Bitcoin Flash Crash on May 19, 2021: What did really happen on Binance?

Tim Baumgartner* Andre Guettler†

November 6, 2022

Bitcoin plunged by 30% on May 19, 2021. We examine the outage the largest crypto exchange Binance experienced during the crash, when it halted trading for retail clients and stopped providing transaction data. We find evidence that Binance back-filled these missing transactions with data that does not conform to Benford's Law. The Bitcoin futures price difference between Binance and other exchanges was seven times larger during the crash period compared to a prior reference period. Data manipulation is a plausible explanation for our findings. These actions are in line with Binance aiming to limit losses for its futures-related insurance fund.

JEL Codes: G10, G12, G14, K22.

Keywords: Cryptocurrency, Bitcoin, Derivatives, Market Crash, Binance, Crypto Exchange, Trading Outage, Fraud, Extreme Volatility, Benford's Law.

---

*Ulm University, Institute of Strategic Management and Finance, Helmholtzstr. 22, 89081 Ulm, Germany, tim.baumgartner@uni-ulm.de (corresponding author)

†Ulm University, Institute of Strategic Management and Finance, and Halle Institute for Economic Research (IWH), Germany, andre.guettler@uni-ulm.de

We thank Carol Alexander, Francis Kim, Laura Kodres, Pete Kyle, Justina Lee, Jiasun Li, Shumiao Ouyang, Roberto Pascual, Daniel Rabetti, Andriy Shkilko, participants at the MIT GCFP Annual Conference, Auckland Centre for Financial Research Conference on Derivative Markets, AEFIN Finance Forum, British Accounting and Finance Association Annual Conference, and seminar participants at Ulm University for helpful comments.

1 Introduction

Reliable trading platforms are integral for functional financial markets. On May 19, 2021, Bitcoin plunged 30% and the largest crypto exchange Binance experienced a significant outage. Besides offering a broad range of crypto spot markets, Binance provides highly leveraged crypto futures. When clients engaging in these futures lose more than their account value due to liquidations and slippage[1], the exchange compensates their counterparties. Binance claims to have set up an insurance fund to overtake these liquidated positions when their value becomes negative.

The Financial Times (Samson and Oliver, 2021) and the Wall Street Journal (Kowsmann and Ostroff, 2021), among others, covered the May 19 outage event on Binance. Harmed clients are alleging that the leading crypto exchange intentionally shut down its platform to have free play when liquidating their clients' futures positions in order to limit losses in its insurance fund. Litigation is also pending on the matter.

In this paper we investigate what really happened at the largest crypto exchange on May 19, 2021. We use a data set comprising 73 million Bitcoin future transactions on and around the crash day.[2] We find that Binance stopped transaction-level reporting in the middle of the crash for 40 minutes. Assessing the time lag between Binance's API ("application programming interface", i.e., data connection) processing and the underlying trades, it seems as if Binance paused its API to process the data in a delayed manner. A part of the data has been (partly) back-filled by Binance later on. We investigate the authenticity of these particular transactions, and find that they do not conform to Benford's Law – a widely used approach to detect fake data which also helped uncovering the Libor manipulation (Abrantes-Metz et al., 2012).[3]

[1] Slippage refers to the difference between the expected price of the trade, when a certain liquidation was initiated, and the actual and less favorable price. It is most prevalent during high volatility market conditions such as in flash crash periods (and subsequent recoveries).

[2] We focus on Bitcoin future prices throughout the paper if not stated otherwise. To facilitate the exposition, we just write Bitcoin price.

[3] More recently, Cong et al. (2021) use this approach to analyze wash trading, i.e., market participants simultaneously selling and buying the same assets to create artificial trading volume, at unregulated crypto exchanges.

We compare Binance to other exchanges during the flash crash, and find that Binance's Bitcoin price deviates substantially from the average price based on other crypto exchanges. Specifically, the Bitcoin price difference between Binance and its peers is seven times larger during the crash period compared to a 10-day reference period (serving as a control group) before the crash.

The trading volume of the insurance fund spiked just before the outage started in line with the sharp decline of the Bitcoin price which liquidated long positions to a large extent. During the outage period, for which no insurance fund transaction data is available, we estimate further significant fund activity. While the front-end of the platform was already unavailable, locking clients out from closing positions or adding collateral to avoid liquidation, the liquidation engine continued to operate in the back-end.

We derive a model suggesting that the insurance fund would have had to trade a volume of around 1 bil. USD during the 40 minute period with missing transaction data, corresponding to a loss of 10 mil. USD when assuming a cost of 1% arising from overtaking liquidated positions. On first sight, these losses seem to be easily manageable given a stated insurance fund volume of 290 mil. USD. Such a reasoning rests on the assumption that the insurance fund actually holds the stated amount in liquid assets. However, at the time of writing this paper, there has been no way to verify (like a publicly known wallet address or external validation by auditors) the insurance fund's assets, neither the absolute level (as stated by Binance) nor the liquidity of the underlying assets.

Given how easy it were to publish the insurance fund's wallet address, concerns about the insurance fund's size and asset structure seem warranted – particularly in light of the many crypto exchanges that have disappeared since the infamous Mt. Gox bankruptcy in 2004, when 650 000 Bitcoin were declared missing at the leading crypto exchange at that time (market value of around 25 bil. USD at the time of writing this article).[4]

[4]More recently, Tether Limited had to pay a 41 mil. USD fine for lying about its USDT's backing: rather than the claimed 100%-USD backing of its stable coin USDT, it was actually only 2.9% (Robinson, 2021).

During the outage period, some limited trading volume remains. Based on the estimated trading volume for the back-filled time interval, our model suggests that the actual trading volume was only approximately 20% of the expected volume in absence of the outage. A potential explanation might be Binance's priority trading capabilities for VIP clients, such that this small group might have made extensive profits at the expense of liquidated traders and those hindered to enter into positions to participate in the rapid Bitcoin price recovery. Such an unequal treatment appears particularly problematic since the SEC investigates trading firms being potentially linked to Binance's founder (Benson, 2022).

We contribute to the literature in several ways. First, there is a broad literature offering means to uncover general numeric fake data. For example, Cho and Gaines (2007), Durtschi et al. (2004), and Drake and Nigrini (2000) propose evaluating the distribution of the numbers' digits to reveal deviations to the law of Benford (1938). We contribute to this strand of the literature by assessing the validity of the back-filled data period on May 19, 2021, at the leading crypto exchange Binance.

An extensive body of literature covers trading halts and their implications (Gerety and Mulherin, 1992; Subrahmanyam, 1994; C. Jiang et al., 2009). Sokolov (2021) examines congestions on the blockchain associated with ransomware attacks. Alike, high volatility events are well investigated (Gatheral et al., 2018; X. F. Jiang et al., 2017). A number of papers also consider the market microstructure of crypto currencies (Bouri et al., 2019; Makarov and Schoar, 2020; Aleti and Mizrach, 2021).[5] We contribute to these considerations with our event study surrounding the May 19, 2021 crash and characterize the consequences of leveraged positions during trading outage periods with ongoing liquidations.

A further line of the literature deals with regulatory issues in crypto currency markets and the downside of crypto exchanges in particular. Foley et al. (2019) estimate the exploitation of Bitcoin for illegal purposes, while Griffin and Shams (2020) analyze whether excessive issuances of the crypto stable coin Tether influences

[5] A very rich literature performs order book analysis in traditional assets and consider high-frequency trading strategies (e.g., Brogaard et al. (2014)).

cryptocurrency prices.[6] Moore and Christin (2013) study the risks associated with fraudulent exchanges and assess the principal-agency problem between brokers and their clients, whereas Gandal et al. (2018) explore suspicious activity on the formally leading crypto exchange Mt. Gox. Amiram et al. (2021), Chen et al. (2022), and Cong et al. (2021) provide evidence for substantial fake trading volumes ("wash trading") of crypto exchanges aiming at pretending higher liquidity. Our paper adds to this literature by highlighting the need for more consumer protection, in particular with respect to providing excessive leverage, in the context of these largely unregulated crypto exchanges.

2 Institutional Setup

2.1 Market Overview

There are more than 600 crypto exchanges, of which Binance has by far the largest trading volume. Exchanges depend on network effects: The trading venue with the highest liquidity attracts further liquidity because this situation corresponds to ceteris paribus lower transaction costs.

Thousands of different cryptocurrencies exist. Bitcoin had the biggest market capitalization of around 1 tril. USD on May 19, 2021, accounting to around 40% of the entire crypto market.[7] The exchanges offer to trade different pairs of cryptocurrencies directly against each other, or to swap cryptocurrencies against traditional currencies. In addition, various exchanges also offer derivative instruments such as futures, leveraged tokens, or options on cryptocurrencies.

2.2 Perpetual Futures

Binance has allowed even small retail clients to engage in futures trading, where the exposure may reach a factor of up to 125 of the pledged collateral when entering the position.[8] On top of traditional futures which expire on a particular day to get

[6] From a broader point of view, Liu and Tsyvinski (2021) investigate the risks and returns of cryptocurrencies.

[7] For instance, Huberman et al. (2021) provides an excellent overview of the Bitcoin protocol.

[8] The exchange has already decreased its derivative products spectrum in various jurisdictions, due to regulatory scrutiny.

settled subsequently, Binance offers perpetual futures. This type of futures enables traders to hold them infinitely as they have no expiry.[9] The exchanges manage this extension with a technical mechanism to replicate a rolling over between different future expiries (i.e., selling one future near to expiry and buying the next one): They deduct (or credit) the so-called "funding premium" every day. This payment compensates for the differences between the two futures contracts and avoids any arbitrage opportunities.

2.3 Liquidations

To maintain a levered position, Binance determines the "maintenance margin" which is typically about half of the initial margin – the collateral required to enter into the position. When the collateral falls below the maintenance margin, Binance liquidates the position in the open market. The exchange does not conduct this check based on its own market prices, but continuously computes a special "mark price". This index stems from a formula including prices of other exchanges, to which Binance applies certain adjustments. Binance notes in its FAQ that it takes additional measures during extreme market conditions.[10] This mechanism gives Binance some free play to influence the liquidation trigger.

For the margin calls, Binance does not offer its clients any grace period, i.e., providing them a chance to add collateral. Binance may even not inform the client of an imminent liquidation such that the client learns about it when it has been executed. Liquidations occur immediately at the prevailing market prices and are at a disadvantage compared to positions closed considerately to market circumstances with some time and in a careful manner, e.g., in the middle of the order book.

2.4 Insurance Fund

Liquidations of futures positions are intended to yield a small positive payoff from the exchange's perspective. Nevertheless, the trade execution (shortly after the

[9]Similar to a contract for difference (CFD).

[10]Binance states: "Please note that due to extreme market conditions or deviations in price sources, which may lead to mark price deviate from the spot price, Binance will take additional protective measures." (https://www.binance.com/en/support/faq/360033525071)

automatic initiation of the liquidation) may take place at a less favorable price if market conditions change fast ("slippage"). Then, the payoff can be negative. In such cases, Binance's insurance fund takes over the position, such that the counterparty (on the profitable side of the transaction) receives the total amount due. Otherwise, there would be a deficit at the expense of the counterparty.

Most traditional futures exchanges maintain such funds. However, in the case of Binance, this insurance fund is an opaque vehicle for which little information is known and almost none is verifiable. Technically, it is no separate (legal) entity with independent audits. Binance also does not publish any specific wallet address which would enable a plausibility check. The fund remains a part of Binance's equity. Therefore, any realized losses negatively impact Binance's bottom line. Vice versa, any profits made constitute a gain for the exchange.

To the best of our knowledge, the insurance fund assumes to-be liquidated futures positions over-the-counter, i.e., these transfers are not reported as trades. The insurance fund may hold the positions for a longer period, cancel opposite positions against one another, or place a limit order to close positions carefully. Such intelligent handling of the forced transactions may increase Binance's profits, whereas the liquidated clients receive nothing. When the insurance fund closes positions in the open market, the corresponding transactions receive the insurance fund flag in Binance's API. Binance publishes the total fund's equity on its website at a daily frequency, although it states that this figure does not consider open positions.[11]

2.5 VIP Clients

Binance has a VIP program depending on the clients' rolling 30 days trading volume to incentivize high trading activity. Traders can reach levels from 1 (lowest) to 9 (highest) by spot or futures trading, and holding Binance's coin BNB. The initial benefits include fee discounts, higher 24-hour withdrawal limits, and sub-accounts. Additional benefits are available starting at level 4. For such a level, one has to

[11] Various exchanges, including Binance, use a mechanism to distribute the losses associated with liquidations amongst other traders ("Auto Deleveraging"). It determines that the profiting counterparty bears the irrecoverable losses, as the exchange reduces the position's profit as if it had lower leverage – as a loss-sharing rule.

reach a 30 days spot trading volume matching at least 120 mil. USD (or 600 mil. USD in futures and hold 500 BNB worth 170 000 USD as of May 19, 2021). These conditions illustrate that only a minority of Binance's clients enjoy such level 4 benefits, that include a priority API with a "[h]igher API order frequency", "VIP Risk management and priority notifications", and "'[p]riority support for technical issues" (see Figure A9).

3 Data

3.1 Data Sources

We retrieve tick-level data from Binance's API which the data provider Tardis.dev has cached through the Binance websocket API (denoted by *Tardis API* data).[12] Therefore, we have an observation for every individual transaction at millisecond-level. We focus on Bitcoin (BTC) transactions as this cryptocurrency has been considered most prominently in the incident on May 19, 2021. We make use of the day of May 19, 2021 itself and the ten days before (i.e., starting on May 9) as reference period (control group), as well as the day after the crash, May 20, 2021. We have transmuted all time specifications into Coordinated Universal Time (UTC).

Data points comprise a timestamp at millisecond-granularity, the Bitcoin amount, and the price for each transaction. In addition, Binance provides an indicator whether the Binance-owned insurance fund has been one of the counterparties of the trade, and another one signaling whether the buyer has been the maker[13]. Notably, Binance has ceased the supply of these two data variables from its API for any new data a few weeks later, after public attention to the incident on May 19 arose.

The Tardis API data contains a transaction data gap roughly between 13:00 and 15:00 on May 19, 2021. We fill these gaps with additional data which became available in March 2022 in Binance's Public Data Collection (denoted by *Binance*

[12]Binance's API allows only a limited look-back period, as for instance 30 days. The caching comprises only saving the data streams on an "as is" basis.

[13]As opposed to being the taker, i.e., the party "taking" the liquidity

website data).[14] This data, however, does not contain a flag on insurance fund trades or any other details besides transaction volume and price.

Beyond, we obtain order book depth of the ten highest bids and the ten lowest asks, as well as the corresponding volume for each price level – also from Tardis. The timestamps of the order book depth data are separate from the transaction data. Therefore, we match the order book snapshot prevailing at each trade.[15]

For any API message, there are two timestamps: First, there is the timestamp of the underlying event (the trade or order book change) and the timestamp when the message has been sent out on the API. We also assess this time difference between the event and sending the information via the API. Similarly to the order book data, we further append the long-short ratio delivered by Binance, which includes separate time series for all accounts, the most significant accounts, and the largest Bitcoin positions. Furthermore, Binance releases data on the open interest, the funding premium of futures, and the mark price which determines the value of positions as collateral and thus whether Binance liquidates a position or not.

Along with the tick-level data of individual transactions, we also aggregate another data set (from the Tardis API data) on the minute level to derive metrics like the number of transactions per minute or the share of insurance fund trades of the total transaction volume within the minute.

Besides, we collect minute-level OHLCV (open, high, low, close, volume) data from TradingView to which Binance supplies this aggregated data (denoted by *TradingView* data). We also obtain minute-level reference Bitcoin prices for other exchanges from Tiingo Inc.

To gauge social sentiment, we include Augmento Bitcoin crowd sentiment data. Augmento fetches posts from Twitter, Reddit, and Bitcointalk by web scraping these social media sources. Twitter is a social network specialized in short text posts, Reddit is a general discussion forum, while Bitcointalk is a crypto-specific discussion forum. The data provider constructs 93 indicators for each of the three sources,

[14] https://data.binance.vision/?prefix=data/futures/um/monthly/trades/BTCUSDT/

[15] Technically, we merge by adding the latest order book price level and volumes to each transaction in a "last observation carry forward" manner.

signaling the intensity of specific categories, such as optimism, pessimism, sadness, panicking, or rumors. These indicators are positive or zero as they separately capture intensities in opposite directions.[16] These categories are identical across the three social networks and total up to 279 time series with hourly values.

3.2 Descriptive Overview

Table 1 reports summary statistics for May 9 to May 20, 2021. Overall, the data covers 73 million tick-level Bitcoin futures transactions. The Bitcoin price averages at around 46 000 USD within the period considered. The similar averages of the bid and ask order book depth indicate a well-filled order book. In the order book, the top of the book's volume averages at 1.08 Bitcoin (bid) and 0.73 Bitcoin (ask). The following averages going deeper into the book have decreasing average quantities. On average, the insurance fund has executed 0.1% of all trades. The indicator variable of maker trades shows that in 50.3% of all trades the maker has been the buyer.[17]

The open interest ranges broadly between 30k and 40k Bitcoin. The long-short ratio of all traders and the most significant accounts show roughly the same statistics, averaging at 2.5 more long than short exposure. In contrast, the top positions have a less pronounced long-short ratio. The funding premium, which Binance deducts from Bitcoin long futures and credits to short positions (for positive premia, and vice versa if negative), levels off at -0.002%. The mark price, used for valuing positions as collateral, exhibits slight differences from the spot price. Finally, the time lag between market events and the API processing has a mean of 0.1 minutes (trades) and 0.09 minutes (order book), while the maximum shows the large delay during the outage period.

[16] For instance, there is no netted "optimism" indicator which is determined as positive minus negative signals, but a separate optimism and pessimism indicator which both contain positive numbers or zero.

[17] The term "maker" does not necessarily refer to a "market maker". Rather, this expression refers literally to the single transaction, although it seems probable that this trader is practically often a market maker. In contrast to some traditional exchanges, there is usually more than *one* market maker on Binance. The barrier to becoming a market maker under Binance's designated Market Maker Program is relatively low, and anyone meeting the conditions can qualify to join the program. However, this implies a specific minimum liquidity provision.

4 Timeline of the Crash

4.1 Price Collapse and Recovery

The crash of May 19, 2021, has the character of a flash crash. Figure 1 depicts the course of events on the crash day, the day before, and the day afterwards. To the best of our knowledge, there has been no particular news justifying such a volatile price development. On the day before, the Bitcoin price experienced a decrease of roughly 5% to 42 500 USD, which further expanded to 40 000 USD in the first twelve hours of May 19 (UTC). The sharp price decline began at around 12:30. Table 2 summarizes the flash crash statistics. Prices reached the maximum drawdown at 13:09, being 33.9% lower than at midnight.

However, prices recovered rapidly and volatility decelerated at about 35 000 USD at 14:00. The Bitcoin price sets a post-crash intra-day high at 16:48. When focusing on the four hours between noon and 16:00, the trading volume on Binance was 357 607 Bitcoin, summing up to 12.7 billion USD in a total of 2.7 Million transactions. The average transaction value of slightly below 5 000 USD suggests that retail traders have been heavily involved in the events. According to Binance's website, the insurance fund had a balance of 290.4 million USD on the day before the crash, and reported a moderate loss (1.7 million) one day after the crash.

4.2 Binance Outage

To the best of our knowledge, Binance's platform outage evolved as follows. Clients first complained at 12:02 about website errors showing the HTTP status code 503 or 504 when depositing funds, referring to a gateway timeout or an unavailable service. More importantly, clients report the first errors connected to trading at 12:14, showing a gateway timeout when, e.g., trying to purchase crypto assets. At 13:05, Binance tweets: "ETH and ERC20 withdrawals are temporarily disabled due to network congestion" (see Figure A8). Further user complaints include screenshots of a blank user interface in Binance's mobile app at 13:07, and still around 17:00 some customers post "server too busy" errors on social media.

In addition to the user interface outage, there seems to have been a disruption of the API, which is even more severe given that most traders do not submit orders manually but rather rely on algorithmic trading based on the API. The Tardis API data contain a transaction data gap (as outlined in Panel C of Figure 8) roughly between 13:00 and 15:00 on May 19, 2021. The red bars in Panel C denote time periods with missing transactions data, while for some minutes Binance provided heavily delayed transaction data.[18] However, additional gap-free transaction volume and price data became available in March 2022 in Binance's Public Data Collection (*Binance website* data)[19]. We use the Tardis API data (and any gaps filled with the Binance website data) to conduct our analysis in Section 5.

