Empirical Evaluation of MakerDAO's Resilience

Martin Kjäer, Monika di Angelo, Gernot Salzer
TU Wien, Vienna, Austria
{martin.kjaeer,monika.di.angelo,gernot.salzer}@tuwien.ac.at

***Abstract*—Stablecoins are cryptocurrencies with the aim to reduce the price volatility by design. This increases their acceptance as an instrument of payment. Centralized approaches facilitate stable organizational structures, at the cost of introducing dependencies. Decentralized projects, on the other hand, face several challenges, one of them at the core of stablecoins: stability.**

In this paper, we assess the stability of the MakerDAO protocol, one of the major decentralized stablecoins. We conduct a measurement study regarding MakerDAO's resilience during the first year of its full protocol, from November 2019 to 2020, including the cryptocurrency crisis in March 2020. Our analysis is based on the publicly available transaction data of Ethereum that documents the activities of MakerDAO's contracts executing the protocol. We state requirements for resilience and introduce empirical metrics, based on which we test the resilience hypotheses against the transaction data. While the events of March 2020 were unfortunate for several users, the resilience at protocol level proved to be good even under stress.

***Index Terms*—case study, DeFi, Ethereum, measurement study, stablecoin, stability, transaction data**

I. Introduction

The price volatility of cryptocurrencies is one of their less desirable properties, hence approaches to stabilize the exchange rates were discussed early on, resulting in *stablecoins*. These crypto tokens have low price volatility and can therefore be used as a reliable medium for the exchange of values.

While different approaches to stabilize cryptocurrencies keep appearing, many aspects and issues still need to be addressed both theoretically and practically, especially with regard to the assessment of the stability provided. When evaluating a stablecoin project, there is a particular focus on the actual price stability itself – the very aspiration of the project. However, other influencing factors need to be considered as well: (i) the stability mechanism, (ii) the project governance, (iii) the implementation of the project including the underlying platform, and (iv) the interplay of the components.

Stability in general is discussed theoretically in works like [1]–[3]. On a more practical side, there are many stablecoin projects in the wild that offer a live experience of their implemented ideas. The project with the most stablecoins issued to date is Tether, centrally operated by Tether Limited for many years. Regarding decentralized stablecoin projects, MakerDAO is by far the largest. It operates on Ethereum and is among the most trusted ones. Remarkably, it demonstrated the ability to keep its stablecoin DAI close to the USD (soft-pegged) during most times while maintaining the project decentralized.

The practical stability of stablecoins is at risk as the *Black Thursday* price crash during March 12 and 13, 2020 showed. This incident provides valuable empirical data for an evaluation of resilience under stress conditions. With a focus on MakerDAO, the authors of [3] explain the Black Thursday using a stochastic model, while the authors of [4] analyze its pricing oracle and decentralized governance mechanism. However, a detailed empirical analysis of the transaction data of MakerDAO with respect to its practical stability, especially in stress situations like Back Thursday is still missing.

Approach. In this work, we evaluate the resilience of MakerDAO from November 2019 to 2020 with a focus on Black Thursday by means of a case study based on [5]. Due to the fact that almost every part of Maker's protocol is managed in a decentralized way, we look for answers in the actual transaction data. To this end, we first derive requirements for resilience and then define metrics that allow us to evaluate hypotheses concerning these requirements by applying the metrics to the public transaction data. This analysis also delivers a detailed picture of the stablecoin DAI in practice and the usefulness of the metrics.

Contribution: We add to the discussion on the stability of stablecoins with an empirical study of resilience to stress. Furthermore, we explore the informative value of empirical metrics for the particular decentralized stablecoin DAI.

Stakeholders. Stablecoins are particularly interesting to examine because they enable numerous other applications and can provide insight into current and future developments in decentralized business models. A detailed discussion seems relevant from an academic perspective as well as for enterprises, publicly funded research institutions and venture capital firms interested in use cases for blockchains in general, and in decentralized finance (DeFi) in particular. Moreover, our findings may be useful for designers of stablecoin projects.

Roadmap: In section II, we summarize the MakerDAO protocol. In section III, we discuss related work. In section IV, we detail our methods for the investigation. In section V, we analyze the Maker protocol regarding its resilience. In section VI, we discuss the findings and conclude in section VII.

II. MakerDAO

The Maker[1] protocol [6] is designed to run as a decentralized autonomous organization (DAO) on a public per-

[1] 'Maker' and 'MakerDAO' are used interchangeably, even by MakerDAO.

978-1-6654-3924-4/21/$31.00 ©2021 IEEE

missionless blockchain. Rules and behavior of the DAO are implemented as a set of smart contracts on the Ethereum mainnet. The protocol is developed and maintained by the Maker Foundation (makerdao.com), while governance and development are gradually handed over to the Maker community to make the project completely decentralized and independent.

The purpose of Maker is to manage the ERC-20 token DAI as a stablecoin. Though soft-pegged to the USD, it is not backed by dollars directly, but by other assets that can be managed on the Ethereum blockchain in a decentralized manner (proxy collateralization). In 2017, Maker launched with ETH as the only collateral, but extended the protocol in November 2019 to handle multiple collaterals, making the so-called Multi-Collateralized DAI (MCD) more stable against fluctuations of collateral prices.

