



Liquidations: DeFi on a Knife-Edge

Daniel Perez^{1(✉)}, Sam M. Werner¹, Jiahua Xu^{2,4}, and Benjamin Livshits^{1,2,3}

¹ Imperial College London, London, UK

daniel.perez@imperial.ac.uk

² University College London, Centre for Blockchain Technologies, London, UK

³ Brave Software, San Francisco, USA

⁴ École polytechnique fédérale de Lausanne, Lausanne, Switzerland

Abstract. The trustless nature of permissionless blockchains renders overcollateralization a key safety component relied upon by decentralized finance (DeFi) protocols. Nonetheless, factors such as price volatility may undermine this mechanism. In order to protect protocols from suffering losses, undercollateralized positions can be *liquidated*. In this paper, we present the first in-depth empirical analysis of liquidations on protocols for loanable funds (PLFs). We examine Compound, one of the most widely used PLFs, for a period starting from its conception to September 2020. We analyze participants' behavior and risk-appetite in particular, to elucidate recent developments in the dynamics of the protocol. Furthermore, we assess how this has changed with a modification in Compound's incentive structure and show that variations of only 3% in an asset's dollar price can result in over 10 m USD becoming liquidable. To further understand the implications of this, we investigate the efficiency of liquidators. We find that liquidators' efficiency has improved significantly over time, with currently over 70% of liquidable positions being immediately liquidated. Lastly, we provide a discussion on how a false sense of security fostered by a misconception of the stability of non-custodial stablecoins, increases the overall liquidation risk faced by Compound participants.

1 Introduction

Decentralized Finance (DeFi) refers to a peer-to-peer, permissionless blockchain-based ecosystem that utilizes the integrity of smart contracts for the advancement and disintermediation of traditional financial primitives [25]. One of the most prominent DeFi applications on the Ethereum blockchain [27] are protocols for loanable funds (PLFs) [13]. On PLFs, markets for loanable funds are established via smart contracts that facilitate borrowing and lending [28]. In the absence of strong identities on Ethereum, creditor protection tends to be ensured through overcollateralization, whereby a borrower must provide collateral worth more than the value of the borrowed amount. In the case where the value of the collateral-to-borrow ratio drops below some liquidation threshold, a borrower defaults on his position and the supplied collateral is sold off at a discount to cover the debt in a process referred to as *liquidation*. However, little is known

about the behavior of agents towards liquidation risk on a PLF. Furthermore, despite liquidators playing a critical role in the DeFi ecosystem, the efficiency with which they liquidate positions has not yet been thoroughly analyzed.

In this paper, we first lay out a framework for quantifying the state of a generic PLF and its markets over time. We subsequently instantiate this framework to all markets on Compound [17], one of the largest PLFs in terms of locked funds. We analyze how liquidation risk has changed over time, specifically after the launch of Compound’s governance token. Furthermore, we seek to quantify this liquidation risk through a price sensitivity analysis. In a discussion, we elaborate on how the interdependence of different DeFi protocols can result in agent behavior undermining the assumptions of the protocols’ incentive structures.

Contributions. This paper makes the following contributions:

- We present an abstract framework to reason about the state of PLFs.
- We provide an open-source implementation¹ of the proposed framework for Compound, one of the largest PLFs in terms of total locked funds.
- We perform an empirical analysis on the historical data for Compound, from May 7, 2019 to September 6, 2020 and make the following observations: (i) despite increases in the number of suppliers and borrowers, the total funds locked are mostly accounted for by a small subset of participants; (ii) the introduction of Compound’s governance token had protocol-wide implications as liquidation risk increased in consequence of higher risk-seeking behavior of participants; (iii) liquidators became significantly more efficient over time, liquidating over 70% of liquidable positions instantly.
- Using our findings, we demonstrate how interaction between protocols’ incentive structures can directly result in unexpected risks to participants.

2 Background

In this section we introduce preliminary concepts about blockchains and smart contracts necessary to the understanding of the rest of the paper.

2.1 Blockchain

A blockchain, such as Bitcoin [19] or Ethereum [27], is in essence a decentralized append-only database. Data is added to the blockchain in the form of transactions that are grouped in blocks. Some rules are enforced by the protocols on both transactions and blocks to ensure its correct working. Blockchains need to be able to maintain consensus of which blocks are included. Both Bitcoin and Ethereum use the Proof-of-Work consensus that requires block producers, often called *miners*, to solve a computationally expensive puzzle to produce a new block [21]. An important point to note is that miners are allowed to choose which transactions to include in a block and in which order to include them. This can potentially allow miners to profit from having a transaction included before another one. This is commonly referred to as *miner-extractable value* [8].

¹ <https://github.com/backdfund/analyzer>.

2.2 Smart Contracts

Ethereum Smart Contracts. On Ethereum, smart contracts are programs written in a Turing-complete language, typically Solidity [11], that define a set of rules that may be invoked by any network participant. These programs rely on the Ethereum Virtual Machine (EVM), a low-level stack machine which executes the compiled EVM bytecode of a smart contract [27]. Each instruction has a fee measured in so-called *gas*, and the total gas cost of a transaction is a fixed base fee plus the sum of all instructions' gas [2, 20]. The sender of a transaction must then set a gas price, the amount of ETH he is willing to pay per unit of gas consumed for executing the transaction. The transaction fee is thus given by the gas price multiplied with the gas cost [22, 26]. Within a transaction, smart contracts can store data in logs, which are metadata specially indexed as part of the transaction. This metadata, commonly referred to as *events*, is typically used to allow users to monitor the activity of a contract externally.

