UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

ESSAYS IN BLOCKCHAIN ECONOMICS

A DISSERTATION

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

DOCTOR OF PHILOSOPHY

By

PETER CRONIN MUELLER

Norman, Oklahoma

2023

ESSAYS IN BLOCKCHAIN ECONOMICS

A DISSERTATION APPROVED FOR THE
MICHAEL F. PRICE COLLEGE OF BUSINESS

BY THE COMMITTEE CONSISTING OF

Dr. William Megginson, Chair

Dr. Bryan Stanhouse, Co-Chair

Dr. Aaron Burt

Dr. Kose John

Dr. Richard Price

Acknowledgements

I would like to express a deep gratitude towards my Dissertation Committee, consisting of William Megginson, Bryan Stanhouse, Aaron Burt, Kose John, and Richard Price, for excellent support and guidance throughout my PhD program. I am also grateful for helpful comments from Farshid Abdi, Carol Alexander, Ehsan Azarmsa, Carsten Bienz, John Paul Broussard, Chitru Fernando, Abhishek Ganguly, Franz Hinzen, Sara Holland, Takdir Hossain, Felix Irresberger , Engin Iyidogan, Jaeho Kim, Scott Linn, Lubomir Litov, Laura Liu, Summer Liu, Yue (Mark) Ma, Asif Malik, Katya Malinova, William Mann, Nick Pan, Leonid Pugachev, Raghu Rau, Thomas Rivera, Sam Rosen, Ioanid Rosu, Fahad Saleh, Christoph Schiller, Donghwa Shin, Jared Stanfield, Larry Wall, Pradeep Yadav, Baozhong Yang, David Yermack, Huming Zhang, Xun Zhong, Peter Zimmerman and seminar participants at University of Oklahoma, University of Alberta, Fordham University, Georgia State University, Rensselaer Polytechnic Institute, SUNY Buffalo, Texas Tech University, Durham University, McGill University, Sussex University, and University of Leeds, and finally conference participants at the Midwest Finance Association conference, Southwest Finance Association conference, 7th International Young Finance Scholars conference., 5th Shanghai-Edinburgh-London Fintech Conference, FMA Consortium on Asset Management and Fintech, Northern Finance Association conference, FMA conference, 2020 Toronto Fintech conference, Bergen Fintech conference, 2nd Yushan conference, Paris Webinar on Financial Technology and Crypto.

Abstract

This dissertation consists of three essays (chapters) concerning blockchain economics. Chapter 1 is a study of a decentralized lending platform. Decentralized Financial (DeFi) lending applications allow users to take leveraged positions in risky cryptoassets without the use of a financial intermediary, mimicking the functionality of a margin trading account without a centralized broker. This decentralization comes at a cost; leveraged traders must pay blockchain transaction fees to miners/validators and, absent an intermediary, third parties must be incentivized to perform liquidation transactions on positions with insufficient collateral. I provide evidence suggesting that volatile blockchain transaction fees can effectively "price out" a significant portion of market participants engaged in DeFi leveraged trading, forcing these traders to temporarily abandon their leveraged positions. Meanwhile, I show that the market liquidation incentive disciplines leveraged traders with riskier positions, inducing them to take more conservative positions. These results highlight the limitations of blockchain/DeFi technologies in their current state and have interesting policy implications concerning the regulation of leveraged traders.

Chapter 2 paper provides an empirical overview of the largely unexplored public blockchain ecosystem, coauthored with Felix Irresberger, Kose John, and Fahad Saleh. Our overview highlights that only a few blockchains dominate the ecosystem although no single blockchain, not even Bitcoin, dominates uniformly. We explain our empirical findings with a simple theoretical framework that establishes three key economic attributes - adoption, scale, and security - as the determinants of blockchain user utility. We examine each blockchain along those dimensions empirically. We then theoretically establish that whether a blockchain could

be welfare enhancing for some user type can be determined by a comparison across our attributes; in turn, applying the comparison empirically yields that only a few blockchains could be optimal for some user type. Our results thus provide an explanation for why only a few blockchains dominate the broader blockchain ecosystem. Moreover, our work provides an empirical framework from which to evaluate progress within the blockchain ecosystem.

In Chapter 3, I model the entry and exit decisions of cryptocurrency miners. I argue that miners using Application Specific Integrated Circuit (ASIC) equipment, with a lower salvage value, have asymmetric reactions to price shocks, while miners using Graphics Processing Unit (GPU) equipment, with a higher salvage value, have more symmetric reactions to price shocks. Estimating the long-run equilibrium relationship between cryptocurrency price and hashrate (the aggregate computing power of miners), I show that Bitcoin miners, who use ASIC equipment, respond only to negative disequilibria (when hashrate is relatively low compared to price). Meanwhile, Ether miners, when using GPU equipment, respond symmetrically to disequilibria. These results have both practical and theoretical implications for cryptocurrency pricing and applications of blockchain technology.

Table of Contents

List of Figures and Tables

Chapter 1: DeFi Leveraged Trading: Inequitable Costs of Decentralization

1. Introduction

The future of DeFi (decentralized finance) and its impact on our traditional financial systems is a topic of great interest. Following the meteoric rise of blockchains such as Bitcoin, other blockchains have emerged offering additional functionality beyond that of a payment system. DeFi technologies existing on these blockchains utilize secure distributed ledgers to offer services mimicking those of traditional financial institutions (i.e., exchanges and brokers). The novelty of these services is they are executed without the use of a centralized intermediary or authority, instead facilitating "permissionless" and "censorship resistant" exchanges between market participants (Harvey, Ramachandran, and Santoro, 2021).

The dominance of DeFi applications is apparent in the growth of the blockchain ecosystem. Irresberger, John, Mueller, and Saleh (2021) note that DeFi applications attract more users and value than other types of decentralized blockchain applications. Much recent work concerns Automated Market Makers (AMM), which exist as intermediary-less exchanges with algorithmic settlement (Lehar and Parlour, 2021; Capponi and Jia, 2021; Angeris, Kao, Chiang, Noyes, and Chitra, 2021; Park, 2022). However, DeFi lending applications remain largely unexplored. These lending applications may be seen as the second half of the DeFi ecosystem, with available liquidity of $15.91 billion, only trailing decentralized exchange applications in this metric ($23.77 billion).[1] DeFi lending protocols allow users to take leveraged positions in risky cryptoassets without the use of a financial intermediary, mimicking the functionality of a broker margin account.

[1] Source: https://defillama.com/categories

This paper presents an empirical study of Compound, one of the leading DeFi lending applications on the Ethereum blockchain. Studies of the behavior of leveraged traders in traditional (broker) markets are hindered by the availability of suitably granular data, as the trades and positions of individual retail investors are not publicly available. In contrast, the Ethereum blockchain is a public ledger containing all transactions and past states of individual accounts (traders) utilizing Compound. This novel dataset allows me to isolate the effects of market frictions on the behavior of heterogeneous leveraged traders. I focus on two frictions faced by leveraged traders that are unique to the DeFi lending space: blockchain transaction fees and Compound's liquidation incentive. I find that blockchain transaction fees effectively "price out" a good portion of leveraged traders, while Compound's liquidation incentive induces highly leveraged traders to take less risky positions.

Blockchain transaction fees can be viewed as a fundamental cost of decentralization. For a transaction to be considered valid/final, it must first be recorded by miners/validators on the blockchain. Under proof-of-work consensus protocol blockchains, such as Bitcoin and Ethereum, miners compete to solve a computationally challenging problem. The winner of this competition earns the right to propose the next block of transactions on the blockchain. However, these blocks have a limited capacity for transactions, so it is the policy of miners to record them in descending order of transaction fees. If there is a surplus of transactions, those with lower fees are not included in the new block of transactions. Thus, users compete to have their transactions recorded first by attaching larger fees to their transactions.

For DeFi leveraged traders, these fees may be characterized as volatile fixed costs. They are fixed costs in the sense that a very low-value transaction is subject to the same fee as a very high-value transaction at a given time. However, over any time series, the average transaction fee varies greatly with the demand for blockchain transactions. Though fixed costs are not

unique to the DeFi leveraged trading market, the realized volatility of transaction fees distinguishes this market friction from any comparable cost in a traditional market for leverage. For example, the daily estimated average cost to fully reverse a leveraged position on Compound ranges between $0.41 and $72.99 in my sample.

Blockchain transaction fees also necessitate a liquidation incentive to be captured by the liquidators. As borrowers have no single counterparty, market mechanisms must incentivize third parties to perform liquidation transactions on accounts that have insufficient collateral. These liquidators must be compensated for bearing both the search cost to find undercollateralized accounts and the transaction fee they pay when executing the liquidation transaction. In a traditional leverage market, brokers bear this cost and have obvious incentives to do so (protecting their exposure to borrower default). In the DeFi leverage market, this cost is shifted to borrowers via the liquidation incentive (or liquidation penalty), which is currently set at 8% of the liquidated amount. In the absence of transaction fees, the liquidation incentive could be arbitrarily small and liquidations would still occur. Given large transactions fees, these penalties must also be large to incentivize liquidators.

I begin my analysis by testing the propensity of leveraged traders to rebalance their positions in the face of high and volatile transaction fees. Utilizing fixed-effect linear probability and logistic specifications, I show that traders with higher-value leveraged positions manage their positions more actively. A one standard deviation increase in the size of a trader's position is associated with an increased daily propensity to rebalance of 1.3% (average marginal effect). This is consistent with the notion that traders treat transaction fees as fixed costs, which are a greater burden on smaller operations. Meanwhile, a one standard deviation increase in transaction fees is associated with a decreased daily propensity to rebalance of 6.0%.

However, this pooled specification does not do justice to the effects of transaction fees in relation to the size of position held by traders. All else equal, the effect of increased transaction fees should be large for a lower-value trader and negligible for a higher-value trader. To capture this differential effect, I separate accounts into deciles for each day based on the total value of their positions, then pool these deciles to form 10 unique subsamples. Repeating the analysis on these subsamples, I show that, as expected, higher value accounts are less sensitive to blockchain transaction fees. From decile 4 to decile 10, the magnitude of the fee coefficient decreases monotonically, becoming entirely insignificant in the highest value decile. However, trader sensitivity to fees *increases* from decile 2 to decile 4 subsamples. This result is puzzling, as fixed costs should be a greater burden on smaller operations.

I argue that this is indicative of non-linear sensitivity to fees based on the relative value of accounts. A general lack of response from lower-value traders to changes in transaction fees is plausible if contemporaneous fee levels entirely price these traders out of the market. Following this logic, a decrease in transaction fees will illicit no response from traders so long as the transaction fee outweighs the transaction benefit. Measuring transactional benefits is subjective for most types of transactions, but the value of reversing a leveraged position is clear. To reverse their position, the trader simply repays the borrowed asset(s) and then claims the collateral asset(s). All loans in the market are overcollateralized, so the benefit of this transaction to the trader is simply the positive difference in these values. Thus, the *Net Reversal Benefit* would be this difference minus the expected transaction fee.

If the transaction benefit when reversing the leveraged position is less than the blockchain transaction fee, *Net Reversal Benefit* will be negative. Such traders are effectively priced out of the market until *Net Reversal Benefit* becomes positive. For example, consider a trader that collateralizes $125 of ETH (a risky cryptoasset) and borrows $100 of USDC (a

stablecoin pegged to the USD). Reversing this leveraged position would lead to a gain in liquidity of $25 before the transaction fee. If the transaction cost of such a reversal is $40, then the trader would never perform the transaction (their *Net Reversal Benefit* would be -$15). Even if the transaction fee fell to $30, this trader would still be priced out of the market (with a *Net Reversal Benefit* of -$5). For expositional purposes, I consider such accounts to be temporarily abandoned.

Consistent with this explanation, I show that a decreasing proportion of observations in the decile 1 to decile 6 subsamples have a negative *Net Reversal Benefit*. By this measure, a significant portion of market participants on Compound are frequently priced out of the market due to high blockchain transaction fees. The steady rise in Ethereum transaction fees observed towards the end of 2020 led to about 30% of leveraged traders falling into the category of "temporary abandonment." Moreover, following a particularly large jump in transaction fees, more than half of traders experienced a negative *Net Reversal Benefit*.