A. Maker Vaults

Vaults are the centerpiece of Maker for taking DAI loans against collateral. To receive DAI, borrowers lock a self-determined amount of collateral (above a minimum prescribed by Maker) as security in the vault. As collaterals differ in price volatility, the liquidation ratio is defined separately for each type of collateral and can be adjusted by Maker governance. For ETH, it is 150 %, which means that the USD value of the ETH in the vault must be at least 1.5 times the number of DAI taken as a loan.

To finance governance and development, Maker charges a stability fee to users who take out loans. The fee depends on the amount of DAI as well as on the loan time. It has to be paid by the vault owner when returning the DAI and unlocking the collateral.

B. Governance

Maker uses the ERC-20 token MKR for its decentralized governance processes. Initially, MKR tokens were distributed among early contributors, advisors, and team members of the Maker foundation. Later, tokens were also sold to venture capital firms, while a strategic supply is still set aside to fund the project until it is fully self-sustaining. MKR tokens grant its holders the right to participate in executive votes on changes to the Maker ecosystem and in governance polls. Executive votes may have on- or off-chain effects, like the employment of software developers in the name of the DAO.

C. Oracles

Maker uses oracles to track the price of collaterals. Figure 1 depicts the infrastructure. Feed providers (individuals, corporate entities, and other DAOs) collect price data from off- and on-chain sources, compute the median value, sign it and distribute it via a peer-to-peer relayer network. The network forwards the collected price data to a contract `Median` on the Ethereum blockchain (there is one for each type of collateral). The contract `Median` checks the signatures as well as the timestamps of all feeds which have to be greater than the values last received. Then, it saves the median of all feeds so that the price can be queried on-chain from the contract `OSM` (Oracle Security Module). The contract `Median` can publish prices at any frequency, whereas the `OSM` can only be called once per hour. At each call, the `OSM` takes the last published price from `Median` and uses it to update itself in the subsequent call. The delay allows (i) vault owners to add collateral or repay DAI debt in the event of a price drop and (ii) Maker to monitor prices and to act in case oracles are manipulated.

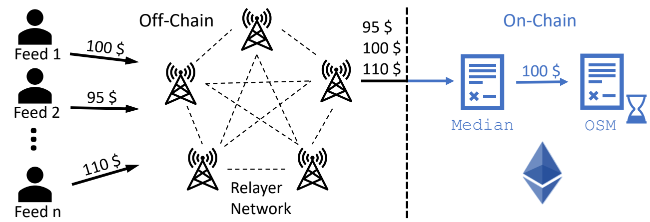

Fig. 1. Oracle Architecture

D. Collateral Auctions

To ensure that the DAI keeps its value even when the price of a collateral drops, locked collateral can be auctioned if a vault becomes under-collateralized. The liquidation of the collateral starts with external actors, *keepers*, calling the function `bite` of the Maker contract `Cat`, which will confiscate collateral. The incentive of keepers is a chance to buy collateral below the market price at the subsequent auction, where the seized collateral is sold to cover the outstanding DAI. The collateral has to be paid in DAI to reduce the supply of the stablecoin to guarantee that all minted DAI is always backed by sufficient collateral.

Auctions consist of two phases. First, the participants bid increasing amounts of DAI calling the function `tend`. Once a bid covers the outstanding DAI debt, the second phase starts, where participants bid via the function `dent` decreasing amounts of collateral they are willing to accept in exchange for the fixed amount of outstanding DAI. The DAI of a bid are locked in the auction contract `Flipper` until a higher bid comes in or the auction ends after the time period `ttl` (set by Maker governance). If the auction reaches the `dent` phase, all collateral that was not auctioned off, will be returned to the Vault at the end of the auction.

III. Related Work

In this section, we first summarize characteristics of stablecoins and their impact. Then, we cover prize stabilization as well as empirical studies and measurements.

A. Characteristics and Impact

In 2014, Buterin put forward design ideas on stablecoins [7]. A few years and many stablecoin projects later, the authors of [8] provide a classification framework for stablecoins with a compact overview of many coins. In a recent survey on decentralized finance (DeFi) [9], the authors describe the ecosystem and the role of stablecoins, and outline open research challenges.

Regarding legal aspects, the authors of [10] "highlight the benefits of stable cryptocurrencies for monetary policy making, overall market stability, and their bilateral impact on the emergence of decentralized commerce."

The G7 Working Group on Stablecoins [11] discusses challenges and risks of global stablecoins mostly from the point of view of public policy and regulation. They stress the importance of assessing a stablecoin system as a whole due to its complexity and "unpredictable interactions between components following a disruption in any individual component." The authors of [12] empirically show, "that (i) stablecoins can serve as safe havens in specific situations, although most act merely as an effective diversifier in normal market conditions."

B. Price Stabilization

The taxonomy of price stabilization in [1] is based on more than 20 stablecoin projects. The authors find that over a third of the projects are vulnerable to "speculative economic attacks – an avoidable design flaw."

In an analysis of price stabilization mechanisms, the authors of [2] claim that the following changes to blockchain-based cryptocurrencies render them stable: "(1) limiting re-adjustment of proof-of-work targets, (2) making mining rewards variable according to the observed over-threshold changes of block intervals, and (3) enforcing negative interests to remove old coins in circulation."

