

Systemic Fragility in Decentralized Markets

Alfred Lehar*
Haskayne School of Business
University of Calgary

Christine A. Parlour†
Haas School of Business
UC Berkeley

May 18, 2023

**Preliminary
Comments Welcome**

Abstract

We analyze a unique data set of collateral liquidations on two Decentralized Finance lending platforms – Compound and Aave. Such liquidations require arbitrageurs to repay the loan in return for the discounted collateral. Using Blockchain transaction data, we observe if arbitrageurs liquidate positions out of their own inventory or obtain “flash loans.” To repay flash loans, arbitrageurs immediately sell the collateral asset. We document the high frequency price impact of such liquidity trades on nine different decentralized exchanges. Consistent with large block trades in equity markets there is a temporary and permanent price impact of collateral asset sales in DeFi. We document the effect of these trades on return distributions. Our work highlights the systemic fragility of decentralized markets.

Keywords: Decentralized Lending, Blockchain, Decentralized Finance, Systemic Risk

*Corresponding author, email: alfred.lehar@haskayne.ucalgary.ca, Tel: (403) 220 4567. Alfred Lehar is grateful to the to the Canadian Securities Institute Research Foundation and SSHRC for financial support.

†email:parlour@berkeley.edu. We thank Terrence Hendershott, Andreas Park, Jonathan Payne, Brandon Rochon, Gustavo Schwenkler, Zhen Shi, Kathy Yuan, and seminar participants at American Economic Association 2022, FIELDS-CFI Workshop on Mathematical Finance and Cryptocurrencies 2022, CryptoAssets and Digital Asset Investment Conference. CICF 2022, OFR, NBIM Transparency Conference, NYU Global Fintech Conference, 2022, Tel Aviv Finance Conference 2022, NUS-SKKU FinTech Seminar Series, CSA systemic risk committee, the HKBU-Monash-NTU Joint Online Digital-Economy Seminar, Warwick Business School Gillmore Centre Conference DeFi & Digital Currencies 2022, Financial innovation: A threat to financial stability conference 2022, and QED for comments. Parlour and Lehar are grateful to the Tel Aviv BlockChain Research Institute and NFI for financial support.

Systemic Fragility in Decentralized Markets

Preliminary and incomplete

Abstract

We analyze a unique data set of collateral liquidations on two Decentralized Finance lending platforms – Compound and Aave. Such liquidations require arbitrageurs to repay the loan in return for the discounted collateral. Using Blockchain transaction data, we observe if arbitrageurs liquidate positions out of their own inventory or obtain “flash loans.” To repay flash loans, arbitrageurs immediately sell the collateral asset. We document the high frequency price impact of such liquidity trades on nine different decentralized exchanges. Consistent with large block trades in equity markets there is a temporary and permanent price impact of collateral asset sales in DeFi. We document the effect of these trades on return distributions. Our work highlights the systemic fragility of decentralized markets.

1 Introduction

Collateral is widely used in financial markets to mitigate credit risk exposure for lenders. However, collateral only fulfills this purpose if it can be efficiently liquidated or seized. This is done in different ways. In Repo markets, ownership of the collateral is directly transferred from the borrower to the lender, whereas in many marketplaces, the lender retains capital, and alerts the borrower with a margin call if the value of collateral falls. In both cases, the lender retains and monitors the collateral. In decentralized finance, collateral is monitored and margins are enforced by third parties. In this paper we document how third party liquidations affect protocol risk, collateral risk and risks to the decentralized finance system.

Collateralized lending in decentralized finance is automated and risk is mitigated in two ways. First, borrowers post collateral against debt, and the loan to value ratio of each individual position is publicly observable. If the LTV rises above a threshold, anyone can liquidate the loan. In this way, the credit risk of each individual position is minimized. Second, all loans are issued at high frequency floating rates which increase as capital is withdrawn from the protocol. Increases in the floating rates add to the loan amount which may trigger liquidations. In this way, protocol run risk is reduced, and transferred to the borrower.

Both of these risk mitigations rely on efficient collateral liquidations. In this paper, we collect a unique data set of collateral liquidations on Compound and Aave, two of the largest DeFi lending protocols. Over our sample, approximately \$9 billion of collateral was locked in Compound, and over \$11 billion locked in Aave.¹ We observe liquidations valued at \$2,487,543,624.

Using these data, we document a temporary and permanent price impact of collateral liquidations: deleveraging leads to lower prices. We observe lower prices on both the exchange where the transaction occurs and then subsequently on other exchanges including off-chain markets. This contagion leads to negative feedback loops: loans are liquidated which leads to downward price pressure on collateral and more loans are then liquidated. Second, these trades have a measurable effect on collateral return distributions: liquidated collateral has heavier tails than unliquidated collateral. Finally, we provide evidence consistent with strategic behavior by liquidators. The contagion, negative feedback loops, strategic behavior by liquidators and measurable effects on collateral return distributions highlight a new form of systemic fragility.

Decentralized finance offers a unique laboratory to investigate the immediate effect and subsequent propagation of collateral liquidations. First, there is no mandated maximum on the amount that can be borrowed to invest. Second, the unique nature of blockchain settlement allows flash loans or loans without credit risk that can be used for arbitrage trades. Thus, arbitrage capital is not constrained. Third, the transparent nature of the blockchain makes it possible to track trades at a high frequency and precisely estimate their impact. Finally, the mechanics of decentralized exchanges allow us to precisely estimate what the price would have been had arbitrageurs not traded to return the price to its equilibrium value.

Next to Decentralized Exchanges, most capital in decentralized finance (Defi) is allocated to collateralized lending protocols. Users can post collateral in one token and take out a loan

¹Approximate figures are available from defipulse.com

in another token. One common use case is to build a levered position in Ether (ETH), the native crypto-currency of the Ethereum blockchain, by posting ETH as collateral, borrowing a USD stablecoin, and then trading the USD for more ETH. In DeFi lending, users interact with a system of smart contracts, computer code – often open source – that is deployed on a blockchain. The smart contracts hold collateral in escrow, approve loans, collect interest, and, most importantly for our study, have a mechanism in place to ensure that the loan is adequately collateralized.

Most lending platforms require collateral to be between 1.2 to 1.5 times the amount borrowed. As soon as the value of the collateral falls below this threshold, the loan is eligible for liquidation. While lending platforms differ in the actual liquidation process, they nonetheless are structured in broadly the same way. To ensure competition, and prompt liquidation, any user can buy the collateral at a discounted price and use the proceeds to repay the loan. Much of this market is automated, and trading algorithms (bots) often called keepers implement these trades. Since the collateral is sold to liquidators at a discount relative to current market prices, liquidators earn a profit which compensates them for their transaction costs and provides an incentive for swift liquidations. The actions of these keepers are instrumental in ensuring the stability and resilience of the lending protocols and eliminating credit risk.

There is a limited but rapidly growing literature on decentralized finance. Various recent papers investigate the properties of decentralized exchanges. These include theory contributions due to Angeris and Chitra (2020), Angeris, Kao, Chiang, and Noyes (2019), Park (2021) and Aoyagi (2020), which characterize automated market maker mechanics and information transmission. Recent empirical contributions by Capponi and Jia (2021), Barbon and Ranaldo (2021) and Lehar and Parlour (2021b). All of these papers note the importance of gas fees.

There is a long literature in Finance that explores the effect of large trades on markets. The seminal paper of Kraus and Stoll (1972), find that block trades on the NYSE lead to permanent price effects that they attribute as recompense for liquidity provision. By contrast, Holthausen, Leftwich, and Mayers (1990) examine the impact of large block trades on the NYSE and find that liquidity effects are reversed after a few trades. We note that in the DeFi swap markets, the liquidity providers are not recompensed for large trades – these benefits accrue to arbitrageurs. The further implication in our context is that there is systemic fragility as liquidations lead to price changes which mechanically trigger further liquidations through oracle updating.

Parallel literatures in economics and finance have considered the effect of leverage on asset prices, returns and risk. In a series of papers, Geanakoplos and Fostel (2015) illustrate how leverage can increase asset prices in incomplete markets. Intuitively, agents with high valuations for an asset will borrow against future claims and so increase the price. The implication is that deleveraging will have permanent price effects. In the finance literature, Gromb and Vayanos (2002) show that if arbitrageurs are financially constrained, prices of assets may diverge even for long periods of time. Similarly, Brunnermeier and Pedersen (2009) show that traders' ability to provide liquidity depends on their capital. We note that in decentralized finance, arbitrage capital is never constrained because of the existence of flash loans.

2 Decentralized Lending

In our analysis, we focus on two DeFi lending platforms, Aave and Compound, both of which are structured in a similar way. These protocols match borrowers and lenders in specific asset pairs or pools. While they are economically similar to banks, they operate as platforms and so do not retain any intermediation risk.

Lenders supply assets that are pooled and then lent out to borrowers. The rate that each lender receives (and borrower pays) is calculated block by block as a function of the ratio of funds lent and borrowed (“the utilization rate”) and a constant. This floating rate ensures that the protocol is not subject to run risk. As lenders withdraw funds, the utilization rate and thus the rate paid by the borrowers increases. This provides an incentive for borrowers to either close out their loan or provides an incentive for liquidators to do it for them. The implication of this high frequency floating rate is that unlike intermediaries, the protocols do not face liquidity transformation risk, rather it is transferred to the borrowers.

In a decentralized system, without the benefit of reputation or identity, lending is collateralized.² Many different tokens are accepted as collateral, but each token differs in the required overcollateralization. The trading price of each token fluctuates and if the relative value of the borrowed token rises sufficiently, the position can be liquidated. The protocols rely on so-called liquidators to monitor the positions and sell the underlying collateral. Liquidators are typically traders who deploy algorithms or ‘trading bots’ that monitor all the collateralized positions. In principle, any Ethereum address may invoke a liquidation function, however in practice this is a specialized activity. We note that expertise is more likely to be the constraint rather than capital because of the existence of flash loans.