Even for the minute-level data delivered to TradingView, there has been a gap initially between 13:16 and 13:56 (highlighted in Figure 8 with the vertical red lines). Notably, Binance has later back-filled this gap. Together with this data delivery, Binance has also overwritten a subset of the previous data points with partly different numbers according to a statement received from TradingView's higher-level support.

5 Has Binance faked Trading Data?

This section focuses on Binance's initial data feed from May 19, 2021, that contained missing data for a critical time frame in the middle of the crash ranging from 13:16 to 13:56 (UTC). As explained above, this gap in the data has been populated with delayed data after about one week, casting doubt on the credibility of these transactions.[20] Potential manipulation of these data could either be due to filtering the actual transactions or fully making up the transaction data. Therefore, we employ forensic methods to investigate the plausibility of the delayed transactions data. We postulate that these transactions should conform to Benford's Law, a well-

[18]There is a red bar if no data has been available at all. If there is *no* red bar, data has been provided (belatedly).

[19]https://data.binance.vision/?prefix=data/futures/um/monthly/trades/BTCUSDT/

[20]We fill any gaps with transaction data Binance is providing on its website (as of March 2022).

respected method which also helped uncovering the Libor scandal (Abrantes-Metz et al., 2012).

Benford (1938) proposes an approach to validate numbers by evaluating them without any need for additional benchmark data. This independence makes this approach reliable and robust against external biases. Benford postulates that certain digits should occur more often than others due to the multiplicative nature of their composition.[21] For instance, the first digit "1" should be more likely than the first digit "9". The observed relative frequency of first digits should converge to their expected value, while substantive deviations speak against the authenticity of the data.

We consider the tick-level transaction data, that were delivered in a heavily delayed manner by Binance for the 40 minutes in question, containing 373 thousand Bitcoin futures transactions. We investigate the individual Bitcoin volumes of these transactions as the Bitcoin prices lack a sufficient dispersion in its first digits (i.e., they were either 2 or 3 in that period) which is necessary for Benford's law to hold.

Figure 2 (a) compares the observed frequencies of first digits (bars) to the expected value according to Benford's law (dashed line). We can see that the observed frequencies are skewed such that lower digits occur more often than expected, and the frequency of higher digits is hence lower. In particular, the first digit "1" appears almost twice more often.

These substantial deviations are consistent with fake data. Therefore, we investigate the data more thoroughly. For this purpose, Benford (1938) provides another theorem on the first digit's probabilities after a second-order transformation of the data. This transformation is carried out by calculating pairwise differences between every two consecutive transactions (sorted by ascending timestamps). Figure 2 (b) presents these results. Under this measure the divergence is even more pronounced, backing up the mistrust in the back-filled transactions.

Moreover, Benford supplies another technique, summing up the digits. He proves that this enforces, under reasonable assumptions, a uniform distribution. We plot

[21] Refer to Fewster (2009) for a helpful introduction of Benford's Law.

the realized distribution and the corresponding differences in Figures 2 (c) and (d), revealing solid variations from the expected outcome again. Figure 2 (e) exhibits the differences belonging to Pearson's χ^2 Test. Differences are most prominent for the digit "1", while in particular the digits "4", "6" and "9" are also notable.

We proceed with statistical tests on the Benford analysis. Table 3 shows the results. Panel A shows the numbers for each digit's observed frequency and the difference to the expected value. For instance, we observe future volumes with a leading digit "1" in 52.7% of all cases, while Benford's Law would suggest only 30.7%. Hence, there is a difference of 22 percentage points (or 71.7% in relative terms). We test the statistical significance of each digit's difference with Pearson's χ^2 Test, which indicates significance at the 1% level in each case.

Panel B summarizes the accompanying descriptive statistics. The distribution exhibits a positive skewness, and the arithmetic mean of absolute differences (to the expected proportion, "MAD") reaches 5%. The distortion factor is negative (-55.5). This magnitude suggests that the numbers appear to be somewhat understated. We should observe this situation if Binance, for instance, had filtered out actual transactions with larger volumes, as opposed to creating invented transactions from scratch. This finding is interesting given that transactions by (larger) VIP clients and related market makers my have been excluded from the transaction data.

Panel C provides test results with respect to the overall conformity of the data with Benford's Law. The canonical way of testing the goodness of fit to Benford's Law, Pearson's χ^2 Test, indicates significance at the 1% level (test-statistic: 93 775). Similarly, the additional tests, the Kolmogorov-Smirnov (test-statistic: 124.98) and Euclidian Distance Test (test-statistic: 132.90), signal significance at the 1% level. Apart from the goodness of fit tests, we apply the Mantissa Arc Test, whose null hypothesis is that the mantissa of the observed distribution is uniformly distributed. The test-statistic of 0.10 rejects this null hypothesis at the 1% level as well.

We run the following robustness check to address potential concerns that Bitcoin futures volumes in general do not conform to Benford Law's distributional assumptions. We hence run the test for reference data comprising the ten day observation period (May 9 to May 18, 2021) before the crash day. Figure A1 presents the Benford analysis for the reference data. In Panel (a), we observe that there are deviations from Benford's distribution, albeit not as pronounced as in the 40-minute interval with delayed transaction data on May 19 (Figure 2).

Panel (f) of Figure A1 shows the resulting difference-in-differences of the digit distribution (relative frequencies) between the 40-minute time period and the reference data (both versus the expected relative frequencies according to Benford's Law). The comparison shows that the abnormal effect seems less pronounced in the reference data compared to the 40-minute interval of interest.

Furthermore, we add a bootstrap test by sampling 20 000 observations, half of which stem from the 40-minute interval and the other half from the reference data (10 000 iterations). Panel (g) of Figure A1 shows bootstrapped difference-in-differences boxplots. The body of each boxplot marks the range between the 25% and 75% quartiles, while whiskers sketch the range of the remaining values.[22] We find statistically significant differences between the reference data and the suspicious 40-minute interval with back-filled transaction data. Our results are in line with the notion that the latter may have been manipulated.[23]

We furthermore repeat this robustness check for the reference period, serving as a control group, to rule out any time-of-the-day effects (**Wood1985**). For this purpose, we filter the reference period data for the same time of the day at which the crash took place. We then perform the same calculations as for Figure A1 and present the results in Figure A2. Results remain qualitatively the same, indicating that time-of-the-day effects are not driving our main results in this section.

Digging deeper into the particular days, we check whether any individual day of the reference period uncouples from the expected distribution. This comparison is also more balanced than aggregating the ten reference days: Matching periods of equal length is the most natural correspondence, weakening any effects stemming

[22]Outliers appear as points. We classify a value as an outlier if the distance between it and the boxplot's body exceeds more than 150% of the boxplot's body length.

[23]We have conducted the same robustness check with only May 18 instead of the period fom May 9 to May 18. Results are qualitatively closely similar and are available upon request from the authors.

from more observations. This evaluation also has a sufficient number of observations because even this subset of the transaction-level data remains well above 100k observations.

Figure A3 shows the results for each day in a separate panel. Independently from the day chosen, results stay vastly similar and robust. This additional evidence alleviates potential concerns that different sample sizes or time-of-the-day effects may drive our results.

6 Can we Reconstruct a Data Filtration?

We choose another crash day prior to May 19 for a simulation analysis in order to reconstruct a potential data filtration. We take the crash of April 18, 2021, since it comes with similar characteristics and takes place in an alike environment, being only a month earlier than May 19, 2021. Figure A4 provides an overview of the course of events around the flash crash. The maximum drawdown of the Bitcoin price reached 12%, while the volume spiked up to 7715 Bitcoin per minute on Binance. But despite the large price drop and the large volume of up to 7500 Bitcoin per minute – in a similar magnitude as on May 19, 2021 – there was no outage at Binance on April 18.

We apply a filtration to the data (i.e., deleting a subset of it) and then compare it to the original data set using the same Benford method. In every minute where the Bitcoin price is at least 0.1% lower on Binance than the average of other major exchanges, we randomly delete 90% of the transactions on Binance. The motivation for this approach is that Binance may have an incentive to reduce the reporting of transactions that differ a lot compared to other exchanges.

Figure A5 shows the results of our filtered data. The Benford method yields qualitatively similar results compared to our study of the suspicious data. Panel (f) of Figure A5 depicts the difference-in-differences between the faked data and the original data set which share the same direction and relation with our previous results. Panel (g) shows the bootstrapped difference-in-differences. This bootstrap suggests that the results hold systematically in the sample, i.e., are not merely driven by some outliers. All in all, our findings are consistent with the notion that Binance also reported filtered data on May 19, 2021.

7 Was Binance Different than Other Exchanges?

In this section, we analyze whether the situation on Binance was different compared to other exchanges during the crash, and if those deviations are notably different during the crash compared to the ten day reference period before the crash. Our analysis extends Makarov and Schoar (2020) that shows evidence for arbitrage opportunities across crypto currency markets.

7.1 Bitcoin Futures

We consider the Bitcoin average of other exchanges as the natural benchmark to investigate any aspect specific to Binance. Generally, Bitcoin prices converge across exchanges as there is little friction limiting arbitrage opportunities between them. Traders can for instance enter into opposite future positions when a divergence occurs and close both positions when the gap has dissolved. As a result, and given the marginal transaction costs for large arbitrageurs, even minor differences between exchanges are noteworthy.

In line with this arbitrage argument, the overall flash crash arose across all exchanges. Figure 3 presents the Bitcoin futures chart from the start of the decline to its recovery. The red line represents the minute-level average futures price on Binance, whereas the blue line shows the average of other exchanges. The ribbons in the corresponding color point out the empirical 95% price range within each minute. The red box highlights the back-fill data as outlined in Section 4.2.

We see that the price fluctuation range on Binance expands substantially during the crash and in particular for the two hours with delayed and back-filled price data. Additionally, the price discrepancy between Binance and other exchanges is considerable. While at the beginning of the outage period Binance tend to be more expensive than its peers, Binance's Bitcoin price was way too low between 13:30 and 14:00. This observation is consistent with a surplus of selling orders at Binance, driven by forced liquidations or voluntarily closed positions to prevent imminent liquidations.

While the fluctuations are most intense in the highlighted time interval, the trading volume reported by Binance in Panel B drops materially. This decrease is in line with locking out a large share of clients because these kind of market environments are typically characterized by above average trading volumes. In Panel D, we see that the volumes at the other exchanges remained high. Hence, despite high trading volume, the other exchanges were able to process the high order flow on their platforms.

Notably, even though trading volume at Binance went down a lot in the time period of interest, volumes were well above zero.[24] A potential explanation for some remaining volume may be the priority API available to VIP clients starting from level 4 as explained in Section 2.5. If this rationale is valid, those VIP clients might have made extensive profits at the expense of liquidated smaller traders, and all those hindered from participating in the steep price recovery. Such a relationship would be particularly noteworthy as there have been recent reports of an SEC investigation of trading firms potentially linked to Binance's founder Changpeng Zhao.[25]

Panels E and F of Figure 3 compare the number of trades on Binance to those of the other exchanges. While the number of transactions reached a high level compared to other days, it does not seem plausible that it justifies the outage from a technical perspective as Binance's platform had already processed a higher number of transactions per minute (e.g., around 12:50 on the same day).[26]

7.2 Leveraged Tokens

Next, we compare different prices for bearish leveraged Bitcoins tokens between exchanges. Binance offers such a product as BTCDOWNUSDT, which traders use to hedge their long positions during time of distress. As such leveraged tokens have individual characteristics, they are unsuitable for constructing an average of other

[24] We estimate realistic trading volume for this time frame in Section 10.

[25] https://decrypt.co/93016/sec-investigating-firms-linked-binance-founder-report

[26] Intra-day trading volume on May 19, 2021 was extraordinary but did not exceed other intense days such as April 18, 2021.

exchanges. We compare it directly to two similar tokens, Poloniex's BEARUSDT and FTX's BEARUSDT. All three tokens are short Bitcoin tokens with a leverage factor of three. We compare their performance defined as the tokens' percentage price change. As benchmark, we use the actual Bitcoin price (average of various exchanges) change which they aim to track. For the bearish tokens, this reference is the three-fold (the leverage factor which the tokens aim to reach) inverted percentage change of Bitcoin.

Figure 4 presents the empirical results for May 19, 2021. The three tokens are tightly correlated in the first hours of the day. Between 12:30 and 15:30, within the crash, an underperformance of Binance's token kicks in, whereas Poloniex's and FTX's tokens remain unaffected and continue to track the Bitcoin performance benchmark closely. Panels B to D show each derivatives' corresponding USD trading volume. Panel B displays that the underperformance of Binance's token coincided with a surge in trading volume. The surge in trading volume was more pronounced on Binance than on other venues.

This underperformance seems to correspond to a one-off tracking error. The associated losses are permanent because the underlying algorithm mechanically carries on replicating a leveraged short exposure on Bitcoin. Hence, any "catch-up" to the benchmark is not feasible since the underlying net assets experienced a lasting reduction. Our results underpin the notion that (most) clients at Binance did not only face an outage period with respect to futures trading, but also experienced substantial disadvantages in other derivative instruments.

8 What Role did Binance's Insurance Fund play?

In this section, we investigate which role Binance's insurance fund played based on the mechanics explained in Section 2.4. The first subsection analyzes reported insurance fund trades while the second subsection provides estimates for the insurance fund's activity during the period with heavily delayed transaction data.

8.1 Reported Insurance Fund Trades

This subsection focuses on the trades flagged as insurance fund trades in the Tardis API transaction data. Panel C of Figure 3 shows the sum of trading volume flagged as insurance fund trades. Due to the nature of reporting, these trades should show the unwinding of positions of the insurance fund, while it is unclear when and at which price the insurance fund received positions. We observe that those trades are associated with the price decline below the thresholds of 35 000 USD and 32 000 USD. Under the assumption that the insurance fund closed the positions immediately after taking them, this relationship is plausible since the traders had enough collateral to keep their (long) positions until the price fell drastically.

There is little insurance fund volume after the slight recovery at 13:00. This observation indicates that either few traders wanted to (or due to the outage few were able to) enter into levered positions at that time, or that the insurance fund was unwilling to close positions at such low prices. However, the maximum Bitcoin drawdown coincides with an economically meaningful volume of the fund, suggesting that the flash crash might also have been a result of insurance fund sell orders.

Considering the consequences of the liquidations and the crash overall, Figure 6 shows the overall open interest of Binance's clients. The open interest is denoted in Bitcoin, meaning that Bitcoin price changes do not affect the graph. We find a sharp decline in open interest around noon of May 19, 2021. Just before the crash, open interest reached a ten-day high. It seems likely that Binance's clients reached this high exposure mainly through increasing leverage, without a sustainable inflow of new capital. In this case, the market environment would have been especially vulnerable to a price collapse. The decrease in open interest of approximately 40% comprises voluntarily closed positions as well as forced liquidations, while we are not able to derive the breakdown into the two reasons from the data.

To further characterize the liquidations, we rely on the fact that the liquidations during the price decline are naturally long positions, making the long-short ratio an interesting metric. Figure 7 shows the long-short ratio on Binance. On top of the long-short ratio of all traders, Binance discloses the long-short ratio of the biggest ("top") accounts, as well as the biggest Bitcoin positions. The chart shows that the long-short ratios of all traders and the top accounts were synchronously dropping on May 19.

In contrast, the time series of top positions detaches entirely from the other two groups. It fluctuates around 1.2, indicating an almost equal long and short exposure. Nevertheless, it also exhibits the negative outlier on May 19 at the same time. Overall, the picture gives the impression that both regular accounts and top accounts were equally affected by the crash, whereas the top positions remained vastly unaffected.

8.2 Model for the Missing Data during the Outage

This subsection derives an insurance fund model to extrapolate its activity, which is necessary because of the missing insurance fund transaction data during the outage. We construct a Tobit model (Tobin, 1958) that is particularly suitable to model the insurance fund's trading volume because these volumes are bound at zero. As discussed in more detail in Section 2.4, the liquidation mechanism aims to close positions with a small positive payoff. The insurance fund only needs to step in once this security margin is insufficient. Hence, we observe zero insurance fund trading volume in most quiet market phases, for which Tobit models are particularly well suited.

We postulate that the main factors for an insurance fund recourse are volatility and order book liquidity. We assume a positive relationship between insurance fund activity and volatility, while the relationship with liquidity is negative (i.e., the insurance fund needs to intervene more often if the order book is relatively thin). The rolling five-minute standard deviation of Bitcoin prices measures volatility. We gauge the order book liquidity on the bid and ask side by summing up the volume available within 0.5% behind the highest bid (or lowest ask) price.

We estimate the model with the data from May 9 to May 19, 13:15 since this is the period before insurance fund transaction data is missing.[27] Table 5 shows the resulting Tobit models in Panel A. The first model shows a positive effect of volatility, in line with the notion that the insurance fund does not step in unless volatility escalates and forced liquidations happens at adverse prices.

For the more saturated second model, we add the order book liquidity measures. The volatility coefficient remains similar albeit slightly more moderate (1.11). In line with intuition, the ask liquidity is negatively associated with the insurance fund trading volume (−24.71). In contrast, the bid liquidity coefficient is positive (1.57) and signifies that insurance fund trading volume is higher in liquid market phases. This result does not confirm the notion that the insurance fund needs to step in less often if there is sufficient liquidity in the order book. The evidence is nevertheless consistent with the insurance fund holding the liquidated positions during the movement to the intra-day low for some time to wait until the market recovered *and* large clients were buying the fire sale-priced positions from the insurance fund. It is important to emphasize again that only VIP clients and large market makers were able to enter these type of very profitable transactions during the outage period.

Figure 5 shows the out-of-sample results for the estimated insurance fund activity. The model gauges the insurance fund's volume adequately and remains conservative as predicted volumes are notably lower than the reported amounts before (or at) 13:15. For the time after 13:15 we estimate the insurance fund's volume for the data gap, where no reported volumes are available. The model predicts considerably higher volumes than earlier on the day. Notably, the estimated insurance fund interventions are nearly at the record high of the day when data becomes unavailable. In the following two hours, there are several waves of predicted interventions, which are then flattening over time. When data availability resumes at 15:15, the reported and estimated interventions converge to zero.

Figure 5's Panel B depicts the cumulative sum of trading volume attributable to the insurance fund. Until 13:16, the volume slightly exceeded 500 mil. USD, being

[27] Note again that the Binance website transaction data does not contain the insurance fund flag.

already a material volume compared to its balance of around 300 mil. USD. More importantly, the model suggests a vast increase reaching somewhat more than 1 bil. USD until 16:00. Panel C further shows that order book liquidity dropped sharply within the crash making it very costly (if possible at all) for the insurance fund to get rid of the overtaken positions. After that, reported volumes seem to stand still with some jumps, speaking for a lack of data delivery.

To check the robustness of the Tobit models, we also estimate OLS regressions. In contrast to the Tobit model, OLS has the disadvantage that the majority of zero insurance fund volumes biases the model predictions to always predict numbers near to zero. Volume predictions can also be negative. We address this obstacle by filtering out the observations of zero (truncation), and further the near-to-zero volumes (below 1 Bitcoin) to make the model reliable for rough market situations. However, this comes at a cost of a relatively low number of observations (93 cases with substantial insurance fund trading volume) and resulting low power of the model.

Panel B of Table 5 shows the results. We obtain similar coefficients as in the Tobit model. Constants are negative and the volatility coefficients are slightly more moderate (0.93 and 0.98 compared to 1.19 and 1.11). Also, the liquidity coefficients remain within in a reasonable span (0.93 and -34.68 compared to 1.57 and -24.71) even though we lose statistical significance due to the power issue mentioned above. Altogether, this robustness confirms the validity of the Tobit models.

From a “back of the envelope” approach, the estimated insurance fund’s trading volume of 1 bil. USD can serve as a factor to reckon the overall cost incurred by the fund. Assuming a spread of 1% between the achievable liquidation price and the market price when the liquidation was triggered, and Binance immediately winding off the overtaken positions, the loss would total 10 mil. USD. That stands in sharp contrast to Binance’s stating a mere reduction of 1.7 mil. USD (see Table 2). Rather than immediately offsetting these insurance fund positions into a thin order book (recall that most clients were not able to place orders), Binance could have held these positions on its own books and at its own risk. But given the volume of 1 bil.