The stochastic modeling approach of stablecoins in [13] enables the authors to derive results about dynamics and liquidity in stablecoin markets. Moreover, they demonstrate how deleveraging spirals can cause illiquidity and collateral loss of value, and suggest design improvements. In [3], they employ the model to explain the Black Thursday crisis and "formally characterize a deflationary deleveraging spiral."

With a modeling framework for formalizing algorithmic stablecoins, the authors of [14] address the question whether algorithmic stablecoins are volatile by design or in practice. Moreover, they claim that "the framework can identify critical conditions under which stablecoins might become volatile." They applied the framework to the stablecoin Basis Cash.

In [15], the authors use the standard deviation of return rates and detect differences between distributions of this measure for various stablecoins. They show that the stablecoins deliver unequally "on the promise to provide stable market value, with tokenised funds being leaders." The analysis is based on data up to September 2019.

C. Empirical Studies and Measurements

In their empirical evaluation [4], the authors focus on the accuracy of the pricing oracle and the robustness of the governance of Maker, and recommend improvements.

In [16], the authors analyze the performance of several stablecoins including Tether and DAI from November 2019 to May 2020. To measure instability, they apply a volatility estimator (based on the exponentially weighted moving average) to the stablecoin exchange rate log-returns.

In their exploration of weaknesses in DeFi protocols [17], the authors present a scenario for lending a stablecoin against ETH (like the Maker protocol) before March 2020. They point out the decline in liquidity of the stablecoin resulting from a rapid price drop of the collateral. Moreover, they demonstrate an attack on Maker governance and point out that the failure of one protocol may spread to others.

D. Our Approach

We address the assessment of price stabilization in the case of a particular stablecoin (DAI) from an empirical perspective. Thus, our approach complements the more theoretical results of [3], [13]. Regarding methods, our work is closely related to [4] since we also analyze historical data using empirical metrics. In contrast to [16], we do not employ methods from quantitative finance as we do not assess market risks. Also, we are not interested in correlating volatility to stability mechanisms in general as in [15]. With respect to the topics covered, we complement [4], [16], [17] by analyzing further aspects of Maker.

IV. Data and Methods

In this section, we first describe the data forming the basis of our investigation. Then we detail our case study including requirements for resilience and empirical metrics.

A. Data

We run an Open Ethereum[2] client that provides transaction data and execution traces. Moreover, we compare this data to AnyBlock's analytics tool[3] for consistency checks. Since Maker's Multi-collateral DAI went live on November 18, 2019, numerous upgrades to its protocol have been released. A change log [18] tracks the addresses of new and re-deployed contracts. To obtain correct and complete data, it is important to select the correct versions of contracts. Our analysis is primarily based on the logs written by the Maker contracts.

B. Case Study

We follow the approach suggested by [5] for case studies, with Maker itself as the case, represented by the transactions recorded on the Ethereum blockchain. We restrict our investigation to the Multi-Collateral DAI (MCD), as it is more advanced than the Single-Collateral DAI (SCD) and implements the full protocol. Moreover, we consider only ETH-A vaults, since they are the most common ones. We limit our analysis to the first operative year of MCD, from November 18, 2019 to November 17, 2020. We link the data to hypotheses by means of empirical metrics defined below. To interprete the findings, we reason about what the hypotheses indicate about the resilience of Maker to stress.

[2] https://github.com/openethereum/openethereum

[3] https://www.anyblockanalytics.com/docs/sql/schema/

C. Resilience

From the design of Maker (cf. section II), we extract the following requirements, which are essential for its resilience.

- *Auctions achieve adequate prices for the collateral.* The idea of a collateral is to be sold when necessary, preferably at market value.
- *Auctions are able to cover outstanding DAI debt.* With the collateralization required by the protocol, most auctions should actually cover the debt.
- *Vault owners maintain a collateralization at double the liquidation on average.* Usual price fluctuations of the collateral should not threaten a vault with liquidation, hence double the collateralization is a safeguard.
- *Vaults threatened with liquidation adjust their collateralization in time.* Liquidation of a vault not only carries the risk of ineffective auctions, oracle failures or smart contract bugs, but also incurs a liquidation penalty of 13 %.
- *Keepers liquidate vaults quickly after they have fallen below the liquidation ratio.* Falling collateral prices require fast reaction so that the outstanding debt can be covered.

We hypothesize that Maker fulfills these requirements for the study period. To test the hypotheses, we define several metrics.

D. Metrics

The following empirical metrics are inspired by observations we made about the course of events on Black Thursday.

M1 *Bite events per week or per hour:* We count the `Bite` events emitted by the contract `Cat` per week (with Monday as start of week), or per hour.

M2 *USD price of ETH:* By $\texttt{Median}(t)$, we denote the value of ETH in USD as reported by the `LogMedianPrice` events of contract `Median`. $\texttt{Median}(t)$ reflects the external value of ETH at time t.

M3 *Collateralization of vaults:* For a time t and a vault v, let $\mathrm{ETH}_v(t)$ be the amount of ETH locked and $\mathrm{DAI}_v(t)$ the amount of DAI taken as a loan. Moreover, let $\texttt{OSM}(t)$ be the ETH value in USD as reported by the event `LogValue` of the contract `OSM`. Then the collateralization of the vault v at time t is given by

$$\mathrm{Coll}_v(t) = \mathrm{ETH}_v(t) \times \texttt{OSM}(t)/\mathrm{DAI}_v(t)$$

To determine $\mathrm{ETH}_v(t)$ and $\mathrm{DAI}_v(t)$, we accumulate all movements of collateral and DAI in and out of the vault, by observing the calls `frob`, `fork` and `grab` to the contract `Vat`.