Liquidation is profitable because a fraction or all of the borrowed amount can be repaid in return for the collateral at the current market price minus a liquidation discount. In other words, the liquidator receives the collateral at a discounted price.

Figure 1 illustrates how a liquidator repays a loan on Aave and seizes collateral. First, the liquidator has to repay the loan. In this case, the borrower has a debt of 1 WBTC and has posted 11 WETH as collateral. The liquidator may either repay the debt out of inventory or obtain it through a flash transaction. On obtaining the 1 WBTC, the liquidator seizes the collateral. She may either keep this in inventory or swap it out on a decentralized exchange (DEX).

To illustrate the mechanics of a liquidation we detail one transaction from block 14759771, that occurred on May 12, 2022 at 7:19 UTC.³ The liquidation was undertaken by a bot.⁴ The liquidator partially repaid an outstanding loan by returning USDT 39,330.04 to the Aave V2 lending pool. In return, the liquidator obtained collateral of 22.80 ETH. The liquidator then swapped ETH 21.50 into USDT 39,330.04 on Sushiswap.

²While overcollateralization is mostly observed, undercollateralization is possible however in these cases the protocol retains control of the lent assets.

³Transaction 0x41eddc70253a40cea41587aab1c46c057a1c7247b9aefd1e79917dd00c6b4715

⁴The address is 0xabcf5d4be599f1c7f71fcbcae4643a2aa849f4c8

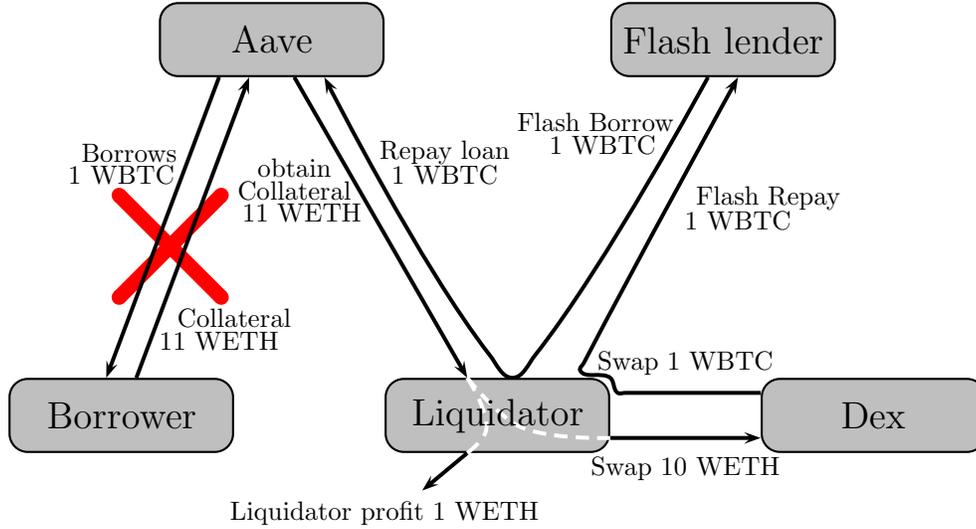


Figure 1. Anatomy of a Liquidation

As the liquidator both repaid USD_T and swapped the collateral for the same amount of USD_T, she had no change in her USD_T position. She made a profit on ETH of 22.80-21.50=1.30 before fees.

The liquidator paid a gas fee of ETH 0.85 to the miner to process the transaction, leaving a net profit of ETH 0.45 or approximately USD 878.76 at the time. On the same day, the same liquidator, liquidated 42 other loans.

Loans are eligible for liquidation on Aave based on a “health factor.” On Aave V2, a health factor H_f is calculated for each wallet. Consider a wallet that has borrowed D (denominated in ETH) against collateral assets C_i $i = 1, N$ also denominated in ETH. Each distinct collateral asset i has a specific liquidation threshold ℓ_i that reflect liquidity risk, volatility etc. The health factor of a position is calculated as

$$H_f = \frac{\sum_{i=1}^N C_i \ell_i}{D}.$$

Any loan with a health factor below 1 can be liquidated. Figure 2 shows the health factor as defined by the Aave lending protocol on a block per block basis around the time of our example liquidation (normalized to block zero).⁵

As is evident from Figure 2 the loan is liquidated as soon as the health factor approaches 1. By reducing the borrower’s position, the liquidation causes his health factor to increase sufficiently so that the remaining loan is adequately collateralized.

Accurate information on prices is crucial to efficient liquidation. Information available on-chain is

⁵We obtain this data by querying the Aave smart contract using an Ethereum archive node to ensure that we have the correct pricing oracles.

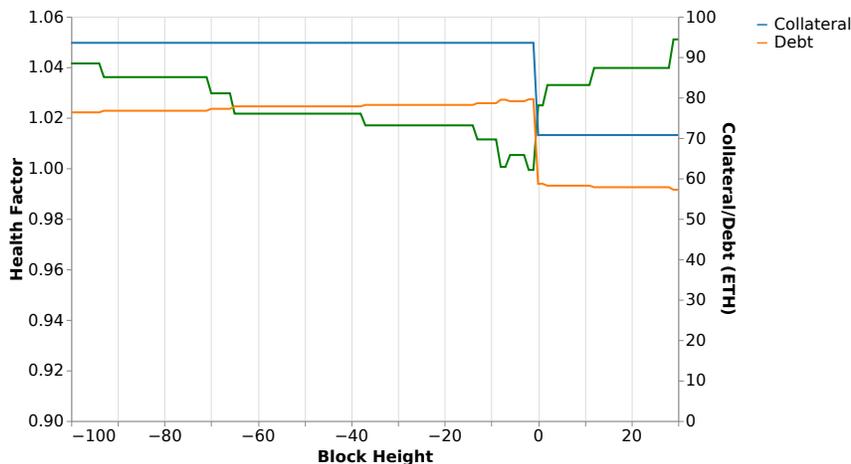


Figure 2. Health factor, debt and collateral around a liquidation. The graph shows the amount borrowed (orange) and the amount of collateral (orange) denominated in ETH as well as the health factor as defined by the Aave protocol (green) in relative blocks around the liquidation.

provided through oracles. Typically these aggregate information across various on-chain sources. To prevent manipulation, the exact mapping between on chain data and the oracle price is not published, however they are based on Dex prices. We note that the liquidation trigger is only a function of public information. Thus, these liquidations are purely liquidity trades and have no informational content.

3 Data and stylized facts

We collect data on collateral liquidations from two of the largest Defi lending protocols, Aave and Compound. We collect data from UniSwap and its most important clones: SushiSwap, ShibaSwap, ZKSwap, SakeSwap, DefiSwap, CitySwap, BTSwap and Equalizer. (For readers unfamiliar with these markets, we present a description in Appendix C.) Interactions with the smart contracts of these decentralized exchanges generate entries on blockchains that run the Ethereum virtual machine. These entries are then stored in the individual blocks that constitute the blockchain. We record the amount and token of the loan as well as the amount and token of the collateral that was liquidated. Tokens are recorded based on the address of the smart contract that governs the token. Using the API from Etherscan.io we identify the name and ticker for each token and the conditions on the trading venue.

We note that our data do not comprise all the liquidations and subsequent sales. First, there could be non-transparent exchanges (i.e., dark trading venues). Second, we do not record information from exchanges such as Bancor and Balancer. In total we observe 42,324 liquidations from September 25, 2018 to May 16, 2022 comprising 27,466 liquidations on Aave and 14,858 liquidations on Compound.

There is no natural numeraire asset in DeFi, as the protocols are international and any assets can be traded against any other. Thus, our data comprise liquidations of 37 distinct collateral tokens. In Table 1 we present the number of liquidations for the top ten collateral tokens. We present values in both USD and ETH. We convert the liquidated collateral to ETH and USD using block by block pricing data from decentralized exchanges such as Uniswap V2 and Sushiswap. Price availability reduces our sample to 38,414 liquidations.⁶ Notice that wrapped ETH, Link and wrapped Bitcoin are the leading collateral tokens.⁷

	Collateral Token	Number Liquidations	Amount USD	Amount ETH
WETH	Wrapped Ether	19,487	1,391,529,875	866,635
WBTC	Wrapped BTC	3,355	402,121,684	186,899
LINK	ChainLink Token	5,583	190,542,689	111,386
USDC	USD Coin	1,502	146,915,160	177,891
DAI	Dai Stablecoin	1,270	101,185,822	167,325
YFI	yearn.finance	611	53,146,479	36,517
AAVE	Aave Token	1,070	47,900,694	23,735
UNI	Uniswap	1,419	34,346,097	19,017
xSushi	SushiBar	286	14,714,561	5,715
COMP	Compound	371	14,437,255	5,991

Table 1. Ten largest collateral tokens sorted by amount liquidated in USD. *Number Liquidations* is the number of liquidation events, *Amount USD* is the sum of collateral liquidated in USD, *Amount ETH* is the sum of collateral liquidated in ETH.

The collateral exhibited in Table 1 was used to issue debt in various 41 debt tokens. The top debt tokens are USD stablecoins. Table 2 matches the collateral against the borrowed tokens. The top four token pairs users borrowed stablecoins against WETH, which is consistent with the widespread belief that lending platforms are used to build levered positions in crypto-currencies such as ETH or Bitcoin.

