

The Decentralized Financial Crisis: Attacking DeFi

Lewis Gudgeon
Imperial College London

Daniel Perez
Imperial College London

Dominik Harz
Imperial College London

Arthur Gervais
Imperial College London

Benjamin Livshits
Imperial College London

Abstract

The Global Financial Crisis of 2008, caused by excessive financial risk, inspired Nakamoto to create Bitcoin. Now, more than ten years later, Decentralized Finance (DeFi), a peer-to-peer financial paradigm which leverages blockchain-based smart contracts to ensure its integrity and security, contains over 1bn USD of capital as of February 2020. Yet as this ecosystem develops, with protocols intertwining and the complexity of financial products increasing, it is at risk of the very sort of financial meltdown it is supposed to be preventing.

In this paper we explore how design weaknesses in DeFi protocols could lead to a DeFi crisis. First, overcollateralized DeFi protocols are vulnerable to exogenous price shocks. We quantify the robustness of DeFi lending protocols in the presence of significant falls in the value of the assets these protocols are based on, showing for a range of parameters the speed at which a DeFi protocol would become undercollateralized. Second, we present a governance attack on Maker—the largest DeFi protocol by market share—that allows an attacker to steal all 0.5bn USD worth of collateral. Moreover, we present a novel strategy that would allow an attacker to steal the Maker collateral within just *two transactions* and without the need to lock any tokens. This paper shows that with the composition of collateralized debt in these DeFi protocols, the failure of one protocol may lead to financial contagion, resulting in losses ranging from 145m USD to in excess of 246m USD.

1 Introduction

Blockchain technology emerged as a response to the Financial Crisis of 2007–8 [52]¹. The perception that banks had ‘misbehaved’ resulted in a deterioration of trust in the traditional financial sector [26]. The causes

of the crisis were several, but arguably chief among them was a lack of transparency regarding the amount of risk major banks were accumulating. When Lehman Brothers filed for bankruptcy, it had debts of 613bn USD, bond debt of 155bn USD and assets of 639bn USD [10]. Central to its bankruptcy was its exposure to subprime (i.e. bad quality) mortgages. This exposure was compounded by the fact that the bank had a leverage ratio² of 30.7x in 2007 [12].

From their inception, blockchain-based cryptocurrencies sought to provide a remedy: facilitating financial transactions without reliance on trusted intermediaries, shifting the power, and therefore, the ability to cause crisis through the construction of opaque and complex financial instruments, away from banks and financial institutions [52]. Ten years later, a complex financial architecture—the architecture of Decentralized Finance (DeFi)—is gradually emerging on top of existing blockchain platforms. Components in this architecture include those that pertain to lending, decentralized exchange of assets, and markets for derivatives (cf. Table 9) [25, 31, 33, 36, 66, 69].

At present, DeFi protocols assume that participating agents have weak identities, since the underlying blockchain operates with weak identities. Consequently DeFi architectures for lending require agents to post security deposits to fully compensate counter-parties for the disappearance of the agent. We will assume that agents are economically rational such that if the agent faces a choice between repayment of a debt or loss of collateral, the agent will choose the least cost option. These security deposits serve to guard against (i) ‘misbehavior’ of agents, where the action that would maximize individual utility does not maximize social welfare, and (ii) external events, such as large exogenous drops in the value of a particular cryptocurrency [34]. Of all DeFi protocols, those with the most locked capital are for lend-

¹The first Bitcoin block famously outlines: “The Times 03/Jan/2009, Chancellor on brink of second bailout for banks”.

²Defined as $\frac{\text{Total assets}}{\text{Equity}}$; the total assets were more than 30 times larger than what shareholders owned, indicating substantial debt.

ing: the biggest, Maker, has (as of January 2020), 57% of all capital locked in DeFi, corresponding to 473.2m USD [61].

Governance is another crucial facet of DeFi protocols and we observe differing degrees of governance decentralization. For example, Maker uses its own token MKR to allow holders to vote on a contract that implements the governance rules. In contrast, Compound, the third largest protocol by market share is centrally governed and a single account can shut down the system in case of a failure [64]. Moreover, as in traditional finance, these protocols do not exist in isolation. Assets that are created in Maker, for example, can be used as collateral in other protocols such as Compound, dY/dX, or in liquidity pools on Uniswap. Indeed, the composability of DeFi — the ability to build a complex, multi component financial system on top of crypto-assets — is a defining property of open finance [62]. However, if the underlying collateral assets fail, all connected protocols will be affected as well: there is the possibility of financial contagion.

This paper. This paper is the first to quantitatively assess the *financial security* of DeFi protocols to the risks posed by a sharp downturn in cryptocurrency prices—including the resulting financial contagion—as well as attacks on their governance architectures.

The first security risk arises when, in the context of lending, security deposits become smaller than the issued debt. Assuming rational economic agents, in such an under-collateralization event, the borrower would default on their debts, since the amount they have borrowed has become worth more than the amount they escrowed. In this paper, we formally model these security constraints for DeFi lending platforms, and simulate their behavior. We find that for plausible parameter ranges, a DeFi lending system could find itself undercollateralized. To the extent that other DeFi protocols allow agents to lend or trade the undercollateralized asset, financial contagion would result.

The second security risk relates to governance concerns, and in particular the possibility of an agent seizing control of a DeFi lending platform and stealing the funds. We introduce two governance attack strategies on Maker that would enable an attacker to steal all funds in the protocol (0.5bn USD). The first attack strategy inspired by [49], covertly executed, is feasible within two blockchain blocks and requires the attacker to lock c. 27.5m USD of collateral. The second attack strategy is a novel contribution and allows an adversary to amass the Maker collateral within two transactions, with an upfront collateral requirement of a few US dollars to pay for gas fees.

Contributions.

- **Formal modeling of DeFi protocols.** We provide definitions for economically-resilient DeFi protocols, introducing overcollateralization, liquidity, and counter-party risk as formal constraints.
- **Stress-testing of DeFi.** We develop a methodology to quantitatively stress-test a DeFi protocol with respect to these security constraints, inspired by risk assessments performed by central banks in traditional financial systems. We simulate a price crash event with our stress-test methodology to a generic DeFi lending protocol that closely resembles the largest DeFi protocols to-date, by volume: Maker, Synthetix, and Compound. We find, for plausible parameter ranges, the manifestation of undercollateralized DeFi protocols.
- **Maker attack.** With specific focus on the largest DeFi project by market share, Maker, we show that it is feasible to successfully steal the funds locked in the protocol covertly and within two blocks or within two transactions. By exploiting recent flash loan pool contracts, we show how an attacker with no capital (besides gas fees), under the assumption of sufficient market liquidity, would be able to execute such an attack.
- **DeFi contagion.** We present two transmission mechanisms for DeFi contagion, which would serve to propagate the effects of under-collateralization or the Maker attack into other protocols.

We have shared our findings with the MakerDAO team in February 2020 and they generally appear to agree with the possibilities described in this paper.

Paper structure. First, Section 2 provides an overview of DeFi. Focusing on DeFi lending, Section 3 presents a formal system model and the constraints that DeFi lending protocols operate under. Section 4 then provides the simulation results of such a lending protocol, and assesses the potential damage due to financial contagion. Turning specifically to Maker, Section 5 considers the feasibility of an attack on the governance mechanism of the largest DeFi protocol. Section 6 considers transmission mechanisms between composable protocols in a DeFi crisis. Section 7 presents related work and Section 8 concludes.

2 Overview of Decentralized Finance

This section provides a working definition of DeFi, before considering three core properties of a DeFi architecture: (i) the types of identities the participating agents have, (ii) the types of agents in the protocol and (iii) the existence of economic cycles.

DeFi is an emergent field, with 1.06bn USD of total value locked into DeFi protocols as of 12 February 2020 [61]. Appendix A.1 Table 9 presents a categorization of DeFi protocols, providing the three largest by locked USD in each case [61]. We observe that Maker dominates the DeFi projects with over 0.5bn locked USD. DeFi protocols mostly emerge for uses such as lending, decentralized exchange and derivatives. We define Decentralized Finance (DeFi) as follows.

Definition 2.1. Decentralized Finance: a peer-to-peer financial system which leverages distributed ledger based smart contracts to ensure its integrity and security.

2.1 Blockchain Model

A DeFi protocol π operates on top of a layer-one blockchain, which provides standard ledger functionality [7, 8, 22, 57]. We assume that the underlying blockchain is able to provide finality [14, 50], construed as a guarantee that once committed to the blockchain a transaction cannot be modified or reversed. We further assume that given n participants in a consensus protocol, the maximum number of Byzantine faults f is such that $f < \frac{n}{3}$. Finally, we assume that the blockchain offers sufficient protection from selfish mining attacks and that fewer than $\frac{1}{3}$ of the miners are malicious [29].

2.2 Identity

We define identity notions as follows³.

Definition 2.2 (Weak identity). Where the mapping between an agent and an online identity is not one-to-one, and in particular may be many-to-one and change through time.

Definition 2.3 (Strong identity). Where the mapping between an agent and an online identity is one-to-one and does not change through time.

In traditional finance, agents have strong identities. The declaration of bankruptcy does reputational damage, affecting an agent’s ability to access credit again in the future as well as affect the interest rates they pay. However, in DeFi as it currently stands, weak-identities prevail⁴. As a result, no such costs are attached to the

³While the authors are unaware of these definitions formally stated elsewhere, credit is due to Rainer Böhme for providing a similar set of definitions [11].

⁴As DeFi develops, it is foreseeable that agents could have strong-identities. This would have significant ramifications, including potentially removing the requirement for overcollateralization, as agents would have reputational incentives to behave in desirable ways. While reputation systems for crypto-economic protocols have been designed (e.g. [34]), strong identities could permit such systems to allow under-collateralization. In the present work, we consider DeFi protocols as they currently are, featuring agents with only weak-identities.

act of defaulting: an agent can simply leave a protocol and rejoin with a new identity⁵.