USD, it is not unlikely that Binance may have run out of buying power at some point because of liquidity issues. Additionally, these numbers only include effects of Bitcoin futures alone and do not comprise the negative effects of futures on other crypto assets such as Ethereum which Binance offers as well.[28]

On first sight, these losses of 10 mil. USD seem to be easily manageable given a stated insurance fund volume of 290 mil. USD (see Table 2). However, as discussed in Section 2.4 neither the absolute level (as stated by Binance) nor the liquidity of the underlying assets can be verified externally. Given how easy it were to publish the insurance fund's wallet address, concerns about the insurance fund's size and asset structure seem warranted. In contrast to centralized (stock) exchanges that are publicly traded, such as Nasdaq or NYSE, for which detailed financial statements are publicly available, nothing is known about the available capital backing operations of crypto exchanges such as Binance. Literally, Binance is reluctant to provide its headquarter's address stating it has none (Baker, 2020).

Clients and regulators should be cautious in light of the many crypto exchanges that have disappeared over time. One outstanding event was the Mt. Gox bankruptcy in 2004, when 650 000 Bitcoin (market value of around 25 bil. USD at the time of writing this article) were declared missing at the leading crypto exchange at that time.

Overall, these arguments, whether immediately offsetting the fund's positions (and realizing losses) or holding the positions in order to avoid short-term losses (but running out of capital), are in line with deliberately "pulling the plug" and shutting down the futures engine because significant and expensive insurance fund interventions were imminent.

9 Did Binance function irregularly during the Crash?

This section sets the delays in Binance's API in relation to the price decline and missing data, and further reports deviations in the mark price which the exchange

[28]The open interest of Ethereum-based perpetual futures account to around 50% of those for Bitcoin.

uses to trigger liquidations. To grasp the chronological sequence of delayed processing in Binance's infrastructure, we assess the time lag between the transactions' (alleged) execution and their announcement in the API stream using Tardis API data. We calculate this lag for both the transaction and order book streams.

Panel A of Figure 8 shows this time difference in minutes for the transactions, and Panel B examines the order book. While there is usually no notable lag, the delay starts at 13:12. The delay increases linearly, climbing up to 32 minutes to transactions executed at 13:20 (and therefore sent via the API at 13:52). Then, the API data gap begins, i.e., the API does not send any transaction data at all, and lasts until 15:10 (with the exception of 13:40 to 13:48). Binance sent the transaction data for 13:40 to 13:48 between 13:55 and 13:57.

Panel C revises the data availability of transaction data: Transaction data that was initially missing between 13:16 and 13:56 has been overwritten, and each red bar signals that Binance has never delivered transaction data via API, even not with a time delay.[29]

The extension of API delays coincides with the overwritten (minute-level) data, which would be in line with an intentional and manually initiated delay. Remarkably, the time lags increase linearly. Each point represents the sending of data in the corresponding stream. All points are partly transparent, but there are multiple points overlapping. Sending multiple messages simultaneously gives the impression that Binance has queued and released data at those points in time. One can see that releases of trade and order book data alternate. After 15:08, the delays come to an end.

Panel D indicates that in every minute liquidations took place. While the data which Binance discloses does not tell the sum of liquidations, it at least signals that Binance triggered liquidations without any disruption even though the majority of clients did neither have access to place orders (or add collateral) via app, website, nor API.

[29]The transaction data is independent of the order book data shown in Panel B.

Next, we evaluate the mark price, which triggers liquidations as it determines the collateral value, resulting in liquidations when crossing a set threshold. Panel A of Figure 9 presents hourly boxplots of the difference (in levels) between the mark price and the spot price. While usually fluctuating closely around zero, the difference widens around noon, hitting its maximum in the three following hours. Notably, the deviations occur in both directions, while the negative direction is more pronounced. As explained in Section 2.2, Binance reserves the right to adjust the mark price, giving them free play to adapt it. The substantial deviations come hand-in-hand with the flash crash (see Panel C) and the liquidations. The fluctuations in both directions result in a maximal extent of liquidations as it can cross thresholds for long and short positions. These liquidations of both types of positions are particularly helpful for the insurance fund to unload its own long positions with liquidations of short positions (and vice versa).

Panel B shows the funding premium (see Section 2.2) of the futures which dropped below -0.1% and recovered subsequently. This development implies a transfer of cash from long to short holders, and characterizes an oversupply of sell orders. Figure A6 in the appendix outlines the funding premium for the reference period, showing that the funding premia never fell below -0.025%.

We proceed with statistical tests to demonstrate that the incident was different from the reference period comprising the ten days before. For each of the following metrics, we consider the arithmetic mean and the median, and test them by a *t*-test and a Wilcoxon signed-rank test. Table 4 shows the results. Regarding the Bitcoin price, we note that the differences between Binance versus other exchanges and the mark price both show firm increases within the crash, being statistically and economically significant. We find similar effects for trading volumes and transaction numbers. Concerning the API time lag, we find noticeable average delays, while the median does not show the immense effect, likely because we have chosen the crash interval so wide.

Finally, we conduct a non-parametric approach to measure the correlation of Bitcoin prices versus other exchanges or the mark price. We compute three correlation coefficients and compare them using Fisher's r-to$z$ approach in Table A2. This approach again finds statistically significant differences. All in all, we find further evidence that the Binance market place did not function regularly on May 19.

10 What Trading Volume would be realistic during the Crash?

This section estimates a realistic trading volume between 13:16 and 13:56 on May 19, 2021, the time period for which Binance was back-filling minute-level data on TradingView (denoted by *back-filled period* in this section). During that time period, Binance reports sharply declining trading volumes (see Figure 3).

Our empirical strategy to estimate a realistic trading volume is threefold: First, we identify predictors for the trading volume on Binance which are independent of the data in question. This independence is necessary as other variables in the data in question, such as the number of reported trades per minute, might also be unreliable. Second, we derive a model of Bitcoin trading volume and feed it with training data before May 19. Third, we apply this model to the back-filled period to better grasp the reported volume difference.

10.1 Model Construction

We run an OLS post-Lasso model to estimate minute-level trading volume using crypto market and sentiment data. As predictors, we first use Bitcoin price volatility, given the well-documented correlation between volatility and trading volume (Karpoff, 1987; Brailsford, 1996). As we consider a short time frame in the order of a few hours, a correspondingly brief look-back period of five minutes seems suitable. Second, we also include the Bitcoin price level (Lee et al., 2000), as market participants may be more inclined to trade at lower price levels (given that the majority of Bitcoin traders are long on average).

Third, we use sentiment indicators that are web-scrapped from social media platforms and were aggregated by the data vendor Augmento. They provide 279 indicators for Twitter, Reddit, and Bitcointalk on an hourly basis, to measure the sentiment relating to Bitcoin. These indicators represent categories of emotions and are derived from linguistic analysis.

To avoid overfitting, we use a Lasso regression in the first step to shrink the number of predictors.[30] Table A1 shows the Lasso regression results for a regularization parameter λ of 21.[31] Besides the volatility and price level, we obtain nine non-zero coefficients for social media indicators. Five of them relate to Twitter, while four stem from Reddit.

In the next step, we perform an OLS regression to estimate a realistic trading volume during the outage period, using these remaining predictors.[32] Table 6 presents the results. The first regression includes the volatility only, which explains already 50.3% of the variation. In line with intuition and the evidence in the literature, the point estimate is positive. The second regression adds the price level. We find a negative coefficient, i.e., there is more trading volume for lower prices. However, the adjusted R^2 increases by only 2.2%. We add the nine sentiment indicators in the third regression. Results are in line with expectation such that a predictor like "correction" is related with higher trading volume. Adding the sentiment data increases the adjusted R^2 to 57.7%.

10.2 Empirical Results

We use the estimated coefficients from the previous subsection to estimate a realistic trading volume on the minute-level for the back-filled period. We report the empirical results of the model in Panel A of Table 7. We observe that errors between the training data and test data (in the morning of May 19 before the crash) increase rather moderately (given that the test data is out-of-sample and thus more prone to deviations), indicating a good performance in the test data set. When assessing

[30] Lasso adds a penalty term to the minimization term of the OLS regression (Tibshirani, 1996). When solving the minimization problem of the sum of squared errors, the penalty term introduces a trade-off between errors and the absolute sum of coefficients. This least absolute shrinkage helps prevent overfitting as it selects variables by setting some coefficients to zero.

[31] We determine λ by ten-fold cross-validation, taking the mean squared error (MSE) as the measure type and a sequence of possible values for λ from 5 to 100. We then chose the largest λ at which the MSE remains within one standard error of the smallest MSE.

[32] Note that we are not interested in the inference of these predictors and hence do not correct their standard errors to account for the two-step estimation procedure (Belloni et al., 2013).

the back-filled period, we find a marked surge in errors, in particular if we compare the increase between columns 1 and 3 to the increase between columns 1 and 2. This finding points out that the volume in the back-filled period of interest is inconsistent with the model. A t-test and a Kruskal-Wallis-Test indicate the statistical significance of the differences.

Figure 10 depicts the predictions (blue) in comparison to the data reported by Binance (red). The first hours display a close correlation between the predicted and observed volumes. For the back-filled period (highlighted by the red box) the model predicts materially higher volumes than Binance reports. Panel B presents the volume difference which is oscillating around zero before the back-filled period and plunges just within the back-filled data to up to -5000 Bitcoin per minute. This estimate is consistent with the reported numbers of other exchanges where volumes remained relatively high at that time (compare Figure 3).

Panel B of Table 7 calculates the resulting differences within the back-filled period. The Bitcoin volume has been reported as 24 935, while the model predicts a sum of 113 152, implying an economically and statistically significant difference of 88 216. Expressed in USD, this value amounts to 3.0 bil. USD.

To sum it up, the model suggests that Binance reported trading volumes far away from the model estimates. These results support the notion that only VIP clients were able to trade (and to add collateral) during the back-filled period.

11 What Impact did the Incident have?

This section estimates the impact Binance's clients experienced during the flash crash as a mere result of being a Binance customer and not one at another exchange. We focus on Bitcoin futures, and any impact in other instruments such as the leveraged tokens comes on top. We report results two-fold: First, based on the volume reported by Binance, and secondly for the more realistic volume estimated by the model in Section 10.

We calculate the impact as the absolute difference of Binance's Bitcoin price minus the reference Bitcoin price of other exchanges as shown in Figure 3. Panel

C of Figure 10 portrays this difference. While staying close to zero before, it spikes within the suspicious time between 13:16 and 13:56, and then slowly converges to zero. This observation speaks in favor of the hypothesis that the suspicious time frame inflicts most of the impact.

Panel C of Table 7 sums up the impact in this forty-minute time frame. Already based on the small volume reported by Binance, the absolute differences add up to a total impact of 13.7. mil. USD. Even more substantial, we arrive at an impact of 64.9 mil. for the estimated volume.

Overall, this price difference of 65 mil. USD may serve as an estimated damage the majority of Binance clients suffered during the flash crash on May 19 because of the front-end and API outage. The beneficiaries in this zero-sum game were likely VIP clients and large Binance market makers. The latter seems particularly noteworthy since the SEC investigates trading firms potentially linked to Binance's founder (Benson, 2022).

12 Conclusion

We have examined the outage the largest crypto exchange Binance experienced during the flash crash of May 19, 2021, when it halted trading for retail clients and stopped providing transaction data. We find evidence that Binance back-filled these missing transactions with data that does not conform to Benford's Law.

During the outage, some trading volume remained. Its marketing of a priority API to VIP clients gives the impression that they had free play while smaller clients were at their mercy. The latter suffered from substantial and possibly unfair losses because they could neither close their open futures positions nor add collateral to avoid liquidations. Still, potential beneficiaries did not even "receive" all the losses since the lack of trading is not necessarily a zero-sum transfer – the total damage is more serious. Our model of a realistic trading volume suggests that the trading volume accounted to only 20 percent of the expected level. Hence, the large majority of clients were not able to act in their interest.

Our study suggests that crypto exchanges require more consumer protection. If an outage occurs, trading needs to stop for everyone. Any front-end (and general API) downtime, where clients submit orders and add collateral, requires a consequent stop of the liquidation engine as the resulting liquidations are otherwise unfair. A prioritization of valuable clients in the trading engine is strictly inappropriate.

The high leverage itself is at the heart of the problem. Offering high leverage seems misguided for (retail) investors in the already highly volatile crypto market. Leveraged products should only be offered to selected clients who have sufficient spot trading experience and enough capital. Finally, the crypto exchanges need to reveal the insurance funds' wallet addresses, put their insurance funds into separate and bankruptcy-remote legal entities, and become audited on a regular basis. Furthermore, implementing a (transparent) circuit breaker would allow time for margin calls to be met and reassess information.

References

Abrantes-Metz, R. M., M. Kraten, A. D. Metz, and G. S. Seow. 2012. "Libor manipulation?" *Journal of Banking & Finance* 36 (1): 136–150.

Aleti, S., and B. Mizrach. 2021. "Bitcoin spot and futures market microstructure." *Journal of Futures Markets* 41 (2): 194–225.

Amiram, D., E. Lyandres, and D. Rabetti. 2021. "Competition and Product Quality: Fake Trading on Crypto Exchanges." *SSRN Electronic Journal.*

Baker, P. 2020. *Binance Doesn't Have a Headquarters Because Bitcoin Doesn't, Says CEO.* Coindesk. www.coindesk.com/markets/2020/05/08/binance-doesnt-have-a-headquarters-because-bitcoin-doesnt-says-ceo/.

Belloni, A., V. Chernozhukov, and C. Hansen. 2013. "Inference on treatment effects after selection among high-dimensional controls." *Review of Economic Studies* 81 (2): 608–650.

Benford, F. 1938. "The Law of Anomalous Numbers." *Proceedings of the American Philosophical Society* 78 (4): 551–572.

Benson, J. 2022. *SEC Is Investigating Trading Firms Linked to Binance Founder.* Decrypt. decrypt.co/93016/sec-investigating-firms-linked-binance-founder-report.

Bouri, E., R. Gupta, and D. Roubaud. 2019. "Herding behaviour in cryptocurrencies." *Finance Research Letters* 29:216–221.

Brailsford, T. J. 1996. "The empirical relationship between trading volume, returns and volatility." *Accounting and Finance* 36 (1): 89–111.

Brogaard, J., T. Hendershott, and R. Riordan. 2014. "High-frequency trading and price discovery." *Review of Financial Studies* 27 (8): 2267–2306.

Chen, J., D. Lin, and J. Wu. 2022. "Do cryptocurrency exchanges fake trading volumes? An empirical analysis of wash trading based on data mining." *Physica A: Statistical Mechanics and its Applications* 586:126405.

Cho, W. K., and B. J. Gaines. 2007. "Breaking the (Benford) law: Statistical fraud detection in campaign finance." *American Statistician* 61 (3): 218–223.

Cong, L. W., X. Li, K. Tang, and Y. Yang. 2021. "Crypto Wash Trading." *SSRN Electronic Journal.*

Drake, P. D., and M. J. Nigrini. 2000. "Computer assisted analytical procedures using Benford's Law." *Journal of Accounting Education* 18 (2): 127–146.

Durtschi, C., W. Hillison, and C. Pacini. 2004. "The effective use of Benford's law to assist in detecting fraud in accounting data." *Journal of Forensic Accounting* 5 (1): 17–43.

Fewster, R. M. 2009. "A simple explanation of benford's law." *American Statistician* 63 (1): 26–32.

Fisher, R. A. 1915. "Frequency Distribution of the Values of the Correlation Coefficient in Samples from an Indefinitely Large Population." *Biometrika* 10 (4): 507–521.

Foley, S., J. R. Karlsen, and T. J. Putnins. 2019. "Sex, Drugs, and Bitcoin: How Much Illegal Activity Is Financed through Cryptocurrencies?" *The Review of Financial Studies* 32 (5): 1798–1853.

Gandal, N., J. T. Hamrick, T. Moore, and T. Oberman. 2018. "Price manipulation in the Bitcoin ecosystem." *Journal of Monetary Economics* 95:86–96.

Gatheral, J., T. Jaisson, and M. Rosenbaum. 2018. "Volatility is rough." *Quantitative Finance* 18 (6): 933–949.

Gerety, M. S., and J. H. Mulherin. 1992. "Trading Halts and Market Activity: An Analysis of Volume at the Open and the Close." *The Journal of Finance* 47 (5): 1765–1784.

Greene, W. H. 1997. *Econometric Analysis.* 373–374. Upper Saddle River: Prentice Hall.

Griffin, J. M., and A. Shams. 2020. "Is Bitcoin Really Untethered?" *The Journal of Finance* 75 (4): 1913–1964.

Huberman, G., J. D. Leshno, and C. Moallemi. 2021. "Monopoly without a Monopolist: An Economic Analysis of the Bitcoin Payment System." *The Review of Economic Studies* 88 (6): 3011–3040.

Jiang, C., T. McInish, and J. Upson. 2009. "The information content of trading halts." *Journal of Financial Markets* 12 (4): 703–726.

Jiang, X. F., B. Zheng, T. Qiu, and F. Ren. 2017. "Extreme-volatility dynamics in crude oil markets." *European Physical Journal B* 90 (2).

Karpoff, J. M. 1987. "The Relation Between Price Changes and Trading Volume: A Survey." *The Journal of Financial and Quantitative Analysis* 22 (1): 109–126.

Kowsmann, P., and C. Ostroff. 2021. *Binance Froze When Bitcoin Crashed. Now Users Want Their Money Back.* The Wall Street Journal. www.wsj.com/articles/binance-froze-when-bitcoin-crashed-now-users-want-their-money-back-11626001202.

Lee, C. M. C., B. Swaminathan, Y. Amihud, H. Bierman, L. Brown, T. Dyckman, D. Easley, et al. 2000. "Price Momentum and Trading Volume." *The Journal of Finance* 55 (5): 2017–2069.

Liu, Y., and A. Tsyvinski. 2021. "Risks and Returns of Cryptocurrency." *The Review of Financial Studies* 34 (6): 2689–2727.

Makarov, I., and A. Schoar. 2020. "Trading and arbitrage in cryptocurrency markets." *Journal of Financial Economics* 135 (2): 293–319.

Moore, T., and N. Christin. 2013. "Beware the Middleman: Empirical Analysis of Bitcoin-Exchange Risk." *Lecture Notes in Computer Science* 7859 LNCS:25–33.

Robinson, M. 2021. *Cryptocurrency Tether is fined $41 million for lying about reserves.* Fortune. fortune.com/2021/10/15/tether-crypto-stablecoin-fined-reserves/.

Samson, O., and J. Oliver. 2021. *'I was panicking': the high-risk bets sparking a backlash at Binance.* Financial Times. www.ft.com/content/f7f7f110-32d5-4931-a7b7-d3cd0f63410a.

Sokolov, K. 2021. "Ransomware activity and blockchain congestion." *Journal of Financial Economics* 141 (2): 771–782.

Subrahmanyam, A. 1994. "Circuit Breakers and Market Volatility: A Theoretical Perspective." *The Journal of Finance* 49 (1): 237–254.

Tibshirani, R. 1996. "Regression Shrinkage and Selection via the Lasso." *Journal of the Royal Statistical Society* 58 (1): 267–288.

Tobin, J. 1958. "Estimation of Relationships for Limited Dependent Variables." *Econometrica* 26 (1): 24–36.

Tables

Table 1: Summary Statistics of Binance Transaction Data

This table reports summary statistics on the tick-level transaction and order book Binance data for the entire time interval between May 9 and May 20, 2021. We use Tardis API data and fill Bitcoin price and volume gaps on May 19 with Binance's website data. There are 73 Mio. observations. The Bitcoin (BTC) price and amount refer to the futures transactions. The order book data states the order book situation prevailing at the time of each transaction. For the bids, Bid 1 refers to the highest bid, whereas Bid 2 is the second-highest bid, and so on (and vice versa for the asks). The volumes relate to the amount available at the corresponding price level. The Insurance Fund dummy variable equals 1 if one counterparty in the underlying transaction was the Binance-owned insurance fund. The funding premium is a mechanism (fee) compensating the term structure of futures to the spot price. The mark price is the Bitcoin price used to trigger liquidations. The time difference refers to the time lag between Binance's API push and the event time, i.e., trade or order book change (depth).

