, except for cUSDT because its collateral factor is set to zero. Among these pools, the crypto-pools (cETH and cWBTC2) are the primary sources of contagion because they account for the bulk of collateral assets. Meanwhile, the stablecoin-pools exhibit a higher likelihood of default due to their elevated utilization rates.

Additionally, the stablecoin-pools also pose systemic risks to other stablecoin-pools, as users partake in liquidity mining by utilizing one stablecoin to borrow another. These strategies are generally considered riskless, unless one stablecoin depegs and, as captured by our simulations, triggers the default of its borrowing pools.

We repeated these simulations for each day in our sample and compiled the results in Figure 6. The left-hand panel summarizes the number of pools affected by the cascades. Consistent with the snapshot reported in Figure 5, cETH and cWBTC/cWBTC2 usually trigger the largest number of defaults. However, by the end of 2022, cUSDC had emerged as a significant threat because users responded to the significant decline in cryptocurrency prices by diversifying their sources of collateral.¹⁵ The middle panel indicates the depth of the cascades, i.e., the number of propagation rounds triggered by the default of a particular pool. It reveals that most cascades unfolded within one or two rounds before reaching a state where no further defaults occurred. By the third round, all cascades had been resolved. The right-hand panel depicts the percentage of total asset loss caused by the cascades. The right-hand chart

¹⁵See Appendix C for further details on the composition of loans over time.

highlights episodes of significant losses, mostly involving the two crypto-pools (cETH and cWBTC). It also indicates significant losses resulting from a default of the cDAI pool towards the end of 2020, when self-borrowing strategies were particularly widespread.¹⁶ It is worth noting that the protocol’s robustness improved over time, as indicated by the diminishing losses in recent years. This improvement is explained by the accumulation of reserves, resulting in a higher net worth for the pools.

The aforementioned simulations follow the cascades that result from the default of a specific pool without providing an explanation for why the default occurred initially. In practice, the initial default can be triggered by a decline in the market price of the pool’s collateral. To evaluate the likelihood of this scenario, we subjected the prices of the main cryptoassets (ETH and BTC) to negative shocks. We use δ to denote the magnitude of the price shock, so that $p'_{\{ETH, BTC\}} = (1 - \delta)p_{\{ETH, BTC\}}$.

The outcomes of these experiments are depicted in Figure 7. It shows that default cascades are more likely to originate from stablecoin pools. This fragility can be primarily attributed to two factors: high utilization rates and the reliance on cryptoassets as collateral for the majority of stablecoin loans. Additionally, we observe that cUSDT is the most prone to default, which rationalizes the protocol’s decision of curtailing the contagion that may originate from cUSDT by setting its collateral factor to zero. Overall, we find that to endanger the protocol, the shocks have to be fairly consequential, involving a decline of 50% or more in market prices.

6 CONCLUSION

Smart contracts reduce counterparty risks, but, as illustrated by the Terra debacle,¹⁷ they do not eliminate contagion risks resulting from flaws in the economic design of their protocol. Fortunately, the transparency of blockchains offers an opportunity to develop sophisticated supervisory tools. The availability of exhaustive and exact data makes it possible to monitor the sources of systemic risks in real time with a precision that far exceeds what is possible in the traditional financial sector. Our paper, by applying classical methods for the analysis of financial networks, provides an example of such an endeavor. It characterizes how contagion might spread through Compound’s network, identifying the pools that are more likely to set off or propagate a domino effect. Moving forward, we aim to further explore this research avenue by delving deeper into the trove of data accumulated during this study. In particular, we intend to study the bipartite structure of the financial network and interact it with the empirical behavior of users in order to identify those that are more likely to endanger the stability of the protocol.

ACKNOWLEDGMENTS

This project has benefited from the financial support of the academic chair Blockchain@Polytechnique. Simon Weidenholzer acknowledges support from the Economic and Social Research Council [grant number ES/T015357/1]. We thank Riho Marten Pallum for helping us develop the code to extract Compound V2 data from The Graph.

¹⁶See Figure 10 in Appendix C for supporting evidence.