2.3 Agent Types

We adopt the BAR model of agents [4] with the extension of private valuations [34] resulting in three agent types:

- **Honest:** An honest agent type will act in the interest of the protocol independent of their private motivations or external payments.
- **Rational:** A rational agent type will choose a course of action based on maximizing their expected utility. However, rational agents generally follow the rules of the protocols.
- **Adversarial:** An adversarial agent type will maximize their payoff in any way possible, including by deviating from the protocol rules as well as adopting their behavior as a response to protocol changes and other agents behaviors.

In DeFi, the combination of weak identities and rational agents (i.e. expected utility maximizers) requires that the borrower must give the lender something of at least equal value—a security deposit A —as the amount they wish to borrow B .

2.4 DeFi Composability

DeFi protocols do not exist in isolation. Their open nature allows developers to create new protocols by composing existing protocols together. Some compare this approach with “Money Lego” [62]. As such, assets created through lending in one protocol can be reused as collateral in other protocols in any kind of fashion. This creates a complex and intertwined system of assets and debt obligations that is hard to understand. Moreover, a failure of a protocol that serves as backing asset to other protocols has a cascading effect on others. Indeed, a hallmark of financial crisis is that such events do not take place in isolation, but rather financial contagion takes place, where a shock affecting a few institutions spreads by contagion to the the rest of the financial sector, before affecting the larger economy [5].

3 DeFi Lending Protocols

This section provides a formal system model for a DeFi lending system and characterizes system constraints.

⁵Identities are also not compositional: if an agent defaults in one protocol, the agent can re-use the same identity in another protocol without suffering any reputation damage. In traditional finance, if an agent is bankrupt, it generally may affect any future interaction with institutions that are required to follow KYC and AML legislation.

3.1 DeFi Lending Protocol Model

Our primary assumption is the existence of a set of protocols for overcollateralized borrowing. Overcollateralized borrowing allows an agent to provide an asset A as collateral to receive or create another asset B, of lower value, in return. The asset B, typically together with the payment of a fee, can be returned and the agent redeems its collateral in return. However, borrowed asset B typically has different properties to asset A: for example, an agent might provide a highly volatile asset A and receive a price-stable asset B in return. Furthermore, a third asset C can serve as a governance mechanism. Holders of asset C are able to influence the rules of the DeFi lending protocol. In absence of a governance asset, DeFi lending protocols typically replace this function with a central privileged operator introducing counter-party risk.

At the agent level, a DeFi lending protocol π permits agents k to escrow units of cryptocurrency i , c_i and borrow (or issue) units of another cryptocurrency d against that value. Both the escrowed and the borrowed cryptocurrency are quoted with reference to a third source of value, e.g. USD. For example, the price of the collateral asset(s) is given by the pair c_i/USD , e.g. 150 USD per unit of c_i , and the price of the borrowed cryptocurrency is given by d/USD . At the system level, a DeFi protocol is the aggregation of these individual acts of borrowing by agents, such that the system collateral of type i is given by $C_i = \sum_{k=1}^K c_{i,k}$ for K agents. We formally define an economically secure DeFi lending protocol as follows:

Definition 3.1 (Economically Secure DeFi lending protocol). Assuming rational agents, a DeFi lending protocol π is economically secure if it ensures that $\forall t$, with reference to a basis of value (e.g. USD), the total value of the system debt D at time t is smaller than the total value of all backing collateral types I ($\sum_{i=1}^I C_i$) at time t .

Given the assumption of rational agents with weak identities, this definition connects to the notion that if the agent is faced with a choice between keeping their borrowed funds d or retrieving their escrowed collateral c , since $d > c$ the agent will renege on their commitment to repay the debt. Whether a DeFi protocol satisfies this definition depends on market conditions, and in particular on the relative values of C_i and D with respect to a common price. In addition:

1. **Volatility.** Since the value of the collateral asset is not fixed in terms of the USD basis of value, the value may fall sharply against the USD.
2. **Liquidity.** In an illiquid market, liquidating the collateral asset may only be possible with a significant *haircut*, where the collateral is sold at a discount. These considerations motivate overcollateralization, which we define as follows.

Definition 3.2 (Over-collateralization). When escrowed collateral c_i has a greater value with respect to a basis of value than the issued loan d .

Denoting the overcollateralization factor as λ and the price and quantity of an asset as $P()$ and $Q()$ respectively, the margin M of overcollateralization at time t at the system level is as follows.

$$M_t = (1 + \lambda) \sum_{k=1}^K \sum_{i=1}^I P_{c_{i,k,t}} Q_{c_{i,k,t}} - \sum_{k=1}^K d_{k,t} \quad (1)$$

Clearly, $M_t \geq 0 \iff \sum_{k=1}^K d_{k,t} \leq (1 + \lambda) \sum_{k=1}^K \sum_{i=1}^I P_{c_{i,k,t}} Q_{c_{i,k,t}}$. Should $M < 0$, then the margin of overcollateralization is negative and therefore the system as a whole is undercollateralized. However, a protocol may have an additional pool of reserve liquidity available, enabling it to act as a lender of last resort. For example, one such pool of collateral could be constituted by governance tokens Π for the protocol itself⁶. In a DeFi protocol, participants can have voting power in proportion to the number of governance tokens they hold. The total value of this pool of collateral is given by $P(\Pi)Q(\Pi)$, and thus adding this into the margin of overcollateralization for the system yields:

$$M_t = (1 + \lambda) \sum_{k=1}^K \sum_{i=1}^I P_{c_{i,k,t}} Q_{c_{i,k,t}} + P_{\Pi,t} Q_{\Pi,t} - \sum_{k=1}^K d_{k,t} \quad (2)$$

Typically, the reserve asset Π would be considered the collateral type with the least risk, similar to a *senior tranche* in traditional finance: the reserve asset is only touched once the riskier tranches have lost their value [30]. A protocol designer faces a trade-off. If the parameter λ is too low, volatile and illiquid markets may mean that the protocol becomes undercollateralized. However, if it is too high, then there is significant capital market inefficiency, with more capital than necessary in escrow, leading to opportunity costs of capital.

3.2 Overcollateralization Constraint

Whether the debt d will be redeemable for $\$d$ depends on the value of the reserve of the collateral asset held. At the system level, the necessary condition for overcollateralization of D is as follows.

$$(1 + \lambda) \sum_{k=1}^K \sum_{i=1}^I P_{c_{i,k,t}} Q_{c_{i,k,t}} + P_{\Pi,t} Q_{\Pi,t} \geq \sum_{k=1}^K d_{k,t} \quad (3)$$

In the event that $(1 + \lambda) \sum_{k=1}^K \sum_{i=1}^I P_{c_{i,k,t}} Q_{c_{i,k,t}} < \sum_{k=1}^K d_{k,t}$, the reserve asset Π of a protocol is used as a “lender of last resort” to buy the collateral value. If equation 3 does not

⁶MKR tokens in the case of Maker.

hold, it means that even liquidation of all of the primary collateral asset and reserve asset would be insufficient to cover the total system debt. This would constitute a catastrophic system failure, since given rational agents with weak identities, the borrowed funds would become worthless as they would no longer be redeemable. Note that in this formulation, P_{c_i} and P_{Π} are expressed as functions of each other, accounting for the possibility of correlation in the price of the collateral asset and the price of the reserve asset ⁷. If collateral and reserve assets are positively correlated, this would serve to limit the ability of the reserve asset to recapitalize the system.

3.3 The Liquidity Constraint

Following [54], we define market liquidity as follows.

Definition 3.3 (Market liquidity). A measure of the extent to which a market can facilitate the trade of an asset at short notice, low cost and with little impact on its price.

The liquidity available in a market implies a security constraint: in expectations, over a certain time horizon, DeFi marketplaces can offer enough liquidity that in the event of a sustained period of negative price shocks, a protocol will be able to liquidate its collateral quickly enough to cover its outstanding debt liabilities.

For a time interval $[0, T]$ this can be expressed as:

$$\int_0^T \mathbb{E}[\Omega] d\Omega < \mathbb{E}[\Omega_{max}] \quad (4)$$

where Ω denotes the total notional traded value, i.e. the (average) price multiplied by the quantity for each trade. For a given trade ω of size q , $\omega = \bar{p}q$; aggregating these trades for a total number of trades J provides $\Omega = \sum_{j \in J} \omega_j$. Ω_{max} denotes the maximum notional value that could be sold off during a period of distress in the financial markets.

3.4 The Counterparty Risk Constraint

In practice, DeFi lending protocols are not fully decentralized on account of, for instance, the possibility of oracle attacks (which could cause a flash-crash), as well as privileged access to the smart contracts. Therefore it is necessary to either assume the ‘‘operator’’ of the protocol is honest, or that the operator only offers the services of the protocol provided they are profitable for them. We formally model this counterparty risk by assuming that its existence in a given protocol creates a risk premium, Ψ , such that for an agent deciding between earning a

⁷There is evidence that crypto-assets display high intra-class correlation, limiting the advantage of diversification [41].

return in a DeFi lending protocol vs elsewhere, the expected return in the DeFi protocol (r_D), once adjusted for the risk premium (Ψ), must be higher than an outside return r_f . Formally, we have participation constraint:

$$r_D - \Psi > r_f \quad (5)$$

There exists an inherent trade-off in counterparty risk. On the one hand, governance mechanisms implemented through voting allow for a certain degree of decentralization whereby multiple protocol participants can influence the future direction of a protocol. Depending on the distribution of tokens, this may reduce the risk of one party becoming malicious. However, it also opens the door to attacks on the voting system, as we introduce in Section 5. On the other hand, a single ‘‘benevolent dictator’’ who controls the governance mechanism can prevent the attacks introduced in Section 5. Yet this requires trusting that this central entity does not lose or expose its private keys controlling access to the smart contracts governing the protocol and that this central party cannot be bribed to behave maliciously.