In total we observe liquidations of USD 2,487,543,624 with a mean liquidation size of USD 64,756 and a median of USD 3,586. The largest loan liquidation in our sample was the liquidation of USD 50,508,256 worth of DAI collateral on Compound on Nov 26, 2020.⁸

4 Liquidations

We are interested in the aggregate effect of liquidations. Figure 3 illustrates a day on which a large number of collateralized debt obligations were liquidated. There was a liquidation “wave.” Specifically, on May 19th, 2021 loans collateralized by the Chain Link network token (LINK)

⁶The reduction is mainly due to changes in token versions (for example imBTC upgraded the token smart contract during our sample period resulting in a new contract address) and due to some liquidations being observed before the deployment of Uniswap V2. These liquidations are in general small.

⁷Cryptos that are not native to the Ethereum blockchain (e.g. BTC) or are not a token (e.g. ETH) are wrapped so that smart contracts can handle them using a standard token interface, called ERC-20.

⁸See transaction 0x53e09adb77d1e3ea593c933a85bd4472371e03da12e3fec853b5bc7fac50f3e4.

	Collateral	Debt Token		Num. Liq.	Amount USD	Amount ETH
WETH	Wrapped Ether	USDC	USD Coin	6,287	518,773,703	267,019
WETH	Wrapped Ether	USDT	Tether USD	3,952	398,853,641	180,786
WETH	Wrapped Ether	DAI	Dai Stablecoin	5,078	333,387,258	306,990
WBTC	Wrapped BTC	USDC	USD Coin	1,211	162,090,711	69,783
WBTC	Wrapped BTC	USDT	Tether USD	697	124,277,693	51,375
LINK	ChainLink Token	USDC	USD Coin	2,351	85,400,856	49,354
WBTC	Wrapped BTC	DAI	Dai Stablecoin	797	58,693,780	33,513
LINK	ChainLink Token	USDT	Tether USD	1,283	52,097,818	28,281
WETH	Wrapped Ether	WBTC	Wrapped BTC	129	49,183,070	30,326
USDC	USD Coin	USDT	Tether USD	120	39,887,094	18,931

Table 2. Ten largest collateral and debt tokens pairs sorted by amount liquidated in USD. *Num Liq* is the number of liquidation events, *Amount USD* is the sum of collateral liquidated in USD, *Amount ETH* is the sum of collateral liquidated in ETH.

were liquidated on various DEXs, with approximately 80% on SushiSwap which at the time had the deepest pool. The gray area shows the aggregate amount of loan liquidations. As is evident from the graph, the liquidity trades affected first SushiSwap and then the other DEXs and even a centralized exchange (Binance). This price pattern illustrates loan liquidation contagion.

As we mentioned previously, the DEX structure allows us to precisely calculate the price impact of each individual trade. In this figure, we identify the trades that are due to collateral liquidations and plot their cumulative price impact – this is the red line. The difference between the red line and the realized prices reflects the important countervailing effect of arbitrageur trades. As we argue, given the closed information system of the blockchain their incentive to do so is important to mitigate systemic fragility.

Of course, as is evident from the Figure the arbitrageurs are neither immediate nor do they completely reverse the liquidity trades. It is interesting to observe that in the first part of the liquidation wave, the drop in prices was mostly driven by the sale of the collateral on DEXs (the red line coincides with the other lines). The big drop in price and the permanent component was thus driven by the sales of the collateral on decentralized exchanges.

This price drop not only spills over among all the decentralized exchanges, it also affects prices on centralized exchanges such as Binance. We emphasize that Binance is off-chain and this demonstrates that the price effect of liquidations spill over to more traditional markets.

In the middle of the liquidation wave, the red line separates from the other prices, which indicates that at this point arbitrageurs step in and trade against the liquidators and push the price back up, although not to the same level that it was before the liquidation wave. This trading pattern results in a higher probability of extreme outcomes and distinctive return properties which we examine carefully in Section 4.1.

Liquidation waves are common. Figure 4 shows the weekly amount of liquidations in USD for our sample period. The red line corresponds to the average price of ETH over the time period. The day with the largest amount of liquidations was May 19, 2021 when 2,007 loans were liquidated

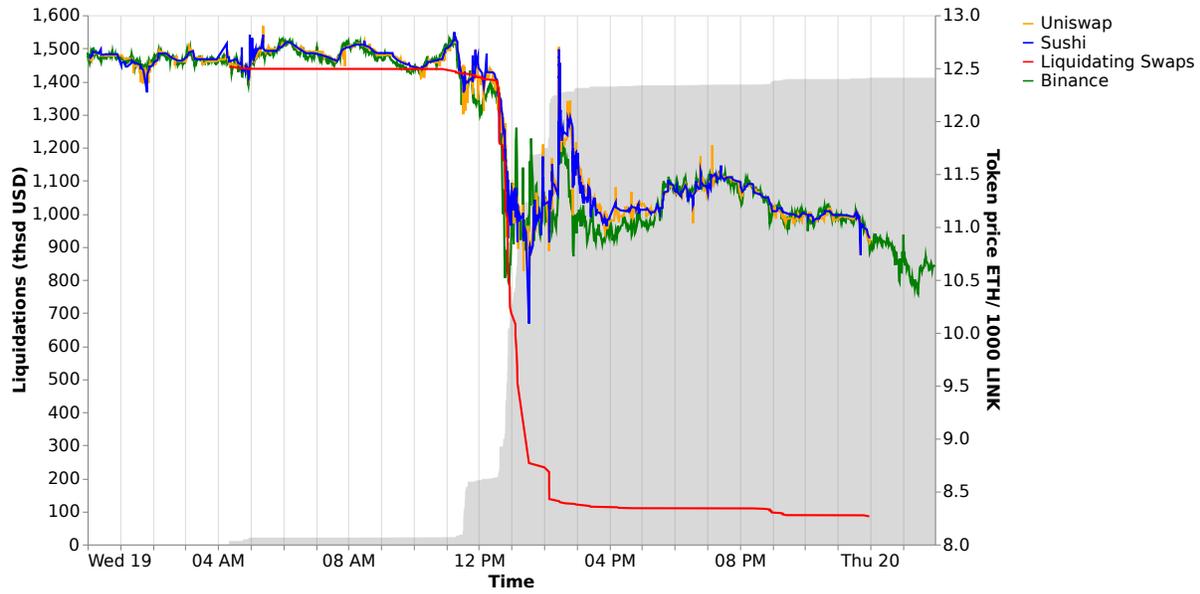


Figure 3. Cumulative return (blue) and cumulative return from loan liquidations (red) for the ETH/LINK exchange rate on May 19th, 2021.

with a total collateral value of USD 503.57 million. On that day ETH dropped from over USD 3,400 to USD 2,014, a 41% decline.

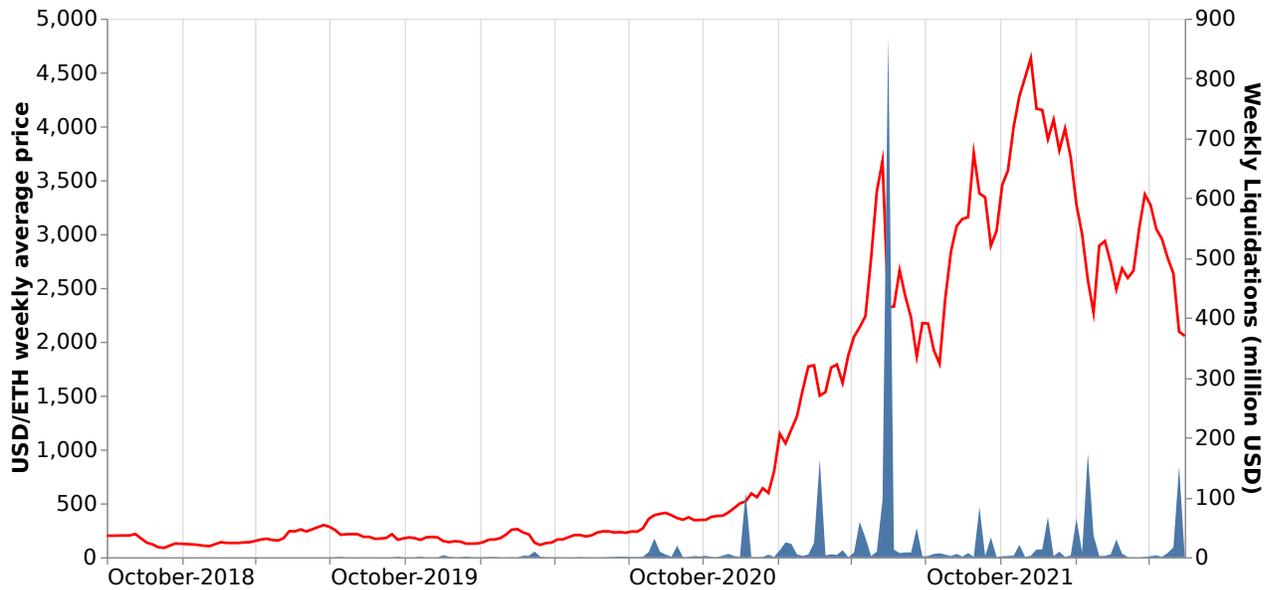


Figure 4. Weekly liquidations in USD depicted in blue (right hand axis). Average Price of Eth in red (left hand axis).

We define a liquidation to be part of a wave if it occurs less one hour after a previous liquidation of the same collateral token. We find 1,028 waves that involve at least 5 liquidations each. In these waves a total of 27,239 loans are liquidated. A total of 19,710 liquidations occur within waves of at least 20 liquidations. In the biggest wave 1,056 loans were liquidated. The average wave with at least 5 liquidations lasts 1.64 hours. A comprehensive examination of cryptocurrency returns appears in Liu, Tsvinski, and Wu (forthcoming) who document momentum at low frequencies. Thus, liquidation waves could reflect a prior increase in the relative value of the collateral asset that led to a cluster of vaults with similar liquidation thresholds.