These findings suggest that blockchain transaction fees can be entirely ignored by the wealthiest participants, and yet at the same time they may price out nearly half of the other market participants. While fixed costs are not a new concept in finance, the extreme volatility in blockchain transaction costs are uniquely damaging to small traders. On the surface, this might simply be priced as an additional market risk. However, it also seems unlikely that half of all leveraged traders expected to have their positions priced out of the market. While the absence of a monopolistic intermediary provides the benefits of inclusivity and lack of censorship to users (Harvey et al., 2021), this decentralization comes at a cost: blockchain transaction fees. My findings suggest that, while not explicitly excluded from transacting, a good portion of leveraged traders are effectively excluded by high and volatile transaction fees.

This highlights the limitations of blockchain technology in its current state, complementing other recent research (Hinzen, John, and Saleh, 2022).

Blockchain transaction fees also necessitate a liquidation incentive such that third parties are induced to perform liquidation transactions on positions with insufficient collateral. This incentive, captured by liquidators from leveraged traders, is a cost unique to the DeFi lending market. To identify the effect of the liquidation incentive on borrower behavior, I exploit a structural change in Compound's liquidation incentive. March 12th, 2020, is often called cryptocurrency's "Black Thursday." Cryptocurrency markets appeared to suddenly collapse, with Bitcoin prices falling by almost 50% before recovering slightly. In response to this realization of systemic risk, the liquidation incentive on Compound was increased from 5% to 8% of the liquidated amount. While this change affected all leveraged traders, it disproportionately impacted traders with riskier positions. Traders with leveraged positions close to the liquidation threshold (i.e., lower margin) were more exposed to this policy change than traders with safer positions, who were unlikely to face a liquidation regardless of the change in liquidation incentive. Utilizing this heterogeneous exposure to the change, I employ a difference-in-differences empirical strategy to isolate the causal effect of the liquidation incentive on borrower behavior.

The difference-in-differences analysis confirms that leveraged traders with riskier positions are disproportionately affected by the increased liquidation incentive. Following the change, traders that utilized more than 50% of their maximum borrowing power (riskier positions) significantly reduced their leverage, taking more conservative positions after the event. On average, such risky traders reduced their borrowing utilization by roughly 8% (a relative reduction of 12%). As the liquidation incentive is a proportional cost to borrowers, this effect is stronger amongst borrowers with larger positions. Confining the sample to traders with

positions valued at over $10,000, risky traders reduced their borrowing utilization by 11% (a relative reduction of 16%) following the change in liquidation incentive.

The behavior of leveraged traders in the face of the liquidation incentive has interesting policy implications. Empirical and theoretical evidence suggests that leveraged traders are a source of market volatility (Hardouvelis, 1990; Hardouvelis and Theodossiou, 2002; Chabakauri, 2013; Rytchkov, 2014). However, the effects of regulatory constraints (such as margin requirements) on systemic risk are unclear.[2] The liquidation incentive in the DeFi lending space induces only leveraged traders with riskier positions to take less leverage, while traders with safer positions are unaffected. A policy such as a liquidation incentive is appealing in this regard as it only constrains highly leveraged traders who are likely the primary contributors to systemic risk, while more conservative traders are not affected.

My findings contribute to a growing literature on blockchain and decentralized finance. Within the DeFi space a number of recent papers study AMMs, primarily focusing on blockchain-related market mechanisms and their efficacy compared to that of centralized exchanges (Lehar and Parlour, 2021; Capponi and Jia, 2021; Barbon and Ranaldo, 2021; Angeris and Chitra, 2020; Angeris et al., 2021; Park, 2022; Aoyagi, 2020). Lehar and Parlour (2022) study DeFi lending markets, documenting the price impact of liquidation transactions occurring on the Compound and Aave platforms. My analysis differs considerably from that of Lehar and Parlour (2022), as I consider the effects of DeFi related market frictions on the behavior of leveraged traders.

[2] The literature on the effect of margin requirements on market volatility provides mixed and often conflicting findings. Older research, such as that of Hsieh and Miller (1990), suggests that lower margin requirements do not lead to increased market volatility, contrasting with the findings of Hardouvelis (1990). Hardouvelis and Theodossiou (2002) reconcile these findings, suggesting that increased margin requirements reduce volatility in normal and bull markets, but have no effect in bear markets. More recent theory suggests that more leverage and lower margin requirements increase market volatility (Chabakauri, 2013; Rytchkov, 2014).

Closely related are papers examining blockchain transaction fees and their effects (both realized and potential) on blockchain markets (Hinzen et al., 2022; Easley, O'Hara, and Basu, 2019; Huberman, Leshno, and Moallemi, 2021; Chiu and Koeppl, 2019; Lehar and Parlour, 2020; Makarov and Schoar, 2020; Pagnotta, 2022). The observed rise in blockchain transaction fees on traditional proof-of-work platforms (i.e., Bitcoin and Ethereum) has likely contributed to the growing popularity of other blockchains with alternative consensus protocols that scale more efficiently (Irresberger et al., 2021). Other consensus protocols, such as proof-of-stake, are studied by Saleh (2021) and John, Rivera, and Saleh (2022). A natural extension of this research would be a study of DeFi lending platforms on blockchains that have a higher transaction capacity than Ethereum (i.e., Solana).

Given that Compound's liquidation incentive effectively disciplines risky leveraged traders, my work also ties into policy related research which examines the effects of leveraged trading on systemic risk. Hsieh and Miller (1990), Hardouvelis (1990), and Hardouvelis and Theodossiou (2002) empirically test the effect of initial margin requirements on market volatility. There is a general consensus that leveraged traders do contribute to the overall systemic risk of the market, though there is some debate as to the effectiveness of regulating margin requirements. Chabakauri (2013) finds a positive relationship between leverage and stock return correlations and volatilities. Similarly, other research finds that tighter leverage constraints lead to a flatter security market line and lower betas (Jylh¨a, 2018; Frazzini and Pedersen, 2014).

The remaining chapter is organized as follows. Section 2 presents relevant institutional details about Compound. Section 3 describes the data and presents summary statistics. Section 4 describes the empirical methodology and results pertaining to the sensitivity of leveraged

traders to blockchain transaction fees. Section 5 presents the difference-in-differences methodology used to examine the effect of Compound's liquidation incentive, along with the results, robustness tests, and alternate specifications. Section 6 contains my concluding remarks.

2. Institutional Details

The term DeFi "lending" application is somewhat misleading here, as borrowers in this market do not gain liquidity. Instead, all loans are overcollateralized; borrowers use a larger value of one asset to borrow a smaller value of another asset. Traders may then achieve a leveraged position by swapping (via AMM or other outside exchange) the borrowed asset for more of the collateral asset. This results in a leveraged long position in the collateral asset and a leveraged short position in the borrowed asset.

2A: Collateral Factor and Leveraged Position

The Compound lending platform is a market of liquidity pools. Users lend to the pools either to earn interest on their lent assets and/or to collateralize the lent assets to gain the right to borrow. Borrowers are thus also lenders in the market, contributing to the overall liquidity of the pool. Lenders who collateralize their assets are allowed to borrow an amount up to the value of their collateral asset times a certain percentage. This percentage, called the *collateral factor*, governs the amount of borrowing power, and thus leverage, the trader may achieve. The *collateral factor* ranges between 0% and 85% across different assets traded on Compound.

To further understand the market mechanisms and leveraged trading, consider the following example. A trader collateralizes \$125 of asset *X* which has a *collateral factor* of 0.8. This allows him to borrow 100 ($\$125 \times 0.8$) of asset *Y* and then proceed to swap asset *Y* for more of asset *X* utilizing an outside exchange (not on the Compound platform). Thus, the trader has \$225 of long exposure to asset *X* and \$100 of short exposure to asset *Y* after a single

iteration of this process. It is important to note that the trader could repeat this process infinitely (collateralizing the \$100 of asset *X* and borrowing another \$80 worth of asset *Y*, swapping this for \$80 of asset *X*, collateralizing the \$80 of asset *X*, ad infinitum) and achieve a theoretical maximum long exposure to asset *X* of $\frac{\$125}{(1-0.8)} = \$625$ and short exposure to asset *Y* of $\frac{\$100}{(1-0.8)} = \$500$. Figure 1 describes how a leveraged trader might take a long (short) position in wrapped Ether, WETH, by borrowing (supplying) a stablecoin, USDC. While technically each position has both a long and short component, the position in USDC should have no variance or expected return assuming stablecoins are actually stable.[3]

2B: Account Health and Liquidation

In the above example, the trader utilizes 100% of their borrowing power. The *health* of this account would be 1. Account *health* is calculated as the borrowing power of the account divided by the account's liabilities (what the trader borrows). Effectively, the account *health* statistic is a minimum maintenance margin normalized to 1. In practice, it is unwise to maintain a position with a *health* of 1, as cryptoasset prices are notoriously volatile. If the value of the collateral asset decreases or the value of the borrowed asset increases, the account *health* will fall below 1 and the trader will immediately be subject to a liquidation transaction. In a traditional market for leverage, brokers execute liquidation transactions when traders fail to meet the minimum margin requirements. Brokers are uniquely situated to perform the liquidation (typically reversing the leveraged trader's position) and have obvious incentives to do so: to mitigate their exposure to counterparty default risk. With no such counterparty in the decentralized lending space, market mechanisms must incentivize third parties to perform this liquidation transaction. This incentive is called the *liquidation incentive.*

[3] The actual "stability" of stablecoins is called into question by several recent papers (Griffin and Shams, 2020; d'Avernas, Maurin, and Vandeweyer, 2022; Li and Mayer, 2022).

Third party liquidators may execute liquidation transactions for any accounts with a *health* value less than 1 (accounts with liabilities exceeding their borrowing limit). For such accounts, the liquidator repays 50% of the liabilities, then claims that value from the account's collateral plus an 8% *liquidation incentive*. As loans are overcollateralized, this transaction increases the account's *health*. The *liquidation incentive* market mechanism is unique to the DeFi leveraged trading environment.

Again, consider the above example in which a trader collateralizes \$125 of asset *X* (with a *collateral factor* of 0.8) and borrows \$100 of asset *Y*. Now suppose the value of asset *X* decreases by 4% while the value of asset *Y* remains constant. The health of the trader's account is $\frac{\$120 \times 0.8}{\$100} = 0.96$ following the change in price, thus the account is subject to a liquidation transaction. A liquidator would then repay 50% of the account's liabilities, \$50 of asset *Y*, and claim that amount plus the 8% *liquidation incentive* from the account's collateral, asset *X*, \$54. Following the liquidation transaction, the trader's account has \$66 worth of asset *X* as collateral, and \$50 worth of asset *Y* as a liability; the liquidation transaction mechanically increases the account health to $\frac{\$66 \times 0.8}{\$50} = 1.056$. Figure 2 illustrates this example. The liquidation incentive in this case is worth \$4 to the liquidator. It is important to note here that the liquidator bears the blockchain transaction fee for initiating the liquidation transaction and would rationally ignore such a liquidation opportunity if the transaction fee is greater than \$4.

2C: Interest Rates and C-Tokens

Apart from blockchain transaction fees and (potential) *liquidation incentives*, leveraged traders also face an interest rate cost. When lending, lenders receive *c-tokens* (or compound-token) corresponding to the underlying lent asset. For example, lending USDC results in the lender receiving cUSDC (compound USDC). The exchange rate between the *c-tokens* and underlying asset increases over time at the pace of the interest rate for lending the underlying

asset. Thus, when lenders reclaim their lent assets by exchanging their *c-tokens* they also claim the accrued interest.