4 Stress-Testing DeFi Lending

This section considers the financial security of a generic DeFi lending protocol, stress-testing the architecture to quantitatively assess its robustness as inspired by central banks [55, 63].

4.1 Stress-Testing Framework

Central banks conduct stress tests of banking systems to test their ability to withstand shocks. For example, in an annual stress test, the Bank of England examines what the potential impact would be of an adverse scenario on the banking system [55]. The hypothetical scenario is a ‘tail-risk’ scenario, which seeks to be broad and severe enough to capture a range of adverse shocks ⁸. Following such best practice, we devise and implement a stress-test of the DeFi architecture.

4.2 Simulation Approach

We leverage the generic DeFi lending protocol architecture as developed in Section 3.1. The results here should be considered relevant to all lending protocols which rely on overcollateralization by volatile collateral and reserve assets. We make the following assumptions about the initial state of the system.

⁸For example, in its most recent stress-test of the UK financial system, the Bank of England considers the impact of a UK specific risk premium shock, which causes sharp falls in UK asset prices alongside a 30% depreciation of sterling [55]

1. The lending protocol allows users to deposit ETH as their single source of collateral c_i .
2. The lending protocol has 1m tokens of a generic reserve asset, which at the start of the simulation has the same price as ETH but with exactly half of the historical standard deviation of ETH taken over the sample period.
3. By arbitrage among borrowers, before the crash the lending protocol as a whole is collateralized to $\lambda + \epsilon$, i.e. just above the minimum collateralization ratio.
4. At the start of the crisis, the protocol has a collateralization ratio of exactly 150%, such that every USD of debt is backed by 1.50 USD of collateral.
5. Each unit of debt d^k maintains a peg of 1:1 to the US dollar, allowing us to abstract from the dynamics of maintaining the peg.

Next, we detail the methodology we follow to obtain our simulation results.

Price simulation. Firstly, we obtain OHLCV data at daily frequency Compare [18], focusing on the period 1 January 2018 to 7 February 2020, incorporating the large fall in the ETH price in early 2018. Taking parameters from this historical data, we use Monte Carlo simulation to capture how the ETH and reserve prices may be expected to evolve over the next 100 days. Monte Carlo simulation leverages randomness to produce a range of outcomes of a stochastic system. We simulate 5,000 randomly generated paths, using a geometric Brownian motion, specified with the following equation.

$$P_{c_{i,t}} = P_{c_{i,0}} \exp\left(\left(\mu_i - \frac{\sigma_i^2}{2}\right)t - \sigma_i W_t\right) \quad (6)$$

where

$$W_t = \mathcal{N}(0, t) \quad (7)$$

W_t denotes a Wiener process [72], μ denotes the mean of the log returns and σ denotes the standard deviation. The shocks are drawn from a standard normal distribution. However, given the possibility of the behavior of the different asset prices being correlated, we explicitly incorporate a strong positive correlation structure between the log returns of ETH and the reserve asset into the simulation by correlating these random shocks as they are generated⁹. Of the 5,000 simulations, our subsequent analysis is focused on the iteration which yields the lowest price after 100 days. By focusing on this worst-case, we test the DeFi lending protocol with a ‘black-swan’ event, representing a severe challenge to its robustness.

System simulation. In the event of such a severe price crash, since by assumption the protocol is collateralized

⁹We consider weak positive and strong negative correlations in the Appendix.

to 150%, the protocol will seize 100% of the debt value from the collateral pool, and seek to sell this collateral as quickly as possible on a market pair to the debt asset. Once a buyer has traded the debt asset d for the collateral, the protocol then burns the debt d , effectively taking it out of circulation, offsetting the liability. Therefore, the impact that our scenario has on the DeFi lending protocol, and how quickly it materializes, depends on liquidity available on all collateral/debt pairs. However, an common feature of crashes in financial markets is the *drying-up* of liquidity. In the event of a liquidity crisis, the demand for liquidity outstrips supply¹⁰, such that the constraint in Section 3.3 is binding. Thus a liquidity crisis occurs when this constraint is binding: there are not enough buyers in the market to buy the ETH that is for sale. We propose a simple model for the decline in liquidity over time as follows.

$$L = L_0 \exp(-\rho t) \quad (8)$$

where L_0 denotes the initial amount of ETH that can be sold per day. Intuitively, this equation captures the notion that in the event that the protocol attempts to sell large volumes each period, the amount of liquidity available in the next period will be lower.

In this simulation approach, we make a simplification by not modeling the impact that selling large volumes of collateral will have on the price of the collateral asset. It is highly likely that in such a sell-off scenario, the selling of large volumes would serve to endogenously push the price lower. Therefore what we present here represents an upper bound on the price behavior: in reality, the price drop may be even worse than the one we examine.

4.3 Simulation Results

We start with the Monte Carlo simulation of the correlated asset paths, before considering the impact this would have on a DeFi lending protocol and an ecosystem of multiple lending protocols.

4.3.1 Monte Carlo Price Simulation

We use data on the price of ETH/USD over the period 1 January 2018 to 7 February 2020¹¹. Perhaps the most notable element is the decline in the ETH/USD price over the course of 2018, with the price of ETH falling from an all-time-high of \$1,432.88 to c. \$220 as of 7 February 2020. With reference to Section A.2, we consider such a downtrend can be a period of reversion after

¹⁰Indeed, such liquidity crises were at the heart of the Financial crisis of 2007-8, as the value of many financial instruments traded by banks fell sharply without buyers [1].

¹¹See corresponding plot in Appendix 12

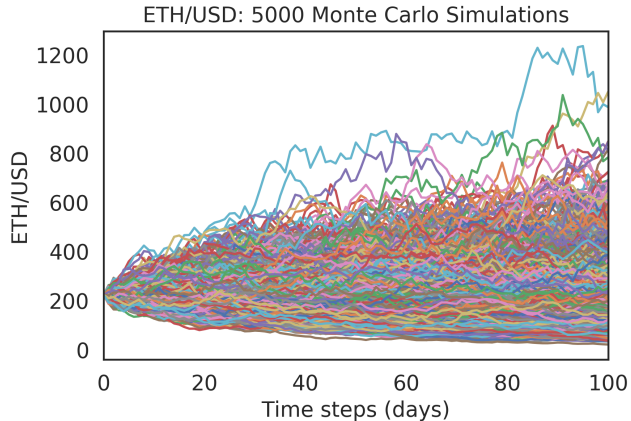


Figure 1: Monte Carlo forecast of ETH prices over the next 100 days from 7 February 2020.

excessive debt has been accumulated in the financial system, culminating in a bubble.

To capture the effects of different correlations between the collateral asset and the reserve asset, we consider three different extents of correlation between the collateral and reserve asset: (i) strong, positive correlation (0.9), (ii) weak, positive correlation (0.1) and (iii) strong negative correlation (-0.9). We then generate correlated asset paths during the Monte Carlo simulation process. In this section we report results for strong correlation, but include those for weak correlation and strong negative correlation in the Appendix.

Figure 1 shows the results of 5,000 runs of the Monte Carlo simulator for the ETH price, and Appendix A.3 Figure 11 shows the results for the reserve asset price in the presence of strong positive correlation in the asset price returns. The starting prices of assets as used in the simulator is the close price of ETH/USD on 7 February 2020.

We isolate the simulation in which the ETH/USD price is the lowest at the end of 100 days¹². In Figure 12 it is clear that in this worst case scenario for the ETH/USD price, the price of the reserve asset similarly falls. This illustrates the risk of using a reserve asset which is positively correlated with the collateral asset: if the price of the collateral asset falls, relative to the same basis of value the reserve asset value is likely to fall, limiting the ability of a DeFi lending protocol to recapitalize itself.

4.3.2 Impact on Collateral Margin

We take the simulation corresponding to the lowest ETH/USD price after 100 days and consider the impact this would have on the collateral margin of a DeFi lending protocol.

¹²We plot the co-evolution of the asset price paths for strong correlation in Appendix A.3 Figure 12.

Negative margins. The main results of this are presented in Figure 2. Plotted with solid lines is the evolution of the total collateral margin (comprising the collateral and the reserve asset) over time as the prices of the collateral asset and reserve asset decline. The dashed lines indicate how the amount of system debt evolves through time, on the assumption that at the start of the 100 day period, the protocol seeks to sell off all of the debt. The speed at which the debt can be liquidated through the sale of its backing collateral in turn depends on the available liquidity in the market for which we consider 3 cases: (i) constant liquidity (such that it is possible to sell a constant amount of ETH every day at the average daily price); (ii) mild illiquidity (where the illiquidity parameter is arbitrarily set to some low level $\rho = 0.005$); and (iii) $\rho = 0.01$, such that the market becomes relatively illiquid.

In particular, this figure makes the following parameter assumptions.

- at the start of the sell-off, it is possible to sell 30,000 ETH per day without having an impact on price¹³.
- the amount of reserve asset Π is fixed at the start of the sell-off at 1m units.
- system debt levels range from \$100m to \$400m, seeking to approximately reflect the levels of capital escrowed in DeFi protocols as in Section 2 Table 9.