M4 *Auction effectiveness:* The effectiveness of an auction a is the ratio of DAI received to ETH sold.

$$\mathrm{AE}_a = \mathrm{DAI}_a/\mathrm{ETH}_a$$

M5 *Scaled auction effectiveness:* The scaled effectiveness of an auction a relates the DAI received to the value of the sold ETH at the time t_a, when the auction ends.

$$\begin{aligned}\mathrm{SAE}_a &= \mathrm{DAI}_a/(\mathrm{ETH}_a \times \texttt{Median}(t_a)) \\ &= \mathrm{AE}_a/\texttt{Median}(t_a)\end{aligned}$$

We determine DAI_a, ETH_a and t_a from the log data of the auction contract `Flipper`, in particular from the calls to the functions `tend` and `dent`.

M6 *Liquidation delay:* The liquidation delay for a vault is the time between the vault becoming under-collateralized and the vault getting bitten. Technically, we measure the time between the block, where the contract `OSM` emits an event `LogValue` with an ETH price that drives the vault below the liquidation ratio, and the block, where the contract `Cat` emits a `Bite` event for this vault, with the vault remaining under-collateralized in-between.

V. Analysis

In this section, we start with an overview of the events on Black Thursday before we investigate Maker in detail. We analyze `Bite` events, vault management and collateralization, liquidation delay, auction effectiveness and phases.

A. Black Thursday

On March 12, 2020, ETH fell by more than 30 % in a matter of hours. With ETH being the collateral backing most of the circulating DAI, this event affected Maker as well. The sharp drop of ETH forced many vaults below the liquidation ratio before their owners could react, resulting in thousands of collateral auctions being kicked off by keepers.

During this time, the Ethereum network was heavily congested due to the drop in prices and the bot transactions triggered by the drop. ETH holders as well as users of various DeFi apps tried to liquidate their positions. This led to an increase in gas prices, which resulted in underpriced transactions not being mined, which in turn prevented vault owners from taking action against the liquidation of their collaterals.

An even bigger problem were auctions ending with almost zero DAI. Due to the congestion of the network and the short auction duration of only 10 minutes, some keepers succeeded in buying large quantities of ETH for virtually no DAI [19].

B. Bite Events

When a keeper calls the function `bite` of the contract `Cat`, the latter logs the action by emitting the event `Bite`. Figure 2

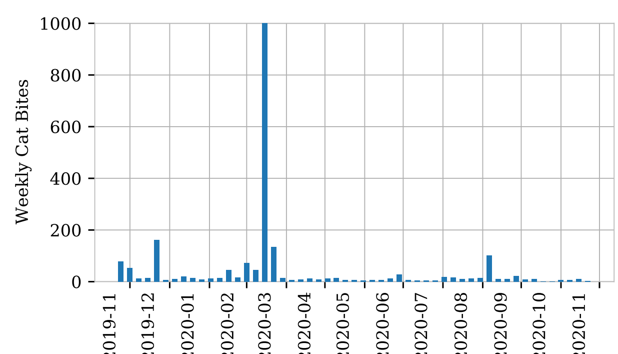

Fig. 2. Bite events per week (metric M1) during the study period. The bar for the week of Black Thursday is clipped, with the actual value being 4574.

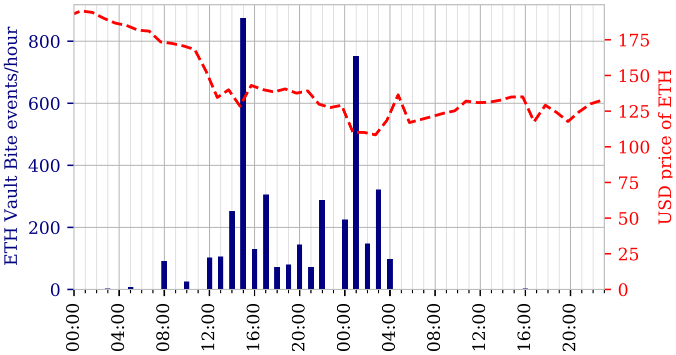

Fig. 3. `Bite` events (metric M1) in comparison to ETH prices in USDs (metric M2) per hour during Black Thursday.

shows the number of weekly `Bite` events (metric M1). Usually, there are just a dozen, but never more than 200 bites per week, except for the week around Black Thursday when they peaked at 4574 bites.

On Black Thursday, 92 % of the bites concerned ETH-A vaults. Figure 3 shows the hourly `Bite` events for this type of vaults and the corresponding ETH price in USD (metric M2). Most `Bite` events took place within the 18 hours between noon of March 12 and the next morning at 6:00, with the maximum after the largest price drop of ETH. The delay of roughly an hour between price drop and rise in liquidations is due to the delay of the price oracle.