4.1 Price Impact of Collateral Liquidations

Collateral liquidations are large liquidity trades, and we should expect a high frequency price impact. This price impact should only be observed on the exchange on which the collateral is liquidated. In the face of a price dislocation on a Dex, arbitrage bots should reverse the trade so that the price reflects the market value of collateral.

To make our investigation of liquidations concrete, we present one liquidation in our sample.⁹ On February 23rd 2021 a liquidator used SushiSwap to trade 12,841.22 ETH for 385.36 WBTC and used the latter to repay an undercollateralized loan. The liquidator then seized the collateral of 14,343.93 Eth, worth over USD 20 million, from Compound.

The Sushi-pool that the liquidator used was deep and had, before the trade, an inventory of 9,353.94 WBTC and 297,957.06 WETH. Because of this liquidation the price in this pool moved from 31.39 WBTC/1000 WETH to 28.85 WBTC/1000 WETH or by 8.08%. Figure 5 shows the price of WBTC per 1,000 WETH around the liquidation event. The price drop caused by the liquidator’s token swap is clearly visible and arbitrageurs brought the price partially back to its fundamental value. In spite of this activity, the trade had a permanent price impact on all exchanges after the liquidation. From 10 blocks after the liquidation to 100 blocks after the liquidation, the average price was 30.64, a 2.38% decrease over the price before the liquidation.

In this example, arbitrageurs partially reversed the trade. In spite of this, the liquidation affected medium term token prices and spilled over to other exchanges. Or, the markets demonstrated high frequency contagion.

Our broader sample constitutes all liquidations in which a swap was used to exchange the collateral for the debt asset in the same transaction where the liquidation was recorded. Notice, this includes both liquidations powered by flash loans and liquidations in which another asset was swapped for the debt asset in order to recover the collateral.

Price impact depends on the depth of the liquidity pools. We find that in the average liquidation wave, which includes many single loan liquidations, 1.08% of the available liquidity on Uniswap-type DEXs gets liquidated. For large liquidation waves with at least 100 liquidations on average 10.93% of available liquidity reserves get liquidated which results in an average price impact of 18.73%. We first investigate the effect of loan liquidations on subsequent high frequency prices.

⁹see transaction 0xd70b42daec5bb9ac6e5df3d25d309f186db50df701f667e1f20b22448ea27d41

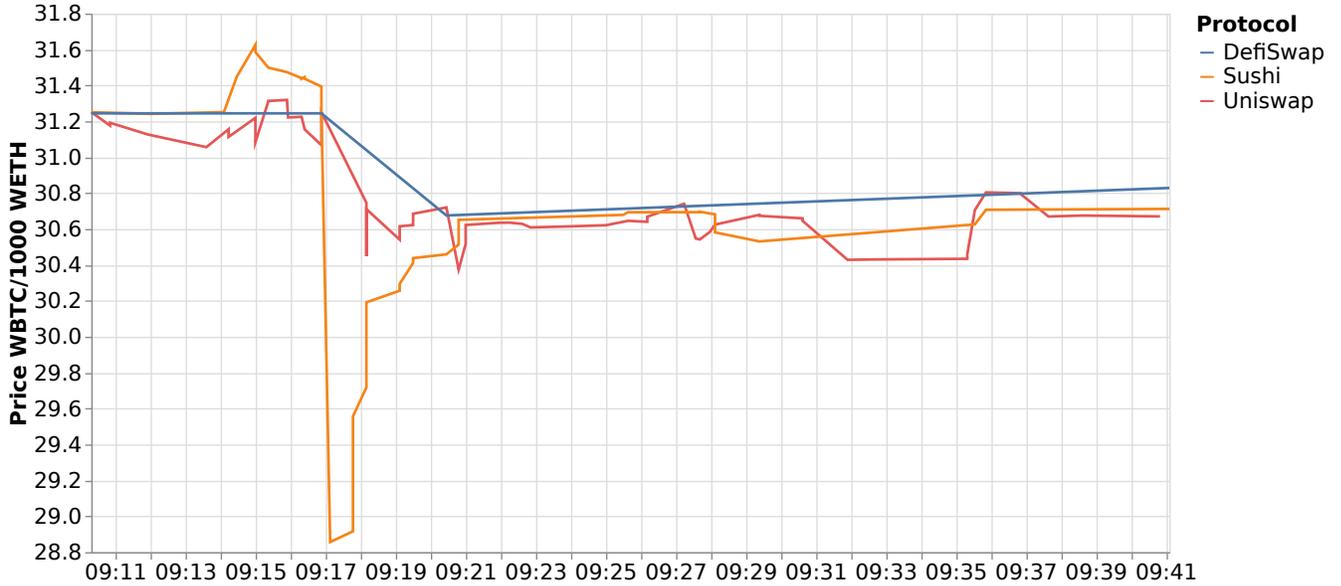


Figure 5. Price reactions to liquidation. Prices of the WETH/WBTC exchange rate on three decentralized exchanges, DefiSwap, Uniswap, and Sushiswap. A liquidator seized over USD 20 million of WBTC from collateral and swapped them immediately for WETH at SuishSwap, diving down the price. The graph illustrates spillover effects to other markets. The graph shows Dex prices from 30 blocks before the liquidation to 100 blocks after the liquidation.

Let $r_5(t)$ denote the 5 block return of each debt assets. Here, t is the block in which a liquidation occurs. We extract the latest recorded price from the Dex on which the liquidation occurs. We determine the price of the last transaction on that Dex five blocks later. The return is calculated from these two prices. We choose a five block window as this is sufficient for arbitrageurs to bring prices back to equilibrium after the liquidation, similar to the effect to the quick reversal in the price on Sushiswap in Figure 5. In addition, we calculate r_t^ℓ , which is the return generated by the liquidation event. Specifically, if a liquidation occurs in block t , we record the Dex price before the liquidation and the Dex price after the liquidation. (Recall, that balances in decentralized exchanges change after each trade and can thus change multiple times within blocks. The exchange rate of a Dex is defined as the ratio of balance of one token over the balance of the other token.) The return is based on these two prices.

Our regression considers the extent to which $r_5(t)$ can be explained by r_t^ℓ and is presented in Table 3. Our control variables include the gas fee associated with the liquidation, and also the wave length in hours and the position (between $[0, 1]$) of the liquidation within the wave. These variables capture a measure of congestion on the blockchain. Columns (3) and (4) present the findings for the exchanges where the collateral was actually liquidated. We find that 38.7% of the price movement of swaps that liquidate collateral persist for the medium term. This finding is important with respect to future loan liquidations. When the liquidation of collateral has a lasting impact on prices, such a liquidation will cause other loans that use the same collateral to be under-collateralized and thus subject to liquidation as well. In columns (1) and (2) we

present results for exchanges that trade the same token pair but which were not involved in the liquidation. We observe strong contagion effects. Selling collateral on one exchange affects prices on other exchanges in the same direction. We find that the price drops upon liquidations are stronger in waves that are shorter and towards the end of a wave.

	Other exchanges		Dex where liquidated	
Return of liquidating Swap	1.302*** (0.273)	1.081*** (0.236)	0.389*** (0.0718)	0.387*** (0.0717)
Gas Price	-1.87e-16* (1.00e-16)	3.93e-17 (7.88e-17)	1.73e-16*** (3.73e-17)	1.62e-16*** (3.84e-17)
Wave Length		-0.00714*** (0.00196)		-0.000118*** (0.0000232)
Position in Wave		0.00760 (0.00955)		0.00395*** (0.000222)
R ²	0.000362	0.00246	0.0916	0.0972
Observations	38,812	38,812	7,786	7,786

Table 3. Regression explaining the return of the debt token/collateral token return around loan liquidations. *LiqCollateral* is the value of the liquidated collateral in million USD. *Wave Size* is the aggregate amount of collateral liquidated in the wave in million USD. *Wave Length* is the length of the wave in hours, *Liquidator Size* is log of the sum of all collateral (in USD) that a specific liquidator has liquidated, and *Gas Price* is the gas price offered on the liquidation transaction in Twei (1 million twei is 1 ETH). Standard errors are clustered by liquidator. One, two, and three stars indicate significance at the 10%, 5%, and 1% level, respectively.

From Table 3 we see that liquidations where the collateral gets immediately swapped have a medium term price impact on all exchanges, whether collateral gets sold on that exchange or not. This feature has implications for systemic stability of the decentralized system. This is because prices across various Dexs are aggregated and used as a price oracle to determine the collateralization of other loans.

Any price changes in the value of collateral across multiple Dexs will cause more loans to be undercollateralized and lead to more liquidations, potentially leading to a liquidation wave. Figure 6 illustrates the feedback effect in the informationally closed blockchain system.

We have demonstrated that liquidations have a price effect both locally on the Dex where the swap occurs and then spread to other Dexs. The price impact of large trades is consistent with other financial markets. The economic difference in decentralized protocols is that, by construction, the blockchain is a closed system which means that information generated on the blockchain is used for other protocols. In particular, the price oracles that inform lending platforms on the value of collateral depend on the prices generated by the Dexs. Thus, there is a natural feedback loop between liquidations and further liquidations.