Where the initial system debt level is 100m USD regardless of the liquidity parameter the collateral margin does not become negative. However, at higher levels of debt, we see that the margin gets closer to 0, and once the debt level reaches \$400m does indeed fall below 0, such that the protocol is undercollateralized overall. In the fourth panel of Figure 2 we see that after just over 50 days of the protocol attempting to liquidate as much debt as possible, due to illiquidity it is unable to liquidate in time and the margin becomes negative. This would constitute a crisis in a DeFi protocol: each unit of debt would not have sufficient collateral backing, and rational agents with weak identities would walk away from the protocol without repaying their debt. Notably, the results presented in the Appendix A show that a weakly correlated reserve asset is able to prevent the collateral margin from becoming negative (see Appendix A Figure 15) while a strongly negative correlation between the assets is actually able to bolster the collateral margin (see Appendix A Figure 17)

The effect of liquidity on margin. Given a set of parameters, where the collateral and reserve assets are strongly positively correlated we consider how quickly a

¹³This assumption is based on the 24-hour volume of ETH/DAI across markets listed on CoinGecko on 7 February 2020, and as such is only a rough proxy for the market liquidity. We use this figure only as a baseline for parameterization and to highlight the theoretical possibility of illiquidity causing default.

A Decentralized Financial Crisis: liquidity and illiquidity causing negative margins

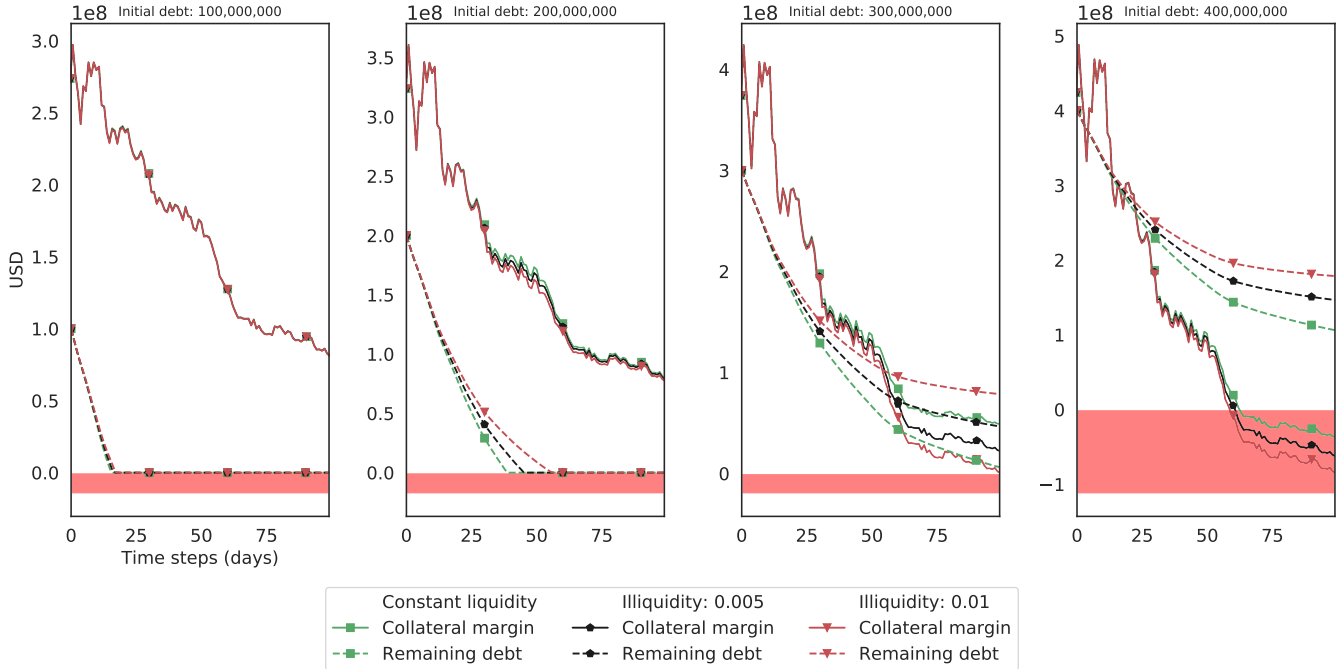


Figure 2: A DeFi lending protocol experiencing a sharp decline in the price of its collateral and reserve assets. Panels correspond to 4 different levels of system debt, with each panel showing the evolution of the collateral margin and the total debt outstanding. Each panel also shows the consequences of different liquidity parameters.

crisis may materialize (i) for varying values of liquidity parameters ρ in Figure 3, and (ii) for varying starting values of the amount of ETH that can be liquidated in Appendix A.3 Figure 13¹⁴.

Figure 3 shows that for a given amount of debt, as the liquidity parameter increases the margin becomes negative more quickly. Vice versa, it also shows for a given liquidity parameter, the more system debt there is, the more quickly the margin will become negative. We find that for a liquidity parameter of 0.025 and a debt of 750m USD, the margin can become negative within 40 days.

5 Governance Attack on Maker

In this section, we turn to a specific attack on an extant DeFi protocol—Maker [33]. In particular, we present and analyze in depth a potential attack on the Maker governance model. Although the basic idea of attack had been briefly presented in a blog post [49], the feasibility of the attack has not been analyzed. We use a representation of the current state of the Ethereum main network and the Maker contract to simulate as realistically as possible how such an attack could take place.

¹⁴The selection of these particular value ranges of debt is on the basis of providing ranges of values from which meaningful insights can be drawn.




Figure 3: Number of days before the collateral margin becomes negative, depending on the amount of system debt and the liquidity parameter.

5.1 Background and Threat Model

The governance process relies on the MKR token, where participants have voting rights proportional to the amount of MKR tokens they hold. MKR can be traded on exchanges [15].

Executive voting. Using executive voting, participants

can elect an *executive contract*, defining a set of rules to govern the system, by staking their token on it. Executive voting is continuous, i.e. participants can change their vote at any time and a contract can be newly elected as soon as it obtains a majority of votes. The elected contract is the only entity allowed to manipulate funds locked as collateral. If a malicious contract were to be elected, it could steal all the funds locked as collateral.

Defense mechanisms. Several defense mechanisms exist to protect executive voting. The *Governance Security Module* encapsulates the successfully elected contract for a certain period of time, after which the elected contract takes control of the system. At the time of writing, this period is set to zero [71] despite efforts to try to increase this delay to 24 hours [46]. Another defense mechanism is the *Emergency Shut Down*, which allows a set of participants holding a sufficient amounts of MKR to halt the system. However, this operation requires a constant pool of 50k MKR tokens, worth 27.5M USD at the time of writing¹⁵.

Threat model. An adversarial executive contract can steal the Maker collateral and mint new MKR tokens. Those can then be traded until the MKR price crashes and effectively destroy the Maker system. We assume the existence of an adversary with at least one (external) account in the Ethereum blockchain. The adversary is rational, i.e. would only engage in the attack if the potential returns are higher than the costs. There are currently circa 150k MKR tokens used for executive voting and the current executive contract has 76k MKR tokens staked. We observe that the amount of stake changes relatively often and the amount of tokens staked on the elected contract often drops below 50k MKR tokens (eq. 27.5M USD). At the time of writing, there are around 470M USD worth of ETH locked as collateral of the DAI supply, which an executive contract can dispose of freely.

5.2 Crowdfunding and Flash Loans

An adversary can choose between the two following options to amass the capital required for the governance attack.

Crowdfunding. Crowdfunding MKR tokens may allow users to lock their tokens in a contract and program the contract so that when the required amount of MKR tokens is reached, it stakes all its funds on a malicious executive contract. This would allow multiple parties to collaborate trustlessly on such an attack, while keeping control of their funds and being assured that they will be compensated for their participation in it¹⁶.

¹⁵For consistency, we use the price of MKR on 2020-02-01, which was 550 USD

¹⁶An (admittedly informal) poll on Twitter from late 2019¹⁷




Figure 4: Evolution of the amount of MKR tokens staked on different executive candidate contracts. We observe that at times the MKR amount of the executive contract drops below 50k MKR.

Liquidity pools and flash loans. A shortcoming of the crowdfunding attack is the required coordination effort between the participants and the likely alerting of benevolent MKR members. Instead, an attack could use liquidity pools offering so called “flash-loans” [3]. A flash-loan is a *non-collateralized* loan that can be valid within one transaction only. In the EVM, a transaction can be reverted entirely if a condition in one part of the transaction is not fulfilled. A flash-loan then operates as follows: A party creates a smart contract that (i) takes out the loan, (ii) executes some actions, and (iii) pays back the loan with interest.

The interesting aspect for our purposes is that if in step (ii) the execution of the actions fails or step (iii) the payment of the loan cannot be completed, the EVM treats this loan as it never took place. Hence, under the assumption there is enough liquidity available in protocols like Aave [2], an attacker could execute the MKR governance attack in step (ii), and, if successful, pay back the flash loan with interest in step (iii). Since the flash-loan requires no collateral, the capital lock up cost for the attacker is significantly reduced. If there is enough liquidity available in these pools, the attacker might even not have to lock any tokens. Furthermore, the liquidity provider also profit from the execution of the attack as they receive the 0.35% fee required in step (iii).

5.3 Practical Attack Viability

In this section, we use empirical data to show how such an attack could potentially take place, and what would be the different shortcomings.

conducted by a user soon after this attack shows that several participants may be interested in such an attack

We first analyze all the transactions received by the current governance contract of Maker DAO: `0x9ef05f7f6deb616fd37ac3c959a2ddd25a54e4f5`. Since the deployment of this contract, in May 2019, there have been a total of 24 different contracts which have been elected as executive contract (cf. Figure 4). When a contract becomes the executive contract, the staked amount is distributed almost equally for a short period of time between the previous and the new executive contract, which reduces the amount of required tokens by more than 50%. This is particularly visible at the end of November 2019 (80k MKR to 40k MKR) or at the middle of January 2020 (120k MKR to 45k MKR) ¹⁸.

5.4 Attack Steps

We note that an attacker is also able to combine crowd-funding and flash loans to achieve a successful attack.