				Quantiles					
Variable	Mean	SD	Min	5%	25%	50%	75%	95%	Max
Traded BTC Price	46229	45214	0	7274	40433	45214	50169	57697	294838
Traded BTC Amount	0.094	0.017	0.001	0.001	0.002	0.017	0.072	0.397	490
Orderbook Bid 1 Price	46226	45212	28737	37270	40432	45212	50168	57696	59652
Orderbook Bid 2 Price	46222	45210	3587	37267	40430	45210	50165	57692	59651
Orderbook Bid 3 Price	46219	45208	0.99	37265	40428	45208	50163	57691	59651
Orderbook Bid 1 Volume	1.05	0.137	0.001	0.002	0.04	0.137	0.496	2.402	2535
Orderbook Bid 2 Volume	0.444	0.1	0.001	0.002	0.022	0.1	0.289	1.288	2543
Orderbook Bid 3 Volume	0.439	0.106	0.001	0.002	0.036	0.106	0.301	1.404	2515
Orderbook Ask 1 Price	46230	45215	28762	37276	40435	45215	50171	57698	62110
Orderbook Ask 2 Price	46233	45217	28780	37278	40438	45217	50173	57700	90818
Orderbook Ask 3 Price	46236	45219	28780	37281	40439	45219	50176	57703	91918
Orderbook Ask 1 Volume	0.909	0.139	0.001	0.002	0.04	0.139	0.5	2.413	1966
Orderbook Ask 2 Volume	0.373	0.1	0.001	0.002	0.023	0.1	0.291	1.271	1952
Orderbook Ask 3 Volume	0.374	0.107	0.001	0.002	0.036	0.107	0.302	1.4	1957
Indicator Insurance Fund	0.001	0	0	0	0	0	0	0	1
Open Interest	34786	34794	23695	24538	32212	34794	39865	43529	46484
Long-Ratio All Traders	2.554	2.698	1.187	1.441	2.271	2.698	2.945	3.419	4.066
Long-Short Ratio Top Accounts	2.591	2.689	1.385	1.57	2.386	2.689	2.911	3.363	3.866
Long-Short Ratio Top Position	1.195	1.191	0.859	1.122	1.158	1.191	1.235	1.288	1.439
Funding Premium (Percent)	-0.002	-0.008	-0.139	-0.02	-0.02	-0.008	0.013	0.045	0.088
Mark Price	46224	45209	29563	37264	40431	45209	50167	57696	59595
Message Time Difference Trades	0.101	0	0	0	0	0	0	0	35.726
Message Time Difference Depth	0.096	0	0	0	0	0	0	0	35.547

Table 2: Statistics of the Crash

This table provides an overview of the May 19, 2021, crash event. It is based on Tardis API data and Bitcoin price and volume gaps filled with Binance website data. It starts at 00:01 on May 19. All times are expressed in UTC. The trading volumes comprise the time frame from 12:00 to 16:00 – the start of the price decline until its stabilization. The insurance fund Balance shown represents the numbers published by Binance.

Variable	Value
Timing	
Date	May 19, 2021
Start Time	≈12:30
BTC Price at 00:00	43 025 USD
BTC Price at 12:00	38 630 USD
BTC Drawdown	
Maximum Price	43 527 USD
Maximum Time	0:13
Minimum Price	28 758 USD
Minimum Time	13:09
Maximum Drawdown	33.9%
BTC Recovery	
Stabilization Time	≈14:00
Stabilization Price	35 000 USD
New High reached at	16:38
New High Price	38 764 USD
Trading Volume	
on Binance in BTC	357 607
on Binance in USD	12.7 bil.
# Transactions	2.7 mil.
Insurance Fund Balance	
Day before Crash	290.4 mil.
Day after Crash	288.7 mil.

Table 3: Benford Statistical Analysis of Added Data

This table presents an analysis of the distribution of the volume of Bitcoin transactions by Benford's Law (Benford, 1938). It considers precisely the data which seems to have been added with a time delay between 13:16 and 13:56 on May 19, 2021 using both Tardis API data and Binance website data (in case of gaps). The analysis has been carried out by focusing on the first (non-zero) digit ("significand").

Panel A reports the Benford analysis of each digit. The *relative frequency* reports the observed occurrences of each digit, while the *Benford Difference* states the difference to the relative frequency which Benford's Law suggests. Statistical significance is provided by Pearson's Chi-squared (χ^2) test.

Panel B shows the descriptive statistics relevant to the Benford analysis. The *Mean Absolute Deviation* (MAD) represents the arithmetic mean of absolute differences to the expected proportion. The *Distortion Factor* estimates whether numbers are overstated (positive number) or understated (negative number).

Panel C states the results of the statistical tests indicating whether the overall distribution adheres to Benford's Law. The null hypothesis of *Pearson's χ^2 Test* and the *Euclidian Distance Test* are that the observed distribution conforms to Benford's law. The null hypothesis of the *Mantissa Arc Test* assumes that the mantissa of the observed distribution is uniformly distributed. *, **, and *** denote significance at the 10%, 5%, and 1% level.

Panel A: Digit Analysis

Digit	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**
Relative Frequency	53.5%	16.9%	11.6%	4.2%	5.5%	1.0%	3.3%	3.3%	0.5%
Benford Difference	23.4%	−0.7%	−0.8%	−5.4%	−2.4%	−5.7%	−2.5%	−1.8%	−4.1%
Significance	***	***	***	***	***	***	***	***	***

Panel B: Descriptive Statistics

Measure	**Value**
Observations	942 924
Mean	0.250
Variance	0.088
Ex. Kurtosis	−0.609
Skewness	0.817
Mean Absolute Deviation	0.052
Distortion Factor	−54.294

Panel C: Goodness of Fit Tests

Test	**Statistic**
Pearson's χ^2 Test	304 183***
Kolmogorov-Smirnov Test	227.38***
Euclidian Distance Test	246.28***
Mantissa Arc Test	0.13***

Table 4: Tests of the Difference between Binance and Other Exchanges

Using TradingView and Tiingo data on the minute-level, this table tests a break in the absolute difference of Bitcoin prices in USD on Binance against the other exchanges (Panel A) as well as its mark price (Panel B). Panel C reports the differences in Bitcoin volume, and Panel D presents the number of transactions differences. Panels E and F show the time difference between transactions and order book changes to the API processing. We control for the crash window from 12.00 to 14.00 (UTC) on May 19, 2021, ("during the crash") and the time before ("before the crash"). *, **, *** denote significance at the 10%, 5%, and 1% level. The mean is assessed with a t-test, while the Wilcoxon signed-rank test evaluates the median. We test the values in the first two columns against zero, whereas the difference constitutes a test of the subtrahend and the minuend.

Panel A: Binance Bitcoin Price vs. Other Exchanges

Absolute Difference	Before the Crash	During the Crash	Difference
Mean	53***	375***	322***
Median	43*	189*	146***

Panel B: Binance Bitcoin Price vs. Mark Price

Absolute Difference	Before the Crash	During the Crash	Difference
Mean	41***	295***	254***
Median	31*	165*	134***

Panel C: Binance Bitcoin Trading Volume vs. Other Exchanges

Absolute Difference	Before the Crash	During the Crash	Difference
Mean	280***	1561***	1281***
Median	168*	960*	791***

Panel D: Binance Bitcoin Transaction Number vs. Other Exchanges

Absolute Difference	Before the Crash	During the Crash	Difference
Mean	2124***	20 460***	18 335***
Median	974*	14 719*	13 744***

Panel E: Binance Trades API Time Lag in Minutes

Absolute Difference	Before the Crash	During the Crash	Difference
Mean	0.00***	2.84***	2.84***
Median	0.00*	0.00*	0.00***

Panel F: Binance Order Book API Time Lag in Minutes

Absolute Difference	Before the Crash	During the Crash	Difference
Mean	0.00***	6.59***	6.59***
Median	0.00*	0.00*	0.00**

Table 5: Insurance Fund Model

This table reports the results of models for the trading volume of Binance's insurance fund, trained with the data from May 9 to May 19, 13:15. Panel A shows results for the Tobit. Panel B presents the corresponding OLS regressions which require insurance fund trades of at least 1 Bitcoin. *, **, *** denote significance at the 10%, 5%, and 1% level.

Panel A: Tobit Models

	Tobit Model 1	**Tobit Model 2**
Dependent Variable	Insurance Fund Volume	Insurance Fund Volume
Constant	−438.45***	−303.21***
Bitcoin Volatility	1.19***	1.11***
Bid Liquidity		1.57***
Ask Liquidity		−24.71***
Observations	15 149	15 149
Log-Likelihood	−4910	−4854
Pseudo R^2	0.11	0.12

Panel B: Truncated OLS Regressions

	OLS Regression 1	**OLS Regression 2**
Dependent Variable	Insurance Fund Volume	Insurance Fund Volume
Constant	−199.55***	−70.05
Bitcoin Volatility	0.93***	0.98***
Bid Liquidity		0.93
Ask Liquidity		−34.68
Observations	93	93
Adjusted R^2	0.27	0.28

Table 6: Regressions of Bitcoin Volume

This table reports the results of three ordinary least squares regressions:

(i) $\text{volume}_{\text{BTC}}(t) = \alpha + \beta \cdot \text{volatility}_{\text{price}}(\,t-4,\, t\,) + \epsilon(t)\,.$

(ii) $\text{volume}_{\text{BTC}}(t) = \alpha + \beta \cdot \text{volatility}_{\text{price}}(\,t-4,\, t\,) + \gamma \cdot \text{price}_{\text{BTC}}(t) + \epsilon(t)\,,$

(iii) $\text{volume}_{\text{BTC}}(t) = \alpha + \beta \cdot \text{volatility}_{\text{price}}(\,t-4,\, t\,) + \gamma \cdot \text{price}_{\text{BTC}}(t) + \delta \cdot \text{sentiment}(t) + \epsilon(t)\,,$

The regressions are conducted on a training data set of minute-level Binance data in 2021 up to May 18, to derive a model for the Bitcoin Volume on May 19, 2021. The dependent variable is the traded volume in Bitcoin in each minute. We define volatility as the five-minute rolling standard deviation of Bitcoin prices. The Bitcoin price reflects the volume-weighted arithmetic mean of transaction prices in each minute. The sentiment is a vector of nine indicators measuring crowd psychology on Twitter and Reddit, which have been web scraped and aggregated by Augmento. Each indicator represents one category of the underlying linguistic analysis, to which we refer by its generic term in quotation marks. We have selected these indicators out of 293 available series using the LASSO regression reported in Table A1. *, **, *** denote significance at the 10%, 5%, and 1% level.

	Regression 1	**Regression 2**	**Regression 3**
Dependent Variable	Bitcoin Volume	Bitcoin Volume	Bitcoin Volume
Constant	−3.207	578.781***	34.597
Bitcoin Volatility	5.739***	5.447***	4.626***
Bitcoin Price		−0.011***	−0.003***
Twitter			
"Bottom"			4.155***
"Technical Analysis"			2.728***
"Listing"			6.267***
"Dip"			−0.604
"Price"			1.146***
Reddit			
"Bottom"			6.002***
"Technical Analysis"			2.421***
"Dip"			0.063
"Correction"			6.793***
Adjusted R^2	0.503	0.525	0.577
F-Statistic	13 926***	7610***	239***

Table 7: Empirical Results of the Bitcoin Volume Model

This table reports the empirical evaluation of the model estimated in Table 6 (Regression 3) for the period presented in Figure 10. The training data comprises the ten reference days before May 19. The test data lies between 00:00 and 10:00 on May 19. The back-filled data is 13:16 to 13:56 on May 19, 2021.

Panel A analyzes the model accuracy. We test both the test data and the back-filled data against the training data, using a *t*-test (means) and the Wilcoxon signed-rank test (medians). The Kruskal-Wallis Test has for each row (except the Entropy Coefficient which we have not tested) the null hypothesis that the three data sets originate from the same distribution. The Entropy Coefficient measures forecast accuracy (Greene, 1997).

Panel B evaluates the results of the model for the back-filled data, i.e., precisely for 13:16 to 13:56 on May 19, 2021. It shows the Bitcoin volume difference in the reported data compared to the model's prediction. The volume in US-Dollar reports the USD volume at market prices of the individual trades.

Panel C shows the sum of volume-weighted absolute price differences resulting from the difference between Binance and other exchanges. The "reported" column states this difference to the trading volumes reported by Binance, while the "predicted" column relies on the volume predicted by the model. *, **, *** denote significance at the 10%, 5%, and 1% level.

Panel A: Model Accuracy

Measure	Train Data Data	Test Data Data	Back-Filled Data	Kruskal-Wallis Test
Mean of Absolute Errors (MAE)	139	356***	2151***	122***
Median of Absolute Errors	85	204***	1828***	122***
Root Mean Square Error (RMSE)	241	665***	2562***	408***
Median of Squared Errors	7281	41 883***	3 344 049***	408***
Entropy Coefficient	0.47	0.58	3.65	

Panel B: Volume Difference of the Back-Filled Data

Measure	Reported	Predicted	Difference
BTC Volume	24 935	113 152	88 216***
Volume in USD	0.86 bil.	3.88 bil.	3.0 bil.***

Panel C: Price Difference to Other Exchanges

Measure	Reported	Predicted	Difference
Difference in USD	13.7 mil.	64.9 mil.	51.2 mil.***

Figures

Figure 1: Timeline of the Crash

This figure depicts a chart of the Bitcoin price of the three days from May 18 to May 20, 2021, such that the crash day May 19 lies in the middle. The adjacent plots show the traded volume in Bitcoin and the underlying number of trades executed, both aggregated on a one-minute level. All trading data stems from Binance. The red box highlights the crash time window between 12:00 and 16:00 on May 19. Binance initially supplied no data for the time frame between 13:16 and 13:56 and filled it up later.

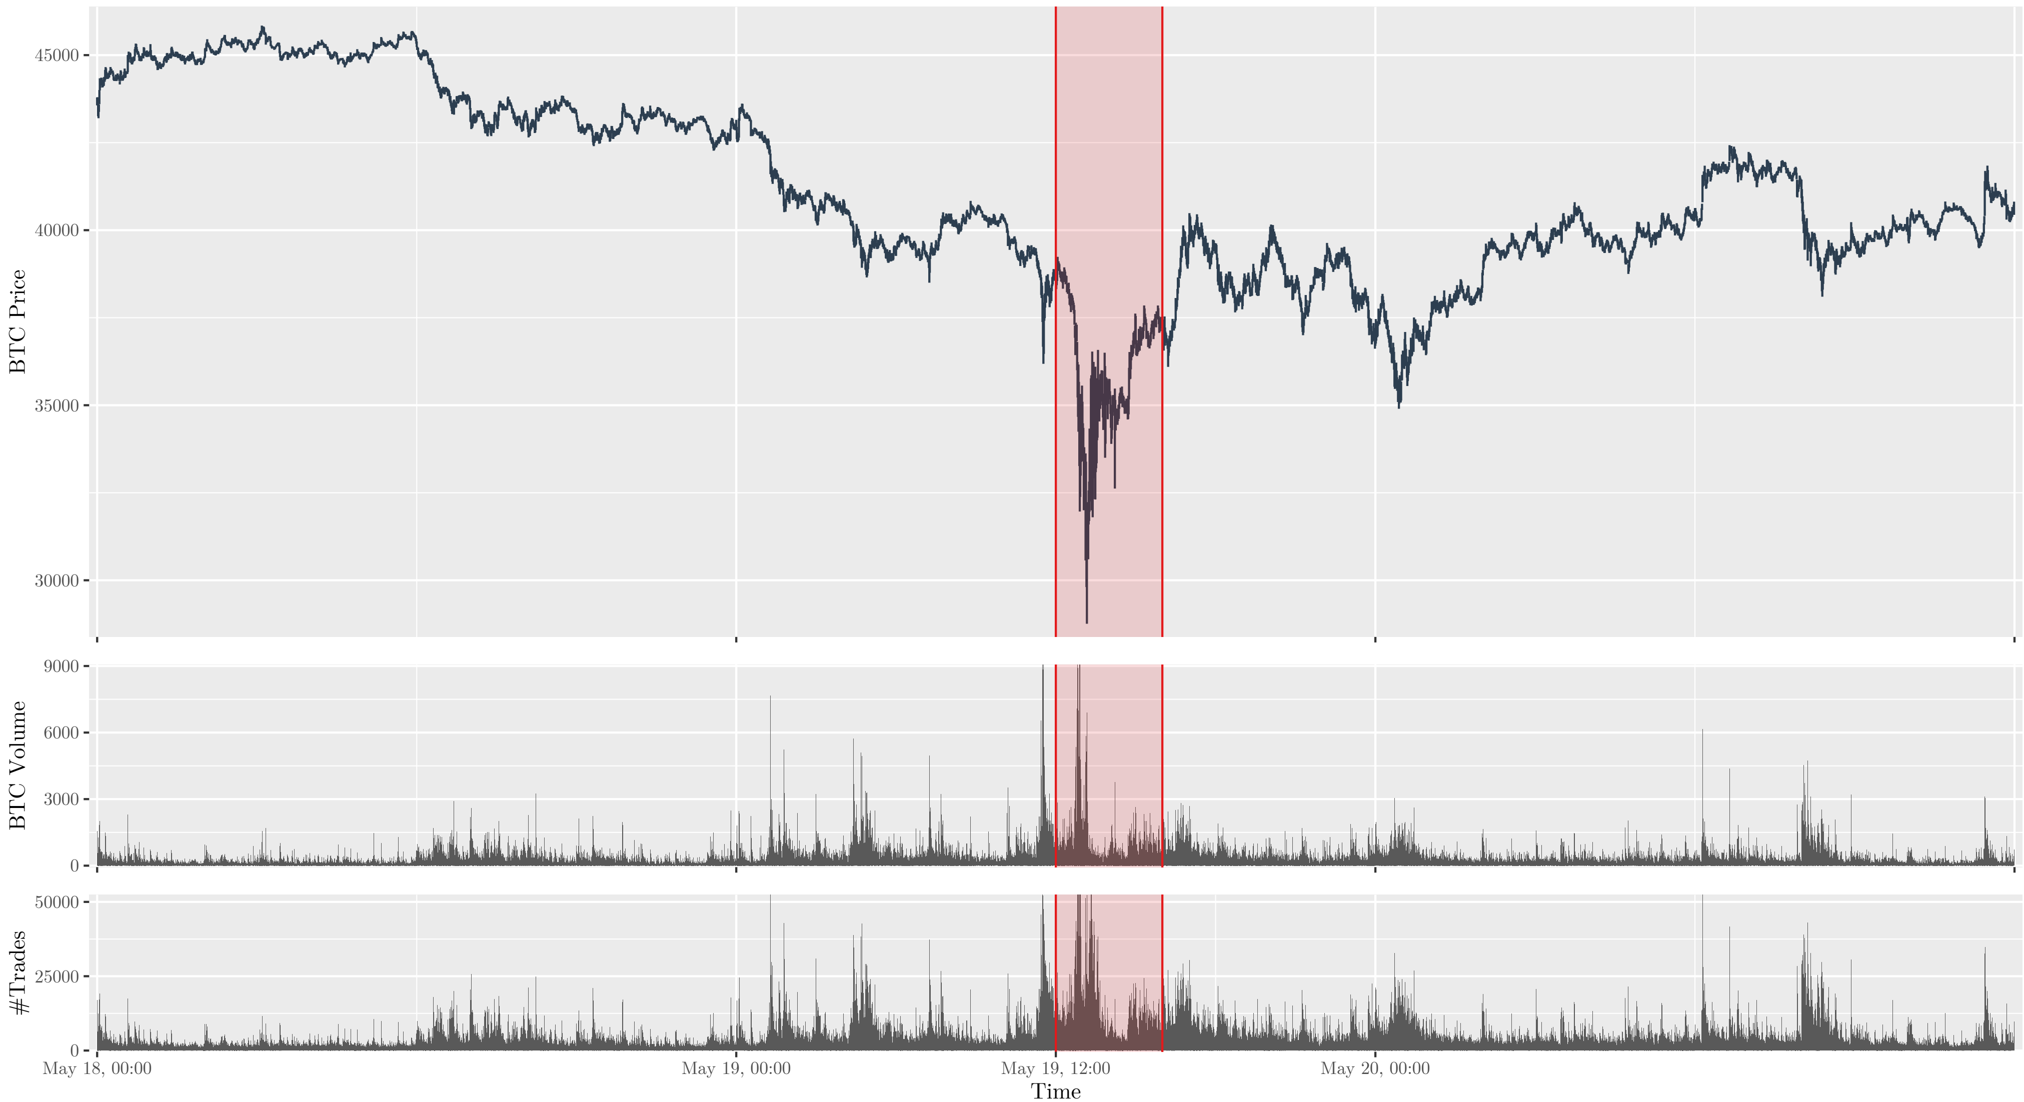

Figure 2: Benford Graphical Analysis of Added Data

This figure presents an analysis of the distribution of the volume of Bitcoin transactions by Benford's Law (Benford, 1938). It considers precisely the data between 13:16 and 13:56 on May 19, 2021, which were back-filled with a time delay. The analysis focus on the first (non-zero) digit ("significand"). The red lines indicate the results to which real data should converge under Benford's Law. Figure (a) depicts the observed frequency of the first digit (significand), whereas Figure (b) shows the pairwise difference between consecutive values. Figures (c) and (d) contain the distribution and difference of the digit summation. Figure (e) plots the χ^2-statistic, pointing out to which digits the discrepancies can be traced back.