The pace at which a *c-tokens*'s exchange rate increases (the interest rate) is determined using a formula with a single input variable: the borrowing *utilization* of the underlying asset. *Utilization* is simply the percent of the total supplied asset that is borrowed by traders. The interest rate formulas are such that higher *utilizations* yield higher interest rates. Each asset has a unique interest rate formula, and most are not monotonic. With the exception of a few assets (i.e., ETH, BAT, and ZRX), interest rate formulas have a "kink" at 80% *utilization*, after which the slope of the interest rate with respect to the *utilization* of the asset increases substantially. For example, the *supply interest rate* for USDC at 60%, 70%, 80%, 90%, and 100% utilizations are 2.10%, 2.87%, 3.77%, 16.77%, and 34.74% APY, respectively.

Borrowers pay a *borrow interest rate* that is slightly higher than the *supply interest rate*. The difference in the borrow and supply interest rate is added to the liquidity pool as permanent liquidity, offsetting the risk that liquidators will ignore certain low-value accounts with health less than 1 in times of high transaction fees. The difference between the borrow and supply interest rates is called the reserve factor, which ranges between 7% and 25% across assets traded on Compound. As borrowers are also lenders given their collateral asset(s), the cost to maintain their leveraged position is simply their realized borrow interest rate minus their realized supply interest rate.

2D: Transaction fees on Ethereum

Compound operates on the Ethereum blockchain, which has a unique method of calculating transaction fees using units of *gas*. *Gas* is a metric of how much "work" is required to process a transaction. A simple transfer of Ether from one account to another will use less

gas than a more complex transaction, which might invoke smart contract code (such as transactions using the Compound application). Whereas a simple transfer uses roughly 21,000 *gas*, different types of transactions on Compound could use an order of magnitude more than that. Table 1 presents the estimated gas usage for different types of transaction on Compound.

Transaction fees are paid in Ether (ETH), the underlying cryptocurrency on the Ethereum blockchain. The price in Ether per single unit of *gas* is called the *gas price*, which is typically measured in GWEI. GWEI is the smallest unit of Ether, one billionth of one ETH. Users may select their own *gas price* when making transactions, however they must do so judiciously as each block has a *gas* limit governing the maximum number/aggregate-complexity of transactions recorded within it. Miners/validators responsible for recording transactions on the blockchain therefore record transactions in descending order of *gas price*. This means that in times of high transaction demand users that post a lower *gas price* will not have their transaction recorded in the current block. Therefore, users compete to have their transactions recorded first by offering higher transaction fees. This competition has led to times of very high transaction fees.

3. Data and Summary Statistics

3A: Data Sources and Sample

I collect account-level data for the population of leveraged traders utilizing the Compound application on the Ethereum blockchain. This is gathered using the Compound Finance public API.[4] Utilizing this API, I determine the "closing block" for each day (the last block recorded before 0:00 UTC) and collect the historical account-state variables associated with this block. For each individual account (trader), I record the borrowing and collateral

[4] https://docs.compound.finance/v2/api/

positions in 18 different assets traded on the Compound platform along with the overall *health* factor of the account.

I supplement this account-level data with market-level data, also from the Compound Finance public API. Similarly, I collect daily asset-pool-level data for the 18 traded assets, including the borrow/supply interest rates, liquidity supplied and borrowed, reserve values, asset to Ether exchange rate, and collateral factors. These are used to calculate important intermediate daily asset-level variables, such as the utilization rate of assets, asset to USD exchange rates, and the realized supply interest rate (the supply interest rate × utilization).

Combining the asset-level and account-level daily data, I calculate additional relevant account-level measures. The dollar value of each asset supplied and borrowed is calculated for each account, along with the net interest rate faced when maintaining the position. I also calculated portfolio weights based on the implied long-short portfolio of each position (with collateral asset representing the long position and the borrowed asset representing the short position). These portfolio weights are used to measure the daily standard deviation (volatility) of the long-short portfolio return, calculated each day using the preceding 30 days of daily asset returns.

The sample of asset-level daily data covers the period from 05/07/2019 to 12/15/2021 (954 days). The account-level daily data covers the period from 05/07/2019 to 10/26/2020 (539 days).[5] I exclude accounts with health values below 0.9, as these are very low-value positions which are effectively abandoned by their traders and ignored by liquidators. I also exclude accounts with health values greater than 10. Such accounts utilize less than 10% of their

[5] The account-level sample ends on 10/26/2020 as the Compound Finance API cannot access historical block-level account-states after this date. This is likely because the number of accounts and size of data exceeded the limitations of the hardware and/or software used to host the public API. To extend the dataset to include more recent periods require access to a local Ethereum archival node (renting access would be prohibitively expensive given the number of requests required and the size of the ideal dataset).

borrowing power and ostensibly represent lenders, not leveraged traders. Following these restrictions, the sample consists of 906,917 account-day observations, corresponding to 12,386 unique accounts.

I supplement these data with the average daily gas price on the Ethereum blockchain, collected from ycharts.com. This is the average fee paid, measured in GWEI, for a single unit of gas based on all recorded transactions within a given day. As GWEI is a unit of Ether, I convert this gas price into the average daily price for one gas in USD (*USD GWEI*). This can be used to estimate the average daily cost of specific types of transactions on Compound, based on the amount of gas they use.

3B: Summary Statistics

The Compound lending platform has experienced enormous growth during the last few years, with the value of supplied assets surpassing $20 billion in 2021, with close to $10 billion in borrowed assets. This growth is documented in Figure 3. Table 2 presents asset-level summary statistics. Historically, the most supplied assets are ETH, USDC, DAI, and WBTC/WBTC2. Of interest is the mean utilization of stablecoins, which is an order of magnitude higher than that of most other assets. The high demand to borrow stablecoins is indicative of the popularity of leveraged long positions in risky cryptoassets, particularly in Ether and Bitcoin, over my sample period. The overwhelming long interest in this period is not surprising given the meteoric rise in Bitcoin and Ethereum prices.

Collateral factors, though generally constant, have occasionally been adjusted. Two assets, Tether and TUSD, maintain a collateral factor of zero; traders cannot gain any borrowing power by using them as collateral. However, this does not mean there is no incentive to lend them. Suppliers of stablecoins have realized interest rates upwards of 22% APY in times of high demand for their assets. In contrast, the maximum interest rate earned for supplying

ETH over this period is roughly 1.2%, also indicative of the overwhelming long ETH interest during this period.

Table 3 presents borrower-level sample summary statistics. Generally, most accounts have aggressively leveraged positions, with a median health of 1.56 (64% of borrowing power utilized). Moreover, the distribution of account values (measured as the sum value of an account's borrowed and supplied assets) is highly skewed, with a mean value of $382,556.50 and a median value of $869.77. This implies that a smaller number of wealthy traders are responsible for the majority of borrowing in the market.

The cost to maintain a leveraged position is the *Net Interest Rate*, which is the realized *borrow interest rate* minus the realized *supply interest rate* for the traders' assets. Both the mean and median *Net Interest Rate* in my sample are relatively small, 1.1% and 0.9% respectively. This suggests that the interest earned on traders' collateral assets generally offsets the cost of interest accrued on their borrowed assets. In fact, the interest earned on collateral assets can entirely outweigh the interest cost from the borrowed assets for some accounts. The minimum Net Interest Rate in the sample is -23.7%, which implies that the trader was earning 23.7% APY while holding their leveraged position. This is not the case for traders who borrow assets in high demand, who face up to a 15.5% APY maximum interest rate in the sample. It is important to remember here that these interest rates are dynamic. Measured daily within my sample, rates typically do not remain at very high levels for long periods. A higher cost to borrow an asset leads to lower borrowing demand and a higher supply (as lenders capture this interest rate).

4. Transaction Fees – Descriptive Analysis

I begin with a descriptive analysis of the effect of blockchain transaction fees on the behavior of leveraged traders. Leveraged traders face volatile fixed costs in the form of

blockchain transaction fees (*gas*). While the amount of *gas* a transaction will use will be fixed over time, the market price for gas is not. As transactions are recorded by miners in descending order of *gas* fees, Ethereum users compete to be recorded first by a higher per-unit price of *gas* for their transactions. This competition is a source of fee volatility, as transactional demand is not constant. These transaction costs are also independent of the value of transactions; a $10 transfer would use the same amount of *gas* as a $10 million transfer. Given this property, I expect that leveraged traders have decreasing sensitivity to transaction costs as the size of their leveraged position increases. If this is the case, I expect to find that higher-valued accounts manage their leveraged positions more actively and take riskier positions (positions with health closer to the liquidation threshold).

4A: LPM and Logistic Specification

To validate this hypothesis, I test the effects of gas prices and borrower account size on the propensity of a leveraged trader to update/rebalance their position. I fit the following linear probability model:

$$\text{update}_{i,t} = \alpha + \beta_1 \text{gas\_price}_t + \beta_2 \text{total\_value}_{i,t} + X'_{i,t}\delta + \eta_i + \epsilon_{i,t} \quad (1)$$

where $\text{update}_{i,t}$ is a dummy variable equal to 1 if trader *i* rebalances their position on day *t*, gas_price_t is the average daily cost in USD of one unit of *gas*, $\text{total\_value}_{i,t}$ is the value (in USD) of trader *i*'s collateral plus borrowed assets, and $X'_{i,t}$ is a vector of control variables, including the health of the trader's account, the volatility of the trader's leveraged portfolio, and the net interest rate paid by the trader for maintaining the position (interest paid on borrowed asset(s) minus interest earned on supplied asset(s)). Account-level fixed effects are denoted by η_i. Standard errors are clustered at the account (trader) level, and all variables are standardized with mean 0 and standard deviation of 1 to allow for more intuitive interpretations.

To address some of the shortcomings of the linear probability model, I also present a logistic regression version of the above specification. Interpreting the coefficients of the linear probability model is difficult, as they are confounded by the account fixed effects. The interpretation of the average marginal effects, calculated from the logistic specification, is more straightforward in this regard as the fixed effects are suppressed. Because the independent variables are standardized, the average marginal effect coefficients can be interpreted as the average predicted change in propensity of the (marginal) trader to rebalance their position given a one standard deviation increase in the associated independent variable. A shortcoming of the fixed effects logistic approach is that groups (accounts) with no variation in the dependent variable (the *update* dummy) must be dropped prior to estimation. Practically, this means that roughly 8% of accounts (representing 0.2% of the observations) are dropped from the sample. The majority of these are very low-value accounts that never rebalance.

Table 4 presents the results from the fixed effects linear probability and logistic specifications. The results from the two specifications are consistent in terms of significance and direction of the predicted effects. A one standard deviation increase in transaction fee cost is associated with a 6% decreased propensity to rebalance a leveraged position. Conversely, a one standard deviation increase in the total value of a position predicts a 1.3% increased propensity to rebalance that position. These results are relatively intuitive and thus far consistent with the idea that traders are generally sensitive to transaction costs and that higher-valued accounts manage their positions more actively.

4B: Decile Analysis

However, the pooled specification does not do justice to the effects of transaction fees in relation to the size of position held by traders. All else equal, the effect of increased transaction fees should be larger for a lower-value trader and negligible for a higher-value trader. The relative sensitivity to fees based on the size of the traders' positions is easily

measured by categorizing traders based on the size of their positions. Thus, I repeat the above analysis after breaking the sample into value deciles. Within each day, I separate accounts into deciles based on the total value of their positions, then pool these deciles to form 10 unique subsamples.

The resulting summary statistics, presented in Figure 3, are illuminating. The unconditional propensity of traders to rebalance their positions increases significantly with the relative value of their leveraged position. Notably, in decile 1 this propensity is almost zero, suggesting that 10% of the market participants almost never rebalance. Conversely, within highest value decile, market participants have a daily propensity to update their positions of 20%, suggesting that they rebalance (on average) once every five days. There is an opposite trend when the mean *health* within each decile is considered. The *health* of lower value accounts is greater than that of higher value accounts, suggesting the wealthier traders take relatively more leverage (riskier positions).

To describe the effect of fees based on the value of traders' positions I perform the same fixed-effect logistic analysis for each subsample, re-standardizing the variables within deciles. I exclude analysis of decile 1, as there is a relative lack of variation in the outcome variable (traders do not update their positions). The results from this specification are presented in Table 5. From decile 4 to decile 10, the sensitivity of leveraged traders to transaction fees monotonically decreases and becomes insignificant within the highest value subsample. This is consistent with the idea that blockchain transaction fees disproportionately affect traders with less capital, while those with more capital more actively manage their leveraged positions. However, from decile 2-4 the sensitivity to fees is increasing with the size of the leveraged position. This is somewhat counterintuitive, as fixed costs (transaction fees) should be a greater burden on smaller operations.

