

An Analysis of the Market Risk to Participants in the Compound Protocol

Hsien-Tang Kao, Tarun Chitra, Rei Chiang, John Morrow
Gauntlet

{hsien-tang, tarun, rei, john}@gauntlet.network

[Research long paper]

ABSTRACT

Compound is widely considered to be the most significant protocol that facilitates participants to trustlessly supply and borrow Ethereum assets. The protocol has grown in popularity in recent years due to the transparency of mechanism, appealing interest rates for borrowers and passive income for suppliers. Despite the increased transparency, understanding the market risk to participate in the protocol is difficult because of the complexity of decentralized market dynamics. The causes of complexity are the multitude of participant behaviors and their interactions with the protocol's endogenous variables and exogenous markets. This paper gives an overview of the market risks, the liquidation mechanism and the implications of protocol and market variables. We rigorously evaluated the economic security of the protocol by using agent-based modeling and simulation to perform stress tests. Our results showed that the Compound protocol is able to scale to a large market size while having a low probability of default under volatile market conditions. Results also revealed that the current protocol parameters for specifying the minimum collateral requirements and liquidation incentives are sufficient for liquid collateral assets.

Keywords

Decentralized finance, agent-based simulation, smart contracts

1. INTRODUCTION

Compound allows participants to trustlessly supply and borrow Ethereum assets, providing appealing interest rates for borrowers and passive income for suppliers. By using collateral and amortizing risk across individual suppliers in a liquidity pool, Compound's Ethereum smart contract has been a profitable place to supply crypto since its inception in 2018. The protocol implemented in Compound's smart contract is detailed in the Compound whitepaper.[1]

However, despite the fact that Compound has grown well

past nine figures (of USD value) without any suppliers losing money, it is still technically possible, under extreme conditions, for borrowers to default on their borrowed assets and suppliers to lose their principal. Understanding when this failure condition can happen boils down to understanding various types of risks associated with the protocol, including protocol security risk,¹ governance risk,² and market risk. This paper focuses on evaluating market risk — the risk of a user experiencing losses due to market fluctuations external to the smart contract itself.

We use a rigorous definition of market risks to construct simulation-based stress tests that evaluate the economic security of the Compound protocol as it scales to underwriting billions of dollars of borrowed assets. These stress tests are trained on historical data and put through a battery of scenarios that represent the expected and worst case economic outcomes for the protocol. Our stress tests are constructed analogously to how transaction-level backtesting is done in high-frequency and algorithmic trading. These techniques are used to estimate the market risk of a systematic trading strategy before it is deployed to the market. As there are over \$1 trillion US dollars of assets managed by funds that use these techniques to provide daily actuarial analyses to risk managers, we believe that these are the best methodologies for evaluating market risk.[2, 3, 4] By modifying these techniques to handle the idiosyncrasies of cryptocurrencies, we are able to provide similar statistical power in these actuarial analyses.

The first portion of this paper will define the set of market risks that users of the Compound protocol face, breaking them down into their principal quantitative components. Subsequently, we will describe the incentives behind the mechanism that the Compound protocol uses to ensure that it is solvent — liquidations. Finally, we will conclude by detailing how liquidators are similar to trading strategies and detail the market impact models that are used to analyze their incentives and expected returns.

The second portion will focus on methodology and results from agent-based simulations of the Compound smart contract. Our methodology utilizes careful simulation to closely replicate the live environment that users interact with in the Compound protocol. This approach and some of our novel technologies, such as a custom Ethereum virtual machine, ensure that our results replicate reality with high fi-

This article is published under a Creative Commons Attribution License (<http://creativecommons.org/licenses/by/3.0/>), which permits distribution and reproduction in any medium as well allowing derivative works, provided that you attribute the original work to the author(s) and FAB 2020.

Third International Symposium on Foundations and Applications of Blockchain (FAB '20) May 1, 2020, Santa Cruz, California, USA.

¹Examined by independent smart contract auditors: Certora, OpenZeppelin, and Trail of Bits

²More broadly, this refers to things like administrator mismanagement, voter participation, etc.

delity. We conclude by detailing the results of these simulations, providing actuarial assurances for the conditions under which the Compound protocol is insolvent.

Our conclusions show that **the Compound protocol can scale to a larger size and handle high volatility scenarios for a variety of collateral types**. In particular, we find statistically significant evidence that even when Ether (ETH) realizes its maximum historical volatility, the Compound system is able to grow total borrowed value by more than 10x while having a sub-1% chance of default.³ Note that in this paper we will refer to the protocol being in ‘default’ as equivalent to being under-collateralized. Moreover, we find that the system stays significantly over-collateralized in extreme scenarios and that current liquidation incentives are sufficient for more liquid collateral types, such as ETH. Finally, a glossary of terms utilized throughout this paper can be found in Appendix 9.1.

2. MARKET RISKS

The decentralized nature of the Compound protocol renders risk assessment both more complex and crucial than similar assessments in traditional markets. The main causes for this increase in complexity are the multitude of participant behaviors in the Compound protocol as well as their interactions with exogenous markets, such as centralized cryptocurrency trading venues. Unlike formal verification and smart contract auditing, which focus on *endogenous* risks within a smart contract, economic analysis of protocols focuses on how *exogenous* shocks affect participant behavior. As the Compound protocol uses a deterministic function of liquidity supply and borrowing demand to determine the interest rates that suppliers and borrowers receive, one need only consider market prices, supplier supply behavior, and borrowing demand to accurately model exogenous risk (see Appendix 9.2). More specifically, the primary sources of exogenous risk stem from the following components:

1. Shocks to market prices of collateral that cause the contract to become insolvent due to under-collateralization
2. Loss of liquidity in an external market place, leading to a liquidator being disincentivized to liquidate defaulted collateral
3. Cascades of liquidations impacting external market prices which in turn lead to further liquidations (i.e. a deflationary spiral)

In order to quantify the effects of these risk components, we first need to delve into the notions of assets and liabilities within the Compound protocol.