Oracles. One of the major challenges smart contracts face concerns access to off-chain information, i.e. data that does not natively exist on-chain. Oracles are data feeds into smart contracts and provide a mechanism for accessing off-chain information through some third party. In DeFi, oracles are commonly used for price feed data to determine the real-time price of assets. For instance, via the Compound Open Price Feed [6], vetted third party reporters sign off on price data using a known public key, where the resulting feed can be relied upon by smart contracts.

Stablecoins. An alternative to volatile cryptoassets is given by stablecoins, which are priced against a peg and can be either custodial or non-custodial. For custodial stablecoins (e.g. USDC [4]), tokens represent a claim of some off-chain reserve asset, such as fiat currency, which has been entrusted to a custodian. Non-custodial stablecoins (e.g. DAI [18]) seek to establish price stability via economic mechanisms specified by smart contracts. For a thorough discussion on stablecoin design, we direct the reader to [15].

3 Protocols for Loanable Funds (PLF)

In this section, we introduce several concepts of Protocols for Loanable Funds (PLFs) necessary for understanding how liquidations function in DeFi on Ethereum.

3.1 Supplying and Borrowing in DeFi

In DeFi, asset supplying and borrowing is achieved via so-called *protocols for loanable funds* (PLFs) [13], where smart contracts act as trustless intermediaries of loanable funds between suppliers and borrowers in markets of different assets. Unlike traditional peer-to-peer lending, deposits are pooled and instantly available to borrowers. On a DeFi protocol, the aggregate of tokens that the PLF smart contracts hold, which equals the difference between supplied funds and borrowed funds, is termed locked funds [9].

3.2 Interest Model

Borrowers are charged interest on the debt at a floating rate determined by a market's underlying interest rate model. A small fraction of the paid interest is allocated to a pool of reserves, which is set aside in case of market illiquidity, while the remainder is paid out to suppliers of loanable funds. Interest in a given market is generally accrued through market-specific, interest-bearing derivative tokens that appreciate against the underlying asset over time. Hence, a supplier of funds receives derivative tokens in exchange for supplied liquidity, representing his share in the total value of the liquidity pool for the underlying asset. The most prominent PLFs are Compound [5] and Aave [1], with 2.5bn USD and 2.7bn USD in total funds locked respectively, at the time of writing [9].

3.3 Collateralization

Given the pseudonymity of agents in Ethereum, borrow positions need to be overcollateralized to reduce the default risk. Thereby, the borrower of an asset is required to supply collateral, where the total value of the supplied collateral exceeds the total value of the borrowed asset. Each asset is associated with a collateralization ratio, namely the minimum collateral-to-borrow ratio when the asset is used to collateralize a new borrow position. For example, in order to borrow 100 USD worth of DAI with ETH as collateral at a collateralization ratio of 125%, a borrower would have to lock 125 USD worth of ETH to collateralize the borrow position. Thus, the protocol limits monetary risk from defaulted borrow positions, as the underlying collateral of a defaulted position can be sold off to recover the debt. The inverse of the collateralization ratio is referred to as the *collateral factor*, which is the amount of a deposit that may be used as collateral. For example, if the collateralization ratio on a PLF for the market of DAI is 125%, the collateral factor would be 0.8, implying that for each \$1 deposit of DAI, the supplier may borrow \$0.8 worth of some other asset.

3.4 Liquidation

The process of selling a borrower's collateral to recover the debt value upon default is referred to as *liquidation*. A borrow position can be liquidated once the value of the collateral falls below some pre-determined liquidation threshold, i.e. the minimum acceptable collateral-to-borrow ratio. Any network participant may liquidate these positions by paying the debt asset to acquire the underlying collateral at a discount. Hence, liquidators are incentivized to actively monitor others' collateral-to-borrow ratios. Note that in practice, the amount of liquidable collateral that a single liquidator can purchase may be capped.

3.5 Leveraging

In finance, leverage refers to borrowed funds being used as the funding source for additional, typically more risky capital. In DeFi, leverage is the fundamental

component of PLFs, as a borrower is required to first take up the role of a supplier and deposit funds which are to be used as leverage for his borrow positions, as we have just seen. The typical aim of leveraging is to generate higher returns through increased exposure to a particular investment. For example, a borrower wanting to gain increased exposure to ETH may:

1. Supply ETH on a PLF.
2. Leverage the deposited ETH to borrow DAI.
3. Sell the purchased DAI for ETH.
4. Repeat steps 1 to 3 as desired.

This behavior essentially enables users to construct so-called *leveraging spirals*, whereby a user repeatedly re-supplies borrowed funds in order to get increased exposure to some cryptoasset. However, increased exposure comes at the cost of higher downside risk, i.e., the risk of the value of the leveraged asset or borrowed asset to decrease due to changing market conditions.

3.6 Use Cases of PLFs

We present the different incentives² an agent may have for borrowing from and/or supplying to a PLF:

Interest Suppliers of funds are incentivized by interest which accrues on a per block basis.

Leveraged long position To take on a long position of an asset refers to purchasing an asset with the expectation that it will appreciate in value. These positions can be taken on a PLF by leveraging the asset on which the long position shall be taken.

Leveraged short position A short position refers to borrowing funds of an asset, which one believes will depreciate in value. Consequently, the taker of a short position sells the borrowed asset, only to repurchase it and pay back the borrower once the price has fallen, while profiting from the price change of the shorted asset. This can be achieved by taking on a leveraged borrow position of a stablecoin, where the locked collateral is the asset to short.