4C: Temporary Abandonment of Accounts

The most obvious explanation for this trend is that the *USD_GWEI* coefficient estimates are very imprecise within the first few value deciles, given the relative lack of variation in the outcome variable *Update*. As the results presented in Table 5 are purely descriptive, there is no obvious solution to this problem. Another possible explanation for this perplexing trend is a non-linear sensitivity to fees based on the value of accounts. A general lack of response within lower-valued accounts to changes in transaction fees is plausible if contemporaneous fee levels entirely price these accounts out of the market. If this is the case, we would rationally expect such traders to have no response to changes in transaction fees below a certain threshold (i.e., a coefficient of 0).

To further illustrate the latter explanation, recall the earlier example of a trader that collateralizes $125 of ETH and borrows $100 of USDC. The gross benefit the trader enjoys when reversing their position is $25, and thus it is arguably rational for this trader to perform the reversal transaction as long as the transaction cost does not exceed the transaction benefit. However, the trader will be entirely insensitive to changes in transaction fees as long as they are above $25, as there is no observable or distinguishing action taken by traders that are "a little out of the money" versus "way out of the money." In such cases, the account can be considered *temporarily abandoned* until the transaction fee falls below $25.

But how pervasive is this problem? In order to explain the rising fee coefficients in Table 5 from deciles 2-4, a decreasing proportion of observations within the decile 2 to decile 4 subsamples should be "priced out" of the market. The point at which a trader is "priced out" is when the cost of transacting exceeds the benefits of transacting. While the cost of transacting is relatively easy to estimate using the average daily cost of gas and the gas usage of different

types of transactions on Compound (outlined in Table 1), the benefit is more difficult to measure for most types of transactions. In this regard, considering only reversal transactions is attractive, as the reversal benefit is simply the difference between the collateral and borrowed value of the traders' assets. Thus, I calculate the *Net Reversal Benefit* for each account day observation in my sample as follows:

$$\text{Net\_Reversal\_Benefit}_{i,t} = \text{Collateral\_Value}_{i,t} - \text{Borrowed\_Value}_{i,t} - \text{Fee}_{\text{t}} \quad (2)$$

A simple reversal transaction uses roughly 340,000 gas (90,000 to repay the borrowed asset plus 250,000 to reclaim the collateral asset). Fee_{t} is then calculated based on the average USD price for one gas in day *t* times 340,000 gas. There is appealing time series variation in the estimated transaction fee for a reversal transaction, ranging from \$0.41 to \$72.99 in the sample.

I consider the fraction of accounts with negative *Net Reversal Benefits* over time. This is a conservative cutoff, as it would be irrational for a trader to reverse a leveraged position if the transaction resulted in a net loss of liquidity, and therefore likely inefficient to rebalance the position in any other way. Figure 5 graphs the proportion of such accounts over time, along with the contemporaneous estimated cost of reversing a leveraged position. The fraction of temporarily abandoned accounts (those with a negative Net Reversal Benefit) remains significantly positive throughout the sample spiking to 50% following a sharp increase in transaction costs. For most of 2020, greater than 20% of accounts fall into this category.

The description provided by Figure 5 suggests that a good portion of lower-value traders are entirely priced out of the market and forced to temporarily abandon their positions. To relate the above figure to my decile analysis, I consider the percent of account-day observations below zero Net Reversal Benefit in each value decile. Within value decile 1, 99.9% of observations fall below this threshold, further validating the choice to exclude this

from the decile logistic analysis. 89.4% of observations in decile 2, 42.2% of observations in decile 3, and 4.0% of observations in decile 4 also have a negative *Net Reversal Benefit*. Finally, 0.3% and 0.1% of observations fall into this category in decile 5 and decile 6 respectively, while the remaining deciles contain no observations with a negative *Net Reversal Benefit*. This trend is consistent with the latter explanation for the results from the decile logistic analysis; a significant proportion of traders are insensitive to fees because they are already entirely priced out of the market.

Though my analysis of blockchain transaction fees offers no causal inference, the results are descriptive of the current state of DeFi technology. The extreme volatility in blockchain transaction costs is uniquely damaging to small traders. Though it is tempting to characterize this as just another market risk, it is hard to justify the argument that roughly half of all market participants expected to have their positions priced out of the market. It is entirely possible that early users of Compound failed to fully anticipate the true volatility of transaction fees on the Ethereum blockchain. Regardless, these findings highlight the limitations of DeFi technology in its current state. Moreover, the implications of temporarily abandoned leveraged trading accounts clashes with several of the fundamental ideals of decentralized blockchain systems: inclusivity and a lack of censorship. While not explicitly excluded from transacting, a good portion of leveraged traders are effectively excluded by high and volatile transaction fees.

5. The Causal Effect of Compound's Liquidation Incentive

Brokers perform liquidation transactions in the traditional market for leveraged trading. Brokers have obvious incentives to monitor leveraged trader accounts and perform these liquidation transactions, as they are the counterparty to the leveraged traders' loans. This cost is shifted to borrowers via the liquidation incentive in the DeFi leverage market. This DeFi market mechanism incentivizes third parties to monitor leveraged trading accounts on

Compound and perform liquidation transactions on accounts with insufficient collateral (health values less than 1).

5A: Difference in Differences Specification

March 12th, 2020 is coined cryptocurrency's "Black Thursday," as Bitcoin prices fell by almost 50% before recovering slightly. In response to this realization of systemic risk, the liquidation incentive on Compound was increased from 5% to 8% within Ethereum block number 9,683,141 on March 16th, 2020[6]. While this change affected all leveraged traders with borrowing positions, it disproportionately impacted traders with riskier positions. For example, a trader with an account health of 1.25 is utilizing 80% of their borrowing capacity and is more likely to face a liquidation (and the subsequent liquidation incentive) than a trader with an account health of 5 who is utilizing only 20% of their borrowing capacity. Given this heterogeneous exposure to liquidation incentive change, I conduct a quasi-natural experiment to test the effect of the policy change on risky borrowers.

I designate all traders who utilize more than 50% of their borrowing capacity as *risky*, creating a dummy variable equal to 1 if account health is less than 2. To test the effect of the liquidation incentive change, I consider the account *health* level at the time of rebalancing. While account *health* will naturally vary over time with borrowed and supplied asset exchange rates, traders implicitly choose a level of *health* at the time of rebalancing. Thus, the subsample including only accounts that have updated their positions in a given day may provide unique insight into the choices and behavior of leveraged traders.

To measure the effect of the change in liquidation incentive, I employ the following difference in differences specification:

[6] See transaction hash on Block 9,683,141: https://etherscan.io/tx/0xe477b9829951883a7e3950f41e312648a057e0cf80451d65a8915b75dbead56c\#eventlog

$$\text{health}_{i,t} = \alpha + \beta \text{Risky}_{i,t-1} \times \text{Post}_{\text{t}} + \text{X}'_{i,t}\delta + \eta_{\text{i}} + \eta_{\text{t}} + \epsilon_{i,t} \quad (3)$$

where *risky* is a dummy variable equal to 1 if the pre-update *health* of the account is less than 2, *Post* is a dummy variable equal to 1 in all periods after the change in liquidation incentive, $\text{X}'_{i,t}$ is a vector of control variables including total account value, volatility of leveraged position, net-interest rate at the account level, and the daily average *gas* price in USD. η_{i} indicates trader fixed effects and η_{t} indicates day fixed effects. Observations with negative *Net Reversal Benefit* are omitted, as are accounts with *health* values less than 1 or greater than 10. Also omitted are observations before August 9th, 2019, as the parallel trends assumption is violated when the sample includes the formative days of the Compound platform. Standard errors are clustered at the account level. The estimates from this specification are outlined in Table 6.

Column 1 includes the above specification with controls, but no fixed effects. Column 2 also includes account fixed effects, while Column 3 includes both account and day fixed effects. The interaction term of interest is positive and significant. Interpreting the most stringent specification (Column 3), traders with riskier positions react to the change in liquidation incentive by increasing their health value by 0.17. Given the average health of traders categorized as risky, this corresponds to a reduction of their borrowing utilization by roughly 8%, or a relative reduction of 12% compared to their pre-update position.

As the liquidation incentive is a proportional cost to borrowers, it is logical to expect that traders with larger positions have a greater response to the change. Column 5 and Column 6 report the difference in differences estimates for accounts that have values greater than $1,000 and $10,000 respectively. Consistent with this notion, traders with larger positions are more affected by the change. Those with positions valued at over $10,000 respond by increasing the

health of their accounts by 0.27, a reduction in their borrowing utilization of 11% (or relative reduction of 16%).

This behavior has interesting policy implications. To the extent that leverage traders are a source of systemic risk, a policy such as a liquidation incentive is appealing as it only constrains highly leveraged traders (i.e., risky traders), while more conservative traders are not affected. In this sense, such a regulation more appropriately targets the source of risk.

5B: Parallel Trends Assumption

For causal interpretation of the above specification, a parallel-trends assumption must be satisfied. This means that, absent the change in liquidation incentive, the *health* trends of traders designated as risky should be parallel to the *health* trends of traders who are not, both before and after the change. As the change in liquidation incentive affects all traders, there is no post-event counterfactual. Despite this, I can still verify no un-parallel pre-trends exist between the treated and untreated groups. This may be accomplished in two ways. First, Figure 6 plots the mean *health* over time of the treated group (*risky*) and the control group, including fitted lines for both subsamples. The trend lines of these two groups are visually parallel. Second, I conduct a version of the difference in differences equation above in which pre-event period dummies are interacted with *risky*. The coefficient estimates are reported in Table 7. Consistent with the parallel trends assumption, these interaction terms are statistically insignificant.

5C: Alternative Specification: Continuous Treatment

The difference in differences estimators described above utilize a discrete categorization of risky vs. non-risky traders to identify the treatment effect. This approach is attractive as it produces results that are intuitive and easy to interpret. Moreover, the methodology is widely used in finance and economics literatures, making the execution and

tests of assumptions relatively straightforward. This said, grouping all accounts with *health* values less than 2 might not be the best approach to indicate their exposure to the change in liquidation incentive.

There are two issues with the discrete categorization. First, the account *health* statistic is dependent on the *collateral factor(s)* of the supplied assets. Ceteris paribus, accounts that supply assets with high *collateral factors* will have a greater *health* value than those that supply assets with a lower *collateral factor*. This means that accounts that lend assets with a higher *collateral factor* might lose more in the event of a liquidation than accounts that lend assets with a lower *collateral factor*. For example, consider two accounts with a *health* value of 2. The first account lends \$100 of an asset with a *collateral factor* of 0.8 and borrows \$40 of another asset. The second account lends \$100 of an asset with a *collateral factor* of 0.6 and borrows \$30 of another asset. Both accounts use 50% of their borrowing power (as denoted by a *health* of 2), yet the first account is clearly more exposed to the cost of liquidation. Given an 8% liquidation incentive, the first account would lose \$1.60 to the liquidation incentive and the second account would lose \$1.20.

The second issue with the discrete difference in differences approach is that all accounts, regardless of their health, could be affected by the change in liquidation incentive. Thus, accounts that are not categorized as risky might not be a good counterfactual group. In reality, it is quite unlikely that accounts that use less than 50% of their borrowing power will face a liquidation. It is more likely that accounts that use slightly more than 50% of their borrowing power (i.e., with *health* factors just below 2) do not have a substantial risk of liquidation. This would bias the interaction coefficient towards zero in the discrete case.