¹⁷See <https://www.coindesk.com/learn/the-fall-of-terra-a-timeline-of-the-meteoritic-rise-and-crash-of-ust-and-luna/>




Figure 6: Daily snapshots of default cascades for the top six pools.




Figure 7: Minimal price shocks triggering default.

REFERENCES

- [1] Daron Acemoglu, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2015. Systemic risk and stability in financial networks. *American Economic Review* 105, 2 (2015), 564–608. <https://doi.org/10.1257/aer.20130456>
- [2] Ziqiao Ao, Gergely Horvath, and Luyao Zhang. 2023. Is decentralized finance actually decentralized? A social network analysis of the Aave protocol on the Ethereum blockchain. arXiv:2206.08401 [econ.GN]
- [3] Massimo Bartoletti, James Chiang, Tommi Junttila, Alberto Lluch Lafuente, Massimiliano Mirelli, and Andrea Vandin. 2022. Formal analysis of lending pools in decentralized finance. In *International Symposium on Leveraging Applications of Formal Methods*. Springer, 335–355. https://doi.org/10.1007/978-3-031-19759-8_21
- [4] Massimo Bartoletti, James Hsin-yu Chiang, and Alberto Lluch Lafuente. 2021. SoK: Lending Pools in Decentralized Finance. In *Financial Cryptography and Data Security: FC 2021 International Workshops*. Springer, 553–578. https://doi.org/10.1007/978-3-662-63958-0_40
- [5] Jonathan Chiu, Emre Ozdenoren, Kathy Yuan, and Shengxing Zhang. 2022. On the Fragility of DeFi Lending. Available at SSRN 4328481 (30 November 2022). <https://ssrn.com/abstract=4328481>
- [6] Michael Darlin, Georgios Palaiokrassas, and Leandros Tassioulas. 2022. Debt-financed collateral and stability risks in the defi ecosystem. In *2022 4th Conference on Blockchain Research & Applications for Innovative Networks and Services (BRAINS)*. IEEE, 5–12. <https://doi.org/10.1109/BRAINS55737.2022.9909090>
- [7] Douglas W Diamond and Philip H Dybvig. 1983. Bank runs, deposit insurance, and liquidity. *Journal of Political Economy* 91, 3 (1983), 401–419. <https://doi.org/10.1086/261155>
- [8] Larry Eisenberg and Thomas H Noe. 2001. Systemic risk in financial systems. *Management Science* 47, 2 (2001), 236–249. <https://doi.org/10.1287/mnsc.47.2.236.9835>
- [9] Matthew Elliott, Benjamin Golub, and Matthew O Jackson. 2014. Financial networks and contagion. *American Economic Review* 104, 10 (2014), 3115–3153. <https://doi.org/10.1257/aer.104.10.3115>
- [10] Prasanna Gai, Andrew Haldane, and Sujit Kapadia. 2011. Complexity, concentration and contagion. *Journal of Monetary Economics* 58, 5 (2011), 453–470. <https://doi.org/10.1016/j.jmoneco.2011.05.005>
- [11] Prasanna Gai and Sujit Kapadia. 2010. Contagion in financial networks. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 466, 2120 (2010), 2401–2423. <https://doi.org/10.1098/rspa.2009.0410>
- [12] Paul Glasserman and H Peyton Young. 2016. Contagion in financial networks. *Journal of Economic Literature* 54, 3 (2016), 779–831. <https://doi.org/10.1257/jel.20151228>
- [13] Lewis Gudgeon, Daniel Perez, Dominik Harz, Benjamin Livshits, and Arthur Gervais. 2020. The decentralized financial crisis. In *2020 Crypto Valley Conference on Blockchain Technology (CVCBT)*. IEEE, 1–15. <https://doi.org/10.1109/CVCBT50464.2020.00005>