(a) Digit Distribution

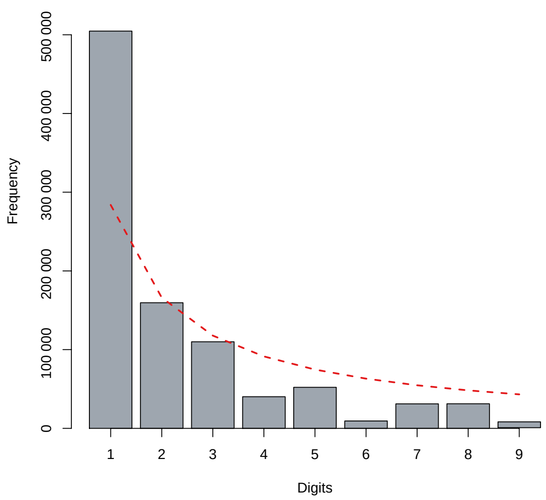

(b) Digit Distribution Second Order Test

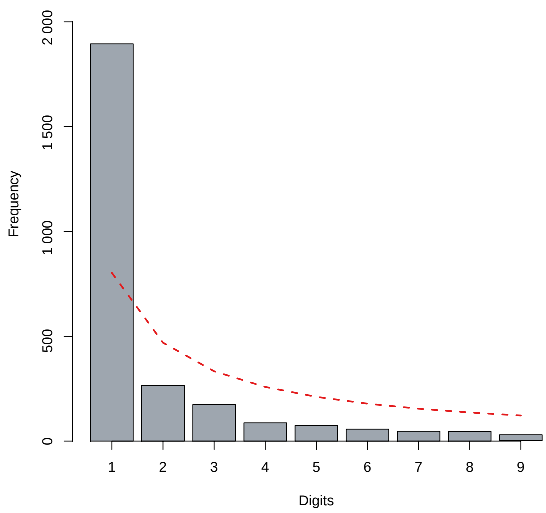

(c) Summation Distribution

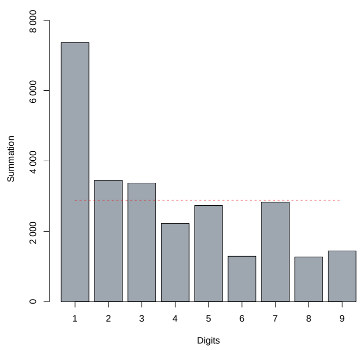

(d) Summation Difference

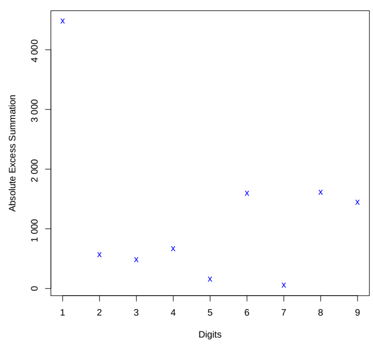

(e) Chi-Squared Difference

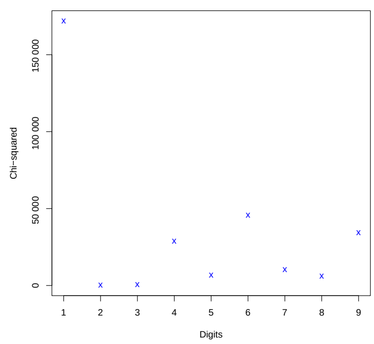

------- Distribution Expected by Benford's Law

Figure 3: Comparison with Other Exchanges

This figure compares the course of the May 19 crash on Binance with other exchanges. The red box highlights the time frame between 13:16 and 13:56 for which Binance back-filled data. Panel A depicts the Bitcoin futures transaction prices. While the lines mark each minute's mean price, the ribbons represent the 95% quantile price range. Panels B and D show the Bitcoin volume traded at Binance and the other exchanges. Panel C presents the Bitcoin volume traded by Binance's insurance fund. Panels E and F plot the number of executed trades at Binance and the other exchanges.

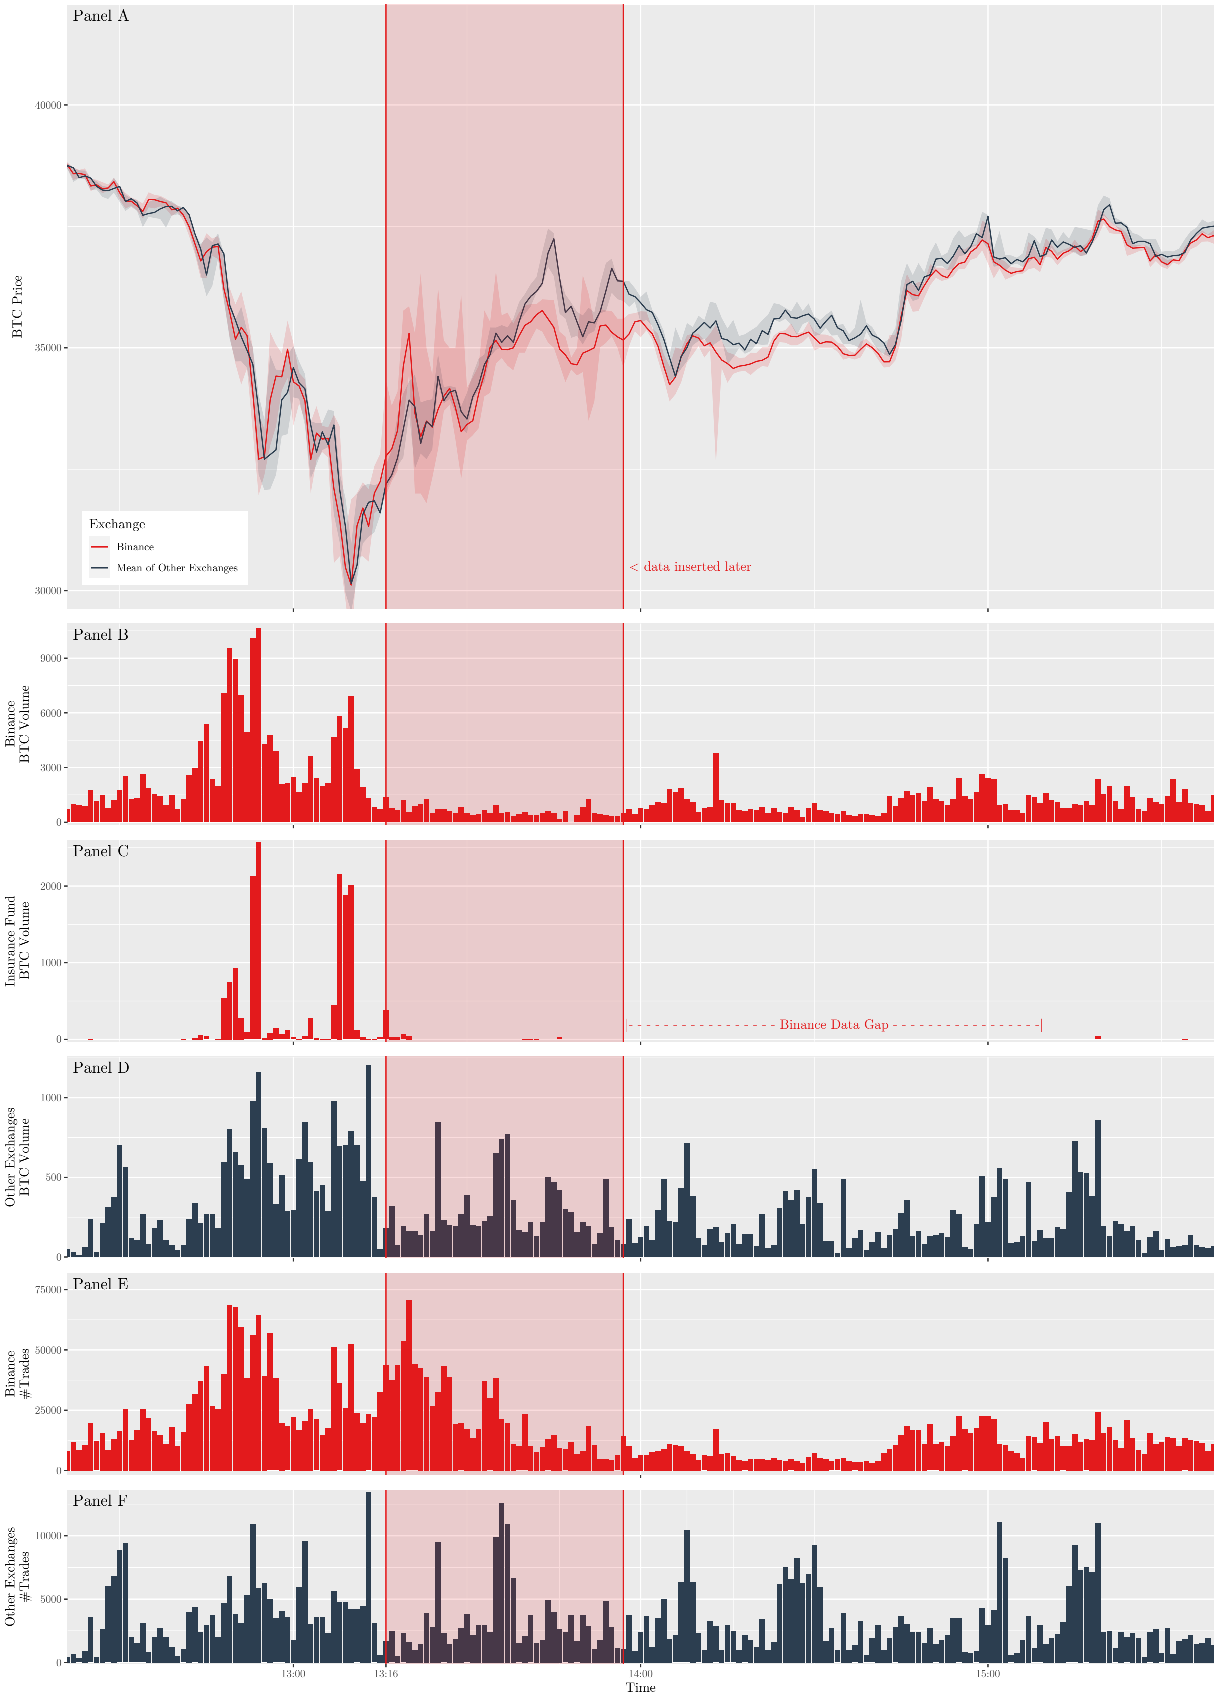

Figure 4: Leveraged Bearish Tokens

This figure presents the performance of bearish leveraged Bitcoin tokens on Binance (BTCDOWNUSDT) and other exchanges (Poloniex's and FTX's BEARUSDT). The red box highlights the crash and recovery window from 12:30 to 15:30. The tokens aim to replicate the inverted (minus) Bitcoin performance (average of various exchanges), levered by the factor 3. Hence, the relevant benchmark is the inverse of the Bitcoin spot price, multiplied by the leverage factor. Volumes are expressed in thousand USD.

Figure 5: Insurance Fund Model

This figure presents the May 19 out-of-sample results of our model for the trading volume of Binance's insurance fund. Table 5's Tobit model 2 yields the insurance fund's trading volume. Panel A compares the model prediction (blue) to the numbers reported by Binance (red). Panel B shows the resulting cumulative sum of insurance fund volume in million USD, based on Binance's reported numbers before 13:16 and the model estimate thereafter as this data is missing. Panel C exhibits the order book bid liquidity, which is one of the regressors. We calculate it as the sum of volume available within no more than 0.5% of the highest bid price. Panel D revises the data availability. Panel E adds the Bitcoin price for comparison.

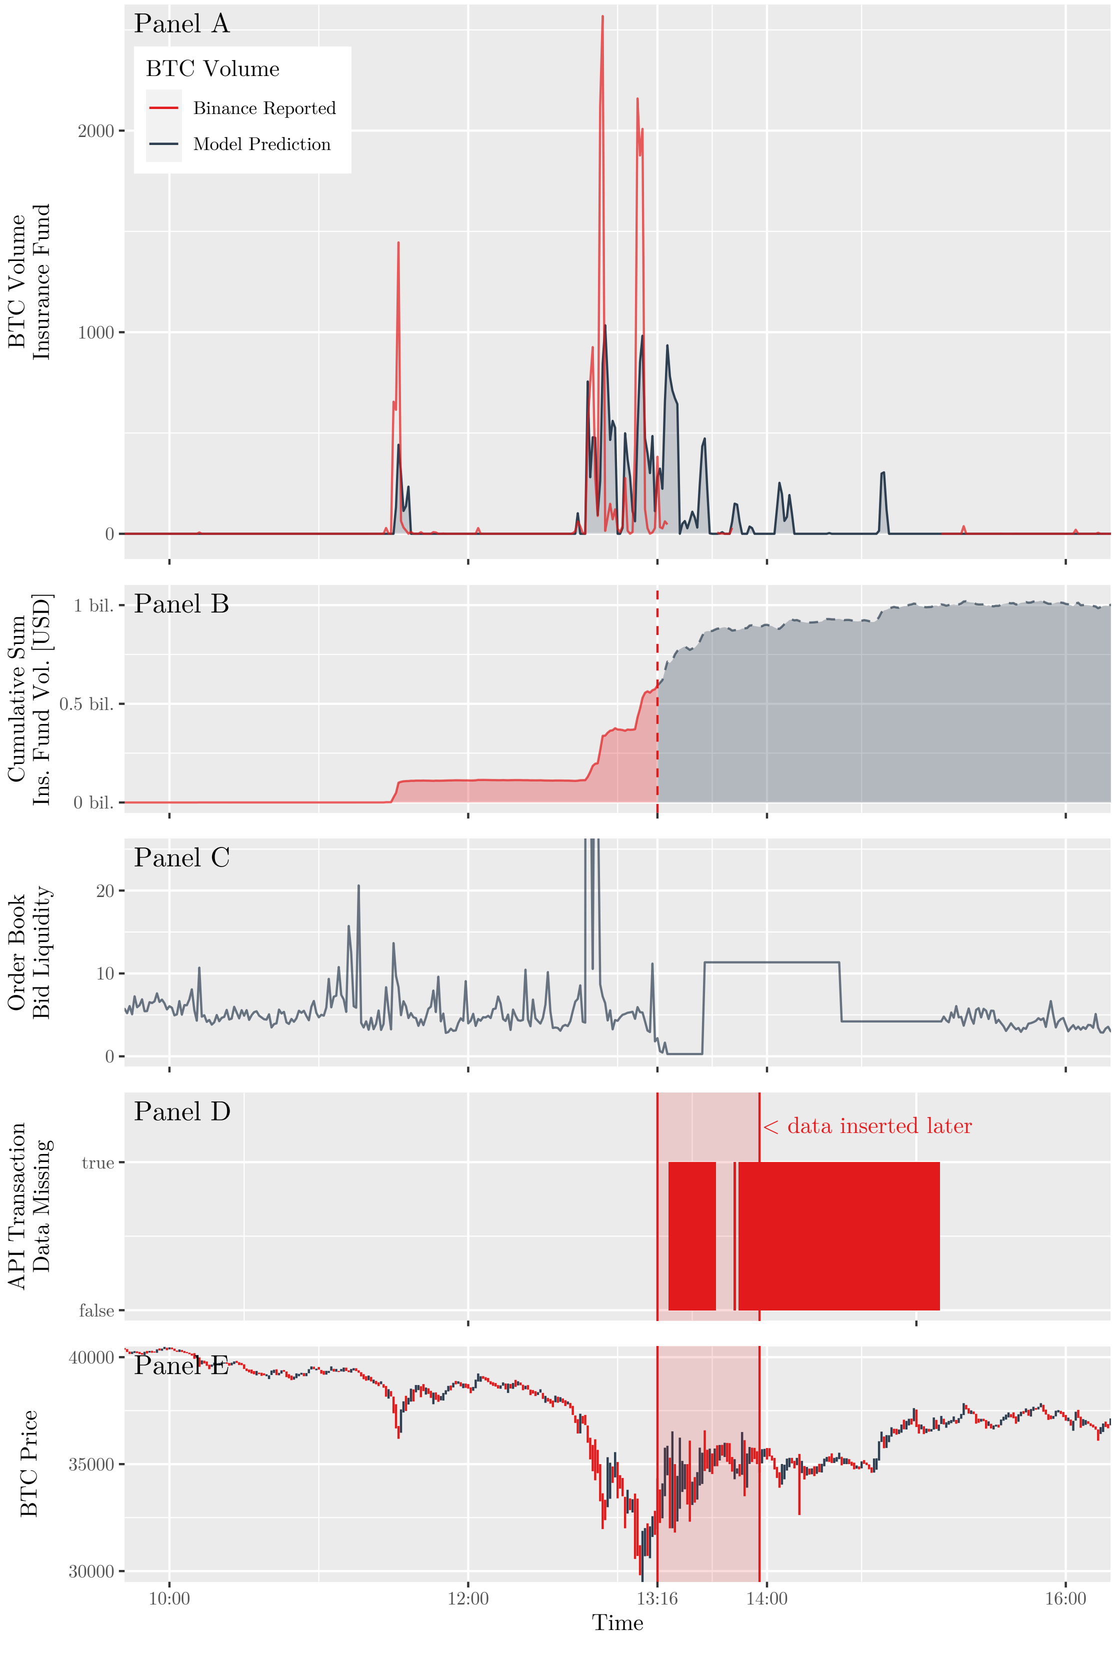

Figure 6: Bitcoin Open Interest on Binance

This figure depicts the development of Bitcoin's open interest on Binance. The chart starts ten days before the crash. The red box highlights the crash day of May 19, 2021. The data stems from Binance's API.

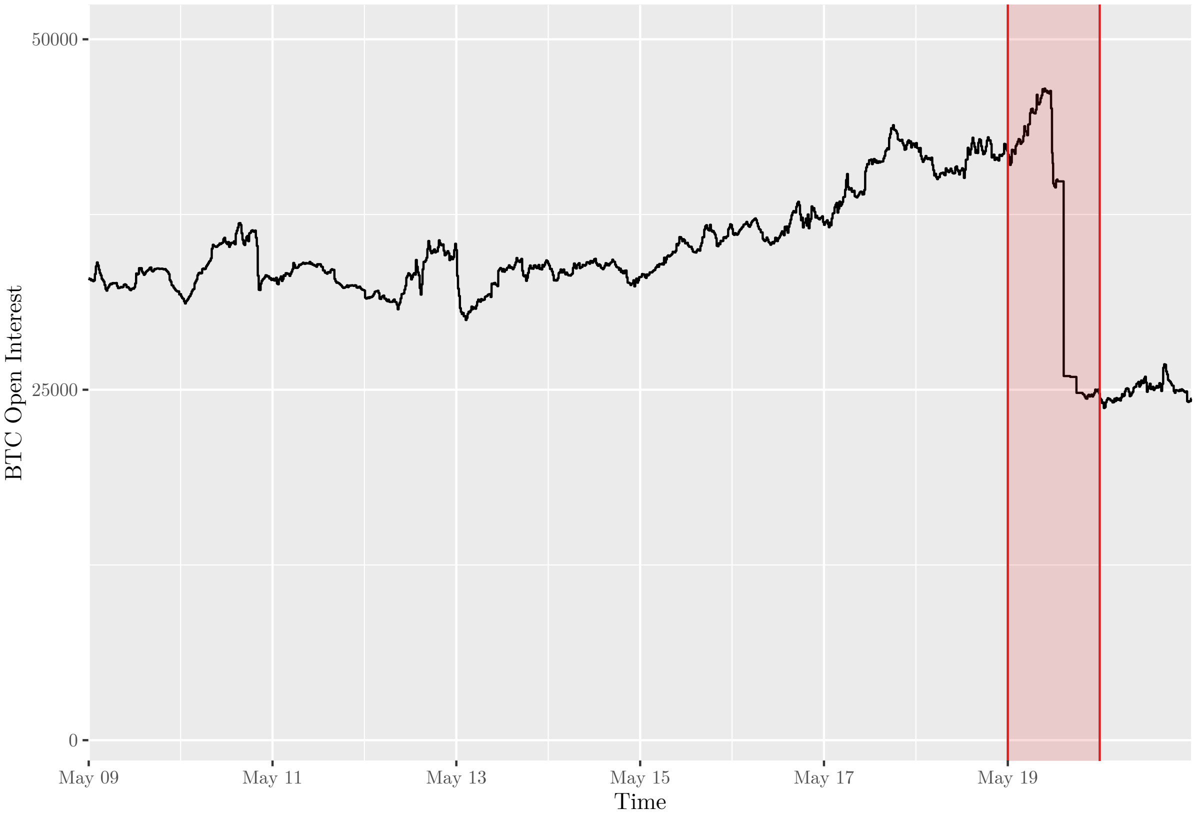

Figure 7: Long-Short Ratio by Account Type

This figure depicts the long-short ratio of Bitcoin traders on Binance. The long-short ratio indicates the volume-weighted quotient of long accounts and short accounts in Bitcoin. For reference, an equilibrium between long and short positions would imply a ratio of 1. Values higher than 1 represent more long than short interest.