2.1 Assets and Liabilities

In the Compound protocol, the main assets are the collateral tokens that suppliers have committed to liquidity pools, whereas the main liabilities are the outstanding borrowed assets. Token holders contribute their ERC-20 assets to a liquidity pool, and are in turn paid a yield on their supplied tokens. Borrowers borrow an asset by first committing collateral before withdrawing up to a certain amount

³This 10x is relative to the size of the Ether market. In the case where that grows a commensurate amount, as it easily could, then Compound could grow even larger.

from the liquidity pool. This amount is controlled by the *collateral factor*, which is the ratio of the maximum outstanding debt to collateral. The system forces borrowers to over-collateralize their borrowed assets (e.g. a fully-secured credit facility), thus enforcing the invariant that assets must always be greater than liabilities. For instance, one can deposit \$100 of ETH⁴ into the contract and withdraw \$75 if the contract has a collateral factor of 75%.

The *net liabilities* of Compound are defined as the asset values less liabilities, so that the system is deemed solvent when the net liabilities are positive. As a decentralized protocol, Compound utilizes a series of economic incentives to ensure that net liabilities are always positive. The borrower’s collateral requirement is the value of outstanding debt divided by the collateral factor. When the market value of the collateral backing a lien falls below the collateral requirement, the collateral position becomes *liquidatable* and the protocol sells the collateral at a discount to a liquidator. This discount, termed the *liquidation incentive*, provides a liquidator with financial incentive to buy the collateral from the protocol, effectively repaying the borrowed asset on behalf of the borrower. With liquidation, the protocol acts much like a bank selling a defaulted asset at a foreclosure auction to increase their net liabilities. In particular, the liquidator acts analogously to the foreclosure auction winner, who is usually able to claim the defaulted asset at a discount.

As an oversimplified example, suppose that the Compound protocol has an borrow position that is in default, with the current collateral amount equal to \$100. If the liquidation incentive is 105% (5% extra bonus), then the liquidator would pay the Compound Smart Contract \$95 for the ETH collateral. Moreover, if the liquidator has low time preference, then they will sell the collateral as soon as possible. In practice, the Compound protocol only lets liquidators liquidate a portion of the borrow amount, and they receive collateral equal to 105% of the borrow value repaid. This has the benefit of increasing the collateralization ratio on the remaining portion of the borrowed asset, while avoiding complicated mechanics of completely closing borrow positions.⁵ In this sense, liquidation in Compound resembles an algorithmic trading strategy, as there is a race to be the first liquidator to claim portions of the collateral and sell it on the market with minimal transaction and slip-page costs.

2.1.1 Synthetic Assets: cTokens

There is a slight nuance in how assets and liabilities are treated — technically, the assets that suppliers and borrowers interact with are cTokens. These tokens, which wrap standard ERC-20 assets, serve as contingent claims on assets and earned interest. Suppliers supply assets as ERC-20 tokens and are returned cTokens, whereas borrowers supply

⁴In this stylized example, we use US Dollars as a numéraire, whereas in reality, one would have to execute this transaction in the Compound protocol against a USD stablecoin. Stablecoins are digital representations of US dollars, with some backed by bank deposits (USDC, TUSD) and others backed by digital collateral (DAI).

⁵Contrast this with the model MakerDAO uses, where there are auctions to liquidate the entire borrowed asset. This can create a delay which adds to market risk as well as unnecessarily closes borrow positions which could be merely reduced to a safe level.

collateral, which is converted to a cToken and used to make outstanding interest payments. Unlike traditional assets, cTokens immediately realize earned interest as payments are paid pro rata to holders on every block update.

Technically, there is a security risk that a cToken cannot be converted back to the underlying asset if the contract has many outstanding borrowed assets that are not being repaid as collateral is redeemed. This would mean that the contract is illiquid, but not necessarily insolvent. This paper focuses on solvency, and liquidity will be considered more deeply in future analysis.

2.2 Risk Sensitive Parameters of the Protocol

The main levers protocol designers can wield in Compound to reduce risk are the collateral factor and liquidation incentive. However, these two levers impact the incentives of the protocol in different ways. The collateral factor controls the riskiness of borrowers — the closer it is to 100%, the more likely risky borrowers will default by borrowing USD stablecoin against collateral that is rapidly decaying in value. On the other hand, the liquidation incentive controls how likely liquidators are to take liabilities off of the smart contract’s balance sheet. The higher the liquidation incentive, the less time a defaulted borrowed asset will be a liability on the Compound protocol. If we dissect how the three risk components of §2 connect to these two parameters, we find the following:

- The risk inherent in the collateral factor is connected to the nature of shocks to the market price of the collateral
- The risks that liquidators with low time preference face is connected to the loss of liquidity in an external market place
- Cascading liquidations affect both the collateral factor and the liquidation incentive because they create a feedback loop between price shocks and a loss of liquidity

This implies that under normal market conditions, when liquidations are independently distributed (e.g. uncorrelated), the collateral factor and liquidation incentive control borrower risk and supplier’s ability to recoup losses, respectively. However, in situations when liquidations have a ‘knock-on’ effect and are correlated, these parameters affect both borrower and supplier behavior. Therefore, to study the true market risk of the system, we need to sample a variety of market and liquidity conditions in order to stress test these scenarios.