Liquidity mining As a means to attract liquidity, PLFs may distribute governance tokens to their liquidity providers. The way these tokens are distributed depends on the PLF. For instance, on Compound, the governance token COMP³ is distributed among users across markets proportionally to the total dollar value of funds borrowed and supplied. This directly incentivizes users to mine liquidity in a market through leveraging in order to receive a larger share of governance tokens. For example, a supplier of funds in market *A* can borrow against his position additional funds of *A*, at the cost of paying the difference between the earned and paid interest. The incentive for pursuing this

² Note that leverage on a PLF in DeFi may in part be motivated by tax benefits, as certain jurisdictions may not tax capital gains on borrowed funds. However, a detailed analysis of this lies outside the scope of this paper.

³ Contract address: [0xc00e94cb662c3520282e6f5717214004a7f26888](https://etherscan.io/address/0xc00e94cb662c3520282e6f5717214004a7f26888).

behaviour exists if the reward (i.e. the governance token) exceeds the cost of borrowing.

Token utility An agent may be able to obtain a token from a PLF which has some desired utility. For example, for the case of governance tokens, the desired token utility could be the right to participate in protocol governance or a claim on protocol earnings.

4 Methodology

In this section, we describe our methodology for the different analyses we perform with regard to leveraging on a PLF. To be able to quantify the extent of leveraged positions over time, we first introduce a state transition framework for tracking the supply and borrow positions across all markets on a given PLF. We then describe how we instantiate this framework on the Compound protocol using on-chain events data.

4.1 Definitions

Throughout the paper, we use the following definitions in the context of PLFs:

Market A smart contract acting as the intermediary of loanable funds for a particular cryptoasset, where users supply and borrow funds.

Supply Funds deposited to a market that can be loaned out to other users and used as collateral against depositors' own borrow positions.

Borrow Funds loaned out to users of a market.

Collateral Funds available to back a user's aggregate borrow positions.

Locked funds Funds remaining in the PLF smart contracts, equal to the difference between supplied and borrowed funds.

Supplier A user who deposits funds to a market.

Borrower A user who borrows funds from a market. Since a borrow position must be collateralized by deposited funds, a borrower must also be a supplier.

Liquidator A user who purchases a borrower's supply in a market when the borrower's collateral-to-borrow ratio falls below some threshold.

4.2 States on a PLF

In this section, we provide a formal definition of the state of a PLF. We denote \mathfrak{P}_t as the global state of a PLF at time t . For brevity, in the following definitions, we assume that all the values are at a given time t . We define the global state for the PLF as

$$\mathfrak{P} = (\mathcal{M}, \Gamma, \mathcal{P}, \Lambda)$$

where \mathcal{M} is the set of states of individual markets, Γ is the price the Oracle used, \mathcal{P} is the set of states of individual participants and $\Lambda \in (0, 1)$ is the close factor of the protocol, which specifies the upper bound on the amount of collateral a liquidator may purchase.

We define the state of an individual market $m \in \mathcal{M}$ as

$$m = (\mathcal{I}, \mathcal{B}, \mathcal{S}, \mathcal{C})$$

where \mathcal{I} is the market's interest rate model, \mathcal{B} is the total borrows, \mathcal{S} is the total supply of deposits, and \mathcal{C} is the collateral factor.

\mathcal{P}^m is the state of all participants in market m and the positions of a participant P in this market is defined as

$$P^m = (B^m, S^m)$$

where B^m and S^m are respectively the total borrow positions and total supplied deposits of a market participant in market m .

For a given market m , the total deposits supplied \mathcal{S}^m is thus given by:

$$\mathcal{S}^m = \sum_{P^m \in \mathcal{P}^m} S^m \quad (1)$$

Similarly, the market's total borrows \mathcal{B}^m is given by:

$$\mathcal{B}^m = \sum_{P^m \in \mathcal{P}^m} B^m \quad (2)$$

The state of a participant P is liquidable if the following holds:

$$\frac{\sum_{m \in \mathcal{M}} \{ [S^m \cdot \mathcal{C} + \mathcal{I}(S^m)] \cdot \Gamma(m) \cdot \mathcal{K}^m \}}{\sum_{m \in \mathcal{M}} \{ [B^m + \mathcal{I}(B^m)] \cdot \Gamma(m) \}} < 1 \quad (3)$$

where $\Gamma(m)$ returns the price of the underlying asset denominated in a predefined numéraire (e.g. USD), $\mathcal{I}(S^m)$ returns the interest earned with supply S^m , $\mathcal{I}(B^m)$ returns the interest accrued with borrow B^m , and $\mathcal{K}^m \leq 1$ denotes the liquidation threshold of market m . In Compound, liquidation threshold \mathcal{K}^m is set to be constant at 100% protocol-wide, whereas with other protocols such as Aave, \mathcal{K}^m is specific to the collateral asset from market m , and can be dynamically adjusted when the risk level of the asset changes.

The transition from a state of a market m from time t to $t + 1$ is given by some state transition σ , such that $m_t \xrightarrow{\sigma} m_{t+1}$.

4.3 Leveraging Spirals on a PLF

Here we examine the workings of leveraging in DeFi using a PLF. We assume a speculator on some volatile asset B , holds initial capital α in B . In order to increase his exposure to B , the speculator may borrow a stable asset A against his α on a PLF at a collateralization ratio $\delta > 1$. For simplicity, we shall assume in this illustrative example that a speculator will leverage his position on the same PLF. Note that the cost of borrowing is given by some floating interest rate γ for the specific asset market. In return for his collateral, the borrower

receives $\frac{\alpha}{\delta}$ in the volatile asset B . As the debt is denominated in units of a stable asset (e.g. DAI), the borrower has an upper limit on his net debt, remaining unaffected by any volatility in the value of asset A . In order to leverage his position, the debt denominated in A may be used to buy⁴ additional units of asset B , which can subsequently be used to collateralize a new borrow position. This process is illustrated in Fig. 1 and can be repeated numerous times, by which the total exposure to asset A , the underlying collateral to the total debt in asset A , increases at a decaying rate.