C. Vault Management

Vault owners are incentivized to keep the vaults over-collateralized at all times. To analyze how timely vault owners react to liquidation threats, we compute the collateralization of the vaults (metric M3) for every `LogMedianPrice` event emitted by the contract `Median`. For vaults that would become under-collateralized if `OSM` took over the ETH price from `Median`, we analyze all user-triggered operations (calls to the functions `frob` and `fork` in the contract `Vat`) as well as collateral liquidations (function `grab` in `Vat`). At the time, when `OSM` copies (or would have copied[4]) the ETH price from `Median`, we check whether the vault owners took action or whether the vault gets liquidated.

Table I accumulates the DAI saved and the DAI liquidated for the time before, during and after Black Thursday. The percentage of DAI saved from threatened vaults was 59 % before Black Thursday, dropped to 54 % during the event, and rose to 89 % afterwards.

For comparison, table I also shows the numbers for the beginning of September, when ETH dropped by more than 30 % within a few days (the second largest decline after Black Thursday within the observation period). With 82 %, the share of saved DAI is only slightly lower than the 89 % for the whole period after Black Thursday.

TABLE I
DEBT SAVED VS. DEBT LIQUIDATED [MILLION DAI]

period	saved	liquidated	% saved
before Black Thursday	1.9	1.3	59 %
March 12–13	21.5	18.1	54 %
after Black Thursday	37.3	4.6	89 %
September 2–6	3.7	0.8	82 %

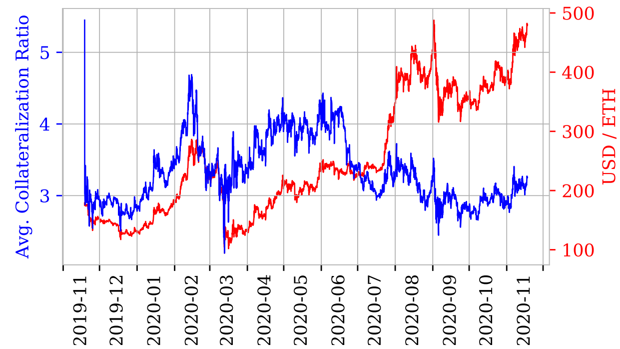

Fig. 4. Collateralization of vaults (metric M3) in comparison to ETH prices in USD (metric M2) for the entire study period

D. Collateralization

We investigate the collateralization (metric M3) over the study period and display its average in figure 4 together with the corresponding value of ETH from the contract `Median`. Since the collateralization depends on ETH, both lines move in sync most of the time. The collateralization finds a first equilibrium at approximately 2.75 before it moves up (together with ETH) to levels above 4.5. During the second half of February, the collateralization declines steadily together with ETH prices and ultimately arrives at its lowest level of 2.2 on March 12. In the following months, the collateralization oscillates at higher rates of about 4, before declining to levels around 3. This decline takes place during a period of low volatility of ETH. Even though the ETH prices then rise to the highest values during the study period, collateralization is kept at 3.

E. Liquidation delay

Next, we examine the length of time from the point a vault becomes under-collateralized to the point its collateral is liquidated (metric M6).

TABLE II
LIQUIDATION DELAYS [MINUTES AND BLOCKS]

period	median time	median blocks	average time	average blocks	maximum time	maximum blocks
March 12–13	2.6	14	9.3	43	57.1	265
before and after	0.2	1	1.7	8	55.0	228

Table II shows the delay for Black Thursday in comparison to all liquidations without Black Thursday by listing the

[4]The ETH price by `Median` may change by further input from the relay network before `OSM` reads the value.

respective average, median and maximum. The minimum time is zero in both cases, since some liquidations start in the block of the `OSM` price update.

F. Effectiveness of Auctions during Black Thursday

Auctions are intended to recapitalize the DAI debt. To measure how well the auctions succeeded, we investigate their effectiveness, i.e., how much DAI they yielded per ETH (metric M4). Figure 5 depicts this data for Black Thursday. For each hour with auctions, it gives the median and the distribution of effectiveness. Bars indicate the number of auctions. We notice two periods with unusual activities, the first one from 9:00 to 16:00 on May 12 and the second one from midnight to 3:00 on May 13. During these times, most auctions took place, and the effectiveness of auctions shows a large variance. For five hours, the median value even drops to almost zero DAI per ETH as the result of a large number of zero bid auctions.

G. Scaled Effectiveness of Auctions

The scaled effectiveness of auctions (metric M5) relates the DAI paid per ETH to the USD value of the ETH purchased. A scaled effectiveness of 95 % means that the auction winner paid 95 % of the value of ETH, getting a discount of 5 %.

Figure 6 compares the scaled effectiveness of all 5178 auctions (including Black Thursday) with the scaled effectiveness of the 1084 auctions not including May 12 and 13. For all auctions, the median of the scaled effectiveness is 76.9 %. The large interquartile range highlights, how much the effectiveness varies over all auctions. The lower quartile of virtually zero reflects the impact of zero bid auctions on Black Thursday. When omitting March 12 and 13, the median rises to 97.4 %, with a lower quartile of 94.4 % and an upper one of 99.2 %.

H. Auctions Reaching the `dent` Phase

Table III lists the number of auctions with respect to the phases they reach. As all auctions start in the `tend` phase, this number coincides with the total number of auctions. Before Black Thursday, 97 % of the auctions reach the second phase. On Black Thursday, their share drops to 30 % and rises afterwards again to 91 %. Beyond Black Thursday, all actions without a `dent` phase take place in the two days afterwards. We find no such auctions in the last eight months of the study period.