- [14] Lewis Gudgeon, Sam Werner, Daniel Perez, and William J Knottenbelt. 2020. Defi protocols for loanable funds: Interest rates, liquidity and market efficiency. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*. 92–112. <https://doi.org/10.1145/3419614.3423254>
- [15] Matthias Hafner, Romain de Luze, Nicolas Greber, Juan Beccuti, Benedetto Biondi, Gidon Katten, Michelangelo Riccobene, and Alberto Arrigoni. 2023. DeFi Lending Platform Liquidity Risk: The Example of Folks Finance. *The Journal of The British Blockchain Association* (2023). [https://doi.org/10.31585/jbba-6-1-\(5\)2023](https://doi.org/10.31585/jbba-6-1-(5)2023)
- [16] Andrew G Haldane and Robert M May. 2011. Systemic risk in banking ecosystems. *Nature* 469, 7330 (2011), 351–355. <https://doi.org/10.1038/nature09659>
- [17] Lioba Heimbach and Wenqian Huang. 2023. DeFi Leverage. Available at SSRN 4459384 (24 May 2023). <https://ssrn.com/abstract=4459384>
- [18] Lioba Heimbach, Eric Schertenleib, and Roger Wattenhofer. 2023. DeFi Lending During The Merge. arXiv:2303.08748 [q-fin.GN]
- [19] Matthew O Jackson and Agathe Pernoud. 2021. Systemic risk in financial networks: A survey. *Annual Review of Economics* 13 (2021), 171–202. <https://doi.org/10.1146/annurev-economics-083120-111540>
- [20] Hsien-Tang Kao, Tarun Chitra, Rei Chiang, and John Morrow. 2020. An analysis of the market risk to participants in the compound protocol. In *Third International Symposium on Foundations and Applications of Blockchains*. https://scfab.github.io/2020/FAB2020_p5.pdf
- [21] Alfred Lehar and Christine A Parlour. 2022. *Systemic fragility in decentralized markets*. BIS Working Papers 1062. Bank for International Settlements. <https://www.bis.org/publ/work1062.pdf>
- [22] Robert Leshner and Geoffrey Hayes. 2019. *Compound: The money market protocol*. Technical Report. <https://compound.finance/documents/Compound.Whitepaper.pdf>
- [23] Daniel Perez, Sam M. Werner, Jiahua Xu, and Benjamin Livshits. 2021. Liquidations: DeFi on a Knife-Edge. In *Financial Cryptography and Data Security*, Nikita Borisov and Claudia Diaz (Eds.). Springer, 457–476. https://doi.org/10.1007/978-3-662-64331-0_24
- [24] Kaihua Qin, Liyi Zhou, Pablo Gamito, Philipp Jovanovic, and Arthur Gervais. 2021. An empirical study of defi liquidations: Incentives, risks, and instabilities. In *Proceedings of the 21st ACM Internet Measurement Conference*. 336–350. <https://doi.org/10.1145/3487552.3487811>
- [25] Kanis Saengchote. 2023. Decentralized lending and its users: Insights from Compound. *Journal of International Financial Markets, Institutions and Money* 87 (2023), 101807. <https://doi.org/10.1016/j.intfin.2023.101807>
- [26] Kanis Saengchote and Carlos Castro-Iragorri. 2023. Network Topology in Decentralized Finance. Available at SSRN 4469783 (5 June 2023). <https://ssrn.com/abstract=4469783>
- [27] Xiaotong Sun, Charalampos Stasinakis, and Georgios Sermpinis. 2023. Liquidity Risks in Lending Protocols: Evidence from Aave Protocol. arXiv:2206.11973 [q-fin.RM]