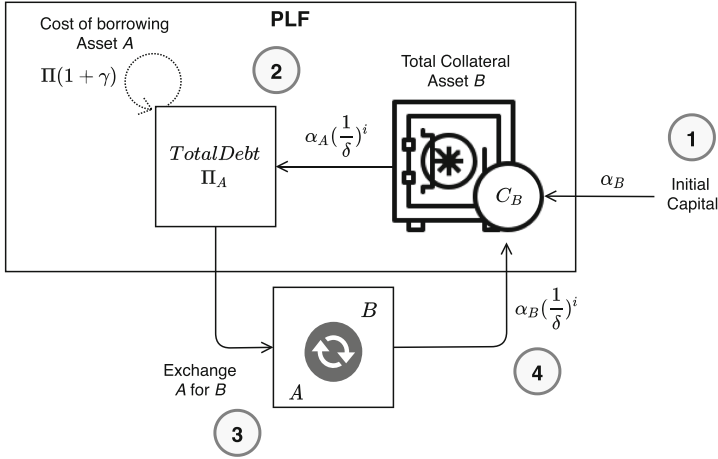


Fig. 1. The steps of leveraging using a PLF. **1.** Initial capital α_B in asset B is deposited as collateral to borrow asset A . **2.** Interest accrues over the debt of the borrow position for asset B . **3.** The borrowed asset A is sold for asset B on the open market. **4.** The newly purchased units of asset B are locked as collateral for a new borrow position of asset A .

The total collateral C a borrower must post through a borrow position with a leverage factor k , a collateralization ratio δ and an initial capital amount α can be expressed as $\sum_{i=0}^k \frac{\alpha}{\delta^i}$. Hence, the total debt Π for the corresponding borrow position is:

$$\Pi = \left(\sum_{i=1}^k \frac{\alpha}{\delta^i} \right) \cdot (1 + \gamma) \tag{4}$$

where γ is the interest rate. Note that Eq. (4) assumes a borrower uses the same collateralization ratio δ for his positions, as well as that all debt is taken out for the same asset on the same PLF and hence the floating interest rate is shared across all borrow positions.

⁴ In practice this may be done via automated market makers [29] (e.g. Uniswap [23]) or via decentralized exchanges [10].

Event	Description	State variables affected
Borrow	A new borrow position is created.	\mathcal{B}
Mint	cTokens are minted for new deposits.	\mathcal{S}
RepayBorrow	A borrow position is partially/fully repaid.	\mathcal{B}
LiquidateBorrow	A borrow position is liquidated.	\mathcal{B}, \mathcal{S}
Redeem	cTokens are used to redeem deposits of the underlying asset.	\mathcal{S}
NewCollateralFactor	The collateral factor for the associated market is updated.	\mathcal{C}
AccrueInterest	Interest has accrued for the associated market and its borrow index is updated.	\mathcal{B}
NewInterestRateModel	The interest rate model for the associated market is updated.	\mathcal{I}
NewInterestParams	The parameters of the interest rate model for the associated market are updated.	\mathcal{I}
NewCloseFactor	The close factor is updated.	\mathcal{A}

Fig. 2. The events emitted by the Compound protocol smart contracts used for initiating state transitions and the states affected by each event.

4.4 States and the Compound PLF

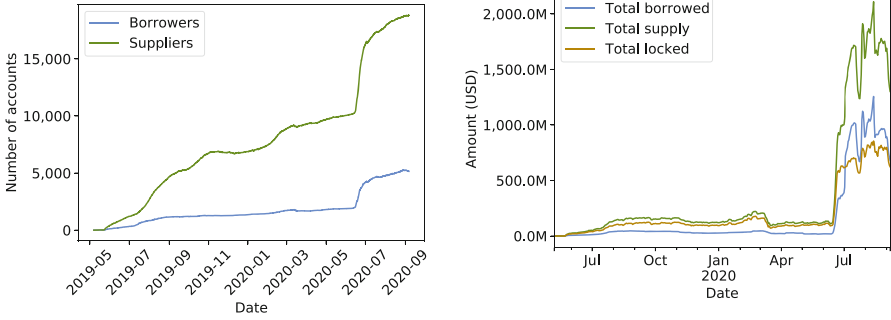
For our analysis, we apply our state transition framework to the Compound PLF. Therefore, we briefly present the workings of Compound in the context of our framework.

State Transitions. We initiate state transitions via events emitted from the Compound protocol smart contracts. We provide an overview of the state variables affected by Compound events in Fig. 2.

Funds Supplied. Every market on Compound has an associated “cToken”, a token that continuously appreciates against the underlying asset as interest accrues. For every deposit in a market, a newly-minted amount of the market’s associated cToken is transferred to the depositor. Therefore, rather than tracking the total amount of the underlying asset supplied, we account the total deposits of an asset supplied by a market participant in the market’s cTokens. Likewise, we account the total supply of deposits in the market in cTokens.