To address these concerns, I implement a continuous version of the differences in differences specification outlined above in Equation 3. Forgoing the *risky* dummy, I create a

new variable to measure the relative exposure of a position to the increased cost of liquidation, *Liq_Harm_Frac*, which is simply the fraction of an account's collateral value that would be lost via liquidation incentive. This is calculated as follows:

$$\text{Liq\_Harm\_Frac}_{i,t} = \frac{0.5 \times borrowed\ value_{i,t} \times\ 8\%}{collateral\ value_{i,t}} \quad (4)$$

Thus, the continuous version of the difference in differences specification is modeled as follows:

$$\text{Health}_{i,t} = \alpha\ + \beta\ \text{Liq\_Harm\_Frac}_{\text{i,t}} \times \text{Post}_{\text{t}} + \ \text{X}'_{\text{i,t}}\delta + \eta_{\text{i}} + \eta_t + \epsilon_{\text{i,t}} \quad (5)$$

The estimates from this specification are presented in Table 8. *Liq_Harm_Frac* is standardized with mean zero and standard deviation one to allow for easier interpretation. Consistent with the discrete estimation, the change in liquidation incentive induces traders that are more exposed to the change (a higher *Liq_Harm_Frac*) to take less leveraged positions. Interpreting the most stringent fixed effects specification (Column 3), a one standard deviation increase in exposure to the liquidation incentive cost is associated with such traders rebalancing their positions after the policy change by increasing the *health* of their accounts by 0.204. This effect is stronger for traders with total account values greater than $1,000 and $10,000, reported in columns 4 and 5 respectively.

While the treatment variable measuring a continuous exposure to the change in liquidation incentive likely better represents reality than the discrete treatment identifier, the estimates from continuous difference in differences specifications face several issues in terms of interpretation and identifying assumptions. These issues are explored by Callaway, Goodman-Bacon, and Sant'Anna (2021). First, interpreting the differences across different values of treatment (*Liq_Harm_Frac*) requires an assumption of a no heterogeneity in relative

treatment effect across different levels of treatment. Absent this assumption, estimates suffer from a heterogeneous treatment effect bias.

It is also more difficult to satisfy the parallel (pre)trends assumption with a continuous treatment variable, as this implicitly means that trends for all levels of treatment should be parallel. This is impossible to display graphically without breaking the continuous variable into discrete groups, which arguably defeats the purpose of the continuous treatment variable. Figure 7 presents this graphical analysis, separating the sample into three subsamples based on levels of *Liq_harm_frac*: within one standard deviation of the mean value of *Liq_harm_frac*, greater than 1 standard deviation above the mean, and less than 1 standard deviation below the mean. Graphically, these three groups appear to have parallel pre-trends in health.

To test the parallel pre-trends assumption in a continuous manner, I also employ a version of the specification outlined in Equation 5 with pre-event period interactions with the treatment variable. These estimates are reported in Table 9. The first and third thirty-day pre-period interactions are significant at the 95% confidence level, violating the parallel pre-trends assumption for the continuous specification. The results of the continuous pre-trends test are highly sensitive to the specification of pre-event periods.

6. Conclusions

The growth of DeFi blockchain applications presents exciting opportunities for economists, as blockchain data are public and transparent. Taking advantage of this, I study how market frictions effect the behavior of leveraged traders utilizing the Compound application on the Etheruem blockchain. Examining extensive trader-level data, I argue that blockchain transaction fees, a necessary feature of decentralized blockchain markets, price out a significant proportion of leveraged traders. These traders are effectively excluded from transacting due to high transaction fees.

I also examine the effect of Compound's liquidation incentive on the behavior of leveraged traders. The liquidation incentive is a necessary market mechanism in the DeFi lending space, as there exists no intermediary to liquidated accounts with insufficient collateral. Thus, third party liquidators capture the liquidation incentive from borrowers in this market. Utilizing a difference in differences specification by exploiting a change in the liquidation incentive, I show that risky borrowers respond to an increased liquidation incentive by taking more conservative leveraged positions.

These findings have several interesting implications. First, they highlight the shortcomings of blockchain and DeFi technologies in their current state. As blockchain transaction fees effectively exclude a portion of market participants from transacting, promising avenues for future research include studies of normative or positive solutions to the blockchain scaling problem. Furthermore, the liquidation incentive market mechanism may be of interest to policymakers, as it induces leveraged traders to take more conservative positions, with a targeted response amongst the riskiest of leveraged traders.

Chapter 2: The Public Blockchain Ecosystem: An Empirical Analysis

Coauthored with Felix Irresberger, Kose John, and Fahad Saleh

1. Introduction

There exist hundreds of public blockchains, but the economics literature has studied only Bitcoin in depth. While each individual blockchain platform might not be worthy of study, it is unclear whether there are blockchains other than Bitcoin that deserve attention. The key purpose of this paper is to examine the blockchain universe to identify which elements might be worthy of study from an economic perspective. Our main finding is that there exists an exclusive set of blockchains, which we term the *blockchain frontier*, that could generate the highest user utility among all blockchains for some user type. Crucially, no single blockchain is optimal for all user types, and most blockchains are not optimal for any user type. Our blockchain frontier, by construction, excludes those blockchains that are not optimal for any user type, thereby highlighting those public blockchains most relevant from an economic welfare perspective.

Our frontier does not necessarily equate with the set of blockchains possessing the highest market capitalization. Rather, our frontier identifies those blockchains that bring the highest utility to at least some users. As an example, a blockchain targeted for a small set of users could appear in our frontier but would be unlikely to possess a large market capital. In particular, the small set of targeted users may find the blockchain preferable to all other blockchains, thereby entering it into our blockchain frontier. At the same time, the small set of targeted users would imply low demand for using the blockchain and thus a low overall market capital.

We begin in Section 2 with a broad overview of the public blockchain ecosystem. Immediately, it becomes clear that the universe of blockchains consists of many small

platforms unlikely to be of broad economic significance. Nonetheless, there is also clear evidence of an evolution since Bitcoin's birth in 2009. In particular, we document an explosion in the launching of platforms that are fundamentally different from Bitcoin, many of which possess nontrivial market capitalization. As an example, the number of blockchains employing protocols different than Bitcoin to generate agreement on ledger contents, commonly referred to as consensus, far outnumber those using the same protocol as Bitcoin. It is noteworthy that Bitcoin remains the largest blockchain in terms of market capitalization, but it is also important to recognize that Bitcoin does not dominate across all relevant economic quantities. Of particular note, Bitcoin has processed far fewer transactions than several other blockchain platforms. Moreover, Bitcoin possesses narrow functionality so that it has limited relevance for important blockchain applications such as Decentralized Finance (DeFi). As a whole, our empirical overview unveils a few dominant blockchain platforms but no monolithic leader, not even Bitcoin. To understand these empirical patterns and the economic relevance of individual blockchains more deeply, we turn to a theoretical framework.

Our theoretical framework, put forth in Section 3, establishes that a blockchain user's welfare depends critically upon three characteristics of the blockchain: adoption, scale, and security. User utility increases in each of these characteristics. Adoption refers to the number of active users on the same blockchain platform equipped to engage in an economic interaction of interest. For example, if Alice would like to exchange assets without an intermediary, then her utility would increase in the number of active users that could serve as her counterparty as more such users would yield improved liquidity. More generally, a larger active user base improves the likelihood of finding a suitable counterparty in any economic interaction so that adoption increases transaction surplus and thereby enhances user utility. Scale refers to the rate at which transactions are added to the blockchain. We demonstrate that a higher scale enables not only more expeditious processing but also endogenously lowers transaction fees, thereby

enhancing user utility. Security refers to the likelihood that a transaction could be reversed. In practice, counterparties impose delays in settlement to compensate for such security risks. Consequently, higher security blockchains generate higher user utility because higher security implies shorter settlement delays.

We empirically implement our theoretical framework in Section 4. Specifically, we compare existing blockchains on the dimensions of adoption, scale, and security. We find that Bitcoin is the most widely adopted among users seeking simple payments, but Bitcoin does not lead on any other dimensions. Of particular note, Bitcoin's lack of functionality renders it without any active user base for all economic interactions beyond payments.

Consequently, Bitcoin has no economic value for users interested in activities such as DeFi transactions, gaming, gambling or data storage. We find that Ethereum leads by a wide margin on DeFi and narrowly over EOS on gaming. TRON leads all other platforms, even Ethereum, with regard to data storage and gambling applications. While Bitcoin, Ethereum and TRON all lead on some aspect of adoption, none lead with regard to scale. The EOS blockchain leads with regard to scale which enables EOS to provide quick transaction processing at low costs relative to all other blockchain platforms. All of the referenced blockchains lag behind on security when compared to Stellar, which is the most secure blockchain.

The purpose of our theoretical model and its empirical implementation is to identify those blockchains that fit within the blockchain frontier. As discussed, the blockchain frontier excludes any blockchain that cannot provide the highest utility among all blockchains for any user type. To identify this set of blockchains, we invoke one of our theoretical results from Section 3, which establishes that a blockchain is in the frontier if and only if it is not dominated by any other blockchain on several pairwise comparisons involving the key blockchain characteristics of adoption, scale, and security. Applying the aforementioned comparisons, we

find a blockchain frontier consisting of seven blockchains. The set includes all five of the previously referenced blockchains that lead on some characteristic: Bitcoin, Ethereum, TRON, EOS and Stellar. It also includes Binance Coin and Bitcoin Satoshi's Vision.

As our theory demonstrates, a blockchain that leads on adoption for users with any particular economic purpose (e.g., payments, DeFi, etc.) is necessarily part of the blockchain frontier. Intuitively, this follows because a user who is insensitive to wait times values only the transaction surplus from her economic interaction, and that transaction surplus depends exclusively on finding a suitable counterparty. Consequently, such a user finds whichever blockchain leads on adoption with respect to her economic purpose to be optimal for her. This means that Bitcoin is best for a user desiring to conduct a simple purchase or sale so long as the user is sufficiently insensitive to wait times. Similarly, a wait insensitive user desiring to conduct a financial transaction without an intermediary (i.e., a DeFi transaction) would prefer Ethereum to all other blockchains, and a wait insensitive user desiring to engage in decentralized data storage would prefer TRON to all other blockchains. Accordingly, Bitcoin, Ethereum and TRON are part of the blockchain frontier.

A blockchain that leads on scale also necessarily enters the blockchain frontier. Intuitively, users that are highly sensitive to waiting prefer the highest scale blockchain. Such users tend to be those that are interested in applications that involve a large number of actions such as some gaming applications. Each action in a game on a blockchain typically requires a transaction for that action. Thus, some gaming applications involve frequent transactions, and frequent transactions necessitate quick and cheap processing for each individual transaction. In turn, as we demonstrate in Section 3, the highest scale blockchain most appropriately suits these needs as it not only provides quick processing of transactions but also endogenously generates low transaction costs. As noted, EOS leads on scale and thus is necessarily also part

of the blockchain frontier. It is also noteworthy that EOS attracts a high level of gaming activity potentially because its scale makes it well suited for such activities.

A blockchain leading on security does not imply its inclusion in the blockchain frontier. This is because counterparties internalize security risks and compensate for those risks by imposing settlement delays upon users. As we demonstrate in Section 3, those settlement delays depend not only upon the blockchain's security but also upon the rate at which blocks are added to the blockchain. In particular, less secure blockchains incur higher settlement delays to compensate for the lack of security. However, our results establish that a given level of security implies a settlement delay of a particular number of blocks to be added to the blockchain, not a specific amount of time. In turn, since users incur disutility with respect to wait times rather than the number of blocks, user disutility associated with settlement delays depends also on the block arrival rate of a blockchain.

While a blockchain leading on security does not imply its inclusion into the blockchain frontier, a blockchain generating the lowest settlement delay does imply such inclusion. Binance Coin exhibits the lowest settlement delay and is therefore a part of the blockchain frontier. Binance Coin is used heavily in trading with the Binance cryptocurrency exchange. The low settlement delay is especially advantageous for users interacting frequently with a cryptocurrency exchange since the low settlement delay enables such users to trade frequently without lengthy interruptions.

Stellar and Bitcoin Satoshi's Vision are also part of the blockchain frontier. These blockchains feature in the frontier because they perform well on several of the key blockchain characteristics of adoption, scale, and security. This overall performance is sufficient such that no single blockchain induces a higher utility than either of them across all user types and hence enables both to enter the blockchain frontier.