- [28] Christian Upper. 2011. Simulation methods to assess the danger of contagion in interbank markets. *Journal of Financial Stability* 7, 3 (2011), 111–125. <https://doi.org/10.1016/j.jfs.2010.12.001>
- [29] Jakub Warmuz, Amit Chaudhary, and Daniele Pinna. 2023. Toxic Liquidation Spirals. arXiv:2212.07306 [econ.GN]
- [30] Aviv Yaish, Maya Dotan, Kaihua Qin, Aviv Zohar, and Arthur Gervais. 2023. Suboptimality in DeFi. Cryptology ePrint Archive, Paper 2023/892. <https://eprint.iacr.org/2023/892>

A DESCRIPTION OF ALGORITHMS

This appendix outlines the logic of the algorithms used to simulate default cascades. We first present in Algorithm 1 how liquidations are simulated. All borrowers with a health factor h^u below one are liquidated up to the maximum amount determined by the close factor (γ). The liquidator repays the borrowed assets and seizes the collateral assets proportionally.¹⁸ Under proportional liquidations, the liquidator acquires collateral worth $(1 + \lambda)\zeta_j^u(p)\gamma v(b^u, p')$ for each asset j used as collateral by user u . Here, $\zeta_j^u(p) \equiv p_j c_j^u / v(c^u, p)$ represents the value of assets j relative to the user's total collateral value.

For each liquidation round, the algorithm updates the new borrow (b^u) and collateral (c^u) positions of the borrower. Moreover, the deposits of the collateral pools (D_j) and tokens of the borrowing pools (D_i) are also increased as a result of liquidation. The algorithm liquidates the borrower's position until $h^u > 1$. We set the gas costs of liquidations equal to the median transaction fee on that snapshot date (i.e., $\text{median}(\text{gasFees}) \times \text{gasUsed}$, where $\text{gasUsed} = 500,000$).

Algorithm 1: Simulation of liquidation process

```

1 for  $u \in \mathcal{U}$  do
2   while  $h^u < 1$  do
3     // 1. Calculate % borrow and collateral.
4      $\beta^u \leftarrow pb^u / v(b^u, p)$ ;
5      $\zeta^u \leftarrow pc^u / v(c^u, p)$ ;
6     // 2. Calculate repaid and seized amounts.
7      $\text{repay}^u \leftarrow \min\{v(b^u, p) \cdot \gamma, v(c^u, p) / (1 + \lambda)\}$ ;
8      $\text{seize}^u \leftarrow (1 + \lambda) * \text{repay}^u$ ;
9     // 3. Stop liquidating if profits  $\leq \text{txFees}$ .
10    if  $\text{seize}^u - \text{repay}^u \leq \text{txFees}$  then
11      break;
12    end
13    // 4. Repay the borrowed assets.
14     $pb^u \leftarrow pb^u - \text{repay}^u * \beta^u$ ;
15     $T \leftarrow T + \text{repay}^u * \beta^u$ ;
16    // 5. Seize the collateral assets.
17     $pc^u \leftarrow pc^u - \text{seize}^u * \zeta^u$ ;
18     $D \leftarrow D + \text{seize}^u * \zeta^u$ ;
19    // 6. Update the health of user  $u$ 
20     $h^u \leftarrow v(c^u, \kappa p) / v(b^u, p)$ ;
21  end
22 end

```

¹⁸We assess the impact of this assumption in Appendix B.

Algorithm 2 explains how we simulate the propagation of a pool's default across the network. Since the cTokens of the defaulting pool are not anymore liquid, we set their collateral factor κ_i to zero. This captures the fact that the cTokens cannot be used to repay the loans. The algorithm calls two procedures. The first one, called *simulate_liquidation*, follows the logic outlined in Algorithm 1. It returns the updated collateral (c) and borrowing (b) matrices defined in Section 2, along with the deposits (D) and token holdings (T) of each pool. The second function, called *update_balance_sheet*, follows the methodology presented in Section 3 to compute the interpool liabilities (L), reserves (R), buffer (B), excess buffer (X), and net worth (V) of each pool. As explained in footnote 14, the excess buffer collects all the cTokens in the buffer that can be redeemed without triggering additional liquidations. The simulation iterates until the cascade reaches a state where no further pools default.