Funds Borrowed. A borrower on Compound must use cTokens as collateral for his borrow position. The borrowing capacity equals the current value of the supply multiplied by the collateral factor for the asset. For example, given an exchange rate of $1 \text{ DAI} = 50 \text{ cDAI}$, a collateral factor of 0.75 for DAI and a price of $1 \text{ DAI} = 1 \text{ USD}$, a holder of 500 cDAI (10 DAI) would be permitted to borrow up to 7.5 USD worth of some other asset on Compound. Therefore, as funds are borrowed, an individual’s total borrow position, as well as the respective market’s total borrows are updated.

Interest. The accrual of interest is tracked per market via a borrow index, which corresponds to the total interest accrued in the market. The borrow index of a market is also used to determine and update the total debt of a borrower in the respective market. When funds are borrowed, the current borrow index for the market is stored with the borrow position. When additional funds are borrowed



(a) Number of suppliers and borrowers. (b) Amount of funds supplied, borrowed and locked.

Fig. 3. Number of active accounts and amount of funds on Compound over time.

or repaid, the latest borrow index is used to compute the difference of accrued interest since the last borrow and added to the total debt.

Liquidation. A borrower on Compound is eligible for liquidation should his total supply of collateral, i.e. the value of the sum of the borrower’s cToken holdings per market, weighted by each market’s collateral factor, be less than the value of the borrower’s aggregate debt (Eq. (3)). The maximum amount of debt a liquidator may pay back in exchange for collateral is specified by the close factor of a market.

5 Analysis

In this section, we present the results of the analysis performed with the framework outlined in Sect. 4. We analyze data from the Compound protocol [17] over a period ranging from May 7, 2019—when the first Compound markets were deployed on the Ethereum main network—to September 6, 2020. The full list of contracts considered for our analysis can be found in Appendix A. When analyzing a single market, we choose the market for DAI, as it is the largest by an order of magnitude.

5.1 Borrowers and Suppliers

We first examine the total number of borrowers and suppliers on Compound by considering any Ethereum account that, at any time within the observation period, either exhibited a non-zero cToken balance or borrowed funds for any Compound market. The change in the number of borrowers and suppliers over time is displayed in Fig. 3a.

We see that the total number of suppliers always exceeds the total number of borrowers. This is because on Compound, one can only borrow against funds




Fig. 4. Cumulative distribution of funds in USD. Accounts are sorted from least to most wealthy and bucketed in bins of 10, i.e. a single bar represents the sum of 10 accounts.

he supplied as collateral, which automatically makes the borrower also a supplier. Interestingly, the number of suppliers has increased relative to the number of borrowers over time. There is notable sudden jump in both the number of suppliers and borrowers in June 2020.

In terms of total deposits, a very similar trend is observable in Fig. 3b, which shows that at the same time, the total supplied deposits increased, while the total borrows followed shortly after. Furthermore, the total funds borrowed exceeded the total funds locked for the first time in July 2020 and have remained so until the end of the examined period. We discuss the reasons behind this in the next part of this section.

Despite the similarly increasing trend for the number of suppliers/borrowers and amount of supplied/borrowed funds, we can see in Fig. 4 that the majority of funds are borrowed and supplied only by a small number of accounts. For instance, for the suppliers in Fig. 4a, the top user and top 10 users supply 27.4% and 49% of total funds, respectively. For the borrowers shown in Fig. 4b, the top user accounts for 37.1%, while the top 10 users account for 59.9% of total borrows. While one could think that this concentration comes from the fact that top accounts are pools receiving money from several participants, only one of the top 10 suppliers and none of the top 10 borrowers fit in this category. We provide a list of the top suppliers and borrowers with a description of the accounts in Figure 10 of Appendix B.

5.2 Leveraging Spirals

As we have seen in Sect. 3, in PLFs, leveraging can be used either to gain more exposure to a particular currency or to gain some incentive provided by the protocol. To understand how leveraging can affect the total amounts borrowed and supplied on Compound, we use the methodology we defined in Sect. 4.3 to measure the existence of leveraging spirals on Compound.

We find that the top supplier deposited a total of 342 million USD and borrowed 247 million. However, after the inspection of leveraging spirals, we find that the user has provided only 16% of the funds, while the rest of the minted funds have been part of leveraging spirals, which means that the user provided a total of roughly 55 million USD to the protocol.

In total, we find a total of 2,141 accounts using this leveraging spiral technique for a total of over 600 million USD, or roughly half of the total amount of funds supplied to the protocol.

5.3 The COMP Governance Token

The sudden jumps exhibited in Figs. 3a and 3b can be explained by the launch of Compound's governance token, COMP, on June 15, 2020. The COMP governance token allows holders to participate in voting, create proposals, as well as delegate voting rights. In order to empower Compound stakeholders, new COMP is minted every block and distributed among borrowers and suppliers in each market.

Initially, COMP was allocated proportionally to the accrued interest per market. However, the COMP distribution model was modified via a governance vote on July 2, 2020, such that the borrowing interest rate was removed as a weighting mechanism in favor of distributing COMP per market on a borrowing demand basis, i.e. per USD borrowed. The distributed COMP per market is shared equally between a market's borrowers and suppliers, who receive COMP proportionally to their borrowed and supplied amounts, respectively. Hence, a Compound user is incentivized to increase his borrow position as long as the borrowing cost does not exceed the value of his COMP earnings. This presumably explains the drop in the degree of collateralization, as the total amount locked is seen surpassed by the total borrows after the COMP launch (Fig. 3b), leading to elevated liquidation risk of borrow positions.