There are two natural countervailing forces to the feedback channel presented in Figure 6. First, as we have observed, if the price on a Dex is dislocated, arbitrageurs have an incentive to trade against the liquidation. We note that if arbitrageurs are also liquidators, these incentives become less clear. The second countervailing force is that of gas fees. Each execution of the

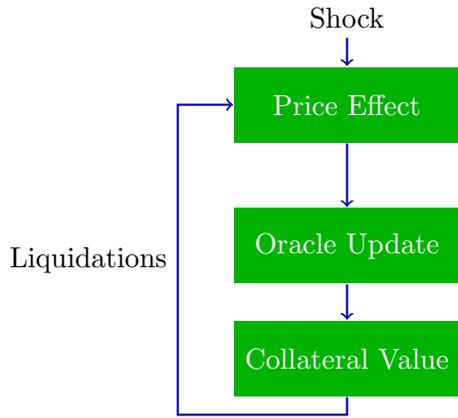


Figure 6. Systemic Fragility Channel

EVM requires a pre-specified amount of gas. In addition, an incentive amount of gas for miners can be added to a transaction. Figure 7 show the average daily gas price in USD for a simple swap. Of course more complex transactions or swaps will require higher fees. Gas fees may affect the stability of the DeFi system in two ways. First, a higher gas fee due to high demand for transaction services ensures that liquidators will require higher payoffs to liquidate positions. This may lead to few liquidations. Second, substantially higher anticipated higher gas fees will provide an incentive for arbitrageurs to trade earlier rather than wait. This will dampen price feedback effects.

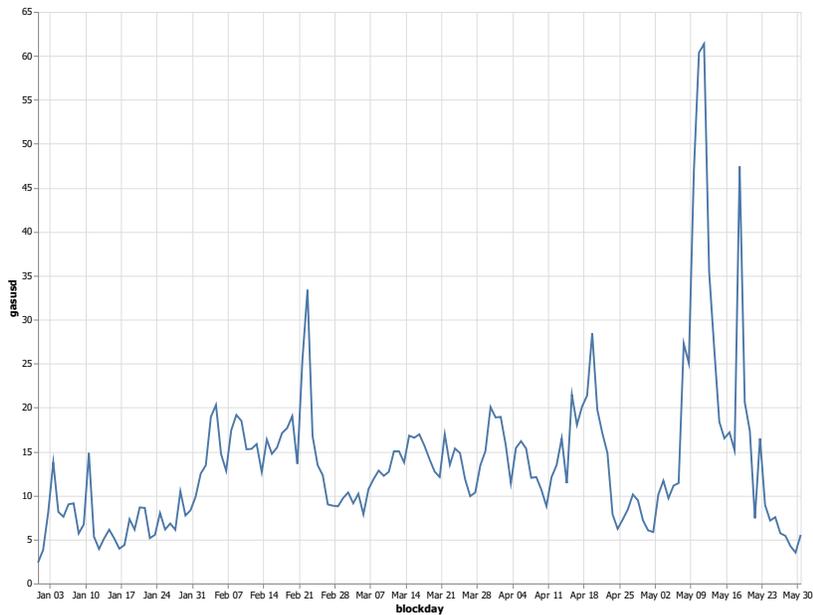


Figure 7. Average Daily Gas fees for a Swap. This calculation is based on a simple swap (50,000 gas units).

In addition to the spillover or contagion effects across multiple exchanges at a high frequency we also document a low frequency price impact of collateral liquidations. Figure 8 illustrates the cumulative return of collateral relative to the start of a liquidation wave. That deleveraging leads to lower prices is consistent with the incomplete markets lending literature.

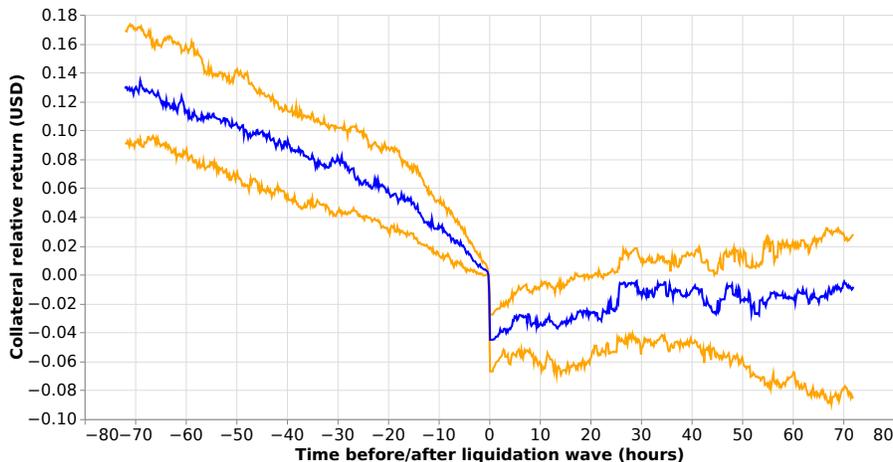


Figure 8. Long term price impact of a liquidation waves. The graph shows the cumulative relative return of the collateral token in USD before a liquidation wave (negative time) and after a liquidation wave (positive time) for a sample of 227 liquidation waves with at least 20 individual liquidations. The blue line is the median return and the orange lines represent the 33% and 66% quantiles of the return distribution.

5 Collateral returns and liquidation

To investigate how the liquidation process affect collateral returns, we collect 2,048,565 5-minute returns from 36 decentralized exchanges for all 16 tokens that serve as collateral at one of the two lending platforms in our sample. We label returns where a position in this specific token was liquidated within a 5 minute interval as the ‘liquidation return’ and contrast these with all other returns for all other tokens.

We choose 5 minute intervals for two reasons: (i) they are sufficiently long in block time. Blocks on Ethereum are generated on average every 15 seconds. If the swap of the liquidated collateral was a pure idiosyncratic event then an arbitrageur has plenty of time to reverse the trade and bring the token price back to its fundamental value. (ii) 5 minute intervals are sufficiently short to separate the effect of liquidations from fundamental movements in token prices. If there is a fundamental reason that makes the price of a token drop over a day we expect to have enough 5 minute observations in our sample to cover both, intervals with and without liquidations.

Figure 9 shows distribution functions for liquidation and other returns. More extreme returns are more likely in intervals where a position in the same token was liquidated. For example, the probability of a return with absolute value greater than 1% is 7.85% when there was a

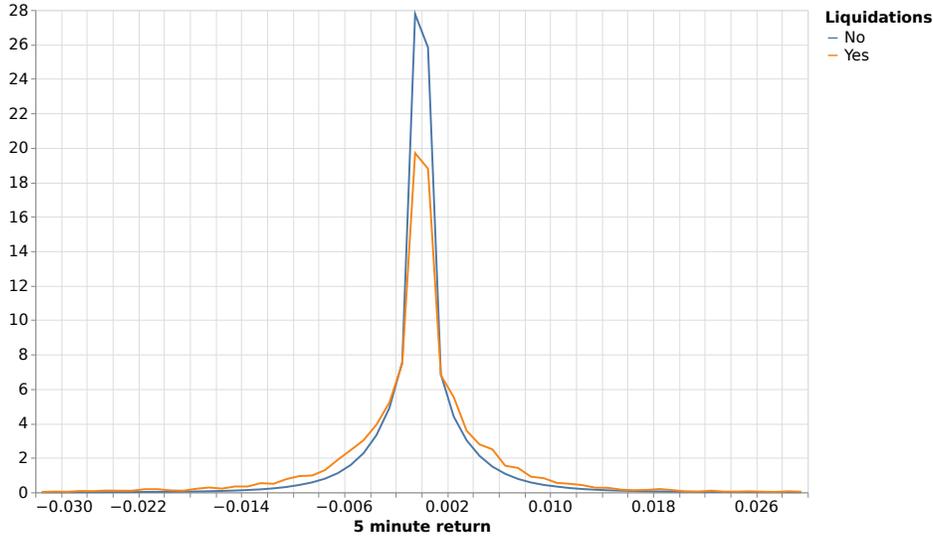


Figure 9. Return distribution for 16 tokens that serve as collateral over 5 minute intervals that coincide with liquidations and ones that do not. For this graph we include two five minute returns before and after liquidations to control for characteristics around liquidations.

liquidation versus 3.25% for cases without a liquidation. Using a Kolmogorov-Smirnov test, we reject that both distributions are equal with a p-value less than 0.000.

To ensure that our findings are not driven by characteristics unique to the time of liquidations, such as changes in the fundamental value of the token, we construct a subsample of both liquidation returns and two 5-period returns before and after the liquidation event. Thus, all returns are observed at the same time and our findings cannot be driven by effects unique to the time period. Our sample is reduced to 47,404 observations.

Figure 10 shows the difference in density function between liquidation returns and other returns. We find again that more extreme returns are more likely to occur with liquidations whereas small returns close to zero are more likely when there is no liquidation.

To examine the impact of swaps on price drops in liquidation waves we compare 1,154 liquidation waves to 1,516 single loan liquidations. The liquidation waves contain on average 11.4 liquidations. For each wave we compute the return of the collateral token from one block before the start of the wave to one block after the end of the liquidation wave. We regress this return on the aggregate return that was caused by all swaps that happened in the same transaction as the liquidations. Our results can be found in Table 4.

For single liquidations (columns (3) and (4)) the price impact of liquidating swaps are inconsequential. This idea is consistent with the fact that arbitrageurs push the price back to its fundamental value after a liquidator’s trade. For liquidation waves (columns (1) and (2)) we find that the liquidation of the collateral on decentralized exchanges makes a significant contribution to the overall price movement of the collateral token throughout the liquidation wave. This finding is consistent with a feedback effect in loan liquidations.

