

Systemic Fragility in Decentralized Markets

Alfred Lehar*
Haskayne School of Business
University of Calgary

Christine A. Parlour†
Haas School of Business
UC Berkeley

June 13, 2022

**Preliminary and Incomplete
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 the Fintech Dauphine Chair in partnership with Mazars and Crédit Agricole CIB for financial support.

†email:parlour@berkeley.edu. We thank Terrence Hendershott and Andreas Park and seminar participants at American Economic Association, FIELDS-CFI Workshop on Mathematical Finance and Cryptocurrencies 2022, CryptoAssets and Digital Asset Investment Conference, NYU Global Fintech Conference, 2022, NUS-SKKU FinTech Seminar Series for comments. Parlour and Lehar are grateful to the Hogege BlockChain Research Institute 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

What effect do liquidity trades have on market stability? Since the Great Depression, the Federal reserve has restricted the amount that investors may borrow to invest in part to prevent large market crashes. Further, organized exchanges employ circuit breakers to mitigate the effects of large trades. In this paper use some of the unique features of decentralized finance protocols to investigate this question.

Decentralized finance offers a unique laboratory to investigate the immediate effect and subsequent propagation of large liquidity trades. 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.

We collect a unique data set of loan liquidations on Compound and Aave, two of the largest DeFi lending protocols. There is approximately \$9 billion of collateral locked in Compound, and over \$11 billion locked in Aave.¹ Over our sample period, we observe liquidations valued at \$2,487,543,097.

Most capital in decentralized finance (DeFi) is allocated to collateralized lending protocols. Users can post collateral in one token and take out a loan 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 when liquidated is sold 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.

Loan liquidations are most often performed by bots, who monitor lending platforms for loans that are in technical default and liquidate them for profit. Such liquidators have no interest in holding the collateral long term and so they immediately sell the collateral on one of the decentralized exchanges. All decentralized exchanges have deterministic price impacts, and there is no strategic incentive to split orders or delay trade. Thus, shocks to the collateral asset

¹Approximate figures are available from defipulse.com

propagate across the DeFi system. There are two ways in which this happens. First, a price dislocation on one DEX will lead to ancillary strategic trading on the others. Second, prices across various DEXs are aggregated and used as a price oracle to determine the collateralization of other loans. Using detailed blockchain data, we can trace liquidated keeper positions to specific DEXs, and then observe the chain effect of this on both other DEXs and contagion effects.

For each of the liquidations in our sample, we are able to trace how they affect prices on one of nine different decentralized exchanges. These liquidations are liquidity trades. We document that liquidations have impacts across multiple blocks that last several minutes. In addition, we find that these liquidity shocks affect not only the exchange on which they occur but also competing venues. This is potentially an important observation as price oracles use average prices across various Dexs. Therefore, in an informationally closed blockchain system liquidations can lead to systemic fragility. Figure 1 illustrates the feedback effect in the informationally closed blockchain system.




Figure 1. Systemic Fragility Channel

To make this argument more concrete, Figure 2 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) 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 Dexes 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.




Figure 2. Cummulative return (blue) and cummulative return from loan liquidations (red) for the ETH/LINK exchange rate on May 19th, 2021.

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. 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.

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 (2021). 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




Figure 3. Lending Protocols provide incentives for Liquidators to monitor asset loans and ensure that they remain over-collateralized

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.

1.1 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. In return for supplying assets, a lender receives a token that is a claim to both the principal and interest earned on the loaned amount. The fungibility of loaned tokens means that the claim is liquid and the protocol design minimizes idiosyncratic or borrower-specific risk.

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. Specifically, a fraction 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.

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.

Information available on-chain is 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. Liquida-

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

tion is triggered based on the oracle price. For our purposes it is important that the liquidation trigger is only a function of public information. Thus, these liquidations are purely liquidity trades.

2 Data and stylized facts

We collect data on liquidations from two of the largest Defi lending protocols, Aave and Compound. We collect data from UniSwap and all its clones: SushiSwap, ShibaSwap, ZKSwap, SakeSwap, DefiSwap, CitySwap, BTSwap and Equalizer. (For readers unfamiliar with these markets, we present a description in the Appendix) 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. Specifically 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. Second, we do not record information from exchanges such as Bancor and Balancer. In total we observe 42,324 liquidations over a time period from September 25, 2018 to May 16, 2022 consisting of 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 40 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 38,409 liquidations.³ Notice that wrapped ETH, Link and wrapped Bitcoin are the leading collateral tokens.