1A: Related Literature

Our paper joins a growing literature on blockchain economics that began with Harvey (2016) and Yermack (2015, 2017). Chen, Cong, and Xiao (2019) provide an overview of that literature. Li, Shin, and Wang (2019), Liu and Tsyvinski (2020), Liu, Tsyvinski, and Wu (2022), Makarov and Schoar (2019), Shams (2020), Griffin and Shams (2020) and H¨ardle, Harvey, and Reule (2020) study the asset pricing dimensions of cryptoasset markets. In contrast, we focus on the economic features of the underlying blockchain. We empirically analyze blockchains employing a variety of consensus protocols, a topic that has been extensively studied in theoretical terms. Most studies consider PoW. Prominent studies of PoW include Arnosti and Weinberg (2018), Basu, Easley, O'Hara, and Sirer (2018), Budish (2018), Chiu and Koeppl (2018), Biais, Bisiere, Bouvard, and Casamatta (2019), Chiu and Koeppl (2019), Benetton, Compiani, and Morse (2019), Easley, O'Hara, and Basu (2019), Hinzen, John, and Saleh (2022), Huberman, Leshno, and Moallemi (2021), Saleh (2019), Alsabah and Capponi (2020), Cong, He, and Li (2021a), and Pagnotta (2022). Prominent studies of PoS include Irresberger (2019), Rosu and Saleh (2021), and Saleh (2021).

This paper proceeds as follows. Section 2 details the major consensus protocols and provides a descriptive analysis of the public blockchain ecosystem. Section 3 puts forth our theoretical framework, while Section 4 implements that framework empirically. Section 5 describes the blockchain frontier resulting from our empirical analysis. Section 6 concludes.

2. The public blockchain ecosystem: An overview

This section provides a descriptive analysis of the public blockchain ecosystem. We document the existence of hundreds of public blockchains, the consensus protocols they use, and a stark heterogeneity in relevance, as indicated by market capital and blockchain usage. We first give an overview of consensus protocols and then provide evidence on the current

state of ecosystem with respect to them. Our results highlight and motivate the need for an economic framework to explain the success of very few blockchains within the ecosystem.

2A: Institutional background on consensus protocols

A blockchain is an electronic ledger that is distributed across a network of agents referred to as validators.[7] Blockchain validators differ from blockchain users in that users submit transactions for processing on the blockchain, whereas validators determine whether those transactions achieve settlement. Blockchain users generally submit transactions via a native platform currency that we hereafter refer to as the native cryptoasset.[8] For the blockchain to be useful for users, there must be a process by which submitted transactions achieve settlement. By definition, a transaction is considered settled on the blockchain only if the validators agree on the transaction being entered on the blockchain. Accordingly, agreement among validators, known as *consensus*, is a key concern for any blockchain. Each blockchain attempts to resolve that concern via a set of rules for updating the blockchain known as a *consensus protocol*. Such protocol is a technical feature that is often used to distinguish between different blockchain designs. For this study, we partition the space of consensus protocols into five groups: (1) PoW, (2) PoS, (3) Hybrid PoW/PoS, (4) DPoS and (5) Nonstandard protocols. The remainder of this subsection provides a brief description of each category.

Proof of Work

PoW was first introduced by Dwork and Naor (1992) as a method to disincentivize spam emails. PoW then gained broad attention when Nakamoto (2008) employed it as the consensus protocol for Bitcoin. PoW was also the most prominent consensus protocol adopted by early public blockchain alternatives to Bitcoin, such as Litecoin and Dogecoin. As of May

[7] For Bitcoin and similar blockchains, these validators are referred to as miners.
[8] For example, Bitcoin has a native cryptoasset called bitcoin, and Ethereum has a native cryptoasset called ether. Some blockchains, such as Ethereum, also allow for other assets to be exchanged on the platform.

2021, the two largest blockchains, as measured by the market value of their native cryptoassets, Bitcoin and Ethereum, employ PoW. Under PoW, participating in the process of validating new blocks requires solving a computational puzzle. Nakamoto (2008) argues that the computational complexity of this puzzle makes it difficult to change the contents of the ledger which in turn helps generate consensus.[9]

In subsequent work, however, Biais et al. (2019) argue that persistent disagreement can arise on PoW blockchains. PoW has further been criticized on the grounds of requiring excessive amounts of resources. Mora, Rollins, Taladay, Kantar, Chock, Shimada, and Franklin (2018) provide empirical evidence in favor of that claim and establish that the PoW structure generates an exorbitant energy expenditure. These and other concerns associated with PoW (see, e.g., Hinzen et al. 2022) have spurred researchers on a quest to uncover alternative protocols.

Proof of Stake

PoS constitutes one of the early alternatives to PoW. King and Nadal (2012) provide the first PoS proposal. PoS overcomes PoW's need for exorbitant energy expenditure by removing PoW's computationally complex puzzle from the consensus process. The first pure PoS protocol was employed on the Nxt blockchain in 2013. Under PoS, agents are randomly conferred the authority to validate the next block. The likelihood of being chosen depends on an agent's "stake" which refers to an agent's overall holdings of the native cryptoasset.[10]

Hybrid PoW/PoS

Some blockchains have attempted to combine both PoW and PoS protocols in an effort to capture the benefits of both protocols. The first such hybrid consensus protocol was utilized on the Peercoin blockchain in 2012. Since then, dozens of other blockchains have followed this approach. Prominent Hybrid blockchains include Dash and Decred. Both blockchains employ

[9] For further context regarding PoW, the interested reader can consult Biais et al. (2019).
[10] For further context regarding PoS, the interested reader can consult Saleh (2021).

PoW to validate blocks, but grant governance rights to agents based on their stake in the native cryptoasset. Such governance rights can include the right to vote on changes to the protocol or confirmation of the validity of blocks.

Delegated Proof of Stake

Both PoW and pure PoS blockchains have faced criticism for their alleged inability to process transactions quickly. A variant of PoS protocols known as Delegated Proof-of-Stake (DPoS) arose in response to this criticism. DPoS was first introduced by Daniel Larimer, the eventual founder of the blockchain EOS. Under DPoS, a fixed number of delegates, also called "witnesses," validate new blocks. Holders of the native cryptoasset vote for delegates. Similar to PoS, the number of votes is assigned in proportion to holdings of the native cryptoasset. The relatively small and fixed number of delegates reduces the validating agent network size compared to traditional PoS protocols. This small network size, in theory, allows for faster transaction rates, as fewer agents have to converge on a common version of the ledger.

Nonstandard Protocols

Research on new consensus protocols is an active field. Numerous proposed blockchain protocols starkly differ from the protocols described above. We group these protocols and refer to them as Nonstandard protocols. Within this group, no single protocol is widely used. The most prominent public blockchains within this category are Ripple and Stellar.

2B: Empirical evidence on public blockchains

To provide an overview of the public blockchain ecosystem, we use cross-sectional data from cryptoslate.com on hundreds of public blockchains (using February 21, 2021 as cutoff date).

For our analysis, we distinguish public blockchains from cryptoassets. Each public blockchain could possess multiple cryptoassets, but each public blockchain typically has only

one native cryptoasset, which is created at the inception of the blockchain.[11] Cryptoslate.com lists a total of 2,400 different cryptoassets. Approximately one third of this set is made up of native cryptoassets with an aggregate market capitalization of ca. $1,594 billion, and two thirds consist of tokens with a total market capitalization of $203 billion. We subsequently focus on the 783 native cryptoassets corresponding to a unique public blockchain.[12]

Figure 8 plots native asset market capitalization (on a log-scale) against blockchain announcement dates for the full set of blockchains listed on cryptoslate.com. The majority of public blockchains carries little to no value in terms of market capitalization. Half of all blockchains (400) exhibit a market capitalization of its native asset below $1 million or carry no value (e.g., as it is not traded on exchanges or has been abandoned). There are a total of 344 blockchains with native asset market capitalizations between $1 million and $1 billion, and only 39 with market capitalization above $1 billion. The latter group consists not only of older blockchains like Bitcoin, which may have a first-mover's advantage, but also of newer blockchains, such as EOS, TRON or Binance Coin.

To distinguish between different types of blockchains, we divide the large number of public blockchains into groups based on the consensus protocol they employ. Table 10 characterizes the public blockchain ecosystem by the number of counts of a particular consensus protocol over time until the end of 2020. Bitcoin was the first blockchain launched in 2009 and employs PoW as a consensus protocol. Over the years, many other public

[11] As an example, Ethereum has the native asset ether (ETH) but allows for the creation of many other nonnative assets, such as Basic Attention Token (BAT) or Maker (MKR). Nonnative assets are generally referred to as tokens and have been extensively studied within the initial coin offerings (ICO) literature (see, e.g., Malinova and Park 2018, Chod and Lyandres 2019, Davydiuk, Gupta, and Rosen 2019, Gan, Tsoukalas, and Netessine 2020, Lee, Li, and Shin 2019, Lyandres, Palazzo, and Rabetti 2019, Howell, Niessner, and Yermack 2020, Li and Mann 2020, and Liu, Sheng, and Wang 2020).

[12] To identify all native cryptoassets, we cross-check the information given by cryptoslate.com with the existence of a block explorer and perform additional background research on the history of the blockchain in case the project was abandoned over time. Whenever the announcement date of a blockchain is not disclosed on cryptoslate.com, we take the timestamped date of the genesis block as the announcement date of the blockchain. We also cross-check the information on the consensus protocol employed by a blockchain disclosed on cryptoslate.com with additional web searches.

blockchains have employed PoW as their consensus protocol, the most prominent ones being Litecoin (LTC), launched in 2011, Ethereum (ETH), launched in 2014, or Bitcoin blockchain forks such as Bitcoin Cash (BCH) and Bitcoin Satoshi's Vision (BSV), launched in 2017 and 2018, respectively. Alternative protocols such as PoS or Hybrid PoW/PoS started to gain significantly in numbers starting from 2014 onwards. Although nonstandard protocols started with blockchains such as Ripple (XRP) in 2013, we have seen the number of nonstandard blockchains increase dramatically from 2016 to 2018 and the new number of nonstandard and PoW blockchains are comparable in 2019 and 2020, respectively. Although much fewer in numbers compared to other protocols, DPoS blockchains such as EOS (EOS) and TRON (TRX) started to emerge more in 2016 and onwards. PoW no longer dominates as the choice for a blockchain's consensus protocol, but no specific protocol has taken PoW's dominant position; rather, recently launched public blockchains frequently employ alternative protocols.

The left panel in Figure 9 shows the distribution of raw counts of consensus protocols. Although the number of blockchains employing PoW is highest among all protocol groups, alternatives such as PoS, Hybrid PoW/PoS, and Nonstandard protocols are not far behind. However, if we view the ecosystem through the lens of market capital share, we see a heavily skewed distribution of relevance. The right panel in Figure 9 shows the market capital share (as of February 21, 2021) of Bitcoin, Ethereum, and native cryptoassets of blockchains using PoW, PoS, Hybrid, DPoS, and Nonstandard protocols, respectively. Bitcoin accounts for two thirds of the total market capitalization of the public blockchain ecosystem, while Ethereum captures a 14.2% market share. DPoS captures the third largest share across these categories with a share of 6.3% of the total native asset market capitalization. Other PoW, PoS, and Nonstandard protocols' market share combined make up around 11.4% and Hybrid PoW/PoS based native cryptoassets have comparatively negligible market capital.

Figure 10 shows the distribution of market shares within each consensus protocol group. Within each group, most of the market capital share is captured by only a handful of blockchains. Other relevant PoW blockchains include Bitcoin Cash, Bitcoin Satoshi's Vision, Litecoin, Monero, and Dogecoin, which together capture over 75% of the market capital attributed to PoW blockchains excluding Bitcoin and Ethereum. Within PoS, Cardano captures over half of the market capital; within the Hybrid group, DASH and Decred are the most relevant. DPoS is highly dominated in terms of market capitalization by Binance Coin and Polkadot, with EOS, Tezos, and TRON together capturing less than a quarter of all market capitalization in this protocol group. The most important nonstandard blockchain is Ripple, with Stellar being a distant second.