5.4.1 Crowdsourcing MKR Liquidity

We inspect the amount of MKR transferred between January 1, 2020 and February 8, 2020 ¹⁹. We observe that on most days the amount is somewhere between 4k and 6k tokens but there are some days when a vastly larger amount of MKR tokens is transferred: more than 18k tokens on February 7, for instance. This results in a mean of roughly 9k MKR tokens traded per day. This corresponds roughly to the estimate provided by other sources [15], and is enough liquidity for an attacker to accumulate tokens in a timely manner.

Given a liquidity of on average 9k MKR per day, an attacker could accumulate enough MKR to perform such an attack in a timely manner. At an average rate of 1k MKR per day, it would take less than 2 months. However, accumulating all the tokens in a single account would likely attract attention. Indeed, from our discussions with the Maker DAO team, the large MKR token holders are known. A potential strategy is to accumulate tokens without perceptibly changing the distribution of MKR tokens. At present there are c. 5k accounts, holding

¹⁸We note the sudden increase of MKR by an executive contract from about 75k to nearly 160k MKR in the beginning of December 2019 (the increase happened one day after a first blog on this topic was published [49]). A token holder [35] injected a very large volume of tokens — around 66k MKR — to potentially help prevent an attack from occurring. It is unclear if the token holder was the Maker Foundation or some other party; in our discussion with Maker they states they knew the identity of the token holder. The holder staked his token to the currently elected contract, making the attack more difficult to execute, before releasing the tokens it staked about one month later. The token holder had vastly more than the necessary amount to execute the attack, and if he would have been malicious could have stolen the funds.

¹⁹See Appendix A.6 Figure 19

a total of slightly more than 272k MKR tokens ²⁰. It is possible for an attacker to accumulate tokens while spreading the wealth in many accounts with this range of wealth. Given that the attack is possible with 50k tokens, an adversary could spread his wealth across 100 accounts, say, with an average of 500 tokens each. One of the drawbacks of this approach is the requirement to vote from these 100 accounts. This vote should happen in the shortest time frame as possible to reduce the chance of MakerDAO preventing the attack by introducing a non-zero delay on new governance contracts becoming effective [46]. Voting for a contract costs on average 69k gas, which means that filling half of a block with voting transactions would allow to vote from more than 72 contracts. Filling half-a-block with transactions costs a mere c. 10 USD [59] at the current Ethereum gas price, which means that an attacker could easily perform the whole attack in two blocks. In the second block, the attacker can finish voting for his malicious contract and execute the attack from the contract, which would leave only one block — or less than 15 seconds — for anyone to react to the attack.

5.4.2 Liquidity pools

To execute the attack without amassing tokens, the attacker can utilize liquidity pools to borrow the required tokens via a flash-loan (e.g. via Aave [3]²¹). The attacker performs the following steps within two transactions (cf. Fig. 5).

Transaction 1: Deploy the malicious governance contract and deploy the attack contract.

Transaction 2: Call the attack contract deployed in step 1 that executes the following steps.

1. Take out a flash loan that allows to convert enough MKR from Aave in the currency with the highest liquidity.
2. Convert the ETH loan into 50k MKR tokens on a decentralized exchange(s) with enough liquidity.
3. Vote with the 50k MKR tokens to replace the current MakerDAO governance contract with the malicious contract deployed in step 1.
4. Take out enough ETH from the MakerDAO system to repay the flash loan with interest.
5. Repay the flash loan with the required 0.35% interest to Aave.

Now, we argue how to realistically construct this attack and give an analysis of the costs involved. Any information given in the steps below with regard to currencies

²⁰Excluding the holders with a low balance (less than 1 MKR token), and a large balance (more than 5k MKR tokens)

²¹Aave is a protocol deployed on the Ethereum mainnet on January 8, 2020, <https://etherscan.io/tx/0x4752f752f5262fb11733e0136033f7d53cdc90971441750f606cf1594a5fde4f>

is based on information obtained on February 14, 2020.

Required funds. Aave’s lending pool prevents reentrancy, i.e. the attacker needs to obtain enough liquidity to obtain the required 50k MKR tokens within a single flash loan. Hence, the attacker should select ETH, as the asset with the largest liquidity. In a naive approach, we could utilize the exchange rate for ETH to MKR to obtain that an attacker requires 114,746 ETH to execute the attack. However, this approach would fall short of the reality of exchange fees: to obtain such large quantities of MKR, the attacker needs to pay a premium. As of 14 February 2020, an attacker could source 50k MKR tokens from three different DEXs with a token split of 38k MKR from Kyber, 11,500 MKR from Uniswap, and 500 MKR from Switcheo for a total of 378,940 ETH²². The attacker thus has to obtain more than an additional 250k ETH from Aave.

Time to attack. As of February 14, 2020 Aave has around 13,670 ETH available in its liquidity pool and is growing with a rate of around 219.5 ETH per day. If we estimate a similar constant growth of the available liquidity pool, it would take about 1,663 days for the pool to be large enough to execute the governance attack *without the attacker requiring to own any tokens*. However, currently the ETH growth rate of Aave is 5.18% per day. Assuming this growth rate continues, it would only take *66 days* until enough liquidity is available in Aave. We also note that if the liquidity of MKR on the observed DEXs increases, the attacker can obtain a better exchange rate. Thus, less ETH needs to be borrowed to attack, decreasing the time until the attack is possible.

5.5 Profitability Analysis

Crowdsourcing. With the crowd funding strategy, the profits from the attack could be split equally between the funders. The only cost are the 20 USD for including the transactions. In return, the attackers can take away the currently 434,873 ETH in collateral in MakerDAO plus the 145m DAI. This amounts to a net profit of 263m USD.

Liquidity pools and flash loans. Assuming Aave’s liquidity pool has accumulated the required 378,940 ETH to execute the attack, we can calculate the profitability as follows. The attacker obtains a total of 434,873 ETH in collateral from Maker as well as the 50k MKR tokens and the 22m DAI currently in circulation. The attacker needs to repay the 378,940 ETH loan with a 0.35% interest (1,326.29 ETH). Furthermore, the attacker needs to pay for the gas fees for the two transactions. The second transaction involves various function calls to other contracts and will cost around 15 USD equivalent of

gas. However, by the end of the attack, the attacker has around 55k ETH, 50k MKR, and 145m DAI. This amounts to a net profit of 191m USD. Moreover, the attacker can design the attack smart contract such that the transaction is reverted if it becomes unprofitable. This makes the attack *risk free* from a cost perspective for the attacker.

6 Composability and Contagion

The possibility of financial contagion in the context of DeFi is of particular importance given the unrestricted composability of protocols [62]. In Sections 4 and 5, we considered an exogenous price drop and a governance attack in each case on a single protocol. Now we argue why these two design weaknesses can lead to contagion to other protocols that might ultimately lead to a decentralized financial crisis.

Both weaknesses, exogenous price drops and governance attacks, result in the debt asset being under-collateralized. Assuming rational agents with weak-identities, it is not individually rational for agents to settle their debts, and thus, the under-collateralized asset eventually reaches a value of 0. Hence, an agent holding such an asset can execute a strategy called *liquidity sweeping* to use his *existing holdings* to buy other assets while the price is not yet 0. A special sub-case occurs in the governance attack scenario where the attacker can additionally mint an unlimited supply of the debt asset to buy up *all* the available liquidity of other assets.

Further, any protocol that uses the now under-collateralized asset as a collateral asset to issue another asset, also becomes under-collateralized. Consider the example of MakerDAO and Compound: MakerDAO uses ETH as a collateral to issue DAI. Compound allows agents to collateralize DAI to issue a cDAI token. As DAI becomes under-collateralized, cDAI loses its debt-backing. We refer to this as *evaporating collateral*. Holders of any composed assets like cDAI then have to essentially also buy up existing liquidity once they are aware that individually rational agents will not settle their debts. We illustrate this cycle in Figure 6 and give a total estimate of the financial damage in Figure 7.

Liquidity sweeping. Contagion occurs when the under-collateralized debt asset is used to “soak-up” as much liquidity as possible, before buyers of the debt asset are able to react. There are two cases: First, if, as detailed in Section 4, the debt asset is under-collateralized due to an exogenous price crash, the agent can trade all their under-collateralized debt (such as DAI) for other assets on DeFi lending protocols which offer liquidity pools such as Compound [31], Uniswap [69], and Aave [3]. This way the agent can dispose of the debt asset if it assumes the debt asset will fall to 0. Second, if, as detailed in

²²For current liquidity and rates see <https://dexindex.io/>.




Figure 5: Example flash loan attack against Maker DAO. All steps can be executed within one transaction, under the assumption that the flash loan pool and DEX have sufficient liquidity available. To execute the attack, the adversary would not need upfront capital, besides the gas fees (estimated to amount to c. 15 USD).

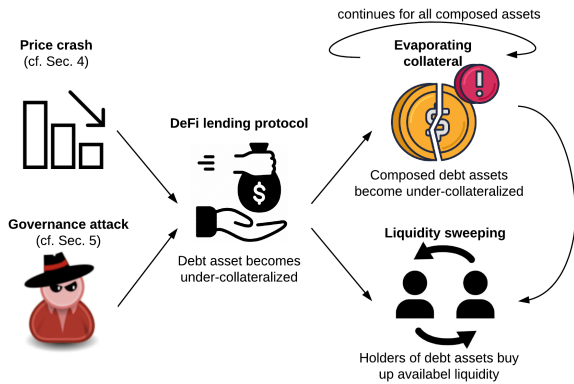


Figure 6: Both the price crash and the governance attack lead to an under-collateralization of a DeFi lending protocol. In turn, agents will try to exit the protocol by buying other assets (liquidity sweeping). Any agent in a protocol using the now under-collateralized asset as debt, will face the same situation and in turn try to exit via buying other assets (evaporating collateral).