3. LIQUIDATION

Akin to foreclosure sale participants in traditional finance, liquidators can repay the outstanding debt with discounts in exchange for the borrower’s cToken collateral. In both foreclosure sales and in Compound liquidations, discounts are used to incentivize purchases of defaulted collateral. The Compound protocol provides a discount by giving liquidators additional collateral as the liquidation incentive to perform liquidation. However, unlike the all-or-nothing transactions of foreclosure sales, an individual liquidator can only repay a portion of the debt. The *close factor* is the protocol parameter that specifies the proportion eligible to be

liquidated by any individual liquidator. When a liquidator finds a profitable trade, she repays a portion of the outstanding debt (determined by the close factor) in return for the borrower’s collateral. Depending on a liquidator’s risk preference, she may sell the collateral immediately to protect against price-fluctuation risk or just hold the received collateral.

Liquidation incentives create an arbitrage opportunity or a price discount for the liquidator in exchange for the reduction of Compound’s risk exposure. The higher the liquidation incentive is, the more liquidators will participate in the liquidation process as they get steeper discounts relative to market prices. In other words, tuning the liquidation incentive is one of the most effective ways to adjust the protocol’s safety boundary. The liquidation incentive also has an influence on a borrower’s decision to borrow asset within the protocol. When a borrower’s lien is liquidated, the liquidation incentive can be viewed as a bonus amount of a borrower’s collateral that is given to the liquidator to compensate for the risk they engender while taking a liability off of the protocol’s balance sheet. If the liquidation incentive is too high, a borrower may be unwilling to borrow assets from Compound in the first place, or she may open a borrowing position and maintain a high collateral factor. In general, one expects that increased liquidation incentives negatively impact borrowing demand.

The collateral factor defines a maximum borrowing capacity for each asset enabled within the protocol. Borrowers must manage their own debt and keep their liens over-collateralized to ensure a certain margin of safety with respect to the maximum borrowing capacity. This margin of safety fluctuates with market conditions and depends on the borrowers’ own risk profile. When the market volatility is high, risk-averse borrowers maintain a high margin of safety to avoid their collateral being liquidated. In contrast, risk-seeking borrowers maintain a low margin of safety and actively refinance their debt to optimize their usage of borrowed capital. Understanding the interaction between collateral factor and the safety margin requires studying the influence of psychology on the participant’s behavior. Randomized controlled trials and other experimental methods are designed to understand this type of causal relationship.

Rational liquidators with short time preference are defined to be participants who purchase collateral from the Compound smart contract and immediately sell it on a centralized venue (e.g. have low risk tolerance). For brevity, we will refer to rational liquidators with short time preference as *greedy liquidators*. To simplify the analysis and simulate the worst-case scenario for Compound, we assume that all liquidators are greedy and sell the collateral immediately to a market, instead of having liquidators that repay the outstanding debt and hold the collateral. This focus on greedy liquidators emulates the worst-case protocol behavior as adverse market and liquidity conditions can cause cascading defaults. Greedy liquidators tend to inflame cascading defaults as they create sell pressure and can cause a deleveraging spiral. [5] The main source of loss for greedy liquidators is the loss due to price impact, or *slippage*, that is caused by selling a large quantity of an asset. Given that greedy liquidators immediately sell, they must optimize the quantity that they are willing to liquidate based on market prices and expectations of slippage.

4. SLIPPAGE

Slippage refers to the expected change in a tradeable asset's price p due to a matched order of size q and is mathematically denoted $\Delta p(q)$. Formally, $\Delta p(q)$ is defined to be the difference between the market midpoint price and the actual average execution price when a market participant executes a trade. Slippage inevitably happens on every trade, and this effect tends to be magnified in thin or high volatility markets. For a liquidation opportunity, slippage is the only cost that can be partially controlled by the liquidator, whereas trading fees and smart contract transaction fees are usually external restrictions. Therefore, slippage is one of the major factors that influence a liquidator's decision-making.

Market impact, which is a synonym for slippage, has been studied extensively in traditional finance.[6, 7] Many market impact models have been proposed and tested for solving optimal order execution problems. In traditional markets, the marginal increase in price impact is usually observed to decrease as a function of trade quantity, which formally corresponds to $\Delta p(q)$ being a concave function.[8, 9] However, this appears to not be true for cryptocurrency markets, where empirical data suggests that $\Delta p(q)$ is linear or even convex (e.g. the marginal cost *increases* with quantity).[10, 11] Despite each type of model having different underlying assumptions and functional forms, a majority of the models comprise trade volume-to-market size, volatility and time variables. Analyzing trade size, volatility and how these variables interact with liquidation is the primary focus of this analysis. The analysis in this paper only considers greedy liquidators that sell repossessed collateral on centralized exchanges with order books, such as Coinbase and Binance. As decentralized exchanges and automated market makers, such as Uniswap,[12] provide an alternative source of liquidity, one might ask why this assumption was enforced. The reasons for this choice are two-fold:

- Order book depth on centralized exchanges is order of magnitudes greater than that of decentralized exchanges for most assets⁶
- Slippage in automated market makers is usually designed to be small for small trades and expensive for large quantities, so greedy liquidators would likely end up going to a centralized exchange during the most volatile times to stay profitable

We will break up the dominant features of slippage into *market variables* that are exogenous to the Compound smart contract state.