The data stems from Binance's API which distinguishes by account type: Most generally, all accounts include everyone. The top accounts take out a subset of the accounts with the highest overall market exposure. The top positions refer to the accounts with the most significant Bitcoin positions. The chart starts ten days before the crash. The red box highlights the crash day of May 19, 2021.

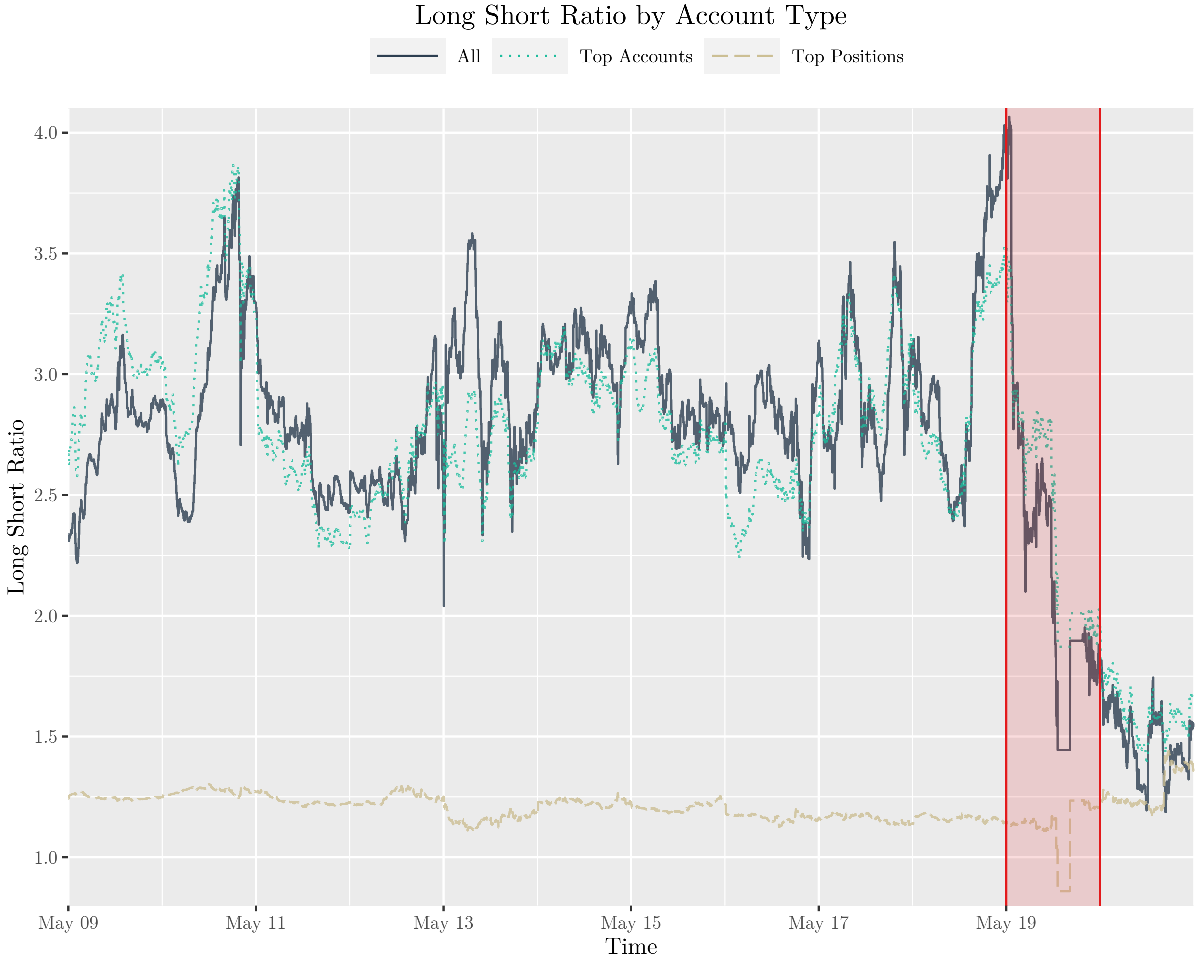

Figure 8: API Time Lag

This figure provides an overview of when Binance sent its API transaction data on May 19, 2021 and whether it has been sent. Panel A and B present the delay of updates sent by Binance on its API on May 19, 2021 for (A) the transaction data, and (B) the order book data stream. The x-axis shows the time at which the transactions have been executed. The delay means that data is sent for a specific point in time later on. Outside May 19's crash, the delay usually remains at seconds. The timestamps of updates sent come directly from Binance.

Panel C shows whether transaction data has been provided at all (no bar shown) or whether it is missing (bar shown). The bars indicate that no transaction data has been delivered, even not (intensely) delayed. However, Binance delivered aggregated data to TradingView, a part of which has been overwritten with other values later, highlighted with the red box. There are accordingly gaps in Panel A, whereas the order book data stream in Panel B is independent of these gaps.

Panel D shows a bar for each minute where at least one liquidation (forced order) has taken place. Binance's data does not permit to consider the volume meaningfully as only the first liquidation per minute is reported.

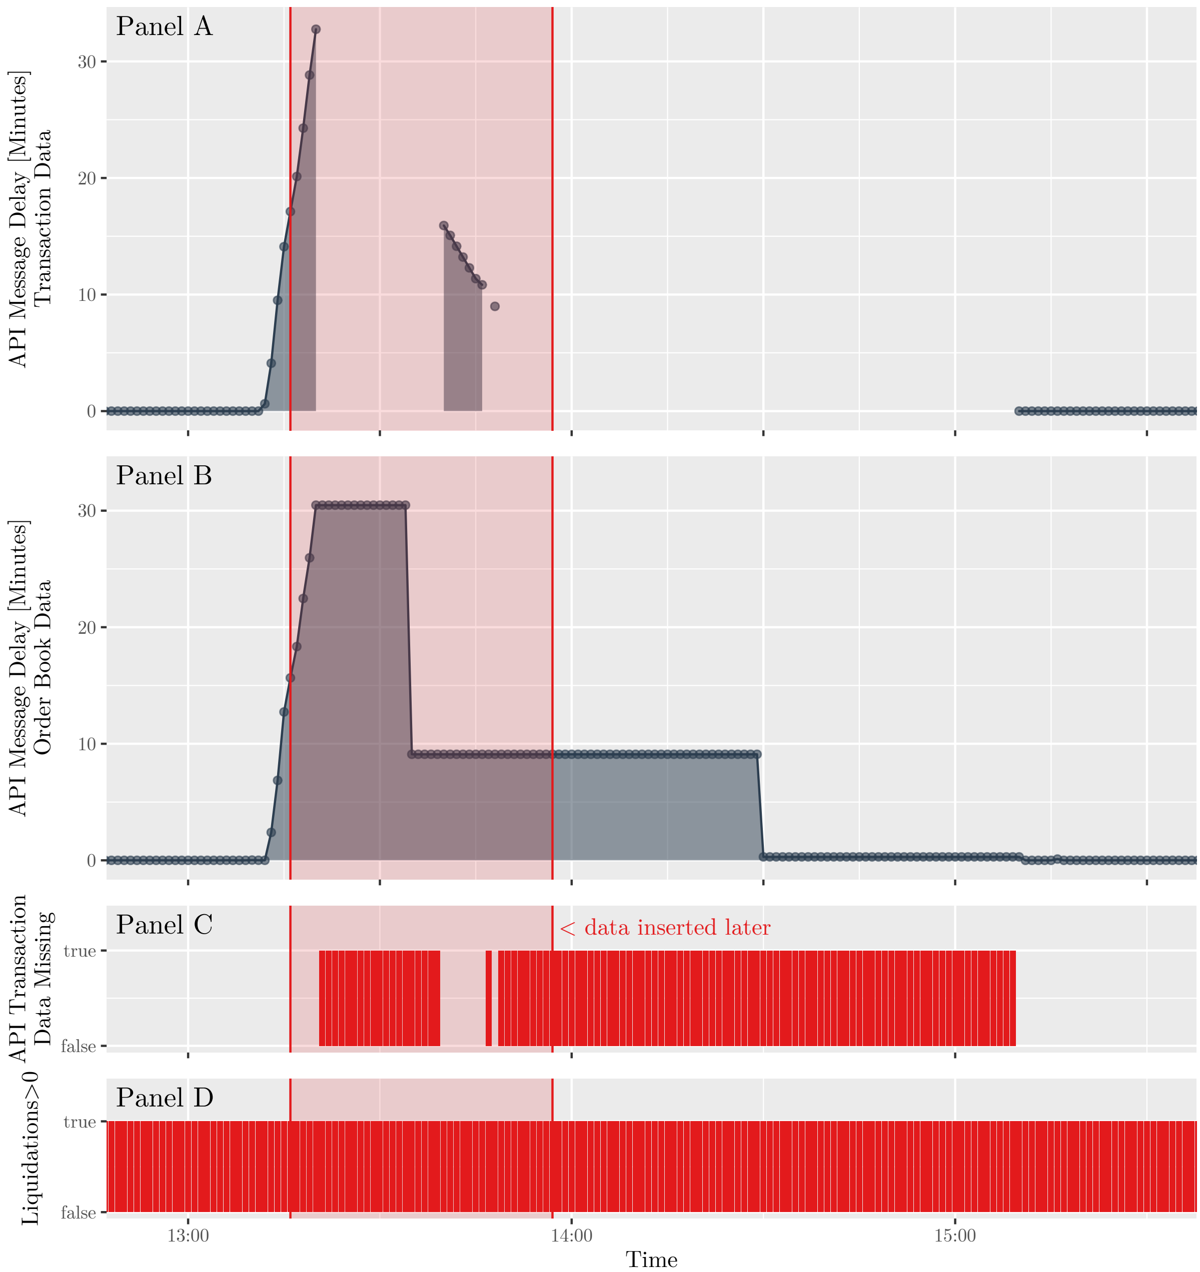

Figure 9: Bitcoin Futures Funding Premium and Mark Price on Binance

This figure compares metrics of Bitcoin Futures on Binance. Panel A shows the difference between the spot price and the mark price which determines liquidations. Panel B shows the futures funding premium. Panel C provides the Bitcoin price for reference. The red box highlights the crash and recovery window from 12:30 to 15:30. Data stems from Binance's API. Figure A6 in the Appendix provides an overview of the funding premium in the ten days before.

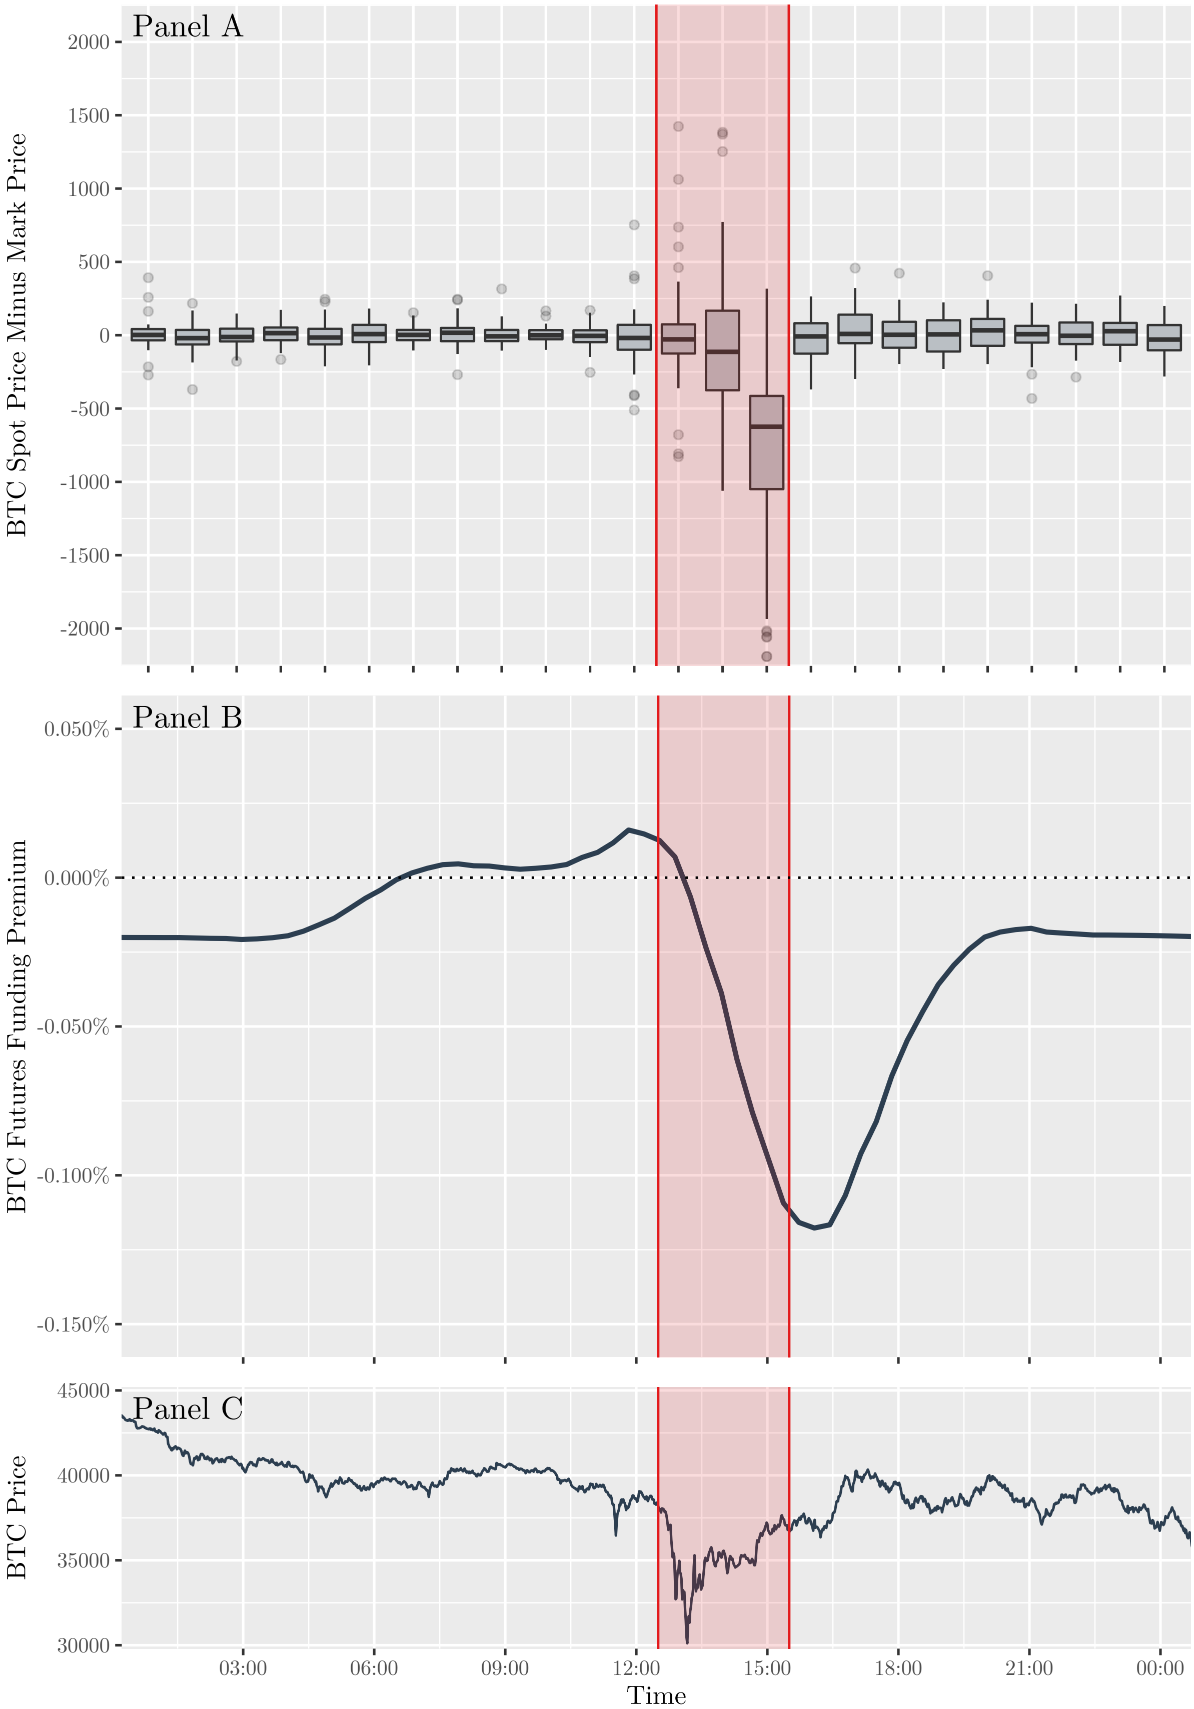

Figure 10: Reported Bitcoin Volume vs. Model Prediction

This figure compares Binance's reported Bitcoin trading volume on May 19, 2021, to the predictions of the model estimated in Table 6 (Regression 3). Panel A shows the volume reported by Binance (red) and the model predictions (blue). Panel B depicts the difference of Binance's reported numbers minus the corresponding prediction. Panel C presents the price difference of one Bitcoin in US-Dollars between the mean of other exchanges minus Binance's price. The red box highlights the back-filled time frame between 13:16 and 13:56 for which Binance initially supplied no data and filled it up later. Table 7 states the corresponding calculations.