Algorithm 2: Simulation of cascading default of the pools

```

Input:  $\text{init\_pool}$ : initial defaulting pool
Output:  $D_r$ : list of defaulting pools at depth  $r$ .
1  $D_0 \leftarrow \{\text{init\_pool}\}$ ; // list of defaulting pools
2  $r \leftarrow 0$ ; // depth of liquidation cascade
// Continue liquidation if any new pool
// defaulted in the previous round
3 while  $r > 0$  and  $|D_r \cap D_{r-1}| > 0$  do
4    $\kappa_{\{d \in D_r\}} \leftarrow 0$ ; // set  $\kappa = 0$  for defaulting pools
5    $r \leftarrow r + 1$ ; // increment the cascade depth
// Liquidate borrowers and update the balance
// sheet
6    $(c, b, D, T) \leftarrow \text{simulate\_liquidation}(c, b, D, T, p, \kappa)$ ;
7    $(L, R, B, X, V) \leftarrow \text{update\_balance\_sheet}(c, b)$ ;
// Add default pools after liquidation
8    $D_r \leftarrow D_{r-1} \cup \{k \in \mathcal{K} \mid V_k < 0 \text{ and } D_k + X_k > T_k + R_k\}$ ;
9 end

```

B ROBUSTNESS CHECK ON SIMULATIONS OF LIQUIDATIONS

In this appendix, we explore an alternative scenario where the liquidator's approach to asset seizure differs from the one used in the main analysis. Instead of seizing assets in proportion to their share of the user's total debt, we employ a *sequential rule* in which the liquidator prioritizes the assets with the highest share of the user's total debt.

To assess the impact of this alternative liquidation strategy, we rely on the Jaccard similarity index to compare the lists of defaulting pools between our benchmark simulations and the simulations based on sequential liquidations. The defaulting pools predicted by the two simulations agree (i.e., Jaccard index = 1) 96.72% of the time. Figure 8 illustrates the evolution over time of the Jaccard index for the top six pools. It indicates that, in general, the outcomes of the two simulations are either identical or very similar.

Despite the stability of our results, we do observe a few short-lived periods of divergence. Particularly, on March 12 and 13, 2022, the Jaccard index for the cDAI pool suddenly dropped to a very low




Figure 8: Jaccard similarity index of defaulting pools between benchmark and sequential simulations of liquidations.

value of 0.1429. To better understand this divergence, we report in Figure 9 a comparison of the default cascades predicted by the two simulation methods on March 13, 2022. It shows that a default of the cDAI pool did not trigger a domino effect in the benchmark simulation. However, in the sequential simulation, it triggered a default of the cUSDC pool, which then propagated to other pools heavily reliant on cUSDC, such as cLINK, cTUSD, cUSDT, and cYFI.




Figure 9: Comparison of cascades triggered by a default of the cDAI pool on March 13, 2022.

This difference in outcomes can be attributed to the fact that liquidators now prioritize USDCs. This concentration of liquidations on a specific asset can exacerbate the fragility of the network, particularly when the utilization rate and borrowing centrality of this asset are high. In summary, our findings are usually robust across the two specifications of the liquidation process. However, there are specific periods where the selection rule followed by the liquidators has an impact on the risk of contagion within the network. An interesting avenue for future research would therefore consist in studying the behavior of liquidators and devising a liquidation algorithm that reflects their actions.

C COMPOSITION OF LOANS

This appendix provides additional information regarding the composition of loans. We categorize them into classes based on the nature of the asset pairs, distinguishing between cryptoassets and stablecoins. Additionally, we differentiate between self-borrowing and interpool loans. Figure 10 confirms that self-borrowing is primarily motivated by liquidity mining since pairs of self-borrowed stablecoins gained prominence shortly after the introduction of Compound’s liquidity mining program on June 16, 2020.

Figure 10 further corroborates the findings presented in Section 4.3, according to which a majority of interpool loans use cryptoassets as collateral for borrowing stablecoins. During the

bear market from late 2021 to mid-2022, the weight of cryptoasset-stablecoin pairs significantly increased, which aligns with our simulation results in Figure 6 where we observe a heightened contagion risk emanating from cETH and cWBTC over this period.




Figure 10: Distribution of nominal liabilities over time.

Subsequently, after the bear market phase, we notice a rise in the proportion of loans associated with stablecoin-stablecoin pairs, as well as an increase in self-borrowing of stablecoins. This trend amplified the borrowing centrality and, consequently, the risk of contagion associated with the cUSDC pool, as indicated in Figure 6.