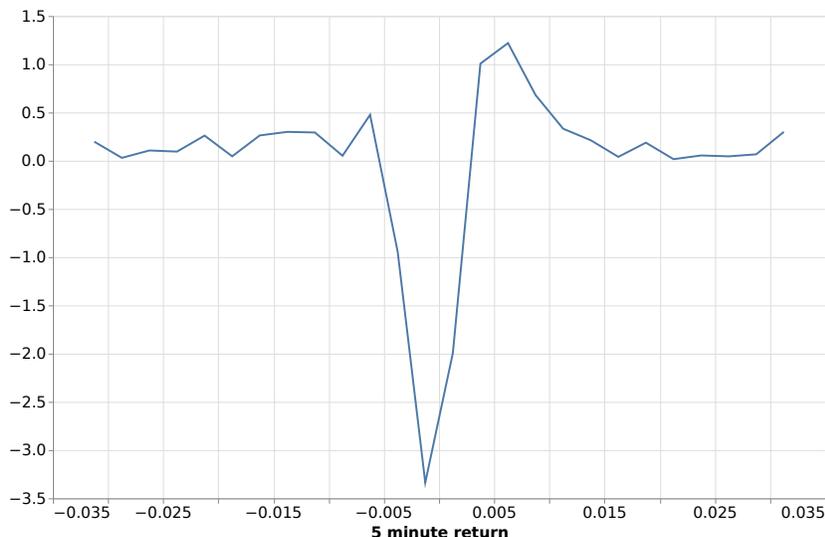


Figure 10. Return distribution for 16 tokens that serve as collateral over 5 minute intervals that coincide with liquidations and ones that do not.

6 Liquidators and Liquidation Incentives

Anecdotal evidence suggests that liquidations are automated and in particular executed by trading bots.¹⁰ We observe 1,007 distinct liquidators, or more precisely liquidator address. Any liquidator may control multiple addresses, so the distinct number of addresses corresponds to an upper bound on the number of liquidators.

We rank liquidators by liquidated collateral amount and present this in Figure 11. Observe that liquidation activity is very concentrated with the top 20 liquidators performing 48.50% of the liquidations and liquidating 75.01% of the collateral. The top liquidator in our sample liquidated 2,732 loans with a total collateral value of USD 427,381,388. We also note from this figure that liquidators concentrate on particular protocols. While similar, different protocols will have slightly different configurations which require different programs to monitor and execute liquidations. There is thus an incentive for specialization.

In order to receive the collateral, the liquidator must repay the loan. This can be done either from available capital, or in the form of a flash loan. Flash loans are uncollateralized and the borrowing and the repayment of the loan happen within the same Ethereum transaction, i.e. at the same instant of time. Because Ethereum transactions are atomic, i.e. they either get executed in whole or not at all, there is no credit risk for the lender because the release of funds to the borrower is conditional on the repayment to the lender within the same transaction.

Upon repaying the debt to the lending platform, the liquidator receives the collateral at a discounted price. She can then choose to keep the collateral token which exposes her to price

¹⁰The Github repository at <https://github.com/haydenshively/Compound-Liquidation-Bot> provides solidity script necessary to program a bot.

	Multiple liquidations		Single liquidations	
	(1)	(2)	(3)	(4)
Return of liquidating Swaps	0.0306*** (0.00261)	0.0306*** (0.00271)	0.00137 (0.00313)	0.00108 (0.00313)
Wave Size		0.0000671 (0.0000686)		-0.00200*** (0.000611)
Wave Length		-0.000169 (0.000425)		
Gas Price		0.000129 (0.000473)		0.0000362 (0.0000823)
R ²	0.106	0.107	0.000126	0.00713
Observations	1,154	1,154	1,516	1,516

Table 4. Regression explaining the return of the debt token/collateral token return through-out a liquidation wave. *Liquidation Returns* is the return that is directly attributable to collateral sales on decentralized exchanges. *Wave Size* is the aggregate amount of collateral liquidated in the wave in million USD. *Wave Length* is the length of the wave in hours, and *Gas Price* is the gas price offered on the liquidation transaction in Twei (1 million twei is 1 ETH). One, two, and three stars indicate significance at the 10%, 5%, and 1% level, respectively.

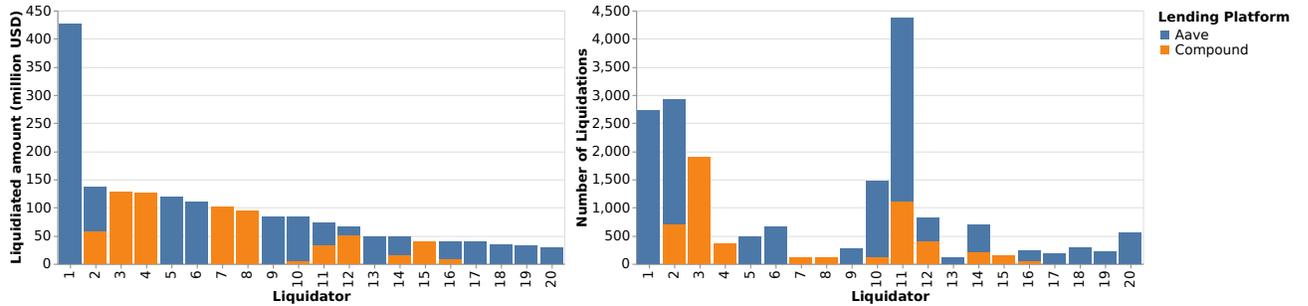


Figure 11. Amounts liquidated by top 20 liquidators, by platform (left panel) and number of loans liquidated by the top 20 liquidators (defined as the liquidators with the largest amount of loans liquidated (right panel)). Each column represents a distinct address, while the dollar volume (left panel) and number (right panel) in Aave is depicted in blue and in Compound is depicted in orange.

risk or she can immediately swap the collateral token on a decentralized exchange. Liquidators that use flash loans typically swap the collateral for the debt token as they have to repay the flash loan.

To get a better understanding when of when swaps are used in liquidations, in Table 5 we regress a dummy that is set to one for all liquidations for which the collateral of a liquidation is immediately swapped on a Dex on liquidation and liquidator specific variables. The results indicate that the collateral of larger liquidations is more likely to be swapped. As alluded to before, this could either be to reduce the exchange rate exposure of the liquidator or because swapping

allows the use of flash loans to overcome capital constraints. We find that swaps often occur in liquidation waves that are large and short. Larger liquidators, probably the more successful bots, are more likely to use swaps. We control for gas fees as swaps incur additional execution cost (gas) for which the liquidator has to pay.

	(1)	(2)	(3)	(4)
Liq.Collateral	0.350*** (0.122)			0.289*** (0.110)
Wave Size		0.00657** (0.00262)		0.00582** (0.00252)
Wave Length		-0.0488*** (0.0152)		-0.0514*** (0.0148)
Liquidator Size			0.0376 (0.0474)	0.0357 (0.0469)
Gas Price			-0.00392 (0.0161)	-0.0157 (0.0161)
R ²				
Observations	38,409	38,409	38,403	38,403

Table 5. Probit regression explaining the use of swaps in loan liquidations. *LiqCollateral* is the value of the liquidated collateral in million USD. *Wave Size* is the aggregate amount of collateral liquidated in the wave in million USD. *Wave Length* is the length of the wave in hours, *Liquidator Size* is log of the sum of all collateral (in USD) that a specific liquidator has liquidated, and *Gas Price* is the gas price offered on the liquidation transaction in Twei (1 million twei is 1 ETH). Standard errors are clustered by liquidator. One, two, and three stars indicate significance at the 10%, 5%, and 1% level, respectively.

6.1 Predatory liquidations

There is anecdotal evidence consistent with predatory liquidations. While this is not surprising for profit maximizing traders, it stands in contrast with the model in which lenders monitor collateral and do not have a profit incentive to liquidate the collateral.

Compound uses a price feed from Coinbase. On November 26, 2020 an attacker pushed the price of DAI (a crypto collateralized stablecoin) up to \$1.30. Compound positions with DAI as a collateral asset appeared undercollateralized. Subsequently, USD 130 million in loans were liquidated.

Note that when a loan is close to the liquidation boundary, a liquidator benefits if it can push the loan over the edge. On May 19, 2021, 147 loans with WBTC as collateral were liquidated amounting to approximately USD 64 million. We trace trades of liquidators 420 blocks before start of liquidation wave

Figure 12 illustrate how the position of the liquidator bots changed around the time of the liquidations. First, as evinced by the blue line the liquidators shorted half a million dollars worth of the collateral asset. This lead to a price drop (orange line). The vaults and collateral

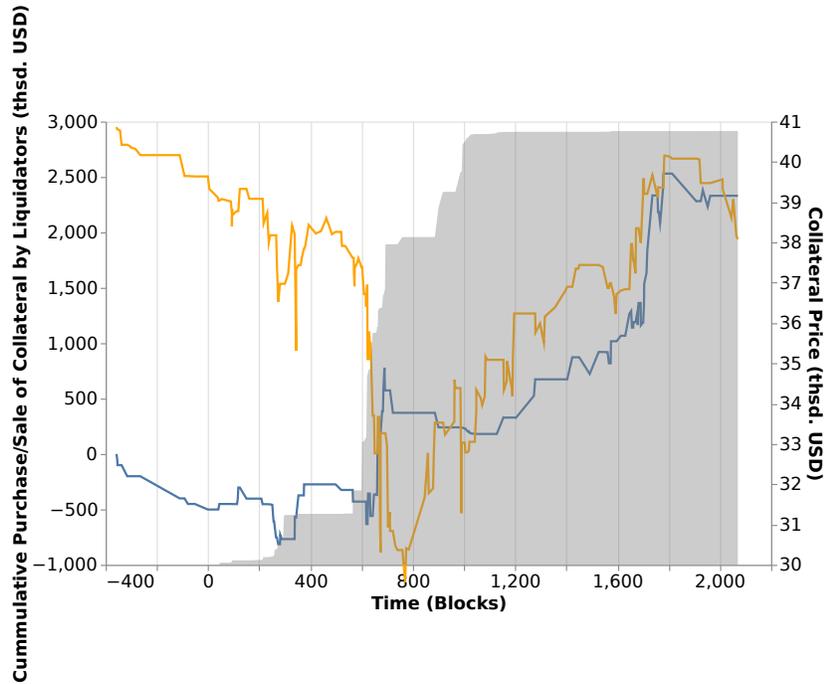


Figure 12. Cumulative liquidated collateral is in grey. Blue (left axis) shows cumulative position in the collateral token by liquidators. Orange: collateral token price in USD (right axis)

were then seized and the liquidation wave unfolded.