4.1 Key Market Variables

4.1.1 Outstanding Debt

The total traded quantity that the protocol will need liquidated in times of net negative liabilities will be a function of the total outstanding debt in the system. Since this quantity is the input to the slippage function $\Delta p(q)$, it is clear that the choice of slippage model needs to be cognizant of

⁶We do note that this is not true for assets such as MKR and SNX, as their primary market is Uniswap. However, for the larger assets that are listed on Compound such as ETH, DAI, and REP, there is far more centralized exchange liquidity.

the amount of outstanding debt. We will define the amount of outstanding debt in this analysis to be the sum of all the borrowers' total outstanding debt value normalized by the average daily trading volume of underlying collateral. This metric captures the size of debt relative to the underlying liquidity, and gives readers a good intuition around how big Compound's market can grow safely relative to the trading markets. Since the trading volume of different assets varies, using unitless metrics (such as the amount of outstanding debt) provides a more intuitive comparison between different assets. The simulation in this paper assumes borrowers borrowing USD stablecoin backed by ETH, as this is the most common use case in the Compound protocol. As an example, suppose that the ETH daily trading volume is 100 million USD, 0.5 total outstanding debt is equivalent to 50 million USD of total outstanding debt value.

Estimating the average daily trading volume of cryptocurrencies is difficult, as wash trading and other market manipulation practices are known issues in the cryptocurrency market.[13] Numerous studies have concluded that the reported volume from various cryptocurrency exchanges may be unrepresentative of the assets' underlying liquidity. For this reason, we aggregated the average daily trading volume from the top 10 exchanges with well-functioning markets identified by Bitwise Investments.[14, 15] This indexing methodology has been adopted as the de facto industry standard, with major brokers and the Securities and Exchange Commission utilizing the Bitwise index for volume estimation.[16]

4.1.2 Asset Volatility

Volatility measures the degree of variation of asset price changes over a given time interval. Historically, it is traditionally defined as the standard deviation of logarithmic returns and is usually denoted σ . [17] Research studies show that volatility is typically a linear coefficient in a market impact model.⁷ Given that asset volatility changes over time and is affected by market microstructure, it's equally important to understand how liquidator behavior changes when the market volatility changes. We assess this by sweeping through a variety of different volatility levels to ensure that we emulate how greedy liquidators interact with a plethora of market environments. Note that we normalize our volatility calculation in a manner akin to what is used by exchanges such as BitMEX.[19]

5. SIMULATED STRESS TESTS

5.1 Agent-Based Simulation

The main tool that we use to perform simulation-based stress tests on Compound's Ethereum smart contracts is agent-based simulation (ABS). ABS has been used in a variety of stress test contexts, including to estimate censorship in cryptocurrency protocols,[20] detect fraudulent trading activity in CFTC exchanges,[21] and in stress testing frameworks from the European Central Bank[22, 23] and the Fed-

⁷Mathematically, this means that there exists a function $f : [0, \infty) \rightarrow \mathbb{R}$ such that $\Delta p(q) = \sigma f(q) + o(1)$; see [17, 18] for theoretical and empirical evidence of this. In particular, note that this appears to hold for many markets in terms of *permanent* impact cost, whereas instantaneous impact cost tends to depend much more on an asset's microstructure details.

eral Reserve.[24, 25] These simulations, while powerful, can be difficult to make both useful and accurate as model complexity can make it hard to match experimental results.[26] Careful design, tuning, and infrastructure architecture can help avoid these pitfalls and has made ABS invaluable in industries such as algorithmic trading and self-driving car deployment.

In such industries, one takes care to ensure that the simulation environment replicates the live environment as closely as possible. This is enforced by having the agent models interact with the same code that is deployed in a live environment in order to minimize errors due to mistranslations or missing minutiae. While the infrastructure overhead of simulating users interacting with a piece of complex software can be heavy, it ensures that errors are limited to those in models of agents as opposed to errors in the models of system dynamics.

As an example, the Compound interest rate curve (Appendix 9.2) is described via a simple mathematical formula. One can simulate agents directly interacting with this formula, without needing to host the Ethereum environment and having the agents generate transactions. However, Ethereum’s 256-bit numerical system and precision differences between different ERC-20 contracts can often lead to disastrous losses due to numerical errors. These cannot be probed without running simulations directly against the Ethereum smart contract and generating the exact same transactions that an agent would if they were a liquidator interacting with the live contract.

5.2 Gauntlet Simulation Environment

The Gauntlet platform, which was used for all simulations and results in this paper, provides a modular, generic ABS interface for running simulations directly against Ethereum smart contracts. In this system, the agent models are specified via a Python domain-specific language (DSL), akin to Facebook’s PyTorch,[27] and interact with a custom-built Ethereum virtual machine that is written in C++. Agents can also interact with non-blockchain modules, such as historical or synthetic market data and/or other off-chain systems. Gauntlet has made significant performance optimizations for interacting with the EVM in Python, resulting in performance gains of 50-100x over the stock tooling. The DSL hides the blockchain-level details from the analyst, allowing the end-user to develop strategies that can migrate from one smart contract to another, should they have similar interfaces. Most of the platform’s design is inspired by similar platforms in algorithmic trading that allow for quantitative researchers to develop strategies that execute over multiple exchanges (with varying order books, wire protocols, slippage models, etc.) without having to know these low-level details. Moreover, the non-blockchain portions of the simulation are analogous to trading back-testing environments,[28] so that agents are interacting with realistic order books and financial data. It should be noted that the strategies emit valid EVM transactions and can be deployed to Ethereum mainnet using the same code path.