5.4 Liquidation Risk

Given the high increase in the number of total funds borrowed and supplied, as well as the decrease in liquidity relative to total borrows, we seek to identify and quantify any changes in liquidation risk on Compound since the launch of COMP. Figure 5 shows the total USD value of collateral on Compound and how close collateral amounts are from liquidation. In addition to the substantial increase in the total value of collateral on Compound since the launch of COMP, the risk-seeking behavior of users has also changed. This can be seen by examining collateral to borrow ratios, where since beginning of July, 2020, a total of approximately 350 m to 600 m USD worth of collateral has been within a 5% price range of becoming liquidable. However, it should be noted that the likelihood of the amount of this collateral becoming liquidable highly depends on the price volatility of the collateral asset.

In order to examine how liquidation risk differs across markets, we measure for the largest market on Compound, namely DAI, the sensitivity of collateral becoming liquidable given a decrease in the price of DAI. Figure 6 shows the




Fig. 5. Collateral locked over time, showing how close the amounts are from being liquidated. Positions can be liquidated when the ratio drops below 100%.

amount of aggregate collateral liquidable at the historic price, as well as at a 3% and 5% decrease relative to the historic price for DAI. We mark the date on which the COMP governance token launched with a dashed line. It can be seen that since the launch of COMP, 3% and 5% price decreases of DAI relative to its peg USD would have resulted in a substantially higher amount of liquidable collateral. In particular, a 3% decrease would have turned collateral worth in excess of 10 million USD liquidable.




Fig. 6. Sensitivity analysis of the liquidable collateral amount given DAI price movement relative to its peg USD. COMP launch date is marked by the dashed vertical line.

5.5 Liquidations and Liquidators

In order to better understand the implications of the increased liquidation risk since the launch of COMP, we examine historical liquidations on Compound and subsequently measure the efficiency of liquidators.

Historical Liquidations. The increased risk-seeking behavior suggested by the low collateral to borrow ratios presented in the previous section are in accordance with the trend of rising amount of liquidated collateral since the introduction of COMP. The total value of collateral liquidated on Compound over time is shown in Fig. 7. It can be seen that the majority of this collateral was liquidated on a few occasions, perhaps most notably on Black Thursday (March 12, 2020), July 29, 2020 (DAI deviating from its peg significantly), and in early September 2020 (ETH price drop).

Liquidation Efficiency. We measure the efficiency of liquidators as the number of blocks elapsed since a borrow position has become liquidable and the position actually being liquidated. The overall historical efficiency of liquidators is shown as a cumulative distribution function in Fig. 8, from which it can be seen that approximately 60% of the total liquidated collateral (35 million USD) was liquidated within the same block as it became liquidable, suggesting that the majority of liquidations occur via bots in a highly efficient fashion. After 2 blocks have elapsed (on average half a minute), 85% of liquidable collateral has been liquidated, and after 16 blocks this value amounts to 95%.

It is worth noting that liquidation efficiency has been skewed by the more recent liquidation activities which were of a much larger scale than when the protocol was first launched. Specifically, in 2019, only about 26% of the liquidations occurred in the block during which the position became liquidable, compared to 70% in 2020. This resulted in some lost opportunities for liquidators as shown in Fig. 6. The account `0xd062eeb318295a09d4262135ef0092979552afe6`, for instance, had more than 3,000,000 USD worth of ETH as collateral exposed at block 8,796,900 for the duration of a single block: the account was roughly 20 USD shy of the liquidation threshold but eventually escaped liquidation. If a liquidator had captured this opportunity, he could have bought half of this collateral (given the close factor of 0.5), at a 10% discount, resulting in a profit of 150,000 USD for a single transaction. It is clear that with such stakes, participants were incentivized to improve liquidation techniques, resulting in a high level of liquidation speed and scale.




Fig. 7. Amount (in USD) of liquidated collateral from May 2019 to August 2020.

5.6 Summary

In this section, we have analyzed the Compound protocol with a focus on liquidations. We have found that despite the increase in number of suppliers and borrowers over time, the total amount of funds supplied and borrowed remain extremely concentrated among a small set of participants.

We have also seen that the introduction of the COMP governance token has changed how users interact with the protocol and the amount of risk that they are willing to take. Users now borrow vastly more than before, with the total amount borrowed surpassing the total amount locked. Due to excessive borrowing without a sufficiently safe amount of supplied funds, borrow positions now face a higher liquidation risk, such that a crash of 3% in the price of DAI could result in an aggregate liquidation value of over 10 million USD.

Finally, we have shown that the liquidators have become more efficient with time, and are currently able to capture a majority of the liquidable funds instantly.

6 Discussion

In this section, we enumerate several points that we deem important for the future development of PLFs and DeFi protocols. We first discuss the influence of governance tokens, by intention or not, on how users behave within a protocol. Subsequently, we discuss potential risks that lie in the use of governance tokens, and the contagion effect that user behavior in a protocol can have on another protocol. Finally, we discuss how miner-extractable value [8] can potentially affect liquidation incentives in such protocols.




Fig. 8. Number of blocks elapsed from the time a position can be liquidated to actual liquidation on Compound from May 7, 2019 to September 6, 2020, shown as a CDF.

6.1 Governance Token Influence

As analyzed in Sect. 5, the distribution of the COMP token has vastly changed the Compound landscape and user behavior. Until the introduction of the token, borrowing was costly due to the payable interest, which implies a negative cash flow for the borrower. Therefore, a borrower would only borrow if he could justify this negative cash flow with some application external to Compound. With the introduction of this token, borrowing could yield a positive cash flow due to the monetary value of the governance token. This creates a situation where both suppliers and borrowers end up with a positive cash flow, inducing users to maximize both their supply and borrow. This model is, however, only sustainable when the price of the COMP token remains sufficiently high to keep this cash flow positive for borrowers. This directly results in users taking increasingly higher risk in an attempt to gain larger monetary rewards, with liquidators ultimately profiting more from their operations.