We analyze a broader sample of 43 large liquidation waves with at least 100 liquidations each. Our sample contains 11,096 liquidations which were liquidated by 155 liquidators.¹¹ We examine these wallets’ trading activity in the collateral token over a 500 block period before the start of the wave, specifically we collect all token transfers, which is a conservative measure of trading activity.¹² For the median liquidation wave liquidators transfer collateral tokens worth 21.80% of the total liquidated amount out of their wallets before the wave starts. We see liquidators giving away collateral tokens in 67.44% of the liquidation waves.

7 Conclusion

Decentralized lending has innovated to mitigate risk. First, by designing a floating rate that increases in the scarcity of lenders to mitigate the chance of “bank runs.” Second by providing incentives to third parties to monitor and liquidate collateral.

¹¹For this analysis we look at wallets that are liquidators in the strict meaning of the protocol, i.e. the wallet addresses that receive the collateral, as well as all wallet addresses that initiate the transactions in which the liquidation call is made as these addresses control the liquidator bots.

¹²We cannot record, for example, trading in derivatives or off-chain positions

We have demonstrated that these liquidity trades can have market wide persistent effects, or in other words we document contagion between decentralized markets. Given that DeFi is designed to be a closed information system it suggests that the system features a systemic fragility. Liquidations engender other liquidations.

We note that transaction fees associated with using the EVM (gas fees) have a nuanced effect on this fragility. On one hand, anticipated future fees will encourage arbitrageurs to trade more rapidly which will quickly reverse the price impact of liquidations. On the other hand, higher fees raise the threshold for liquidators to finance their transactions through flash loans, which restricts the amount of capital available to monitor loans and renders the liquidator market less competitive.

Providing incentives to third parties to liquidate positions is effective in that profit maximizing liquidators carefully monitor positions. This reduces credit risk. However, they also have an incentive to maximize the number of positions that they liquidate, which as we documented can lead to market-wide permanent price impacts. The incentive to liquidate multiple positions is not there for traditional intermediaries who only retain collateral as insurance against bad states. Traditional intermediaries therefore have more exposure to credit risk, but have less incentive for widespread liquidations.

Currently, FTX US has a proposal before the CFTC for “auto-liquidations.” Different from the decentralized lending protocols, these liquidations would be managed directly by FTX, but would immediately liquidate collateral through limit orders.

References

- Angeris, Guillermo, and Tarun Chitra, 2020, Improved Price Oracles: Constant Function Market Makers, *Working Paper*.
- Angeris, Guillermo, Hsien-Tang Kao, Rei Chiang, and Charlie Noyes, 2019, An Analysis of Uniswap markets, *Working Paper*.
- Aoyagi, Jun, 2020, Liquidity Provision by Automated Market Makers, *Working Paper*.
- Barbon, Andreas, and Angelo Ranaldo, 2021, On The quality of Cryptocurrency Markets, *working paper*.
- Brunnermeier, Markus K., and Lasse Heje Pedersen, 2009, Market Liquidity and Funding Liquidity, *Review of Financial Studies* 22, 2201–2238 Market liquidity and the funding of traders are mutually reinforcing, giving rise to "liquidity phenomena" like fragility, commonality and flight to quality.
- Capponi, Agostino, and Ruizhe Jia, 2021, The Adoption of Blockchain-based Decentralized Exchanges, *working paper*.
- Geanakoplos, John, and Ana Fostel, 2015, Leverage and default in binomial economies: a complete characterization, *Econometrica* 83, 2192–2229.
- Gromb, Denis, and Dimitri Vayanos, 2002, Equilibrium and welfare in markets with financially constrained arbitrageurs, *Journal of Financial Economics* 66, 361–407 Limits on Arbitrage.
- Holthausen, Robert W., Richard W. Leftwich, and David Mayers, 1990, Large-block transactions, the speed of response, and temporary and permanent stock-price effects, *Journal of Financial Economics* 26, 71–95.
- Kraus, Alan, and Hans Stoll, 1972, Price Impacts of Block Trading on the New York Stock Exchange, *Journal of Finance* 27, 569–588.
- Lehar, Alfred, and Christine Parlour, 2021a, Battle of the Bots: Flash loans, Miner Extractable Value and Efficient Settlement, *working paper*.
- Lehar, Alfred, and Christine Parlour, 2021b, Decentralized Exchanges, *working paper*.
- Liu, Yukon, Aleh Tsvinski, and Xi Wu, forthcoming, Common Risk Factors in Cryptocurrency, *Journal of Finance*.
- Park, Andreas, 2021, The Conceptual Flaws of Constant Product Automated Market Making, *Available at SSRN 3805750*.

Number	Liquidator Address	Amount Liquidated	Aave	Compound
1	b7990f251451a89728eb2aa7b0a529f51d127478	427.38	427.38	0.00
2	3909336de913344701c6f096502d26208210b39f	137.81	81.11	56.70
3	681bd23f6128db3f9b8914595d1a63830a6212fa	128.84	0.00	128.84
4	3333a5a9d331f0c7b2c3626a1088fe6ee0b69e67	127.10	0.00	127.10
5	2ca158422b940c6722640ac7fa726e8201cccd33	119.10	119.10	0.00
6	b2b3d5b6215d4fb23bf8dd642d385c4b44aad2a	109.83	109.83	0.00
7	333388b6bf358d441cde209f1d58db27f28446e1	100.78	0.00	100.78
8	e8468f05550563aa5bfc5fbc344bf87aa2f6b84	93.94	0.00	93.94
9	19256c009781bc2d1545db745af6dfd30c7e9cfa	84.89	84.89	0.00
10	5af7f71c7747fb0eceb2eef115c3fa34dd4998d3	83.28	78.26	5.03
11	d911560979b78821d7b045c79c36e9cbfc2f6c6f	73.95	41.86	32.09
12	6780846518290724038e86c98a1e903888338875	67.16	17.47	49.69
13	872e0cc1606840bdd532e2f7e09a85cdf95f04bf	47.90	47.90	0.00
14	80d4230c0a68fc59cb264329d3a717fcaa472a13	47.85	32.61	15.24
15	33334570f7e1df34a09377c7f327feb65e2b3faf	39.89	0.00	39.89
16	b206ebe579be55f5b57119bb2e7cc63708eda1aa	39.84	31.77	8.07
17	645e93859ec63abe0c7fe74f17c07c236ee58799	38.96	38.96	0.00
18	3302ea0ddf59a8beb73779555cb2e86d350bfabc	34.96	34.96	0.00
19	08565d290208ea253875efe20d756dc42ae37612	32.56	32.56	0.00
20	d80d99ddad88c35a585e0ed3d287c49988b1e0e5	29.96	29.96	0.00

Table 6. Twenty largest liquidators sorted by amount liquidated in USD. *Number* corresponds to the label in Figure 4, *Amount Liquidated* is the sum of collateral liquidated in million USD by this address, *Aave* is the sum of collateral liquidated on Aave, *Compound* is the sum of collateral liquidated on Compound.

A Top Liquidators

Table 6 lists the wallet addresses of the top 20 liquidators in our sample.

B Detailed Description of the use of swaps and flash loans

We collect detailed log data for all Ethereum transactions where a liquidation events occurs in our sample and use these logs to see identify other smart contracts liquidators interacted with. For analysis in this section we examine if liquidators also interacted with a broad set of decentralized exchanges, including Uniswap and its many clones, Bancor, and Kyber. We mark these transactions as *Swap*. We collect data on flash loan usage by identifying lash loan specific log event patterns for Uniswap, Aave, and DyDx, three leading flash loan providers. We mark liquidations that also have a flash loan in the same Ethereum transaction as *Flash Loan*. All other loan liquidations are labeled as *Capital* as they are funded with the liquidator’s capital.

Figure 13 shows the density of the log loan liquidation amounts in USD by group. We find that the collateral of larger loan liquidations is swapped immediately, the average liquidation in this group is USD 114,763.52. Capital seems not to be a constraint for most liquidators as flash loans are not used that often. Liquidations that use them are on average larger with USD