TABLE III
AUCTIONS REACHING THE DENT PHASE

period	auctions total	dent	
before Black Thursday	626	606	97 %
March 12–13	4094	1210	30 %
after Black Thursday	458	419	91 %
March 14–16	102	63	62 %
March 17–Nov 17	356	356	100 %

VI. Discussion

In this section, we evaluate the resilience-related hypotheses in order to assess the resilience of Maker to stress and discuss Maker's response to the incident.

A. Assessment of Resilience

Since we hypothesize that Maker fulfills the requirements for resilience laid out section in IV, we evaluate the following hypotheses.

1. *Auctions achieve adequate prices for the collateral.*
Test: We check whether the scaled effectiveness (metric M5) is close to one. Since the DAI is soft-pegged to the USD, a scaled effectiveness of one means that the DAI paid for the ETH corresponds to the latter's value in USD.

Assessment: When considering all auctions, the hypothesis has to be rejected. Figure 6 shows the median at 77 %, meaning that in half of the auctions, ETH was sold at 77 % of its value or below, with the center quartiles ranging from 0 to 90 %. When excluding March 12–13, the hypothesis can be retained. Under regular conditions (that include fluctuations in ETH value as well), the median is at 97 % with tight lower and upper quartiles, which indicates that ETH is sold close to its value and auction winners do not receive substantial discounts. Outliers above 100 % indicate the inability of bidders to determine an appropriate ETH price, while outliers below 100 % indicate that even without stress, some auctions may close with a large discount.

2. *Auctions are able to cover outstanding DAI debt.*
Test: We determine the share of auctions receiving at least one `dent` bid, since this implies that the preceding bid was high enough to cover the DAI debt.

Assessment: Apart from Black Thursday, 95 % of the auctions reach the second phase (Table III). In fact, we find that this share rises to 100 % two days after the event and stays there, which indicates that the measures taken immediately after Black Thursday (see below) were effective. Thus, the auctions cover the DAI debt in general. The main exception is Black Thursday, where the rate drops to 30 %, meaning that more than two thirds of the auctions fail to fulfill their purpose. Hence, the hypothesis must be rejected for Black Thursday, while it can be retained for the remainder of the study period.

3. *Vaults are collateralized at twice the liquidation ratio on average.*
Test: Analysis of collateralization over time (metric M3).

Assessment: During the first months of MCD, the over-collateralization was at 2.75 on average, which suggests that the owners considered this value sufficient. Until Black Thursday, owners seem to manage the vaults passively, letting the collateralization rise with the ETH value. The collateralization of 4 after the incident suggests that vault owners became more risk aware and averse. The decline to 3 during a phase of low volatility of ETH looks like a deliberate decision. Vaults start to be managed more actively, as the collateralization stays at 3 even though ETH rises again. We assume that this behavior

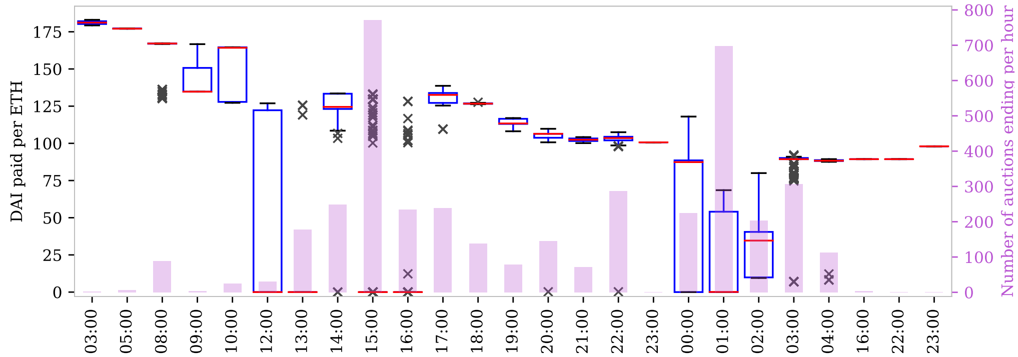

Fig. 5. Auction effectiveness on March 12 and 13. The box plot indicates the price median (red) and the range between upper and lower quartiles (blue) in DAI per ETH per hour. Black crosses mark outliers. The number of ended auctions per hour are displayed as violet bars on the same time line. For space reasons, the plot omits hours without ended auctions.

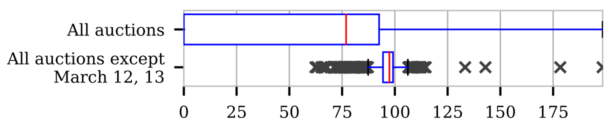

Fig. 6. Comparison of the scaled effectiveness of all auctions (top) to the scaled effectiveness of all auctions without March 12 and 13 (bottom). We omit one outlier at 900 %.

is connected to the boom of yield farming, promising a better use of DAI and ETH than in an overly collateralized vault. Overall, the hypothesis can be retained.

4. *Owners adjust vaults in time when collateral prices drop.*
Test: For vaults that will become under-collateralized as soon as the ETH price announced by `Median` propagates to `OSM`, we determine the ratio of DAI in adjusted vaults to DAI in bitten ones using metrics M1 and M3.