6.2 Governance Token Risks

The increased use of governance tokens across DeFi protocols (e.g. YFI on Yearn Finance, AAVE on Aave, UNI on Uniswap) can be seen as a promising step towards achieving a higher degree of decentralization in terms of protocol governance. However, despite the increased usage of governance tokens, to the best of our knowledge there is still a dearth of academic research examining the different governance models and specifically the relation between their security assumptions and the employed governance token. For instance, the option to aggregate governance tokens via flash loans [24] can pose a significant security risk to DeFi protocols should an attacker attempt to propose and execute malicious protocol

updates. Furthermore, even in the case of flash loan resistant governance models, the relationship between the financial value of a protocol’s governance token and the economically secure regions of the protocol remains unexamined and serves as a further risk that designers of governance models have to take into account. Despite the existence of protective mechanisms against governance attacks on some protocols (e.g. multi-sig approvals or selected “guardians” that are able to halt the governance process), it remains questionable which of such mechanisms are indeed desirable from a decentralized governance perspective and whether there might be more suitable alternatives.

6.3 Contagion Effects

This behavior also indirectly affected other protocols, in particular DAI. The price of DAI is aimed to be pegged to 1 USD resting on an arbitrage mechanism, whereby token holders are incentivized to buy or sell DAI as soon as the price moves below or above 1 USD, respectively. However, a rational user seeking to maximize profit will not sell his DAI if holding it somewhere else would yield higher profits. This was precisely what was happening with Compound, whose users locking their DAI received higher yields in the form of COMP, than from selling DAI at a premium, thereby resulting in upward price pressure [7]. Interestingly, DAI deviating from its peg also has a negative effect for Compound users. Indeed, as we saw in Sect. 5, many Compound users might have been overconfident about the price stability of DAI and thus only collateralize marginally above the threshold when they borrow DAI. This has resulted in large amounts being liquidated due to the actual, higher extent of the volatility in the DAI price.

6.4 Miner-Extractable Value

In the context of PLFs, liquidations can be seen as miner-extractable value. Indeed, it is easy for the miner to check whether a position is liquidable or not after each processed transaction and to add a transaction to liquidate the position immediately after the transaction making it liquidable. In our analysis of the Compound protocol, we have not found any sign of miners participating in liquidations, directly or indirectly. We show the top miners and the top liquidators for each miner in Fig. 11 of Appendix C. Although we found no correlation between miners and liquidators, this is a real risk that could make the role of liquidator, which is essential for the protocol security, less interesting for those who are not collaborating with miners.

7 Related Work

In this section we briefly discuss existing work related to this paper.

A thorough analysis of the Compound protocol with respect to market risks faced by participants was done by [14]. The authors employ agent-based modeling and simulation to perform stress tests in order to show that Compound

remains safe under high volatility scenarios and high levels of outstanding debt. Furthermore, the authors demonstrate the potential of Compound to scale to accommodate a larger borrow market while maintaining a low default probability. This differs from our work as we conduct a detailed empirical analysis on Compound, focusing on how agent behavior under different incentive structures on Compound has affected the protocol’s state with regard to liquidation risk.

A first in-depth analysis on PLFs is given by [13]. The authors provide a taxonomy on interest rate models employed by PLFs, while also discussing market liquidity, efficiency and interconnectedness across PLFs. As part of their analysis, the authors examine the cumulative percentage of locked funds solely for the Compound markets DAI, ETH, and USDC.

In [3], the authors provide a formal state transition model of PLFs⁵ and prove fundamental behavioural properties of PLFs, which had previously only been presented informally in the literature. Additionally, the authors examine attack vectors and risks, such as utilization attacks and interest bearing derivative token risk. This work differs to our work, as the authors of [3] formalize the properties of PLFs through an abstract model, while we provide a thorough empirical analysis with a focus on liquidations and risks brought upon by governance tokens, such as for Compound and the COMP token.

In [16], the authors show how markets for stablecoins are exposed to deleveraging feedback effects, which can cause periods of illiquidity during crisis.

The authors of [12] demonstrate how various DeFi lending protocols are subject to different attack vectors such as governance attacks and undercollateralization. In the context of the proposed governance attack, the lending protocol the authors focus on is Maker [18].

8 Conclusion

In this paper, we presented the first in-depth empirical analysis of liquidations on Compound, one of the largest PLFs in terms of total locked funds, from May 7, 2019 to September 6, 2020. We analyzed agents’ behavior and in particular how much risk they are willing to take within the protocol. Furthermore, we assessed how this has changed with the launch of the Compound governance token COMP, where we found that agents take notably higher risks in anticipation of higher earnings. This resulted in variations as little as 3% in an asset’s price being able to turn over 10 million USD worth of collateral liquidable. In order to better understand the potential consequences, we then measured the efficiency of liquidators, namely how quickly new liquidation opportunities are captured. Liquidators’ efficiency was found to have improved significantly over time, reaching 70% of instant liquidations. Lastly, we demonstrated how overconfidence in the price stability of DAI, increased the overall liquidation risk faced by Compound users. Rather ironically, many users wishing to make the most of the new incentive scheme ended up causing higher volatility in DAI—a dominant

⁵ Note that in [3], PLFs are referred to as lending pools.

asset of the platform, resulting in liquidation of their own assets. This is not Compound's misdoing, but rather highlights the to date unknown dynamics of incentive structures across different DeFi protocols.