The collateral exhibited in Table 1 was used to issue debt in various 44 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 various tokens such as ETH or Bitcoin.

In total we observe liquidations of USD 2,487,543,097 with a mean liquidation size of USD 64,765 and a median of USD 3,587. The largest loan liquidation in our sample was the liquidation of USD 50,508,256 worth of DAI collateral on Compound on Nov 26, 2020.⁴

³The 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.

⁴See transaction 0x53e09adb77d1e3ea593c933a85bd4472371e03da12e3fec853b5bc7fac50f3e4.

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.

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.

3 Liquidations

Figure 6 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 with a total collateral value of USD 503,572,311. On that day Eth dropped from over USD 3,400 to USD 2,014, a 41% decline.

Most liquidations occur in waves. 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 at 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




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

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.

3.1 Liquidators

Anecdotal evidence suggests that liquidations are automated and in particular executed by trading bots.⁵ We observe 1,004 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. A schema of the steps to a loan liquidation is presented in Figure 5.

Consider a borrower who borrows 1 wrapped Bitcoin (WBtc) against collateral of 11 wrapped Eth (WEth). Suppose that the loan has become undercollateralized. This is indicated by the red cross. At this point, a liquidator can profitably repay the debt and liquidate the loan to ensure that Aave does not incur a loss. To do so, she borrows 1 WBtc from a flash lender and repays the outstanding debt. In the same transaction, she seizes the collateral of 11 WEth and swaps 10 of these tokens for 1 WBtc on a decentralized exchange. She then uses the 1 WBtc to repay the flash loan and pockets a profit of 1 WEth.

We rank liquidators by liquidated collateral amount and present this in Figure 6. Observe

⁵The Github repository at <https://github.com/haydenshively/Compound-Liquidation-Bot> provides solidity script necessary to program a bot.




Figure 5. Loan Liquidation and Repayment process.

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 427,368,917. 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.




Figure 6. Amounts liquidated by top 20 liquidators, by platform. Each column represents a distinct address, while the dollar volume in Aave is depicted in blue and the dollar volume in Compound is depicted in orange.

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




Figure 7. Number of Loans liquidated by the top 20 liquidators (defined as the liquidators with the largest amount of loans liquidated). The labels of the liquidators are the same as in Figure 6

discounted price. She can then choose to keep the collateral token which exposes her to price 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 3 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 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.

4 Price impact of liquidations

We have presented stylized facts on liquidations and liquidators. These liquidations are pure liquidity trades, and because of flash loans the size of liquidations are not constrained by liquidator capital. In the Defi system, prices are maintained by arbitrageurs who equalize them across markets. The logic of decentralized exchanges requires that arbitrageurs reverse liquidity trades so the prices quoted by the DEXs are current. Successive liquidations will only have large price impact if either they occur in waves or if arbitrageurs do not trade to reverse the liquidity impact of the liquidations. In such cases, liquidations may trigger other liquidations, and if liquidations spill over into other markets and if prices revert slowly to equilibrium levels then

	(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 3. 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.

these liquidity trades can lead to systemic fragility. We first present an example of a liquidation and the subsequent price impact. We then consider the incentives of an arbitrageur to reverse a liquidity trade and then empirically examine the effect of liquidations on returns.

4.1 A Motivating Example

To make our investigation of liquidations concrete, we present one liquidation in our sample.⁶ 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 8 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

⁶see transaction 0xd70b42daec5bb9ac6e5df3d25d309f186db50df701f667e1f20b22448ea27d41




Figure 8. Price reactions to the liquidation in Subsection 4.1. 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.

contagion.