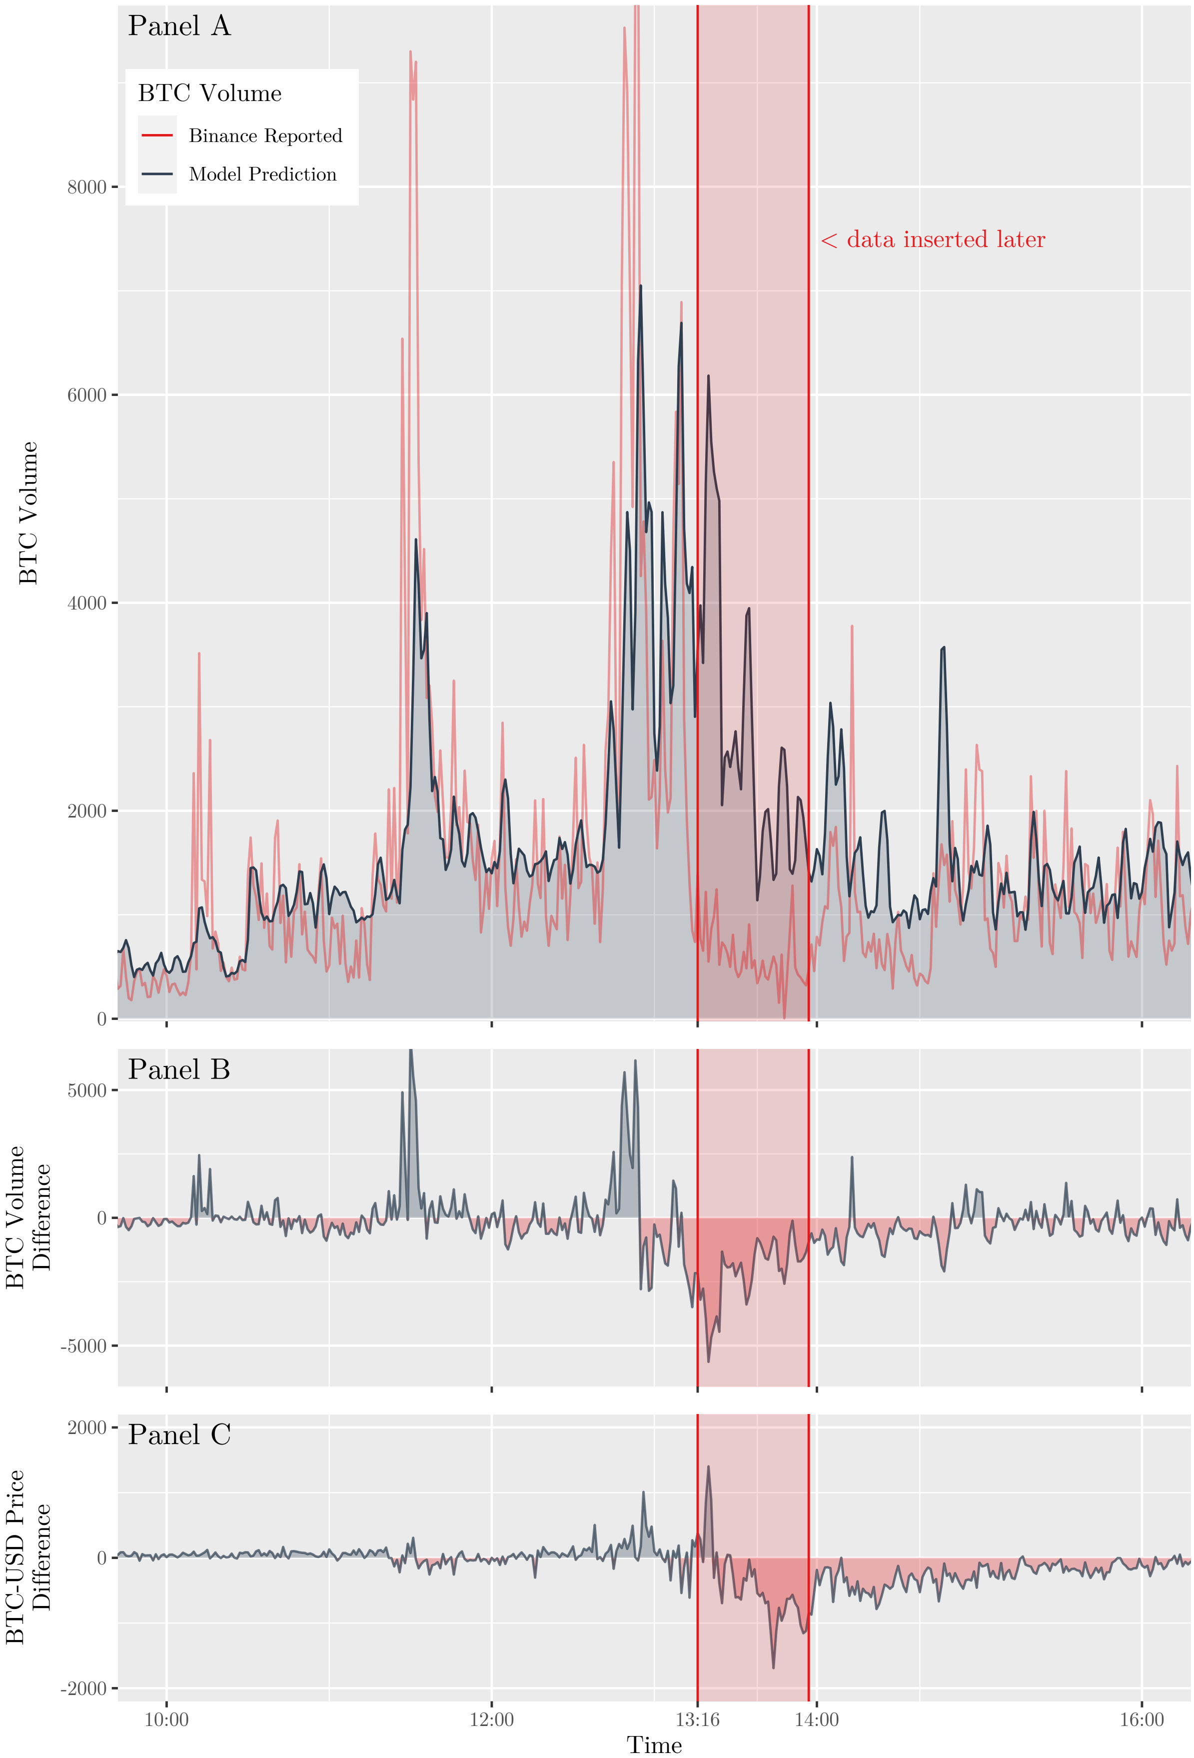

Internet Appendix

Table A1: Lasso Regression of Sentiment associated with Bitcoin Volume

This table reports the results of the lasso regression:

$$\text{volume}_{\text{BTC}}(t) = \alpha + \beta \cdot \text{volatility}_{\text{price}}(\,t-4,\,t\,) + \gamma \cdot \text{price}_{\text{BTC}}(t) + \delta \cdot \text{sentiment}(t) + \epsilon(t)$$

The regression is conducted on a training data set of minutely Binance data in 2021 up to May 18, to select the variables to model the Bitcoin Volume on May 19, 2021 with the regression reported in Table 6. The dependent variable is the traded volume in Bitcoin in each minute. The Bitcoin price reflects the volume-weighted arithmetic mean of transaction prices in each minute. The sentiment is a vector of 279 indicators measuring crowd psychology on Twitter, Reddit, and Bitcointalk, which have been web scraped and aggregated by Augmento. Each indicator represents one category of the underlying linguistic analysis, to which we refer by its generic term in quotation marks. We have chosen a regularization parameter λ of 21. Variables with coefficients of zero are not shown.

	Dependent Variable
	Bitcoin Volume
Constant	12.881
Bitcoin Volatility	4.4436
Bitcoin Price	−0.002
Twitter	
“Bottom”	0.310
“Technical Analysis”	2.743
“Listing”	0.348
“Dip”	0.125
“Price”	1.079
Reddit	
“Bottom”	2.308
“Technical Analysis”	3.264
“Dip”	0.101
“Correcion”	2.253
R^2	0.561
λ	21

Table A2: Tests of a Correlation Break During the Crash

This table tests the correlation of Binance's Bitcoin Price against the other exchanges (Panel A) as well as its mark price (Panel B). We control for the crash window from 12.00 to 14.00 (UTC) on May 19, 2021, ("during the crash") and the time before ("before the crash"). *, **, *** denote significance at the 10%, 5%, and 1% level. Significance refers to the individual coefficients in the first two columns, calculated with the corresponding tests. The third column contains the test statistic and indicates the significance of the difference between the "before the crash" and "during the crash" window based on the transformation (r-to-z approach) of Fisher (1915).

Panel A: Binance Bitcoin Price vs. Other Exchanges

Correlation	Before the Crash	During the Crash	Difference
Pearson	0.9999***	0.9735***	22.59***
Kendall	0.9896***	0.8707***	9.94***
Spearman	0.9998***	0.9744***	19.42***

Panel B: Binance Bitcoin Price vs. Mark Price

Correlation	Before the Crash	During the Crash	Difference
Pearson	0.9999***	0.9820***	22.03***
Kendall	0.9908***	0.8879***	9.80***
Spearman	0.9998***	0.9789***	19.43***

Figure A1: Benford Graphical Analysis of Reference Period Data

Panels (a) to (e) of this figure replicate the analysis of Figure 2 for the reference period (May 9 to May 18). Panel (f) depicts the resulting difference-in-differences (DD) between the back-filled and reference period data of the digit's relative frequency (both versus the expected relative frequencies according to Benford's Law). Panel (g) shows bootstrapped DD boxplots.

(a) Digit Distribution

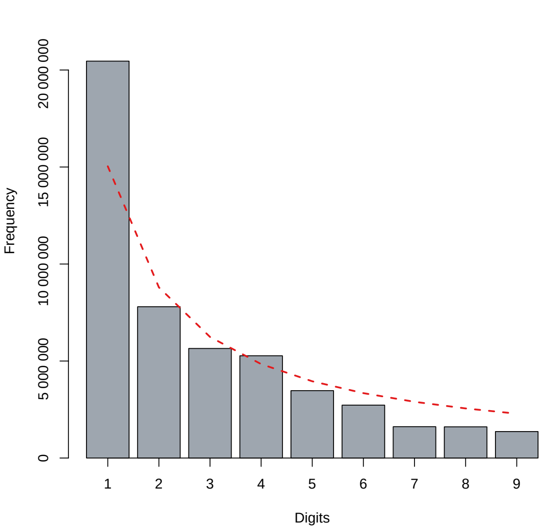

(b) Digit Distribution Second Order Test

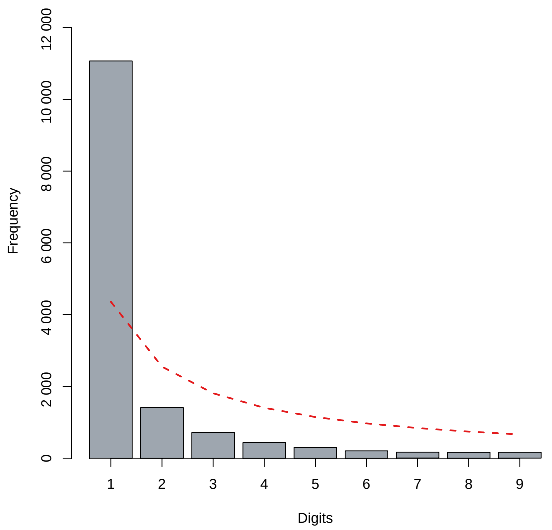

(c) Summation Distribution

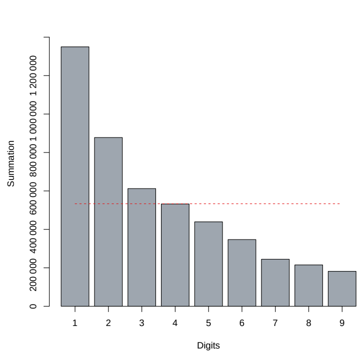

(d) Summation Difference

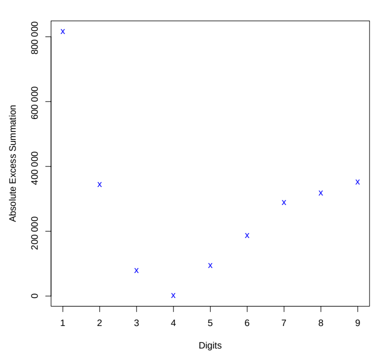

(e) Chi-Squared Difference

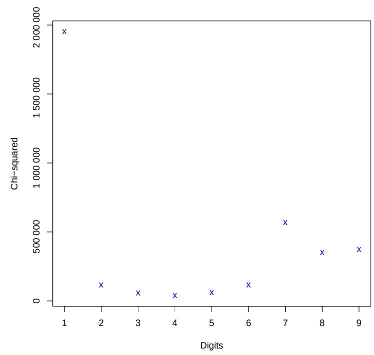

(f) Difference-in-Differences (DD)

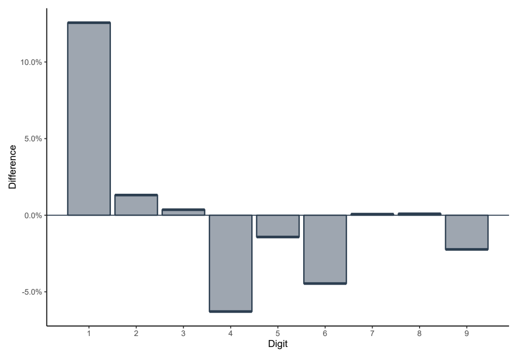

(g) Bootstrapped DD

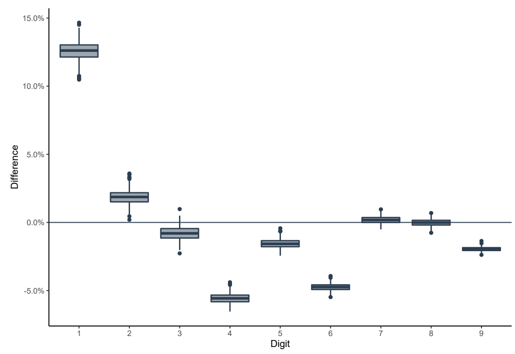

------- Distribution Expected by Benford's Law

Figure A2: Benford Graphical Analysis of Reference Period Data with Time Match

This figure replicates the robustness check of figure A1 but only with the data of the 40 minute time period from 13:16 to 13:56, such that any time-of-the-day effects are ruled out.

(a) Digit Distribution

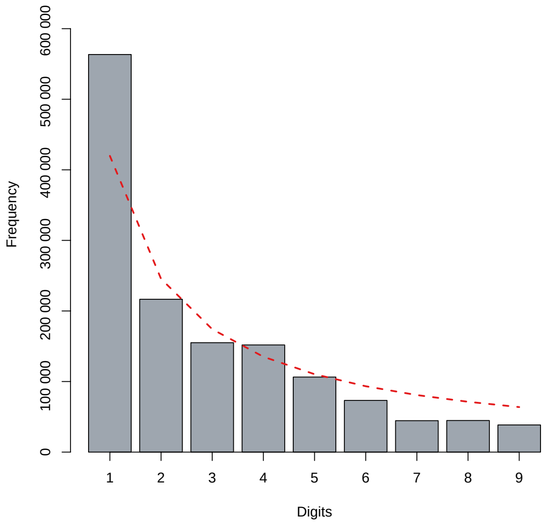

(b) Digit Distribution Second Order Test

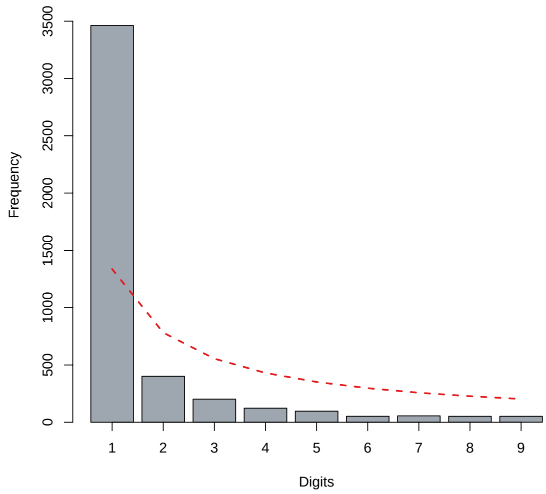

(c) Summation Distribution

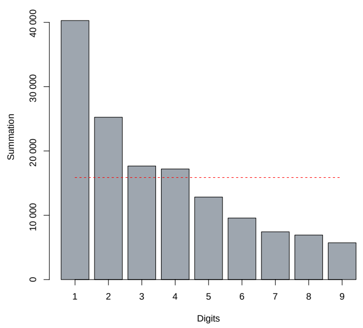

(d) Summation Difference

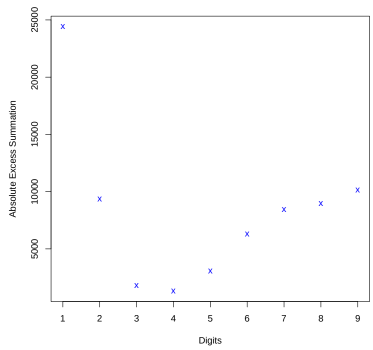

(e) Chi-Squared Difference

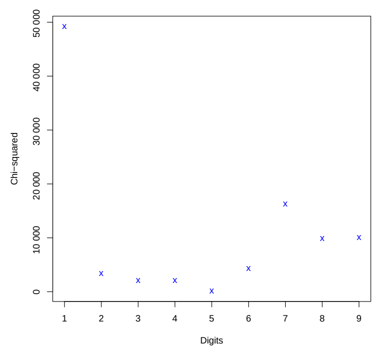

(f) Difference-in-Differences (DD)

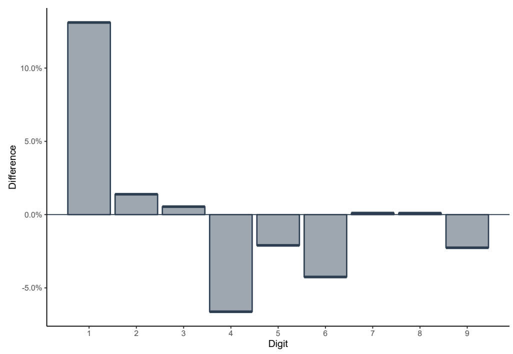

(g) Bootstrapped DD

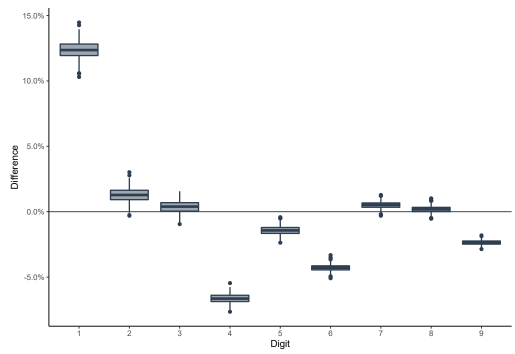

------- Distribution Expected by Benford's Law

Figure A3: Benford Analyses of Individual Reference Days' with Time Match

This figure replicates the analysis of Figure A1 for subsets of the reference data: We evaluate every day individually, and consider only the 40 minute time period from 13:16 to 13:56, to rule of time-of-the-day effects.