Section 5, an attacker gains the ability to create new tokens without collateral backing, the attacker can use these tokens to buy *all existing liquidity* on available exchanges.

However, both cases assume that the price of the under-collateralized debt asset is not 0 at least as long to complete the trades. We support this assumption with the halting of the IOTA cryptocurrency network that occurred on February 13–15, 2020 [38]. Following a security incident, the IOTA foundation stopped the centralized coordinator such that the IOTA network was not able

Weakness	Financial damage
Under-collateralization (price crash) of MakerDAO	\$145m
Under-collateralization (governance attack) of MakerDAO	\$211m
Contagious under-collateralization (price crash) of MakerDAO	\$180+m
Contagious under-collateralization (governance attack) of MakerDAO	\$246+m

Figure 7: Financial damage for an initial failure of the MakerDAO protocol either through a price crash or a governance attack.

to process transactions containing value and advised its users to not access their funds in their wallets [37]. On February 15, the attack is still ongoing [38]. During the incident, the trading volume on exchanges of IOTA fell from 41m USD to 21m USD. Notably, the price of IOTA only moved from 0.34 USD to 0.31 USD within 14 hours of the announcement and recovered back to around 0.32 USD after another two hours. Hence, the inability to transfer the IOTA cryptocurrency and access funds on the IOTA network for an extended period of time, caused only a maximum price drop of 9%.

To quantify the impact of liquidity sweeping, we took an instantaneous snapshot of major marketplaces offering a DAI pair at approximately noon (GMT+8) on 15 February, as presented in Figure 20. In the price crash case, agents could trade the whole of their \$145m holdings. However, an agent who had successfully undertaken the governance attack would be able to soak up all of the order book liquidity, since with in effect an unlimited DAI supply they can pay any price to acquire other assets. We estimate the total USD value, aggregating across markets and pairs, to be c. \$211m²³.

Evaporating collateral. Liquidity sweeping is not contained to agents that participate in a single lending protocol if the assets of the protocol are used in others as well. Rather, a vicious cycle occurs where any debt-backed asset issued in another protocols, e.g. where DAI would be used to issue cDAI on Compound, loaned in Aave, and used as a margin trading collateral in dY/dX. If we consider the DAI example in the price crash case, \$145m of DAI would be exchanged as well as \$35m of cDAI and any other protocol that uses DAI as backing collateral.

Collateral composition. As a stylized example, assume that there are N protocols, with different minimum collateralization ratios, that each use D/N as collateral. In the event of this collateral itself becoming less than 100% collateralized, we assume that the value of the debt would quickly fall to 0. We further assume that agents choose a particular protocol for reasons exogenous to the system, and that each agent obtains the maximum leverage they can on a certain protocol by using their issued debt to purchase additional collateral, which in

²³See Appendix A.7 Figure 20

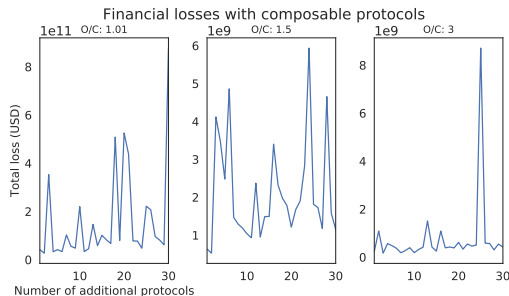


Figure 8: Financial losses for 30 DeFi protocols, which use the debt asset as collateral. “O/C” stands for overcollateralization.

turn is used to issue more debt. To reflect heterogeneity in collateralization ratios among protocols, we assign each of the N protocols a collateralization ratio based on randomly sampling from a uniform distribution over a specified range. The maximum possible leverage for a given protocol π when an agent allocates D/N units of collateral to it is given by $D/N_{\pi}/\lambda_{\pi}$. Therefore the maximum systemic loss is given by $\sum_{\pi=1}^{\Pi} D/N_{\pi}/\lambda_{\pi}$.

Figure 8 demonstrates how for three different maximum overcollateralization parameter ranges, namely (101–105%), (101–150%), and (101–300%), what the maximum system wide losses could be, assuming that the protocol in which the crash occurs has a debt of 400m USD, and the liquidity parameter is $\rho = 0.01$. As the minimum overcollateralization ratio falls, the maximum potential losses for the system grow: from right to left, as the minimum overcollateralization requirement falls from $3\times$ (or 300%) to $1.01\times$ (101%), the maximum possible loss increases. Given our uniform splitting of debt, what also emerges is that the result is driven by the minimum collateralization ratio.

Crisis. Agents seek to exit their under-collateralized asset positions by buying other available assets. Individually rational agents should seek assets that are uncorrelated with the asset they are disposing of. However, this leads to a spread of the initial DeFi lending protocol failure to any other asset that is available for trading on exchanges that accept the initial lending asset. Hence, the DeFi crisis can spread across multiple blockchains and also affect centrally-backed assets like USDT [67].

7 Related Work

There is a paucity of directly related work. However, existing work can be divided into the following categories.

Deleveraging spirals and pegged assets. A series of fundamental results in relation to the ability of non-custodial stablecoins to maintain their peg is provided in [40]. It is shown that stablecoins face deleveraging spirals which cause illiquidity during crises, and that

stablecoins have ‘stable’ and ‘unstable’ domains. The model primarily involves the assumption of two types of agents in the marketplace: the stablecoin holder (who wants stability), and the speculator (who seeks leverage). The authors further demonstrate that such systems are susceptible to tail volatility. While unpublished, [13] uses option pricing theory to design dual-class structures that offer fixed income stable coins that are pegged to fiat currency. [58] considers how one might build an asset-backed cryptocurrency through the use of hedging techniques.

Identity and rationality. In [32] consideration is given as to how a system that is designed to be secure only against rational adversaries cannot be considered secure more generally, these systems only become so if they remain secure even when participating parties behave in a fully Byzantine way. In building the argument, the authors rely on a distinction between ‘weak’ and ‘strong’ identities, similar to our work.

8 Conclusions

This paper presents two mechanisms through which a decentralized financial crisis could manifest. After providing formal constraints on the robust operation of a DeFi lending protocol, we use Monte Carlo simulation to show how, for a range of parameters, a DeFi lending protocol may find itself critically under-collateralized.

We then consider a governance attack on Maker, and show that provided an attacker is able to lock 27.5m USD of governance tokens, they can steal all 0.5bn USD worth of collateral within two blocks. We present a novel strategy that would enable an attacker to steal the collateral within two transactions, without the need to even escrow any assets.

These two sources of weakness in a DeFi protocol are inter-related, and could serve to reinforce each other. In the event that the collateral and reserve assets of a DeFi lending protocol experience a sharp decline in price, tending towards under-collateralization, the cost of acquiring enough governance tokens to undertake the governance attack would also likely fall. Conversely, should an actor undertake a governance attack, this would plausibly send shockwaves through the DeFi ecosystem, serving to reduce the price of the collateral asset, in turn making under-collateralization more likely.

A clear lesson from the Global Financial Crisis of 2008 is the role that financial contagion can play [5]. We consider transmission mechanisms for contagion in DeFi, mechanisms which, in the event of any of the DeFi weaknesses we identify manifesting, could precipitate a decentralized financial crisis.

References

- [1] Three myths that sustain the economic crisis. *The Guardian*, Aug 2012.
- [2] Aave. Aave Liquidity Pools. <https://app.aave.com/borrow>, 2020.
- [3] Aave. Aave Protocol. <https://github.com/aave/aave-protocol>, 2020.
- [4] Amitanand S Aiyer, Lorenzo Alvisi, Allen Clement, Mike Dahlin, Jean-Philippe Martin, and Carl Porth. BAR fault tolerance for cooperative services. *ACM SIGOPS Operating Systems Review*, 39(5):45, 2005.
- [5] Franklin Allen and Douglas Gale. Financial contagion. *Journal of political economy*, 108(1):1–33, 2000.
- [6] Augur. Augur. <https://www.augur.net/>, 2020.
- [7] Christian Badertscher, Peter Gaži, Aggelos Kiayias, Alexander Russell, and Vassilis Zikas. Ouroboros genesis: Composable proof-of-stake blockchains with dynamic availability. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 913–930, 2018.
- [8] Christian Badertscher, Ueli Maurer, Daniel Tschudi, and Vassilis Zikas. Bitcoin as a transaction ledger: A composable treatment. In *Annual International Cryptology Conference*, pages 324–356. Springer, 2017.
- [9] Bancor. Bancor. <https://www.bancor.network/>, 2020.
- [10] BBC. BBC World Service, Aftershock Timeline. http://www.bbc.co.uk/worldservice/business/2009/09/090902_aftershock_timeline_noflash.shtml, 2009.
- [11] Rainer Böhme. A primer on economics for cryptocurrencies, 2019.
- [12] Lehman Brothers. Lehman brothers holdings inc. plan trust – ‘10-k’ for 11/30/07, 2007.
- [13] Yizhou Cao, Min Dai, Steven Kou, Lewei Li, and Chen Yang. Designing stable coins. 2018.
- [14] Miguel Castro and Barbara Liskov. Practical Byzantine Fault Tolerance. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation*, number February, pages 173–186, 1999.
- [15] CoinMarketCap. Maker (MKR) price, charts, market cap, and other metrics. <https://coinmarketcap.com/currencies/maker/>, 2020. [Online; accessed 9-February-2020].
- [16] Coinut. Coin ultimate trading. <https://coinut.com/>, 2020.
- [17] Connex. Connex. <https://connex.network/>, 2020.
- [18] CryptoCompare. Cryptocompare. <https://www.cryptocompare.com/>, 2020.
- [19] Ray Dalio. How the economic machine works, 2013.
- [20] Ray Dalio. Principles for navigating big debt crises. *Westport: Bridgewater*, 2018.
- [21] Cryptowatch Market Data. Cryptowatch dai pair market data. <https://cryptowat.ch/assets/dai>, 2020.
- [22] Bernardo David, Peter Gaži, Aggelos Kiayias, and Alexander Russell. Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 66–98. Springer, 2018.
- [23] DDEX. Ddex: Decentralized margin trading. <https://ddex.io/>, 2020.
- [24] DexIndex. Dexindex. <https://dexindex.io/>, 2020.
- [25] dYdX. dYdX. <https://dydx.exchange/>, 2020.
- [26] Timothy C Earle. Trust, confidence, and the 2008 global financial crisis. *Risk Analysis: An International Journal*, 29(6):785–792, 2009.
- [27] VCC Exchange. Vccexchange. <https://vcc.exchange/>, 2020.
- [28] EXMO. Exmo. <https://exmo.com/>, 2020.
- [29] Ittay Eyal and Emin Gün Sirer. Majority is not enough: Bitcoin mining is vulnerable. In *International conference on financial cryptography and data security*, pages 436–454. Springer, 2014.
- [30] Ingo Fender and Janet Mitchell. Structured finance: complexity, risk and the use of rating. *BIS Quarterly Review*, June, 2005.
- [31] Compound Finance. Compound finance, 2019.
- [32] Bryan Ford and Rainer Böhme. Rationality is self-defeating in permissionless systems. *arXiv preprint arXiv:1910.08820*, 2019.