Assessment: Up to and including Black Thursday, the data in table I does not support this hypothesis. We observe that only 54–59 % of the threatened DAI is saved. Before Black Thursday, the number of `bites` is small (fig. 2), so the penalty for lazy vault management is low. On Black Thursday, network congestion is an additional impediment to re-collateralizing the vaults in time. After Black Thursday, the hypothesis holds. Vault owners start to react faster, probably by refined vault management with more vigilant algorithms, and save close to 90 % of the DAI threatened with liquidation.

5. *Under-collateralized vaults are liquidated fast.*
Test: We determine the time that keepers need to `bite` newly under-collateralized vaults after the contract `OSM` announces the lower price (metric M6).

Assessment: Under regular condition, half of the liquidations are initiated in the same block or the one following the price drop (table II), which is fast. The average delay of eight blocks indicates that some liquidations take longer, with the maximal time being almost an hour. On Black Thursday, the median rises to 14 blocks (three minutes) and the average to 43 blocks (nine minutes), which can no longer be considered fast, in particular in times of stress, when stabilization requires immediate responses.
Hence, the hypothesis has to be rejected for Black Thursday, while it holds during other times. One reason may be that the protocol does not offer incentives for liquidations other than the possibility to participate in auctions. However, this incentive diminished during Black Thursday as assets were locked in the numerous auctions, thus increasing the DAI illiquidity of potential keepers.

B. *Maker's response to the Black Thursday incident*

Due to the zero bid auctions of `Flipper`, the protocol accrued several millions in bad debt, which had to be covered by the sale of new MKR tokens.

As both, vault owners and keepers, sought to buy large quantities of DAI to repay vault debt or to bid on `Flipper` auctions, the demand of DAI soared during this phase, culminating in DAI prices of USD 1.11 on March 14, 2020.[5]

The Maker community and governance facilitators had to act quickly and organized several extraordinary governance and risk calls as well as executive votes. As early as March 13, an executive vote [20] adjusted several of Maker's risk parameters, extending e.g. `ttl` in the `Flipper` contract from 10 minutes to 6 hours. Another vote on March 15 [21] reduced the stability fee from 4 to 0.5 %, in the hope of incentivizing vault owners to generate more DAI, which would restore the exchange rate of DAI to USD. The executive vote on March 17 [22] added the custodial stablecoin USDC [23] as a new collateral type. It was considered the quickest means of providing keepers with a tool that allows them to acquire sufficient liquidity when large amounts of DAI are needed during an event like Black Thursday, with many auctions happening simultaneously. In addition, it offered the opportunity to re-peg the (still overvalued) DAI to the USD, as arbitrageurs could use USDC to mint new DAI and earn a premium by selling them.

The Black Thursday incident sparked a discussion on the design of Maker's collateral auctions, which resulted in two upgrades (Liquidations 1.1 and Liquidations 1.2) that were

[5]https://coinmarketcap.com/currencies/multi-collateral-dai/

implemented in July and August 2020. Both upgrades aim at reducing the amount of assets locked in auctions and thereby freeing resources to maintain more effective auctions in parallel. While Liquidations 1.1 introduced just a minor change, Liquidations 1.2 defines a ceiling of 20 million DAI for the total debt that can be auctioned off simultaneously, which reduces the liquidity needed overall. Currently, a complete redesign of Maker's auctions towards a Dutch auction model [24] is underway. With this approach, the collateral on sale is announced for a high price, which incrementally drops until a bidder is willing to pay the current price. This type of auctions eliminates the problem of locked assets, as keepers can borrow a DAI flash loan on platforms like Aave, buy ETH for DAI from a Maker auction, exchange the ETH back to DAI, e.g. on Uniswap, and repay the flash loan, all in a single atomic transaction.

VII. Conclusion

In this work, we assessed MakerDAO during the first year of the full protocol, from Nov. 2019 to Nov. 2020, by evaluating several stability-related hypotheses. Applying empirical metrics to the public transaction data, we could draw a comprehensive picture of how resilient Maker is to stress and which aspects would benefit from improvements. As we focused on vault management and liquidations, we complement other empirical evaluations of the Maker protocol regarding governance or oracles.

Our findings indicate that vault management has improved a lot over the observation period, which we consider a cornerstone of Maker's resilience to stress situations like Black Thursday. We also found Maker's liquidations to work well outside of stress situations. Regardless of whether the Black Thursday incident was a well-planned attack [25], the resilience of the system is particularly important under stress.

Maker has identified and most recently also resolved most of the core issues that led to the Black Thursday incident by introducing a complete redesign of its liquidation logic. The new Dutch auction system removes many of the core disadvantages of the old system, most notably the need to lock capital during auctions. This improvement greatly reduces requirements on keepers and even allows the use of flash loans for auctions. We therefore strongly support the rollout of Maker's new liquidation system.

Future Work

It would be interesting to study the liquidity of keepers during regular conditions and under stress. Moreover, for assessing the protocol changes since Black Thursday, one could model the incident and replay it with the current parameters. Finally, for the liquidation 2.0 system further metrics could be derived to evaluate its performance and compare the new to the old liquidation system.

References

[1] I. G. Pernice, S. Henningsen, R. Proskalovich, M. Florian, H. Elendner, and B. Scheuermann, "Monetary stabilization in cryptocurrencies – design approaches and open questions," in *2019 Crypto Valley Conference on Blockchain Technology (CVCBT)*. IEEE, jun 2019, pp. 47–59.