(a) Digit Distribution (b) Bootstrapped DD

Panel A: May 18

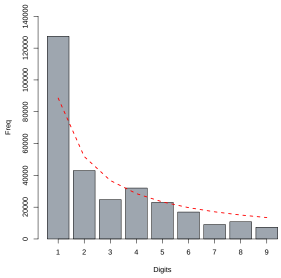

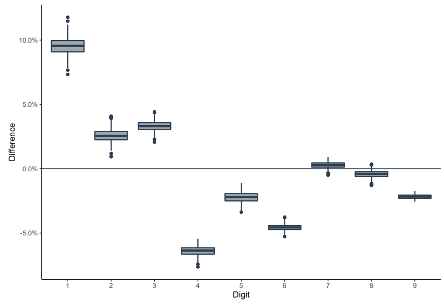

Panel B: May 17

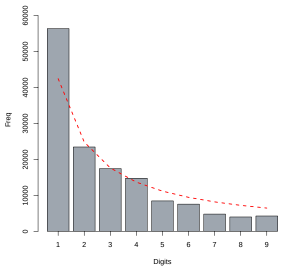

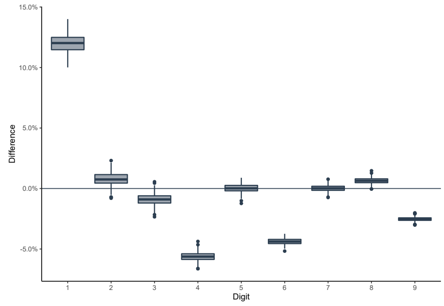

Panel C: May 16

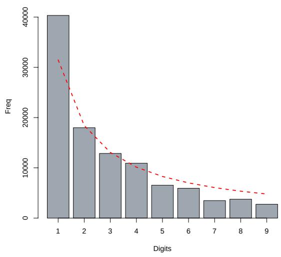

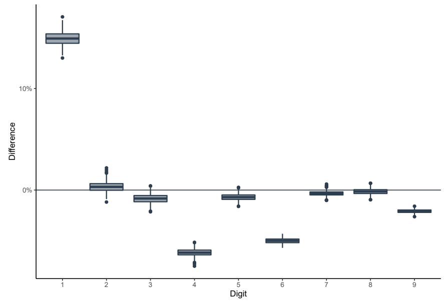

Panel D: May 15

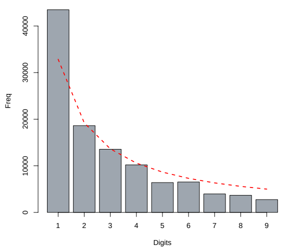

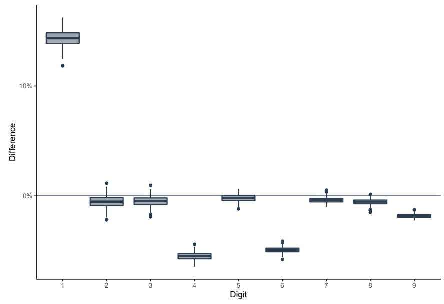

Panel E: May 14

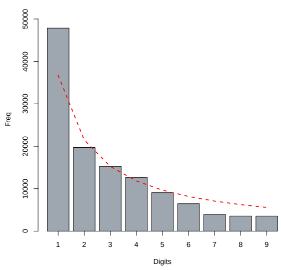

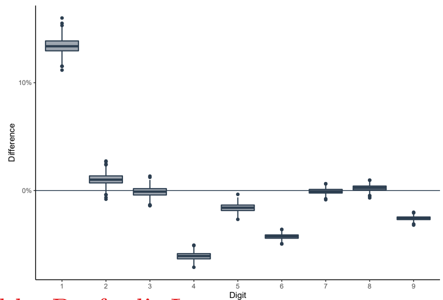

------- Distribution Expected by Benford's Law

Figure A3 continued

(a) Digit Distribution (b) Bootstrapped DD

Panel F: May 13

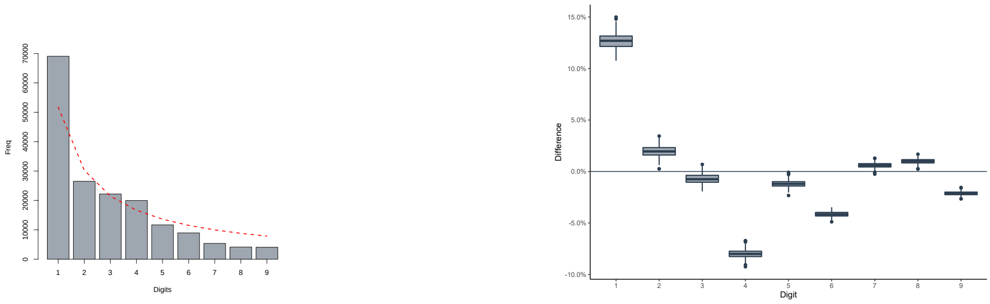

Panel G: May 12

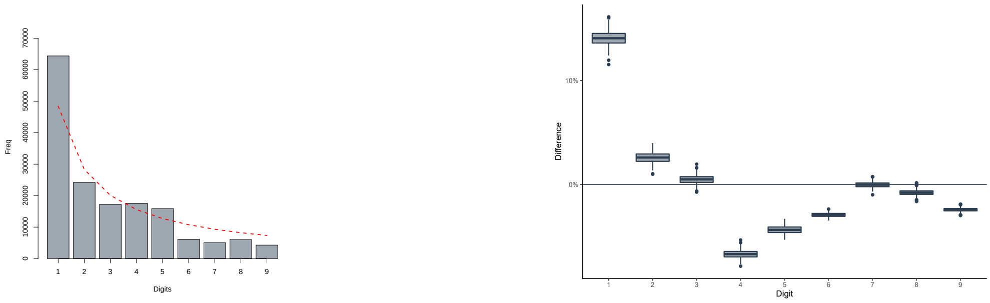

Panel H: May 11

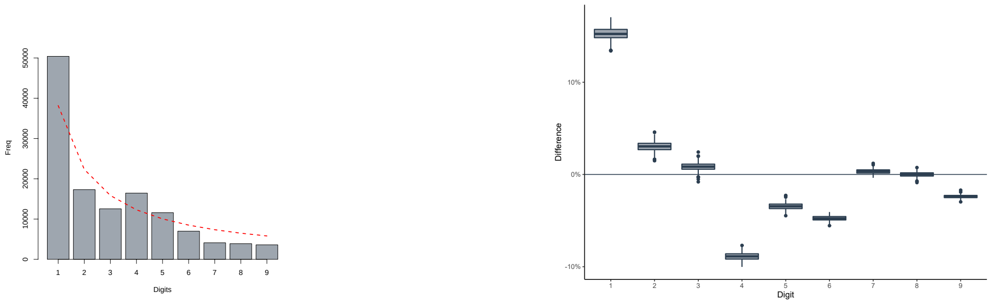

Panel I: May 10

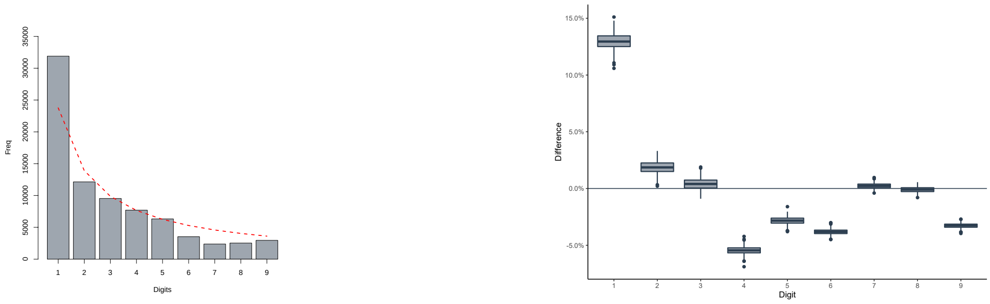

Panel J: May 09

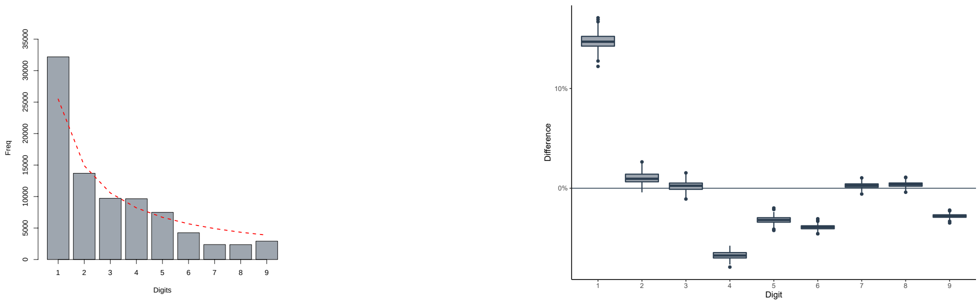

------- Distribution Expected by Benford's Law

Figure A4: April 18, 2021 Overview

This figure compares the course of the April 18, 2021 crash on Binance with other exchanges. The red box highlights the crash time frame between 03:00 and 04:00. Panel A depicts the Bitcoin futures transaction prices. While the lines mark each minute's mean price, the ribbons represent the 95% quantile price range. Panels B and C show the Bitcoin volume traded at Binance and the other exchanges. Panels D and E plot the number of executed trades at Binance and the other exchanges. Panel F presents the Sortino volatility.

Figure A5: Outage Simulation of the April 18, 2021 Crash

Panels (a) to (e) present a Benford analysis for which we have simulated manipulating data by deleting a subset of it for another crash on April 18, 2021. We have deleted 90% of transactions randomly in all minutes where the Bitcoin price of Binance was at least 0.1% lower than on other exchanges. Panel (f) depicts the resulting difference-in-differences (DD) between our manipulated and the original data of the digit's relative frequency. Panel (g) shows bootstrapped DD boxplots.

(a) Digit Distribution

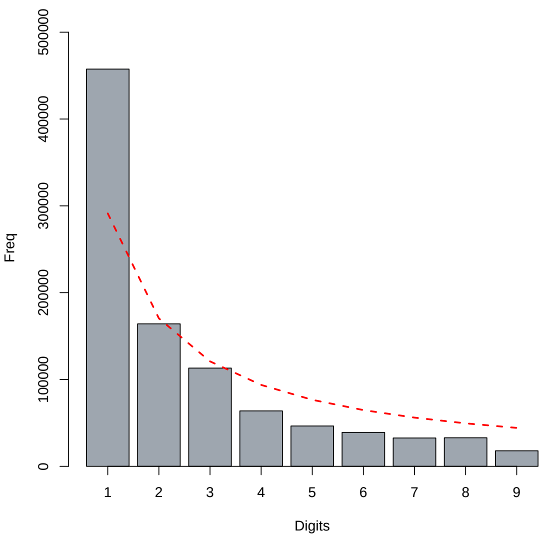

(b) Digit Distribution Second Order Test

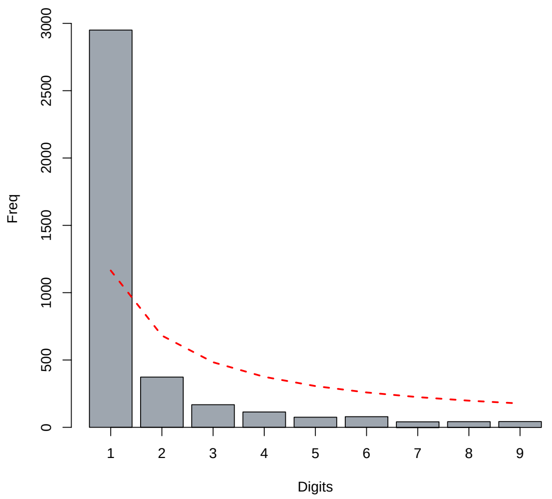

(c) Summation Distribution

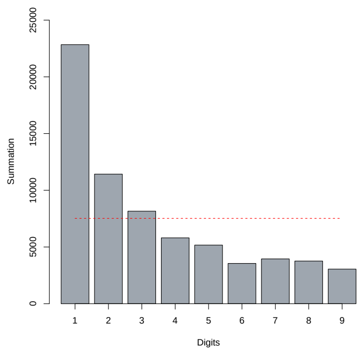

(d) Summation Difference

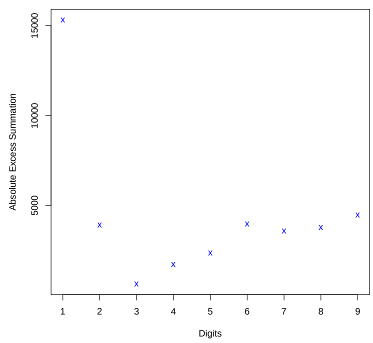

(e) Chi-Squared Difference

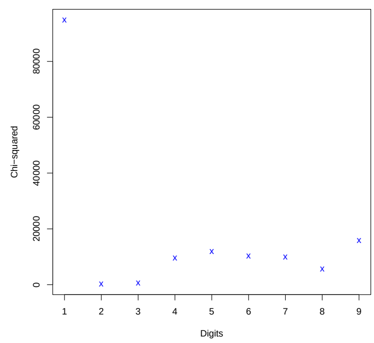

(f) Difference-in-Differences (DD)

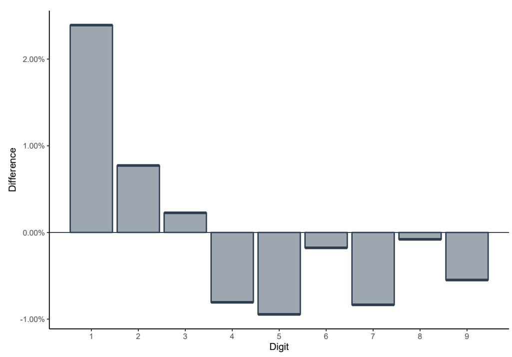

(g) Bootstrapped DD

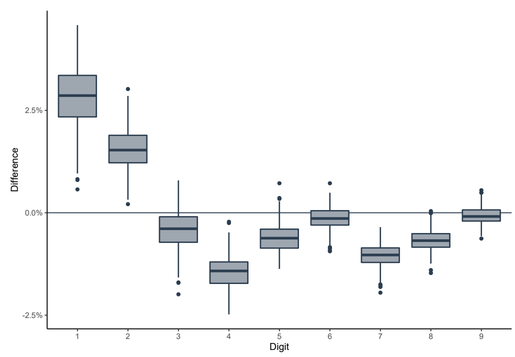

------- Distribution Expected by Benford's Law

Figure A6: Bitcoin Futures Funding Premium on Binance

This figure depicts the funding premium of Bitcoin Futures on Binance. The chart starts ten days before the sharp Bitcoin price decline. The red box highlights the crash day of May 19, 2021. The data stems from Binance's API.

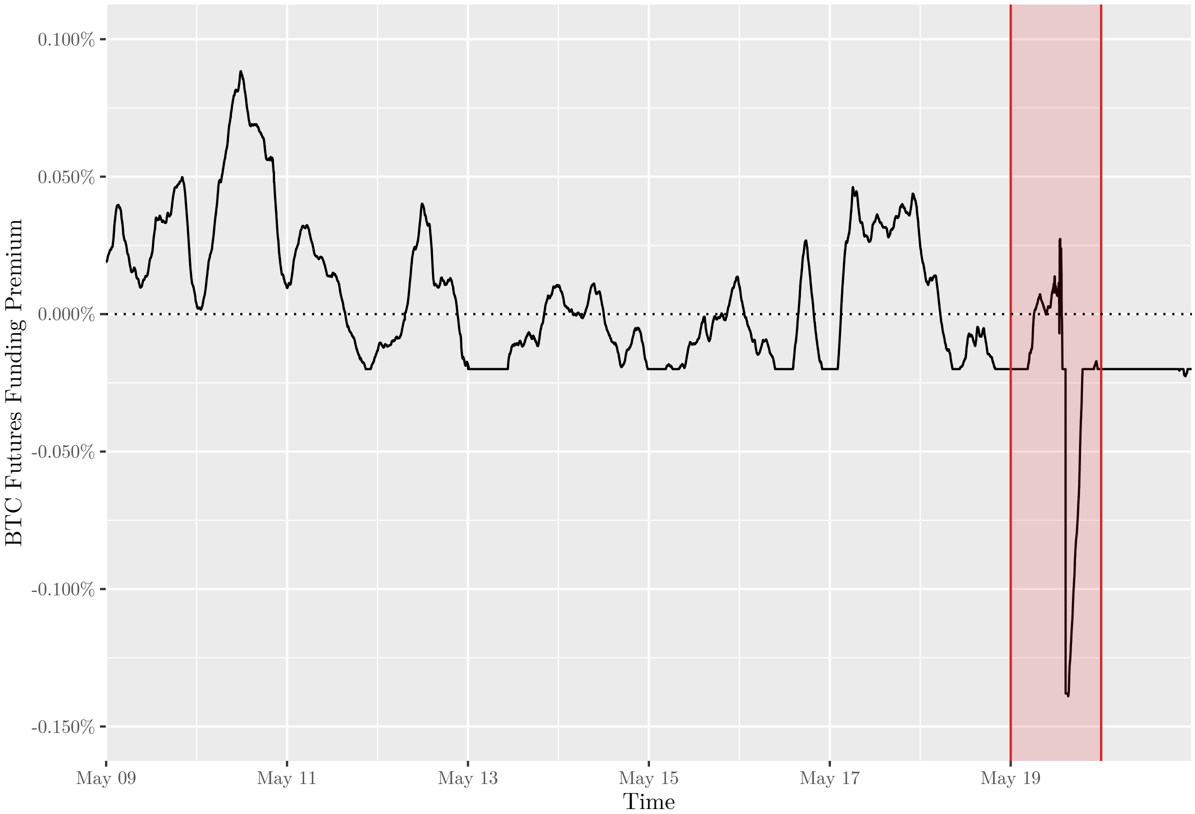

Figure A7: Filled Data Gap

This figure shows screenshots of Bitcoin charts of the crash day May 19, 2021. The screenshots have been taken from TradingView Inc., a data services provider to which Binance directly supplies live data under a cooperation agreement. Panel A depicts the data which Binance sent first, in which a gap is left empty. Panel B presents the newer data version with which Binance overwrote the initial gap. Note that time differences in the screenshots to other paper parts are due to different time zones.

Panel A: Initial Chart with Gap

Panel B: Novel Chart with Filled Gap

Figure A8: Binance Tweet

This screenshot shows Binance's tweet announcing a temporary halt of ETH and ERC20 withdrawals. Source: Twitter.

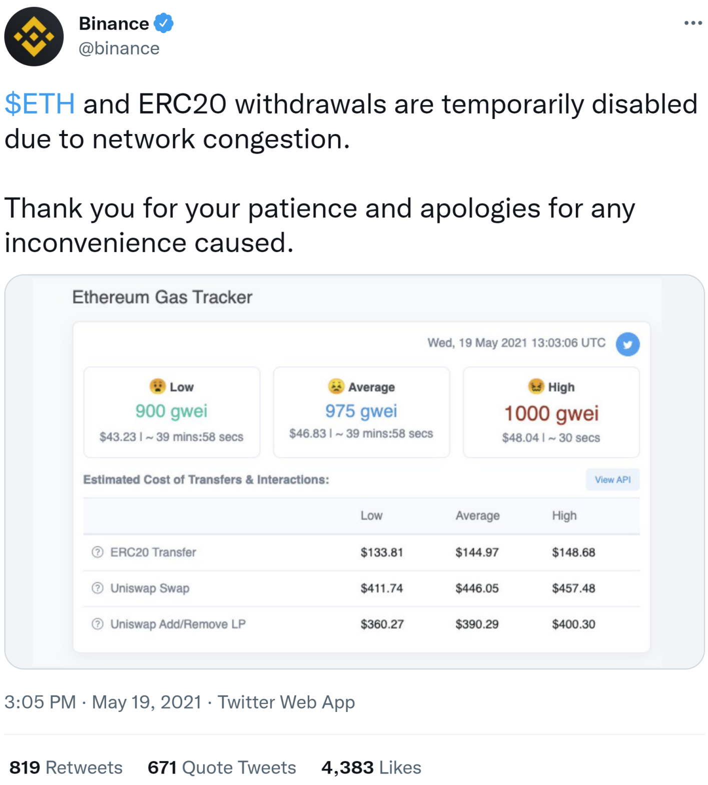
Binance
@binance

$ETH and ERC20 withdrawals are temporarily disabled due to network congestion.

Thank you for your patience and apologies for any inconvenience caused.

Ethereum Gas Tracker

Wed, 19 May 2021 13:03:06 UTC

Low: 900 gwei, $43.23 | ~ 39 mins:58 secs

Average: 975 gwei, $46.83 | ~ 39 mins:58 secs

High: 1000 gwei, $48.04 | ~ 30 secs

Estimated Cost of Transfers & Interactions:

View API

	Low	Average	High
ERC20 Transfer	$133.81	$144.97	$148.68
Uniswap Swap	$411.74	$446.05	$457.48
Uniswap Add/Remove LP	$360.27	$390.29	$400.30

3:05 PM · May 19, 2021 · Twitter Web App

819 Retweets 671 Quote Tweets 4,383 Likes

Figure A9: Binance VIP Client Program

This screenshot shows Binance's overview of its VIP benefits. Source: Binance.

Binance VIP Benefits		VIP 1-3	VIP 4-9
Trading Equity	Exclusive VIP fees (trade, loans)	✓	✓
	Higher withdrawal limits	✓	✓
	Higher API order frequency		✓
	Sub-account / Asset Sub-account / Managed Sub-account	✓	✓
	Higher borrowing rate and trading pair limit for margin trading	✓	✓
	Exclusively with more coins as Multi-Assets	✓	✓
	Higher Futures and BLVT position limit		✓
	Free tier upgrade for VIPs from other exchanges	✓	✓
	Customized services for Spot/Derivative trading		✓
Service Benefits	VIP quotes for Block Trading (OTC)	✓	✓
	Spot/Futures/Fiat MM benefits		✓
	Dedicated customer service fast channel	✓	✓
	Personal 1-to-1 correspondence with key account managers	✓	✓
	VIP Risk management and priority notifications		✓
	Priority support for technical issues		✓
	Priority beta access for new Binance products	✓	✓

Binance VIP Benefits		VIP 1-3	VIP 4-9
Ecosystem Benefits	Exclusive rates for Binance Liquidity Swap	✓	✓
	Exclusive rates for Binance Mining Pool	✓	✓
	Join the Binance x VIP CLUB Group	✓	✓
	Priority for Binance Industry Seminars (online/offline)	✓	✓
	Priority docking for all Binance ecosystem resources.	✓	✓
	The latest industry reports from Binance research institute.	✓	✓
VIP Privileges	Binance VIP private meetups		✓
	Priority registration for offline industry conferences		✓
	Priority sign-up for Global Binance fan meetup.	✓	✓
	Customized Binance Global VIP holiday gifts		✓