- [33] The Maker Foundation. Makerdao. <https://makerdao.com/en/>, 2019.
- [34] Dominik Harz, Lewis Gudgeon, Arthur Gervais, and William J Knottenbelt. Balance: Dynamic adjustment of cryptocurrency deposits. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 1485–1502. ACM, 2019.
- [35] i3nikolai. Note on MKR voting and distribution. https://www.reddit.com/r/MakerDAO/comments/enpad1/note_on_mkr_voting_and_distribution/, 2020. [Online; accessed 9-February-2020].
- [36] InstaDApp. InstaDApp: Trustless smart wallet for DeFi. <https://instadapp.io/>, 2020.
- [37] IOTA. Iota coordinator halted, 2020.
- [38] IOTA. Iota status, 2020.
- [39] Christoph Kern, Anita Kesavan, and Neil Daswani. *Foundations of security: what every programmer needs to know*. Springer, 2007.
- [40] Ariah Klages-Mundt and Andreea Minca. (in) stability for the blockchain: Deleveraging spirals and stablecoin attacks. *arXiv preprint arXiv:1906.02152*, 2019.
- [41] Aikaterina Koutsouri, Francesco Poli, Elise Alfieri, Michael Petch, Walter Distaso, and William Knottenbelt. Balancing cryptoassets and gold: A weighted-risk contribution index for the alternative asset space. *Mathematical Research for Blockchain Economy, Forthcoming*, 2019.
- [42] KuCoin. Kucoin. <https://trade.kucoin.com/spot>, 2020.
- [43] Kyber. Kyber. <https://kyber.network/>, 2020.
- [44] Probit Global Service Limited. Probit. <https://www.probit.com/en-us/>, 2020.
- [45] Maker. MakerDAO Governance Risk Framework (Part 3). <https://blog.makerdao.com/makerdao-governance-risk-framework-part-3/>, 2018. [Online; accessed 9-February-2020].
- [46] Maker. The Governance Security Module (GSM). <https://blog.makerdao.com/governance-security-module-gsm/>, 2019. [Online; accessed 9-February-2020].
- [47] Inc. Maker Ecosystem Growth Holdings. Oasis trade. <https://oasis.app/>, 2020.
- [48] MelonProject. Melon Project. <https://github.com/melonproject/melon-lab/releases>, 2020.
- [49] Micah Zoltu. How to turn \$20M into \$340M in 15 seconds. <https://medium.com/coinmonks/how-to-turn-20m-into-340m-in-15-seconds-48d161a42311>, 2019. [Online; accessed 9-February-2020].
- [50] Andrew Miller, Yu Xia, Kyle Croman, Elaine Shi, and Dawn Song. The Honey Badger of BFT Protocols. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security - CCS'16*, (Section 3):31–42, 2016.
- [51] Nexus Mutual. Nexus Mutual. <https://nexusmutual.io/>, 2020.
- [52] Satoshi Nakamoto et al. Bitcoin: A peer-to-peer electronic cash system. 2008.
- [53] Lightning Network. Lightning Network. <https://lightning.network/>, 2020.
- [54] Kleopatra Nikolaou. Liquidity (risk) concepts: definitions and interactions. 2009.
- [55] Bank of England. Stress testing the uk banking system: key elements of the 2019 annual cyclical scenario, 2019.
- [56] OKEX. Okex. <https://www.okex.com/en/>, 2020.
- [57] Rafael Pass, Lior Seeman, and Abhi Shelat. Analysis of the blockchain protocol in asynchronous networks. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 643–673. Springer, 2017.
- [58] Alex Lipton Thomas Hardjono Alex Pentland. Digital trade coin (dtc): Towards a more stable digital currency. 2018.
- [59] Daniel Perez and Benjamin Livshits. Broken metre: Attacking resource metering in evm. *arXiv preprint arXiv:1909.07220*, 2019.
- [60] Marc Pilkington. Blockchain technology: principles and applications. In *Research handbook on digital transformations*. Edward Elgar Publishing, 2016.
- [61] DeFi Pulse. The DeFi Leaderboard. <https://defipulse.com/>, 2019.
- [62] DeFi Pulse. What is defi?, 2019.
- [63] The Federal Reserve. Dodd-frank act stress tests. <https://www.federalreserve.gov/supervisionreg/dfa-stress-tests.htm>, 2019.

- [64] Ameen Soleimani. What you should know before putting half a million dai in compound, 2019.
- [65] Switcheo. SwitcheoNetwork. <https://switcheo.network/>, 2020.
- [66] Synthetix. Synthetix: Decentralized synthetic assets. <https://www.synthetix.io/>, 2020.
- [67] Tether. Digital money for a digital age, 2020.
- [68] TokenSets. TokenSets. <https://www.tokensets.com/>, 2020.
- [69] Uniswap. Uniswap. <https://uniswap.io/>, 2020.
- [70] WBTC. Wbtc. <https://www.wbtc.network/>, 2020.
- [71] Week in Ethereum News. Week in Ethereum News December 28, 2019. <https://weekinethereumnews.com/week-in-ethereum-news-december-28-2019/>, 2020. [Online; accessed 9-February-2020].
- [72] Norbert Wiener. Collected works, vol. 1, p. masani, 1976.
- [73] xDAI. xDAI. <https://www.xdaichain.com/>, 2020.

A Appendix

A.1 Existing DeFi protocols

	Project	Capital (USD)	Blockchain
Lending	Maker [33]	642.3m	Ethereum
	Compound [31]	133.7m	Ethereum
	InstaDApp [36]	82.9m	Ethereum
DEX	Uniswap [69]	62.8m	Ethereum
	Bancor [9]	11.9m	Ethereum
	Kyber [43]	4.8m	Ethereum
Derivatives	Synthetix [66]	143.3m	Ethereum
	Nexus [51]	3.1m	Ethereum
	Augur [6]	0.6m	Ethereum
Payments	Lightning [53]	8.9m	Bitcoin
	xDAI [73]	39k	Ethereum
	Connex [17]	12.3k	Ethereum
Assets	WBTC [70]	7.2m	Ethereum
	token Sets [68]	4.8m	Ethereum
	Melon [48]	0.3m	Ethereum

Figure 9: Existing DeFi projects [61] (12 February 2020).

A.2 Economic Cycles and Credit

Whether in the context of DeFi or traditional finance, a core feature of economies is a fluctuation between periods of economic expansion and contraction. One of the main

reasons for such cycles is credit. To spur healthy growth, markets allow borrowers to take on debt from lenders in the form of credit. Credit is generally considered positive, if the borrowed money is used productively, and the lender avoids granting excessive credit [20]. While credit can thus be used for good purposes, the market may also abuse such capital injection. To (dis)incentivize borrowing, central banks and policy makers can adjust the *interest rate*. A higher interest rate renders borrowing more expensive, while a lower interest rate encourages lenders to give out credit. An abuse of excessive credit typically leads to debt defaults—i.e. a borrower no longer is able to repay its debt. Lending creates a self-reinforcing upward movement: the more credit is available and affordable, the more lenders take on credit to expand their business and thus the economy expands as a whole. This upwards trend however, necessarily must reverse at some point to repay and equate debt. Due to this dynamic, we observe cycles, with upward and downward trends. The literature differentiates among short-term (roughly every 5 – 8 years) and long-term (every 75 – 100 years) debt cycles [19]. A debt crisis, such as the Financial Crisis of 2007-8, commonly occurs when debt and interest payments rise faster than the income that is supposed to pay back the debt. Once debt is no longer paid back and too much debt accumulated, policy makers (e.g. central banks) should intervene.

A.3 Strong positive correlation




Figure 10: Close prices for ETH/USD over the period 1 January 2018 to 7 February 2020.

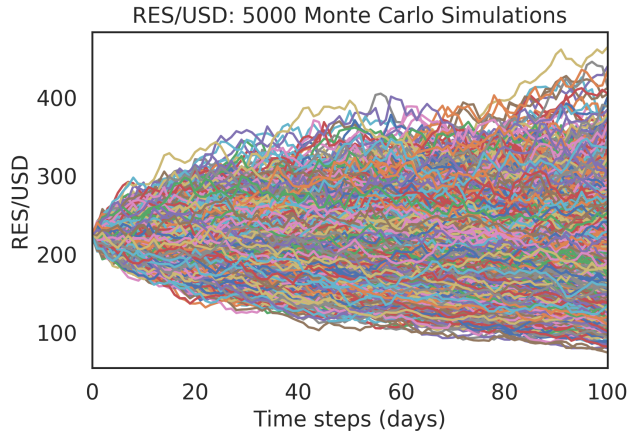


Figure 11: Monte Carlo forecast of the reserve asset price over the next 100 days from 7 February 2020.