Appendix

Appendix is available online at <https://fc21.ifca.ai/papers/144.pdf>.

References

1. AAVE: AAVE (2020). <https://aave.com/>. Accessed 17 Aug 2020
2. Albert, E., Correas, J., Gordillo, P., Román-Díez, G., Rubio, A.: GASOL: gas analysis and optimization for ethereum smart contracts. In: TACAS 2020. LNCS, vol. 12079, pp. 118–125. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45237-7_7
3. Bartoletti, M., Chiang, J.H.Y., Lluch-Lafuente, A.: SoK: lending pools in decentralized finance. arXiv preprint [arXiv:2012.13230](https://arxiv.org/abs/2012.13230) (2020)
4. Circle: USDC (2020). <https://www.circle.com/en/usdc>
5. Compound: Compound (2019). <https://compound.finance/>. Accessed 17 Aug 2020
6. Compound: Open price feed (2020). <https://compound.finance/prices>. Accessed 15 Sept 2020
7. Cyrus: Upcoming comp farming change could impact the dai peg (2020). <https://forum.makerdao.com/t/upcoming-comp-farming-change-could-impact-the-dai-peg/2965>. Accessed 27 Aug 2020
8. Daian, P., et al.: Flash boys 2.0: frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In: 2020 IEEE Symposium on Security and Privacy (SP), pp. 910–927. IEEE (2020)
9. DeFi Pulse: How do we calculate total value locked (TVL)? (2020). <https://defipulse.com/>
10. dYdX: dydx (2019). <https://dydx.exchange/>
11. Solidity v0.8.0 documentation (2020). Accessed 12 Jan 2020
12. Gudgeon, L., Perez, D., Harz, D., Livshits, B., Gervais, A.: The decentralized financial crisis. In: Crypto Valley Conference on Blockchain Technology, pp. 1–15 (2020). <https://doi.org/10.1109/CVCBT50464.2020.00005>
13. Gudgeon, L., Werner, S.M., Perez, D., Knottenbelt, W.J.: DeFi protocols for loanable funds: interest rates, liquidity and market efficiency (2020)
14. Kao, H.T., Chitra, T., Chiang, R., Morrow, J.: An Analysis of the Market Risk to Participants in the Compound Protocol (2020). https://scfab.github.io/2020/FAB2020_p5.pdf
15. Klages-Mundt, A., Harz, D., Gudgeon, L., Liu, J.Y., Minca, A.: Stablecoins 2.0: economic foundations and risk-based models. In: 2nd ACM Conference on Advances in Financial Technologies (AFT 2020), New York (2020). <https://doi.org/10.1145/3419614.3423261>
16. Klages-Mundt, A., Minca, A.: (In) stability for the blockchain: Deleveraging spirals and stablecoin attacks. arXiv preprint [arXiv:1906.02152](https://arxiv.org/abs/1906.02152) (2019)
17. Leshner, R., Hayes, G.: Compound: The Money Market Protocol. Technical report (2018)

18. Maker: The maker protocol: MakerDAO's multi-collateral Dai (MCD) system. <https://makerdao.com/en/whitepaper/>. Accessed 08 June 2020
19. Nakamoto, S.: Bitcoin: A Peer-to-Peer Electronic Cash System (2008). www.bitcoin.org
20. Perez, D., Livshits, B.: Broken metre: attacking resource metering in EVM. In: Network and Distributed System Security Symposium. Internet Society, Reston, VA (2020). <https://doi.org/10.14722/ndss.2020.24267>
21. Perez, D., Xu, J., Livshits, B.: Revisiting transactional statistics of high-scalability blockchains. In: ACM Internet Measurement Conference, vol. 16, pp. 535–550. ACM, New York (2020). <https://dl.acm.org/doi/10.1145/3419394.3423628>
22. Pierro, G.A., Rocha, H.: The influence factors on ethereum transaction fees. In: 2019 IEEE/ACM 2nd International Workshop on Emerging Trends in Software Engineering for Blockchain (WETSEB), pp. 24–31. IEEE (2019)
23. Uniswap: Uniswap whitepaper (2020). <https://hackmd.io/@HaydenAdams/HJ9jLsftz#%F0%9F%A6%84-Uniswap-Whitepaper>. Accessed 26 Aug 2020
24. Wang, D., et al.: Towards understanding flash loan and its applications in DeFi ecosystem (2020). <http://arxiv.org/abs/2010.12252>
25. Werner, S.M., Perez, D., Gudgeon, L., Klages-Mundt, A., Harz, D., Knottenbelt, W.J.: SoK: Decentralized Finance (DeFi) (2021). <http://arxiv.org/abs/2101.08778>
26. Werner, S.M., Pritz, P.J., Perez, D.: Step on the gas? A better approach for recommending the ethereum gas price. In: Pardalos, P., Kotsireas, I., Guo, Y., Knottenbelt, W. (eds.) Mathematical Research for Blockchain Economy. SPBE, pp. 161–177. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-53356-4_10
27. Wood, G., et al.: Ethereum: a secure decentralised generalised transaction ledger. Ethereum Project Yellow Paper **151**(2014), 1–32 (2014)
28. Xu, J., Vadgama, N.: From banks to DeFi: the evolution of the lending market (2021)
29. Xu, J., Vavryk, N., Paruch, K., Cousaert, S.: SoK: Decentralized Exchanges (DEX) with Automated Market Maker (AMM) protocols (2021). <http://arxiv.org/abs/2103.12732>