5.3 Compound Simulation Overview

For the simulations in this paper, we deployed the Compound Solidity contracts within the Gauntlet platform. The simulation environment tracks all the gas used by the transactions, as if users paying transaction fees for submitting

transactions to the Ethereum blockchain. In particular, we use constant gas costs throughout all simulations detailed in this paper to calculate the transaction fees. We use a standard Geometric Brownian motion to simulate price trajectories. This stochastic process S_t obeys the Itô stochastic differential equation, $dS_t = \mu S_t dt + \sigma S_t dW_t$, where W_t is the standard Wiener process and μ is the percentage drift. To understand the sensitivity of the protocol safety to the volatility of the underlying price shock, we ran a set of simulations by varying the volatility σ , as discussed in Section 4.1.2.

We implemented liquidator strategies in our DSL, which allowed for a variety of liquidators with different risk and time preferences to interact directly with the Compound contracts and with simulated order books. These strategies also include optimization components so that liquidators can optimize the amount of collateral purchased based on their slippage estimates. Our simulation uses a linear slippage model: $\Delta p(q) = I\sigma q$, where the intensity I is a model parameter estimated by fitting the Coinbase Pro order book data. We also wrote strategies for borrowers in the Compound protocol using the DSL and fit their risk preferences based on Compound’s historical on-chain data, specifically the collateralization ratio and collateral value.


The simulation assumes that borrowers use ETH as collateral and are borrowing the stablecoin DAI from the Compound protocol. We calculated the initial collateral value and the size of borrowing based on the total outstanding debt, as explained in Section 4.1.1, and the historical collateralization ratio data. Since the price shocks are affecting the collateral value, each liquidator continuously evaluates all borrower’s collateralization ratio and repays DAI on the borrower’s behalf when an arbitrage opportunity exists. Once a liquidator receives the borrower’s collateral, she immediately sells it on a centralized exchange to arbitrage. The arbitrage opportunity exists when the liquidation incentive is greater than the sum of total costs, including slippage, transaction fee, and centralized exchange trading fee. We stress tested a wide range of plausible shocks, as well as market sizes, and analyzed the simulation outcomes.

6. KEY QUESTIONS

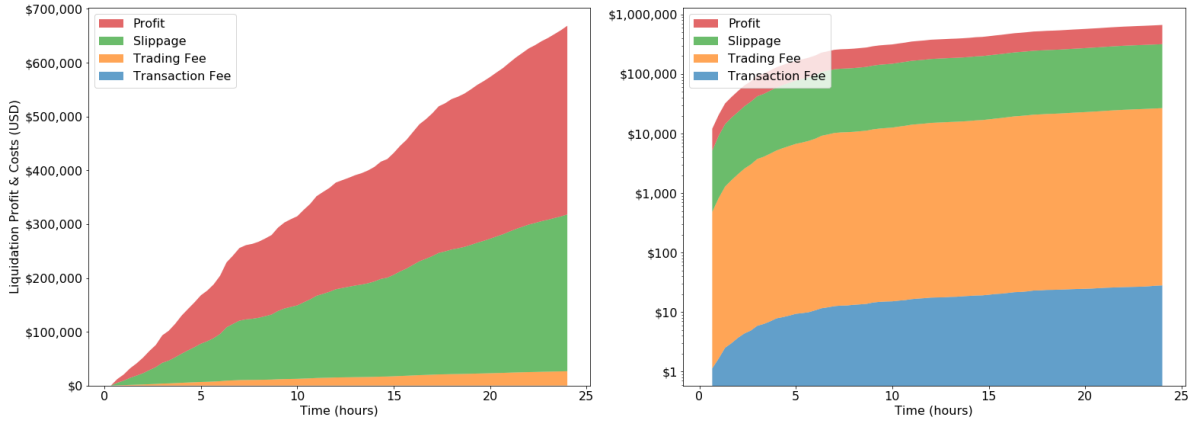
From a liquidity supplier’s perspective, the protocol is safe only if the supplied assets can be safely withdrawn. A functioning liquidation mechanism is critical to the safe operation of the Compound market. When an asset price drops and no liquidators have an incentive to repay the borrower’s outstanding debt, the system fails and some suppliers cannot withdraw their assets. Recall that a rational liquidator’s goal is to make a profit in each liquidation opportunity, which depends on the liquidation incentive and slippage (this is dependent on the trade size and volatility). In light of this, the main questions that we focus on answering are the following:

- Is the protocol safe when the total outstanding debt is high?
- Is the protocol safe under volatile market conditions?

We will first define some metrics that will help us answer these questions in a quantitative manner. An *undercollateralized run* is a simulation run that ends with $>1\%$ of the value of the market’s total outstanding debt that is



(a) 3% daily volatility



(b) 50% daily volatility

Figure 1: Mean cumulative liquidator profit and costs over 30 simulation runs. Note that the y -axis on the left-hand side is using a linear scale (in dollars), whereas the right-hand side is using a logarithmic scale. The simulation assumes \$100MM USD of ETH daily trading volume and the sum of the total outstanding debt value is \$50MM USD.

under-collateralized. Let the *under-collateralized run percentage* be defined as the percentage of simulation runs that are under-collateralized runs. This metric is used to quantify the safety of the system, as the system will be at risk if borrowers with a large amount of outstanding debt are under-collateralized. As we want to ensure that the system is never under-collateralized, we use a strict 1% debt threshold to define the failure criteria.