Figure 12: For the simulation yielding the worst ETH price after 100 days, the co-evolution of the ETH and reserve asset prices where the asset price returns are strongly positively correlated.

be sold within 24 hours), the faster a negative margin precipitates. Similarly, for a fixed initial starting liquidity, the more debt there is in the system the faster the margin will become negative, down to below 40 days.

A.4 Weak positive correlation

Figure 14 shows the evolution of the ETH and reserve asset price to USD when the correlation of these assets is positive but weak (0.1).




Figure 14: For the simulation yielding the worst ETH price after 100 days, the co-evolution of the ETH and reserve asset prices where the asset price returns are weakly positively correlated.

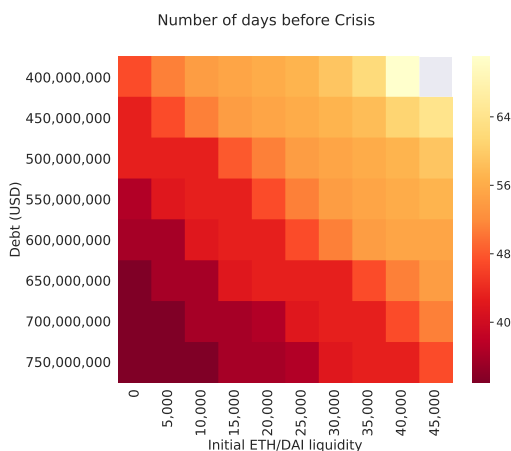


Figure 13: Number of days before the collateral margin becomes negative, depending on the amount of system debt and the initial amount of ETH sellable within 24 hours.

Figure 13 shows that for a given amount of debt, the lower the starting liquidity (i.e. the amount that can

Figure 15 shows the evolution of the collateral margin and debt in the presence of weak positive (0.1) correlation between the returns of the collateral asset and the reserve asset.

A.5 Strong negative correlation

Figure 16 shows the evolution of the ETH and reserve asset price to USD when the correlation of these assets is strong and negative (-0.9).

A Decentralized Financial Crisis: liquidity and illiquidity causing negative margins




Figure 15: A DeFi lending protocol experiencing a sharp decline in the price of its collateral and reserve assets. Panels correspond to 4 different levels of system debt, with each panel showing the evolution of the collateral margin and the total debt outstanding. Each panel also shows the consequences of different liquidity parameters.

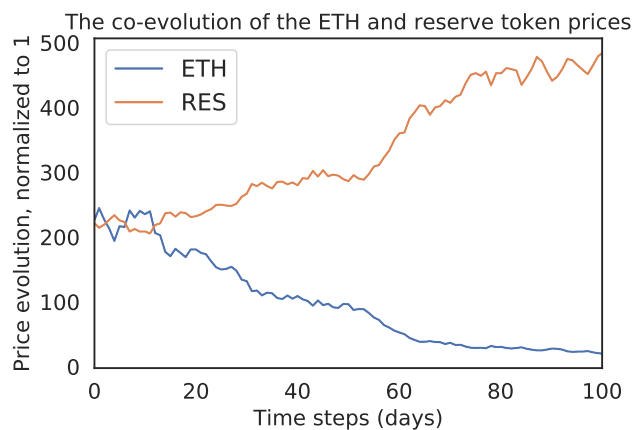


Figure 16: For the simulation yielding the worst ETH price after 100 days, the co-evolution of the ETH and reserve asset prices where the asset price returns are strongly negatively correlated.

Figure 17 shows the evolution of the collateral margin and debt in the presence of strong negative (-0.9) correlation between the returns of the collateral asset and the reserve asset.

A.6 Governance Attack on Maker




Figure 19: Daily traded volume of MKR tokens between 2020-01-30 and 2020-02-08.

A.6.1 Possible Remedies

Although some remedies have been discussed, no improvement has been deployed yet to the network. In this section, we discuss fixes that have been discussed and other potential way to improve the current state.

Governance security module. One of the major issues is that there is currently absolutely no delay between the moment a new contract is elected and the moment it

A Decentralized Financial Crisis: liquidity and illiquidity causing negative margins

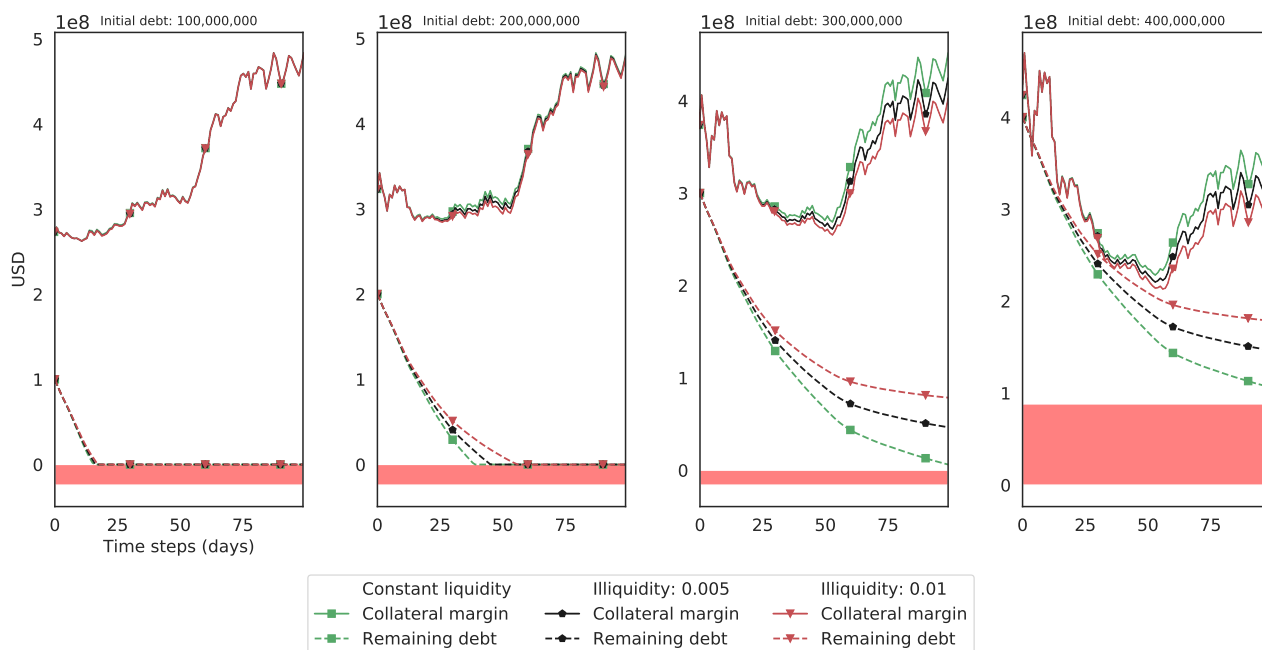


Figure 17: A DeFi lending protocol experiencing a sharp decline in the price of its collateral and reserve assets. Panels correspond to 4 different levels of system debt, with each panel showing the evolution of the collateral margin and the total debt outstanding. Each panel also shows the consequences of different liquidity parameters.



Figure 18: Twitter poll for of a crowd-funding attack on MKR governance.

can control all the locked ETH. The most straightforward fix for this issue is to add a delay, during which participants could review the code of the new contract and eventually trigger an emergency shutdown if the contract looks clearly malicious or vulnerable to attacks. However, this results in another problem: security patches currently go through the same executive vote process and the Governance Security Module, which means that if a 24-hour delay were to be added, the vulnerability would be exploitable for the whole time. Given that both the patched source code would be available as well, it would make it very easy for an attacker to exploit the contract. This creates a large trade-off between protect-

ing against governance attack and keeping the contract secure in case vulnerabilities are discovered.

Principle of least privilege. Another cause of this major issue is the excess of privilege that the executive contract is granted. The contract is currently designed to represent the state of the *whole* system [45]. Although this makes the governance process much simpler, as a user simply needs to vote on a single contract, it also gives an immense level of privilege to this contract, most of which should never be necessary under normal circumstances. This goes against the principle least privilege which has been crucial principle to secure software systems [39].

Although there are technical challenges to dividing privileges in such a system, adding a layer of control, which could obey different voting rules, on the locked collateral could make such an attack vastly more difficult to orchestrate. Many multi-signature wallets [60] have the notion of threshold, over which more parties need to give their approval in order to allow a withdrawal. A similar approach could be taken to ensure that locked collateral cannot be stolen at once.

A.7 Contagion

Marketplace	Total orders for DAI (USD)
CoinbasePro [21]	7,343,000
OasisDEX [47]	2,483,000
Kraken [21]	6,085,000
Bitfinex [21]	2,661,000
Bittrex [21]	39,957,000
KuCoin [42]	2,442,000
Probit [44]	10,000
Switcheroo [65]	176,000
VCC Exchange [27]	21,000
OkEx [56]	11,000
Exmo [28]	61,000
DDEX [23]	46,000
Coinut [16]	6,000
Compound [31]	143,283,000
Uniswap [24]	2,176,000
Kyber [24]	2,176,000
Bamboo Relay [24]	979,000
Eth2DAI [24]	971,000
Airswap [24]	99,000
TOTAL	210,986,000

Figure 20: Assuming DAI:USD peg maintained at 1:1, total USD liquidity on all DAI pairs by marketplace. Rounded to nearest thousand.