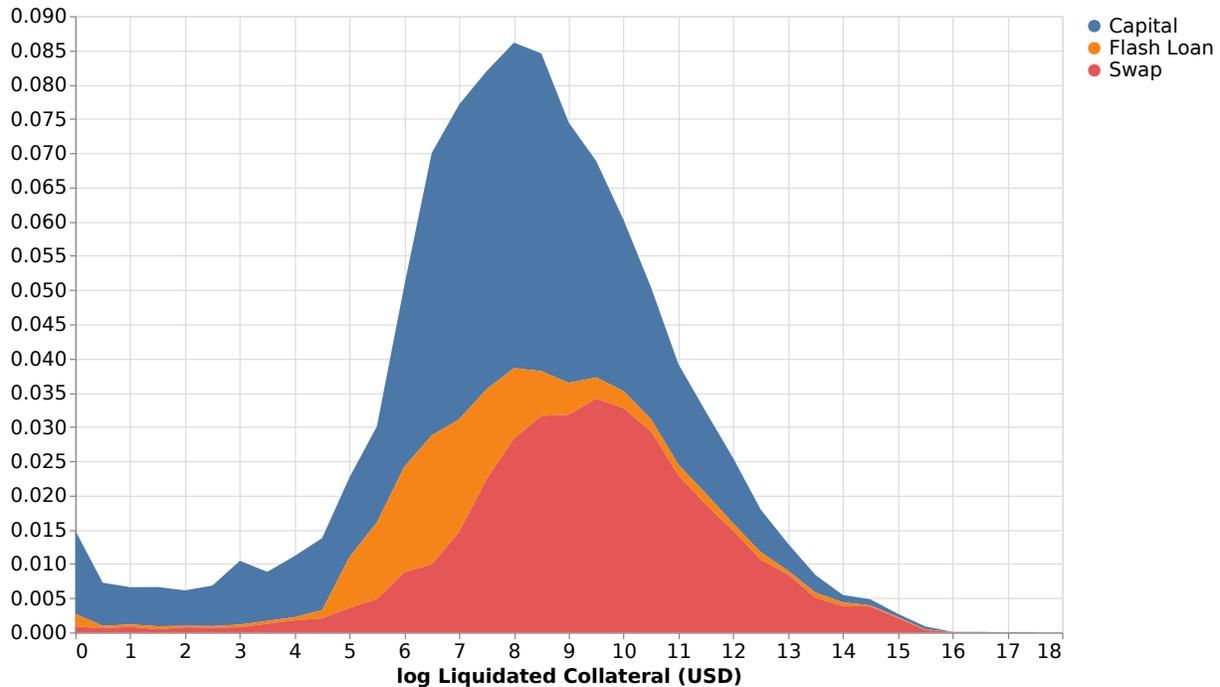


Figure 13. Distribution of log liquidated collateral amounts (in USD) broken down by liquidations funded with the liquidator’s capital for which the collateral was not swapped in the same transaction (blue), liquidations funded with the liquidator’s capital for which the collateral was immediately swapped (red), and liquidations that were funded with flash loans (orange).

40,938.97 compared to the liquidations that use neither swaps nor loans with an average size of USD 36,610.10.

Our analysis most likely underestimates the actual use of swaps and loans as liquidators might have used exchanges or loan providers that we could identify. There is no universal naming system on Ethereum and there is no universal method to identify a smart contract’s purpose. Addresses of regular wallets and smart contract addresses are indistinguishable.

C Detailed Description of Constant Product, Automated Market Making

This appendix is excerpted from Lehar and Parlour (2021b). In it, we describe the market making mechanics on UniSwap V2 for readers unfamiliar with this protocol.

Providing Liquidity: Each swap pool comprises a pair of cryptocurrencies. Most frequently, as we document below, one of the currencies is Eth, the native cryptocurrency on the Ethereum Blockchain. We will typically use Eth as the numeraire, and refer to the other generic coin as the ‘token.’ An agent wishing to provide liquidity to their preferred pool deposits both Eth and

the token into the pool. The deposit ratio of Eth to token is determined by the existing ratio in the pool, which implicitly defines the Eth price of the token.

An agent who makes such a deposit receives a proportional amount of a liquidity token. This third token is specific to the pool and represents an individual liquidity provider’s share of the total liquidity pool. As the pool trades with users, the value of the liquidity pool may rise or fall. Liquidity providers can redeem their liquidity tokens at any time and get their share of the current liquidity pool paid out in equal value of Eth and tokens. Changes in the composition of the pool from the time a liquidity token is minted to when it is cashed in, potentially constitute adverse selection. However, providing liquidity is potentially profitable because each trade faces a fee of 30bps which is redeposited into the pool.

Consummating Trade: Suppose a trader wishes to buy the token. In this case, he will deposit Eth into the pool, and withdraw the token. The amount that he has to deposit or withdraw depends on the bonding curve which is illustrated in Figure 14. Before the trade, there are E_0 Eth and T_0 tokens. The ratio of Eth to tokens is the implied price quoted by the pool. Someone who is interested in selling an arbitrarily small amount of the Token, would pay or receive E_0 . To trade a larger quantity, consider someone who wishes to sell some of the Token. This would mean that the trader deposits some amount $T_1 - T_0$ of the token into the pool. In return, he would receive $E_0 - E_1$. Thus, the amount of Eth in the pool drops.

If the seller was a liquidity trader, the post trade price in the pool ($\frac{E_1}{T_1}$) is now too low, and a potential arbitrageur would enter the market and trade in the opposite direction to return the ratio of Eth to tokens to equilibrium.

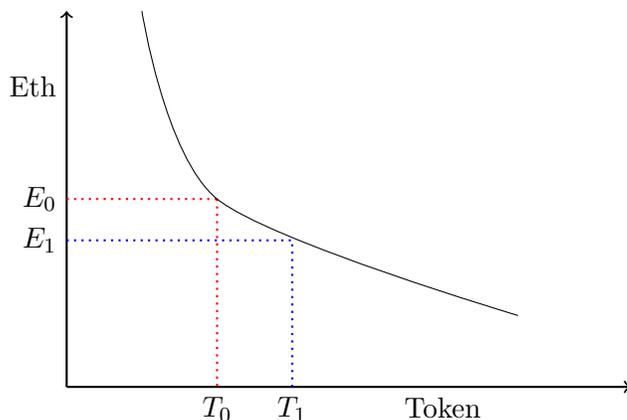


Figure 14. A bonding curve. From an initial amount of Eth and Tokens of E_0 and T_0 respectively, a trader deposits $T_1 - T_0$ tokens (sells) in exchange for $E_0 - E_1$ Eth. the price impact of this trade is determined by the bonding curve.

Specifically, if T_0 is the amount of tokens and E_0 the amount of Eth in the contract’s liquidity pool, then the terms of trade are such that for any post trade quantities before any fee revenue T_1, E_1

$$k := T_1 \cdot E_1 = T_0 \cdot E_0. \tag{1}$$

In other words, the product of the Token and Eth quantities is always on the bonding curve. For each pool, the constant k , depends on the amount of liquidity that has been deposited in the pool up to this point. We note that if more liquidity is posted, the constant changes. This is the mechanism through which the market equilibrates.

Assessing Liquidity Fees: The previous clarifies the terms of trade absent the liquidity fee. Of course, remuneration is important for the liquidity providers. To see how the fee affects trades and prices, suppose that an agent wants to trade e Eth in exchange for tokens. The exchange collects a fee τ , which benefits liquidity holders.¹³ Thus the effective amount of Eth that gets traded is $(1 - \tau)e$. This leads to a post trade, but before fee revenue liquidity pool balance of $E' = E + (1 - \tau)e$. Following the logic of the bonding curve (1), the post trade token balance must be

$$T' = \frac{T \cdot E}{E'} = \frac{T \cdot E}{E + (1 - \tau)e}. \quad (2)$$

The smart contract which executes the trade accepts the e ETH and returns the difference between the pre and post trade token balances. Or, the amount of token t that the trader receives is given by

$$t = T - T' = \frac{(1 - \tau)eT}{(1 - \tau)e + E}. \quad (3)$$

Therefore, the terms of trade expressed in Eth/token is given by

$$p^{tot} = \frac{e}{t} = \frac{e}{T} + \frac{E}{(1 - \tau)T}. \quad (4)$$

The terms of trade have a natural interpretation as a spread. Suppose that the fundamental value of the token denominated in Eth is p_0 . If the pool is in equilibrium then $p_0 = \frac{E}{T}$. The liquidity fee generates what is essentially a tick size that is distinct from the volume-induced price impact that the trader pays when he moves long the bonding curve, then

$$\lim_{e \rightarrow 0} \frac{p^{tot}}{p_0} = \frac{ET}{ET(1 - \tau)} = \frac{1}{1 - \tau} \quad (5)$$

That is, when buying tokens, traders have to pay a fixed spread of $\frac{1}{1 - \tau}p_0$. Similarly for token sales traders have to pay a fixed spread of $(1 - \tau)p_0$.

Pool size: The price that a trader gets is determined by the bonding curve and the volume of posted liquidity. In particular, the price impact of a marginal increase in the order is $\partial p / \partial e = 1/T$. As the liquidity pool grows, the price impact of a fixed order size decreases. Thus, understanding the payoff to liquidity provision is an important determinant of AMM market quality.

¹³Uniswap collects a fee of 30bps per trade.

D Detailed Description of Flash Loans

This appendix is excerpted from Lehar and Parlour (2021a). In it, we describe the mechanics of flash loans for readers unfamiliar with this feature.

Flash loans have neither maturity nor credit risk. They were were invented in July 2018 by Marble, an open source lending platform on the Ethereum blockchain and combine the lending of funds.¹⁴ Flash loans grown rapidly with loans worth on average 1.17 billion USD borrowed per day in the first quarter of 2021 compared to USD 500,000 for the same period a year earlier.

The most common use for flash loans is arbitrage. Decentralized exchanges, which trade tokens worth billions of dollars each day, purposely rely on arbitrageurs to keep prices aligned with markets and consistent with each other. Flash loans provide cheap capital to arbitrageurs to execute their trading strategies. Other use cases for flash loans include swapping collateral for secured loans, loan liquidations, and exploits of weaknesses in other DeFi protocols.

Flash Loans are typically used as one component of more complex transactions on the Ethereum blockchain that interact with numerous Decentralized Finance (DeFi) platforms. One Ethereum transaction can interact with several smart contracts and call functions of these smart contracts to trigger economic actions such as borrowing, lending, conversion between tokens using a decentralized exchange, or transferring tokens between wallets. In a flash loan a borrower takes a loan at the beginning of a transaction and repays the loan at the end of the same transaction, thus repaying the loan at the same time as it was borrowed. Blockchain transactions are atomic, meaning that they either get executed in their entirety or not at all. Therefore borrowers cannot default during a transaction and the loan is only processed or taken out when it is also repaid. Lenders therefore have no credit risk. The atomic nature of transaction also generates an option type payoff for the borrower. A transaction can require to leave a profit for the sender, the person initiating the transaction. Thus if the transaction is not profitable it fails and the loan does not get taken out and the sender is left to pay is the fee for processing the transaction on the blockchain (i.e. the gas cost).

¹⁴Marble never became widely used and is today insignificant.