7. RESULTS

As was discussed in section 4.1, the collateral asset’s quantity to be traded and the asset’s volatility are two major market variables causing slippage, and slippage is one of the main factors influencing a liquidator’s behavior. This suggests that the protocol’s safety heavily depends on the total outstanding debt and the collateral asset’s volatility.

In our simulation, the total outstanding debt is defined as the asset pool’s total outstanding stablecoin debt value normalized by the collateral asset’s daily trading volume. Considering that different collateral assets have different orders of magnitude of trading volume, normalizing the total outstanding debt enables us to intuitively compare the debt

(relative to the collateral asset’s liquidity) between different collateral assets. The simulation time duration is a day, hence we use daily volatility instead of commonly used annualized volatility to make it straightforward to understand.

There is no strong agreement on the daily trading volume of most crypto tokens. Centralized exchanges are susceptible to wash trading, and decentralized exchanges are dwarfed by their centralized counterparts. As the ability to sell collateral quickly is one of the driving factors of safety, this creates an uncertainty that is addressed via simulation. By varying the ratio of outstanding debt to market size widely in our simulations, we cover a broad swath of scenarios that you might see in the practice. If you have very conservative assumptions on the total market depth of the collateral order book, you can assume a higher ratio of debt. Our assumptions on ETH market size are fairly conservative (\$100mm), falling on the lower end of Messari’s daily trading volumes for the beginning of 2020.

In Figure 1, we see liquidator profit and loss charts broken up into transaction fee, trading fee, slippage, and profit. There are more arbitrage opportunities when the asset’s volatility is high (Figure 1b), and subsequently the liquida-

tor’s total revenue (the sum of profit and costs) is higher than the revenue in the low volatility regime (Figure 1a). The chart demonstrates that price slippage is the major cost of arbitrage. In the high volatility scenario, slippage represents more than 30% of the liquidator’s revenue. Even in the low volatility scenario, the liquidator still has to pay more than 10% of the revenue for slippage. In both scenarios, the trading fee takes a fixed percentage of revenue and the on-chain transaction fee is insignificant at this level of the total outstanding debt.






Figure 2: Total liquidated debt amounts over 24 hour period. The simulation assumes \$100MM USD of ETH daily trading volume. A 75% total outstanding debt is equivalent to \$75MM USD worth of the total outstanding debt value. To explain an example cell here: \$20mm in liquidations corresponds to 20% of the Total Outstanding Debt in the 100% case, which is intuitively a worrying number liquidations. However, this does match intuition, because you only see this happen when ETH has a worrying level volatility (close to 50%)

Figure 2 shows the total liquidated debt amount with different initial total outstanding debt and ETH volatility. The results match our intuition: the total liquidated debt amount is proportional to both total outstanding debt and volatility. In the high volatility scenario, a borrower’s collateral value has a high chance to fall below the collateral requirement and, as a consequence, the collateral will get liquidated. Though liquidations are a necessary part of the Compound protocol, they can also serve as a leading indicator of under-collateralization.

Figure 3 demonstrates the liquidatable and under-collateralized run percentages heatmap. Recall that an account becomes liquidatable if the collateral value falls below the collateral requirement and the collateral is available to be liquidated. When the price of the collateral asset drops further and the collateral value is below the outstanding debt value, an account becomes under-collateralized. Here we set a strict 1% debt threshold to define the failure criteria, i.e., the simulation run fails when over 1% of the outstanding debt is liquidatable/under-collateralized. For each cell in the heatmap, we aggregate the results from 30 simulation runs with the same market variables and calculate the percentage of the runs that fail. The lighter the data point is, the fewer simulation runs fail. If the data point is white, the



(a) Percentages of simulation runs that end with > 1% of liquidatable debt



(b) Percentages of simulation runs that end with > 1% of under-collateralized debt

Figure 3: The Compound contracts are deployed with the default parameters (75% ETH collateral factor and 1.05 liquidation incentive). According to BitMEX weekly historical ETH volatility index, the current daily volatility is around 3% and the highest historical daily volatility is around 20%. The simulation assumes \$100MM USD of ETH daily trading volume. Compound’s current total outstanding stablecoin debt value is around \$25MM USD. The current total outstanding debt is around 25%, which is the total outstanding debt value normalized by the daily trading volume of the collateral asset.

protocol is safe and none of the simulation runs have more than 1% of the under-collateralized debt. We use this metric to quantify the safety of the protocol.

The heatmaps demonstrate how large the protocol can scale under a reasonable volatility assumption. The BitMEX weekly historical ETH volatility index reports that the highest ETH weekly volatility in history happened in August 2017 and peaked at around 20% daily volatility⁸. [19] Assuming that the ETH market capitalization will grow over time and the volatility will decrease, we consider daily volatility < 20% as reasonable. Figure 3a shows that when the daily volatility is 20% and the total outstanding debt is greater than or equal to \$100 MM USD, a few risky bor-

⁸The daily volatility is converted from the realized weekly volatility.

rowing positions will not be fully liquidated at the end of the simulation runs. However, with the same 20% daily volatility assumption, none of the borrowers are under-collateralized, and the protocol can scale to at least 10x the current borrow size, as shown in Figure 3b.