For the remainder of the paper, we restrict our empirical analysis to the set of blockchains with daily time-series of network characteristics (e.g., number of on-chain transactions, active addresses, block counts) available on coinmetrics.io. This subsample is small in numbers but already represents the majority of the overall public blockchain universe (ca. 99% of market capital of *native cryptoassets* and 88.6% of *all* cryptoassets, including tokens, listed on cryptoslate.com). The sample blockchains' names and announcement year, market capitalization, the sum of all transaction counts since inception and share of all transaction counts across the sample are documented in Table 11.

Finally, we consider blockchain usage as another economic quantity of interest. Bitcoin (BTC), the oldest PoW blockchain with by far the highest market capitalization, contributes to only 5.5% of all transactions. Ethereum (ETH) captures only 9.1% of the total transaction count. The vast majority of the total transaction count is attributable to EOS and TRON, two blockchains with DPoS protocols. However, Ripple (XRP), a nonstandard blockchain, captures around 15.7% of all transactions. Thus, while Bitcoin remains the most studied blockchain in

the economics literature and the one with the highest market capitalization, it is far from the most used blockchain in terms of on-chain transaction counts.

In summary, we find that there are only a few public blockchains of particular relevance, as measured by native asset market capitalization and transaction counts, and that these blockchains employ a variety of consensus protocols. None of the consensus protocol groups seem to dominate the others, but rather we observe a stark heterogeneity in relevance even within protocol groups. In the following, we put forward a framework that helps explain the heterogeneity in blockchain relevance by demonstrating that only a small set of blockchains are optimal for users to employ.

3. Theoretical Framework

Our model examines a setting in which users arrive randomly to a blockchain and seek to transact on that blockchain. We examine how the welfare of each user varies as a function of the blockchain's characteristics. We demonstrate that user utility depends upon three key blockchain attributes: adoption, scale, and security. Subsequently, we use these three attributes as the attributes of our empirical framework.

We model an infinite-horizon continuous time setting with time indexed by $t \geq 0$. Users with unit transaction demand arrive randomly according to a Poisson Process with rate $a > 0$. We assume that users are heterogeneous in terms of their economic purpose for using the blockchain, the value they place on the economic exchange and the disutility per unit time they incur from waiting for their transaction to be settled. With regard to heterogeneity of economic purpose, we assume that each User i possesses an economic type $\tau_i \in \mathrm{T}$. This economic type, τ_i, reflects that users have different purposes for using the blockchain in practice. For example, some users transact via the blockchain to engage in regular purchases (i.e., payments) whereas others transact via the blockchain to engage in more complex financial

transactions that bypass traditional intermediaries (i.e., Decentralized Finance). Accordingly, two elements of T include payments and Decentralized Finance (DeFi), but there are other economic purposes for using a blockchain as we discuss in Section 4. With regard to how a user values the economic exchange, we follow Cong, Li, and Wang (2021b) and assume that each User *i* possesses a type $u_i \sim F[0, \overline{u}]$ which is proportional to her transaction surplus. With regard to heterogeneity of wait disutility, we assume that each User *i* incurs a disutility of waiting per unit time of $c_i \sim G[0, \overline{c}]$ with *G* being strictly increasing and twice continuously differentiable. We allow that a user may ameliorate her total disutility from waiting by paying a fee but users also dislike paying fees. This modelling choice reflects the fact that blockchain users are processed in descending fee order in practice so that higher fees induce lower wait times. Formally, User *i* selects fee $f_i \geq 0$ by solving:

$$f_i = \arg\max_{f \geq 0} \underbrace{u_i \cdot \Psi(N_{\tau_i})}_{Transaction\ Surplus} - \underbrace{c_i \cdot W(f)}_{Wait\ Disutility} - \underbrace{f}_{Fee} \quad (6)$$

where N_{τ_i} denotes the number of active users on the blockchain with economic purpose τ_i and $W(f)$ denotes the expected wait of User *i* if she pays fee $f \geq 0$. Our specification of transaction surplus follows Cong et al. (2021b) in that we assume that $\Psi' > 0$ so that user transaction surplus depends positively upon the number of users on the blockchain with the same economic purpose. This assumption reflects the fact that it is easier to find a suitable counterparty in a larger community. We refer to N_{τ_i} as adoption hereafter and emphasize that user welfare increases in adoption.

A novelty of our analysis is that we expand the expected wait time, $W(f)$, to incorporate a novel feature relative to prior literature. In particular, our wait time term decomposes into two components with the latter term being novel relative to prior works:

$$W(f) = \underbrace{W_P(f)}_{Processing\ Time} + \underbrace{\kappa}_{Confirmation\ Time} \quad (7)$$

Prior papers consider exclusively that which we refer to here as processing time, $W_P(f)$. This term corresponds to the expected time from the point at which User *i* sends her transaction into the blockchain network for processing to the point at which the transaction appears on the blockchain's ledger. User *i*'s processing time, $W_P(f)$, depends upon the fee that User *i* pays, *f*, and the fees from the set of transactions awaiting inclusion to the blockchain ledger. The set of transactions awaiting inclusion on the blockchain ledger is commonly referred to as the memory pool, which we assume follows its stationary distribution. Recall that the blockchain processes transactions in descending fee order so that User *i*'s transaction is added to the blockchain only after all transactions in the memory pool with higher fees are processed. Moreover, a transaction with a higher fee arriving after User *i*'s transaction would be prioritized ahead of User *i*'s transaction if User *i*'s transaction has not already been added to the blockchain at that time. Thus, User *i* may alleviate her waiting time by paying a higher fee, thereby reducing the number of transactions that are prioritized ahead of her transaction.

Following prior literature, we assume that blocks are created according to a Poisson Process with rate $\Lambda > a$. For exposition, we normalize block sizes to one transaction per block. Under these assumptions, we have the following result:

Proposition 3.1. *User i's fee, f_i, and her Processing Time, $W_P(f_i)$, both decrease in Scale, Λ*

In equilibrium, User i's fee choice is given by:

$$f_i = 2 \cdot a \cdot \Lambda \int_0^{c_i} \frac{g(x)}{\left(\Lambda - a \cdot \overline{G}(x)\right)^3} dx$$

with $\overline{G}(x) := 1 - G(x)$. Moreover, the equilibrium processing time, $W_P(f_i)$, is given by:

$$W_P(f_i) = \frac{\Lambda}{(\Lambda - a \cdot \overline{G}(c_i))\text{^}2}$$

Both the equilibrium fee, f_i, and the equilibrium processing time, $W_P(f_i)$, decrease in the scale of the blockchain.

Proposition 3.1 establishes that fees depend negatively upon scale. The intuition behind this finding is that the value of priority to get on the blockchain decreases with scale so that the value of paying fees decreases as scale increases. Hence, as scale increases, fees endogenously fall.

To understand the aforementioned point, recall that users are processed in descending order of fees so that the number of users for which User *i* must wait depends upon the relative order of her fee among all fees. For concreteness, suppose that User *i* may pay $\varepsilon > 0$ extra to gain priority over one user in expectation. Then, the utility value for her to gain such priority equals $c_i \cdot \frac{1}{\Lambda}$ because c_i denotes her wait disutility per unit time and $\frac{1}{\Lambda}$ denotes the amount of time by which her expected wait falls if she gains priority over one user. Thus, since $c_i \cdot \frac{1}{\Lambda}$ decreases in scale, Λ, so too does User *i*'s benefit from paying the additional $\varepsilon > 0$ in fees. Accordingly, as the scale increases, User *i* becomes less willing to pay higher fees and thus her fee choice also falls.

Proposition 3.1 also establishes that User *i*'s processing time, $W_P(f_i)$, decreases with scale, Λ. This result follows for two reasons. First, holding fixed the number of users that receive service before User *i*, a higher scale directly reduces the expected time required to process those higher priority users. Second, the number of users that receive service before User *i* itself decreases as the scale increases. To see this last point, recall that User *i* must wait behind not only users already in the memory pool that pay higher fees but also later arriving users who pay a higher fee than User *i* if those users arrive before User *i* receives service. A

higher scale reduces the probability that any new user arrives before any user already in the memory pool receives service; in turn, the number of users that receive service before User i receives service decreases in scale thereby also contributing to a reduced wait time for User i.

Our analysis of the processing time, $W_P(f_i)$, is standard (see, e.g., Huberman et al. 2021), but we expand on previous work by explicitly incorporating the confirmation time, κ, into the user wait time. The confirmation time refers to the time in between User i's transaction being entered on the blockchain and the time that User i receives delivery of the other leg of the exchange. For example, if User i is selling cryptocurrency units in return for fiat or some physical good then User i's transaction being posted to the blockchain does not imply that the receiver releases the fiat or goods to User i. Rather, when selling fiat or a physical good for cryptocurrency units, the seller employs a confirmation rule to reduce the risk that the transfer of the cryptocurrency units could be reversed. A confirmation rule requires that the seller of the fiat or physical good release the item to User i only after a set number of confirmations. A confirmation refers to an additional block being placed on top of the same chain that includes the transaction in which User i sends her cryptocurrency units to the other party. Accordingly, we refer to the expected time for a user to receive her goods starting from the time that her transaction is added to the blockchain as confirmation time, κ, which is given by:

$$\kappa = \frac{k}{\Lambda} \tag{8}$$

with k denoting the number of confirmations. The number of confirmations is a choice variable for the seller of the fiat or physical good, and we subsequently derive its relationship to a blockchain's security level.

To derive the relationship between the number of confirmations and a blockchain's security level, we require some minimal theoretical structure. As such, we let $p(s,k)$ denote

the probability that a transaction entered on the blockchain is reversed with *s* denoting the security level of the blockchain, and *k* denoting the number of confirmations required by the counterparty receiving the cryptocurrency. We assume that $\frac{\partial \mathrm{p}}{\partial s} > 0$ so that a more secure blockchain has a lower probability of a transaction being reversed, i.e., *s* is consistent with our conceptualization of security. Further, we assume $p(s,k)$ decreases in *k*, i.e., a higher confirmation number reduces the probability of a transaction being reversed. This relationship arises endogenously within Nakamoto (2008) and Saleh (2021) for PoW and PoS blockchains, respectively. We also assume that $\lim_{k\to\infty} p(s,k) = 0$ so that the probability a transaction is reversed vanishes as the number of confirmations diverges. Since arbitrarily large confirmation rules require arbitrarily large wait times and thus are undesirable, we assume that the counterparty receiving the cryptocurrency selects *k* to be as low as possible, subject to a risk constraint. Formally, we assume that the counterparty receiving the cryptocurrency requires that the probability a transaction is reversed is no more than some threshold, $\alpha > 0$, and determines the confirmation number by $k = k(s) := min\{ k \in \mathbb{N}: p(s,k) \leq \alpha\}$. Note that this technique resembles the widely-used Value-at-Risk (VaR) approach.

The following result gives the relationship between the endogenous confirmation number and the blockchain's security:

Proposition 3.2. *Confirmation number decreases in Security*

$k = k(s) := min\{ k \in \mathbb{N}: p(s,k) \leq \alpha\}$ *decreases in s for any* $\alpha > 0$.

Proposition 3.1 establishes that the endogenous confirmation number, *k*, decreases in a blockchain's security level, *s*. This result arises because the risk level for the transaction is targeted to be a fixed value, α. The transaction's risk level depends on both the underlying security of the blockchain, *s*, and the number of confirmations, *k*. When the blockchain is relatively secure (i.e., when *s* is high), then even a low confirmation number (i.e., low *k*) implies

a low risk. However, when the blockchain is relatively insecure (i.e., when s is low), then a higher confirmation number (i.e., high k) is necessary to achieve the same level of risk. Thus, $k = k(s) := min\{ k \in \mathbb{N}: p(s,k) \leq \alpha\}$ decreases in s, and we have the following corollary:

Corollary 3.3. *Confirmation Time, κ, decreases in Security, s, and Scale, Λ*

User i's confirmation time, κ, equals $\frac{k(s)}{\Lambda}$ *and thus decreases in both security, s, and scale, Λ.*

Corollary 3.3 highlights that a more secure blockchain generates a higher utility for a user. This result arises because a more secure blockchain has a low probability of having a transaction reversed even with a low confirmation number and thus the confirmation number is set low. Then, since the user wait time increases in the confirmation number and users dislike waiting, a more secure blockchain translates to a higher user utility.