4.2 Arbitrage Incentives

To relate changes in the relative value of a token pair to trading in that pair, consider an asset whose value evolves discretely according to

$$p_t = p_{t-1} + \epsilon_t.$$

We briefly review the impact of trade on an automated market maker pool that comprises T tokens against E Eth. Let $k = ET$ denote the AMM constant, then if y tokens are sold to the pool, the amount of Eth removed from the pool, Δ_E^y , is given by

$$\begin{aligned} (E - \Delta_E)(T + y) &= k \\ E - \frac{k}{T + y} &= \Delta_E^y. \end{aligned}$$

Thus, after y tokens are liquidated (sold) to the pool, the pool comprises $T + y$ tokens against $E - (E - \frac{k}{T+y}) = \frac{T}{T+y}$ Eth.

Similarly, if w is bought from the pool, we know that the amount of Eth deposited into the pool, Δ_E^w , is given by

$$\begin{aligned}(E + \Delta_E)(T - w) &= k \\ \frac{k}{T - w} - E &= \Delta_E^w\end{aligned}$$

We are now in the position to consider the payoff to an arbitrageur. Suppose that an amount q_t has been liquidated (sold) to the pool. If they trade x right away, they face a pool of size $(T + q_t, \frac{k}{T + q_t})$. Further, if they buy x tokens, they pay $\frac{k}{T + q_t - x} - \frac{k}{T + q_t}$. Thus, the profit to buying x tokens after a liquidation of q_t tokens is

$$\pi^{now} = p_t x - \left[\frac{k}{T + q_t - x} - \frac{k}{T + q_t} \right]. \quad (1)$$

Clearly, the arbitrageur optimally reverses the trade, so $x = q_t$, if the size of the pool is sufficiently large.⁷

Suppose instead that the arbitrageur waits in anticipation of further liquidations. In this case, his payoff at time $t + 1$ will be

$$\pi^{wait} = p_{t+1} x - \left[\frac{k}{T + q_t + q_{t+1} - x} - \frac{k}{T + q_t + q_{t+1}} \right] \quad (2)$$

Again, he will optimally choose $x = q_t + q_{t+1}$ for a profit of

$$\pi^{wait} = p_{t+1}(q_t + q_{t+1}) - \left[\frac{k}{T} - \frac{k}{T + q_t + q_{t+1}} \right] \quad (3)$$

Comparing Equations 3) and 1), the arbitrageur will optimally wait to reverse the liquidation if

$$\underbrace{E[(p_t + \epsilon_{t+1})q_{t+1} | p_t, q_t]}_{\text{Value of Future Liquidations}} \geq \underbrace{\frac{k}{T + q_t} - E\left[\frac{k}{T + q_t + q_{t+1}} | q_t, p_t\right]}_{\text{Cost of Trading}} \quad (4)$$

In this framework, capital will be slower to allocate if the value of future collateral is higher. This comprises both the price of the collateral and the quantity. The latter of course, is determined by the price path of the collateral. Further, the higher the pool constant, k , the more likely the arbitrageur is to immediately reverse the trade. This implies that smaller and less liquid pools (or collateral) are ceteris paribus less likely to have the liquidation liquidity trades rapidly reversed. There will therefore be cross-sectional differences in collateral tokens and the effect of the liquidity trade will be amplified.

⁷Recall, the price at any point in time is $p = \frac{E}{T}$, and given that $k = ET$, trade is optimal if $T > q_t^2 - q_t$.

4.3 Analysis of the entire sample of Collateral Tokens

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. We first investigate the effect of loan liquidations on subsequent high frequency prices.

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 8. 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 4. 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 41.6% 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.

For the results in Table 4 we see that liquidations where the collateral gets immediately swapped have a medium term price impact on all exchanges, whether a collateral gets sold on that exchange or not. This effect is important because the prices from the exchanges are used as inputs, or oracles, to determine if other loans are sufficiently collateralized. The lending platforms often do not publicly disclose which oracles are used for security reasons but all oracles, more or less rely on the prices that are quoted by decentralized exchanges. Any price changes in the value of collateral will cause more loans to be undercollateralized and lead to more liquidations, potentially leading to a liquidation wave.

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

	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 4. 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.

contrast these with all other returns for all other tokens. Figure 9 shows distribution functions for liquidation and other returns. We can see that 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 liquidation versus 3.25% for cases without a liquidation.

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.