Figure 3b highlights the safe operating space of the protocol. The protocol is safe when the volatility is below 35% and the liquidity pool’s total outstanding debt value is below ETH’s daily trading volume. As the volatility reaches 45%, some suppliers may be unable to withdraw their supplied assets.

8. CONCLUSIONS

In this paper we conducted a market-risk assessment of the Compound protocol via agent-based simulations run against the Compound contracts. We stress-tested the liquidation mechanism under a wide range of market volatility and sizing scenarios to ensure that the protocol can prevent borrowers from becoming under-collateralized in most of these cases. We also used historical market data from centralized cryptocurrency exchanges to ensure that assumptions about volatility and slippage are representative of real-world conditions.

We found that the protocol, as currently parameterized, should be robust enough to scale to at least 10x the current borrow size as long as ETH price volatility does not exceed historical highs. Our methodology can also be applied to other collateral types on Compound with significantly different liquidity profiles, such as REP. This work informs the Compound community on how to evaluate market risk for new assets as they are added to the protocol.

9. APPENDIX

9.1 Glossary

- Debt: Amount of asset borrowed from an asset pool.
- Under-collateralized: An account is under-collateralized if the value of an account’s debt exceeds the value of the collateral.
- Collateral factor: Maximum debt-to-collateral ratio of an asset a user may borrow. When the debt-to-collateral ratio exceeds the collateral factor, the collateral is available to be liquidated.
- Collateralization ratio: The ratio of collateral-to-debt, usually reported in percentage points. For instance, a collateralization ratio of 200% means that one needs two times as much collateral deposited into the contract as the maximum borrow quantity. Concretely, this would mean that one must deposit \$200 worth of ETH in order to borrow \$100 of a stablecoin.
- Borrowing capacity: Current value of collateral deposited into the contract multiplied by the collateral factor.
- Collateral requirement: Value of debt divided by the collateral factor.
- Liquidatable: An account is liquidatable if the account’s value of debt exceeds its borrowing capacity. In other words, an account is liquidatable if the account’s collateral value falls below the collateral requirement.

- Slippage: The amount of price impact that a liquidator engenders when trying to sell collateral. Slippage is denoted $\Delta p(q)$ and is formally defined as the difference between the midpoint price at time t , $p_{\text{mid}}(t)$ and the execution price, $p_{\text{exec}}(q, t)$ for a traded quantity q at time t , $\Delta p(q, t) = p_{\text{mid}}(t) - p_{\text{exec}}(q, t)$. This quantity is usually a function of other variables, such as implied and realized volatilities. Slippage is also known as market impact within academic literature.

9.2 Interest Rate Curves

Within the cryptocurrency space, *bonding curves* are deterministic functions of smart contract state that determine bid and ask spreads. Bonding curves are known as *pricing rules* within the algorithmic game theory literature and were first introduced by Hansen[29] in the study of automated market makers.[30] These were first introduced to Ethereum smart contracts by de la Rouviere[31] as a way to create tokenized markets whose buy and sell prices were determined algorithmically. Instead of using a bonding curve to provide bids and offers, the Compound protocol utilizes a bonding curve to compute the spread between the supply and borrowing interest rates. One that think of this as an analogue of the traditional yield curve from finance, albeit computed in a different manner.

The contract also uses the bonding curve to enforce the no-arbitrage condition that the supply interest rate must be strictly lower than the borrowing interest rate. If this were not true, then an arbitrageur could break the system by borrowing cTokens from the contract and adding liened tokens to the liquidity supply, leading to net negative liabilities. Moreover the contract also enforces softer constraints that control the difference between the supply and borrowing interest rates. The main idea behind the curve used in Compound is that if there is more liquidity supply than borrowing demand, then the interest rate to supply liquidity should be significantly lower than the interest rate to borrow.

Formally, the Compound V2 smart contract[1] constructs the bonding curve as a function of the *utilization rate* at block height t , $U_t \in [0, 1]$. If we denote the borrowing demand at height t (in tokens) as B_t and the liquidity supply at height t as L_t , then the utilization rate is defined as

$$U_t = \frac{B_t}{L_t + B_t}$$

We compute the borrowing interest rate, β_t and the supply interest rate, ℓ_t , using the following formulas, where $\beta_0, \beta_1 \in (0, 1)$ are interest-rate parameters and $\gamma_0 \in (0, 1)$ is a measure of the spread between supply and borrowing (i.e. $1 - \gamma_0$ is the relative spread).

$$\beta_t = U_t(\beta_0 + \beta_1 U_t) \tag{1}$$

$$\ell_t = (1 - \gamma_0)\beta_t \tag{2}$$

For reference, the Compound V2 contract uses the values $\beta_0 = 5\%$ and $\beta_1 = 45\%$. The choice of quadratic bonding curve has a variety of benefits that have been profiled in a number of articles and papers.[32, 33]