[2] K. Saito and M. Iwamura, "How to make a digital currency on a blockchain stable," *Future Generation Computer Systems*, vol. 100, pp. 58–69, nov 2019.

[3] A. Klages-Mundt and A. Minca, "While stability lasts: A stochastic model of stablecoins," *CoRR*, vol. 2004.01304, 2020. [Online]. Available: https://arxiv.org/abs/2004.01304

[4] W. Gu, A. Raghuvanshi, and D. Boneh, "Empirical measurements on pricing oracles and decentralized governance for stablecoins," *SSRN Electronic Journal*, p. 17, 2020.

[5] R. K. Yin, *Case study research and applications : design and methods*, 6th ed. Los Angeles: SAGE, 2018.

[6] Maker Foundation, "The Maker protocol: MakerDAO's multi-collateral Dai (MCD) system," accessed 2021-04-09. [Online]. Available: https://makerdao.com/en/whitepaper

[7] V. Buterin, "The search for a stable cryptocurrency," pp. 1–15, 2014, accessed 2021-04-11. [Online]. Available: https://blog.ethereum.org/2014/11/11/search-stable-cryptocurrency

[8] A. Moin, K. Sekniqi, and E. G. Sirer, "SoK: A Classification Framework for Stablecoin Designs," in *Financial Cryptography and Data Security (FC)*. Springer, 2020, vol. LNCS, pp. 174–197.

[9] S. M. Werner, D. Perez, L. Gudgeon, A. Klages-Mundt, D. Harz, and W. J. Knottenbelt, "SoK: Decentralized finance (DeFi)," *CoRR*, vol. 2101.08778, 2021. [Online]. Available: https://arxiv.org/abs/2101.08778

[10] C. Calcaterra, W. A. Kaal, and V. Rao, "Stable cryptocurrencies: First order principles," *Stan. J. Blockchain L. & Pol'y*, vol. 3, p. 62, 2020.

[11] G7 Working Group on Stablecoins, "Investigating the impact of global stablecoins," Bank for International Settlements, Tech. Rep. October, 2019, accessed 2021-04-11. [Online]. Available: https://www.bis.org/cpmi/publ/d187.pdf

[12] G.-J. Wang, X.-Y. Ma, and H.-Y. Wu, "Are stablecoins truly diversifiers, hedges, or safe havens against traditional cryptocurrencies as their name suggests?" *Research in International Business and Finance*, vol. 54, no. 101225, 2020.

[13] A. Klages-Mundt and A. Minca, "(In)stability for the blockchain: Deleveraging spirals and stablecoin attacks," *CoRR*, vol. 1906.02152, 2021. [Online]. Available: https://arxiv.org/abs/1906.02152

[14] W. Zhao, H. Li, and Y. Yuan, "Understand volatility of algorithmic stablecoin: Modeling, verification and empirical analysis," *arXiv: 2101.08423*, 2021.

[15] K. Jarno and H. Kołodziejczyk, "Does the design of stablecoins impact their volatility?" *Journal of Risk and Financial Management*, vol. 14, no. 2, p. 42, 2021.

[16] C. Jeger, B. Rodrigues, E. Scheid, and B. Stiller, "Analysis of stablecoins during the global covid-19 pandemic," *2nd Int. Conf. on Blockchain Computing and Applications (BCCA)*, pp. 30–37, 2020.

[17] L. Gudgeon, D. Perez, D. Harz, B. Livshits, and A. Gervais, "The decentralized financial crisis," in *2020 Crypto Valley Conference on Blockchain Technology (CVCBT)*, 2020, pp. 1–15.

[18] Maker Foundation, "Maker foundation multi-collateral dai public releases page," 2020, accessed 2021-04-11. [Online]. Available: https://changelog.makerdao.com/

[19] cyrus, "Black thursday response thread - governance plan and forum response," 2020, accessed 2021-04-11. [Online]. Available: https://forum.makerdao.com/t/black-thursday-response-thread/1433/82

[20] Protofire, "Maker Governance Dashboard – Executive Votes," 2020. [Online]. Available: https://mkrgov.science/executive/0x529b8b4b62b5f32bd47412988a0a66d72f86ba00

[21] ——, "Maker Governance Dashboard – Executive Votes," 2020. [Online]. Available: https://mkrgov.science/executive/0xd77ad957fcf536d13a17f5d1fffa3987f83376cf

[22] ——, "Maker Governance Dashboard – Executive Votes," 2020, accessed 2021-04-11. [Online]. Available: https://mkrgov.science/executive/0x049e4d10c1b7280cfed5b0d990e39f9c54529a32

[23] Coinbase, "USD Coin (USDC) - Stablecoin by Coinbase," 2020, accessed 2021-04-11. [Online]. Available: https://www.coinbase.com/usdc

[24] Smart Contracts Domain Team, "Mip45: Liquidations 2.0 (liq-2.0) - liquidation system redesign," 2021, accessed 2021-04-11. [Online]. Available: https://forum.makerdao.com/t/mip45-liquidations-2-0-liq-2-0-liquidation-system-redesign/6352

[25] B. Corporation, "Evidence of mempool manipulation on black thursday: Hammerbots, mempool compression, and spontaneous stuck transactions," 2020, accessed 2021-04-11. [Online]. Available: https://www.blocknative.com/blog/mempool-forensics