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 that just contains the 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 887 liquidation waves to 1336 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




Figure 9. Return distribution for 16 tokens that serve as collateral over 5 minute intervals that coincide with liquidations and ones that do not.




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

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 5.

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.

	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 5. Regression explaining the return of the debt token/collateral token return throughout 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.

4.4 Systemic Fragility and the DeFi 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. We illustrated the feedback effects that generate this systemic fragility in Figure 1 in the introduction. The feedback between Dex prices and liquidations also presents the possibility for strategic liquidations.

There are two natural countervailing forces to the feedback channel presented in Figure 1. First, 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 EVM requires a pre-specified amount of gas. In addition, an incentive amount of gas for miners can be added to a transaction. Figure 11 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 11. Average Daily Gas fees for a Swap. This calculation is based on a simple swap (50,000 gas units).

5 Conclusion

We have demonstrated that liquidity trades can have market wide persistent effects. Further, given that DeFi is designed to be a closed information system it suggests that the system features a systemic fragility. 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.

As the DeFi ecosystem evolves, different blockchains that run the EVM are being developed. Given that blockchains are a closed informational system, differences in costs across blockchains will liquidation, arbitrage activity and hence systemic fragility.

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*.
- Capponi, Agostino, and Ruizhe Jia, 2021, The Adoption of Blockchain-based Decentralized Exchanges, *working paper*.
- 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, 2021, 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	427.37	427.37	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	119.15	119.15	0.00
6	b2b3d5b6215d4fb23bf8dd642d385c4b44aad2a	109.82	109.82	0.00
7	333388b6bf358d441cde209f1d58db27f28446e1	100.78	0.00	100.78
8	e8468f05550563aa5bfc5fbc344bf87aa2f6b84	93.94	0.00	93.94
9	19256c009781bc2d1545db745af6dfd30c7e9cfa	84.78	84.78	0.00
10	5af7f71c7747fb0eceb2eef115c3fa34dd4998d3	83.29	78.27	5.03
11	6780846518290724038e86c98a1e903888338875	67.16	17.47	49.69
12	872e0cc1606840bdd532e2f7e09a85cdf95f04bf	47.93	47.93	0.00
13	b206ebe579be55f5b57119bb2e7cc63708eda1aa	39.32	31.26	8.07
14	645e93859ec63abe0c7fe74f17c07c236ee58799	38.96	38.96	0.00
15	33334570f7e1df34a09377c7f327feb65e2b3faf	38.31	0.00	38.31
16	3302ea0ddf59a8beb73779555cb2e86d350bfabc	34.97	34.97	0.00
17	08565d290208ea253875efe20d756dc42ae37612	32.56	32.56	0.00
18	d80d99ddad88c35a585e0ed3d287c49988b1e0e5	29.94	29.94	0.00
19	80d4230c0a68fc59cb264329d3a717fcaa472a13	26.69	21.03	5.67
20	bf3f6477dbd514ef85b7d3ec6ac2205fd0962039	24.00	20.27	3.72

Table 6. Twenty largest liquidators sorted by amount liquidated in USD. *Number* corresponds to the label in Figure 6, *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.

Detailed Description of Constant Product, Automated Market Making

This subsection is excerpted from Lehar and Parlour (2021). 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 12. 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 12. 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

$$k := T_1 \cdot E_1 = T_0 \cdot E_0. \quad (5)$$

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.⁸ 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 (5), the post trade token

⁸Uniswap collects a fee of 30bps per trade.

balance must be

$$T' = \frac{T \cdot E}{E'} = \frac{T \cdot E}{E + (1 - \tau)e}. \quad (6)$$

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

$$t = T - T' = \frac{(1 - \tau)eT}{(1 - \tau)e + E}. \quad (7)$$

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

$$p^{tot} = \frac{e}{t} = \frac{e}{T} + \frac{E}{(1 - \tau)T}. \quad (8)$$

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

$$\lim_{e \rightarrow 0} \frac{p^{tot}}{p_0} = \frac{ET}{ET(1 - \tau)} = \frac{1}{1 - \tau} \quad (9)$$

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.




Figure 13. Long term price impact of a liquidation waves. The graph shows the commulative return of the collateral token relative to 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.