References

- [1] Robert Leshner and Geoffrey Hayes. Compound: The money market protocol. Technical report, February 2019.
- [2] Robert D Arnott, Jason Hsu, and Philip Moore. Fundamental indexation. *Financial Analysts Journal*, 61(2):83–99, 2005.
- [3] Stephanie E Curcuro, Charles P Thomas, Francis E Warnock, and Jon Wongswan. Uncovered equity parity and rebalancing in international portfolios. *Journal of International Money and Finance*, 47:86–99, 2014.
- [4] Jason C Hsu. Cap-weighted portfolios are sub-optimal portfolios. *Journal of Investment Management*, 4(3), 2004.
- [5] Arian Klages-Mundt and Andreea Minca. (in) stability for the blockchain: Deleveraging spirals and stablecoin attacks. *arXiv preprint arXiv:1906.02152*, 2019.
- [6] Jim Gatheral and Alexander Schied. Optimal trade execution under geometric brownian motion in the almgren and chris framework. *International Journal of Theoretical and Applied Finance*, 14(03):353–368, 2011.
- [7] Bence Tóth, Yves Lempereire, Cyril Deremble, Joachim De Lataillade, Julien Kockelkoren, and J-P Bouchaud. Anomalous price impact and the critical nature of liquidity in financial markets. *Physical Review X*, 1(2):021006, 2011.
- [8] Zoltan Eisler, Jean-Philippe Bouchaud, and Julien Kockelkoren. The price impact of order book events: market orders, limit orders and cancellations. *Quantitative Finance*, 12(9):1395–1419, 2012.
- [9] Jim Gatheral. No-dynamic-arbitrage and market impact. *Quantitative finance*, 10(7):749–759, 2010.
- [10] Igor Makarov and Antoinette Schoar. Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, 2019.
- [11] Wang Chun Wei. Liquidity and market efficiency in cryptocurrencies. *Economics Letters*, 168:21–24, 2018.
- [12] Guillermo Angeris, Hsien-Tang Kao, Rei Chiang, Charlie Noyes, and Tarun Chitra. An analysis of uniswap markets. *arXiv preprint arXiv:1911.03380*, 2019.
- [13] Alameda Research. Investigation into the legitimacy of reported cryptocurrency exchange volume, Jul 2019.
- [14] Matthew Hougan, Hong Kim, Micah Lerner, and Bitwise Asset Management. Economic and non-economic trading in bitcoin: Exploring the real spot market for the world’s first digital commodity. *Bitwise Asset Management*, May 2019.
- [15] Bitwise Asset Management. Bitwise crypto index methodology, Mar 2019.
- [16] Securities and Exchange Commission. Self-regulatory organizations; nyse arca, inc.; order disapproving a proposed rule change, as modified by amendment no. 1, relating to the listing and trading of shares of the bitwise bitcoin etf trust under nyse arca rule 8.201-e, Oct 2019.
- [17] John Hull. *Introduction to futures and options markets*. Prentice Hall Englewood Cliffs, NJ, 1991.
- [18] Robert Almgren, Chee Thum, Emmanuel Hauptmann, and Hong Li. Direct estimation of equity market impact. *Risk*, 18(7):58–62, 2005.
- [19] BitMEX. .evol7d: Weekly historical ether volatility index, 2020.
- [20] T. Chitra, M. Quaintance, S. Haber, and W. Martino. Agent-based simulations of blockchain protocols illustrated via kadana’s chainweb. In *2019 IEEE European Symposium on Security and Privacy Workshops (EuroS PW)*, pages 386–395, June 2019.
- [21] Steve Yang, Mark Paddrik, Roy Hayes, Andrew Todd, Andrei Kirilenko, Peter Beling, and William Scherer. Behavior based learning in identifying high frequency trading strategies. In *2012 IEEE Conference on Computational Intelligence for Financial Engineering & Economics (CIFER)*, pages 1–8. IEEE, 2012.
- [22] Grzegorz Hala. Agent-based model of system-wide implications of funding risk. 2018.
- [23] Anqi Liu, Mark Paddrik, Steve Y Yang, and Xingjia Zhang. Interbank contagion: An agent-based model approach to endogenously formed networks. *Journal of Banking & Finance*, 2017.
- [24] Richard Bookstaber, Mark Paddrik, and Brian Tivnan. An agent-based model for financial vulnerability. *Journal of Economic Interaction and Coordination*, 13(2):433–466, 2018.
- [25] John Geanakoplos, Robert Axtell, J Doyne Farmer, Peter Howitt, Benjamin Conlee, Jonathan Goldstein, Matthew Hendrey, Nathan M Palmer, and Chun-Yi Yang. Getting at systemic risk via an agent-based model of the housing market. *American Economic Review*, 102(3):53–58, 2012.
- [26] Giorgio Fagiolo and Andrea Roventini. Macroeconomic policy in dsge and agent-based models redux: New developments and challenges ahead. *Available at SSRN 2763735*, 2016.
- [27] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pages 8024–8035, 2019.
- [28] Peter Nystrup, Stephen Boyd, Erik Lindström, and Henrik Madsen. Multi-period portfolio selection with drawdown control. *Annals of Operations Research*, 282(1-2):245–271, 2019.
- [29] Robin Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, 2003.
- [30] Abraham Othman, David M Pennock, Daniel M Reeves, and Tuomas Sandholm. A practical liquidity-sensitive automated market maker. *ACM Transactions on Economics and Computation*, 1(3):14, 2013.

- [31] Simon de la Rouviere. Tokens 2.0: Curved token bonding in curation markets, Nov 2017.
- [32] Tarun Chitra. Competitive equilibria between staking and on-chain lending. *arXiv preprint arXiv:2001.00919*, 2019.
- [33] Alexander Obadia. Exploring defi trading strategies: Arbitrage in defi, Nov 2019.