Our preceding results collectively imply the following corollary:

Corollary 3.4. *User utility increases in Adoption,* N_{τ_i}*, Scale, Λ, and Security, s*

User i's utility depends upon the blockchain characteristics of adoption, scale, and security. Moreover, User i's utility increases in each such characteristic.

which establishes that the user utility from using a blockchain depends directly upon three blockchain characteristics: adoption, scale, and security. This result highlights that the welfare implications of blockchain stems from these three characteristics. Accordingly, we use these characteristics as the attributes for our empirical framework in Section 4.

While the utility of each user increases in the referenced three blockchain characteristics, it is important to recognize that users are not homogeneous so that different users may prefer different blockchains. In particular, there is no unique rank of all blockchains for arbitrary users, so we should expect several blockchains to co-exist simultaneously even when each user only transacts on the blockchain that is optimal for her. An important purpose

of our analysis is to determine the subset of blockchains that are optimal for *some* user thereby enabling researchers to focus research efforts on this subset of blockchains. To enable straightforward determination of this subset of blockchains in our subsequent empirical analysis, we offer the following proposition:

Proposition 3.5. *Blockchain Dominance*

Let $N_{\tau,i}$ *refer to the adoption level of Blockchain i for user type* $\tau \in T$, Λ_i *refer to the scale of Blockchain i and* s_i *refer to the security level of Blockchain i. Then Blockchain 1 is preferred to Blockchain 2 for each user under all circumstances if and only if* $N_{\tau,1} \geq N_{\tau,2}$ *for all* $\tau \in T$, $\Lambda_1 \geq \Lambda_2$, *and* $\frac{\Lambda_1}{k(s_1)} \geq \frac{\Lambda_2}{k(s_2)}$.

Proposition 3.5 provides a set of conditions to determine the set of relevant blockchains. In particular, we establish that if Blockchain 1 dominates Blockchain 2 on the stated comparisons (i.e., $N_{\tau,1} \geq N_{\tau,2}$ for all $\tau \in T$, $\Lambda_1 \geq \Lambda_2$, and $\frac{\Lambda_1}{k(s_1)} \geq \frac{\Lambda_2}{k(s_2)}$) then Blockchain 1 is always preferred to Blockchain 2. In turn, Blockchain 2 is then not optimal for any user and thus is also not part of the optimal blockchain ecosystem from the user perspective. In Section 5, we apply these comparisons empirically to eliminate blockchains that would not be optimal for any user, thereby highlighting the public blockchains most relevant for further study.

4. Empirical Framework

In this section, we apply our theoretical framework to data by establishing empirical metrics for each of our three attributes: adoption, scale, and security. We examine and discuss in detail how these attributes vary across a sample of major public blockchains. We find that the leading blockchains, highlighted within Section 2, perform well on at least one of these metrics, emphasizing the underlying economic basis for their relevance.

4A: Adoption

Our theoretical model takes adoption as the number of users of the same type that can transact with each other on a particular blockchain. Intuitively, the probability of a given user finding a suitable counterparty increases with the number of potential such counterparties. Consequently, the more users that transact on a particular blockchain and are active within the relevant user type community, the higher the transaction surplus generated. Accordingly, in our empirical analysis, we consider adoption metrics for multiple user types that each have their own economic purpose for transacting.

First, we use the number of users that have received or initiated a transaction with a blockchain's native cryptoasset on a given day (i.e., we define being active on a day as having been involved in a transaction on a day). We refer to this specific kind of user as a `payment type' user. Second, we consider multiple types of users that seek to engage in more complex transactions than that of the payment type. Users of certain blockchains can make conditional transactions using self-executing programs, i.e., smart contracts, which are coded and run on the blockchain. This functionality allows developers to create decentralized applications (dApps) that mimic relevant economic interactions between users beyond simple payments without the need for a central intermediary. From the perspective of users, dApps represent access to unique networks of other users, bringing efficiency to the matching process. For example, if a user wants to exchange one cryptoasset for another, she searches for a matching set of other users that can provide the liquidity for this exchange. Blockchains with smart contract functionality facilitate the marketplace for such specific economic interactions.

Empirically, for payment type users, we employ as metric the 30 day average of the number of active addresses up to the end of the sample period (February 21, 2021) and rank all blockchains from the coinmetrics.io sample accordingly. Results using this metric are illustrated in Figure 11. Bitcoin (BTC) is the most adopted blockchain with a number of daily

active addresses above one million. There are three other blockchains among the five most adopted blockchains that employ the PoW consensus protocol: Litecoin (LTC), Ethereum (ETH), and Bitcoin SV (BSV). However, the second most adopted blockchain, with almost 600,000 daily active users, is TRON (TRX), which employs DPoS as its consensus protocol.

Figure 11 highlights that there is stark heterogeneity in adoption not just across blockchains but also within consensus protocol groups. For example, PoW protocols' strength in adoption is driven by the dominance of BTC and ETH, while older PoW blockchains such as ZCash (ZEC), Dogecoin (DOGE), or Ethereum Classic (ETC), are lagging behind by orders of magnitude. Nevertheless, some PoW blockchains such as BSV or Bitcoin Cash (BCH), which are lagging far behind BTC in terms of user adoption, are ahead of many other non-PoW blockchains in our sample. Some of these blockchains have been in existence for many years and thus, may possess early-mover advantages, but the success of TRX and ETH on the adoption metric may stem from other factors such as their enhanced functionalities that enable an ecosystem with different types of blockchain users.

To empirically account for the added utility decentralized applications bring to users of different types, we retrieve information on the number of active users of dApps in economically relevant categories from stateofthedapps.com and dappradar.com. We deem a blockchain user address as active on a particular dApp when it has received or initiated a transaction within the last 30 days, i.e., the metric 'Monthly Active Users' (MAU) is the sum of *unique* monthly active users on each dApp across each blockchain.

Table 12 presents summary statistics for dApps across five blockchains with smart contract functionalities and available dApp user data (ETH, EOS, TRON, NEO and

WAVES).[13] There are almost 2,000 dApps hosted on ETH, over 200 on EOS, and less than one hundred each on TRON, NEO, and WAVES, respectively. However, most of these dApps do not have significant user adoption. Less than ten percent of ETH dApps have over 100 MAU and the other four blockchains combined host only 53 dApps with a user base above 100 MAU. With respect to the overall MAU, Ethereum dApp adoption is an order of magnitude higher than all other blockchains, including TRON and EOS.

In terms of different user types, we consider a number of categories that are worth highlighting in terms of the size of their active user base and the distinct economic interactions that take place on respective platforms. *Decentralized Finance (DeFi)* applications bring together users that seek to engage in more complex financial interactions than payments. This includes matching of borrowers and investors as liquidity providers, trading cryptoassets on order books, and the creation of stablecoins and financial derivatives tied to other cryptoassets (see Harvey, Ramachandran, and Santoro, 2021, for an overview of DeFi applications). We combine all users of dApps in the 'Finance' and 'Exchange' categories from stateofthedapps.com to be of the *DeFi* user type. *Gambling* dApps attract types of users that wish to interact within lottery-type structures or tournaments (e.g., dices, sports betting, virtual casino slot machines). Users of the *Gambling* type are typically matched via smart contracts that collate and redistribute cryptoasset funds according to the application-specific (quasi-)random payoff. *Games* dApps serve the purpose of user entertainment and do not necessarily have financial rewards. Users of the *Games* type interact directly with each other in online games, the rules of which are written in smart contracts and any within-game transaction and user interaction is recorded on the respective blockchain. *Storage* dApps provide blockchain-based data services by matching a unique set of consumers with producers of such services.

[13] Note that there are other blockchains that possess some form of smart contract functionality, such as Stellar (XLM) or Ripple (XRP), but none of these blockchains host widely-adopted, decentralized applications and are thus, not included in this subsample.

Users of the *Storage* type interact with each other on dApps that facilitate, e.g., digital file sharing or provision of cloud storage, but in a decentralized setting where incentives to participate in respective peer-to-peer marketplaces are defined in smart contracts. We combine all remaining dApp categories under the umbrella category *Other*.

Figure 12 shows the distribution of the number of dApps and MAU across these categories (combined for all five blockchains) and reveals that the majority of dApp adoption stems from the growing relevance of DeFi within the blockchain ecosystem. Although *DeFi* dApps constitute only 22.4% of the total number of dApps in our sample, the user base of these represent 70.2% of all dApp users. The disparity between the proportion of total dApps and the proportion of total users for each category exists not only for *DeFi*, but also for other categories. *Storage* dApps account for only 1.6% of all dApps but the users of these make up 3.8% of total users. In contrast, *Games* and *Gambling* dApps account for 30.1% of total dApps but only 6.4% of all users have adopted them.

Table 13 provides an overview of the distribution of MAUs within each of the five user types (dApp categories) across the five blockchains. While ETH dominates the *DeFi*, *Games*, and *Other* categories, it does not lead in all categories. TRX leads in the *Gambling* and *Storage* categories, capturing more than 50% and 90% of the total user base within each of those categories, respectively. EOS falls slightly behind ETH in the *Games* category and is otherwise dominated by either ETH or TRX in terms of users in every other category. The user bases of NEO and WAVES dApps fall far behind that of the other three blockchains in all relevant categories.

4B: Scale

Scale determines user welfare both through lowering the processing time as well as the costs users incur in the form of fees paid (see Proposition 3.1). Within our theoretical analysis,

scale corresponds to the rate at which a blockchain can process users' submitted transactions. We estimate a blockchain's transaction rate in terms of recorded transactions in a given time period (denominated in seconds) and employ the maximum of historical daily transaction counts as our scale metric. Our scale metric provides us with a conservative estimate of the transaction rate a blockchain can realistically sustain over longer periods (e.g., more than just one block), is easy to implement empirically, and allows straightforward comparison across blockchains and protocols.[14] Figure 13 displays the scale of each blockchain in the sample.

Figure 13 highlights that, although BTC leads on adoption by *Payment* type users, it does not perform well on scale as it records fewer than 10 transactions per second (TPS). Consequently, users that require higher scale for fast processing times or the ability to send micro-transactions at low transaction costs would opt for alternatives to BTC. The DPoS blockchain EOS performs substantially better than all other blockchains, even other DPoS blockchains, and can sustain an estimated transaction processing rate of 122.6 TPS. Stellar (XLM), a nonstandard blockchain, is second on scale with almost 95 TPS, while the blockchains ranked third to fifth on scale can process only at approximately half of the rate of EOS. These findings highlight that public blockchains have made significant progress since the birth of Bitcoin when it comes to scale, but remain inferior to traditional alternatives (e.g., payment systems like VISA can process over 1,000 TPS).

4C: Security and Confirmation Time

For our security metric, we rely upon Proposition 3.2, which implies that the blockchain ranking by security is the reverse of the blockchain ranking by confirmation numbers (i.e., if $k \leq k'$, then $s \geq s'$). While we cannot directly observe the security of a blockchain, we can

[14] Although there is a theoretical risk of underestimating scale for those blockchains that have high capacity to process transactions quickly but have never experienced large demand for transacting, this risk is unlikely to significantly affect our main results. In particular, the only blockchain in our sample with little transaction history is Polkadot, and the ranking of our blockchain sample according to our scale metric is largely consistent with reported figures in the crypto press.

infer its level by considering the confirmation rules imposed by major users of blockchains. In particular, we collect information on confirmation numbers disclosed by the largest institutions that trade via blockchains: cryptoasset exchanges. We collate a list of major cryptoasset exchanges from cryptocompare.com. We consider all exchanges with daily trading volume above $1 million (n=122; as of March 25, 2021). For each exchange, we search for and browse through its FAQ and support web pages and are able to extract the required number of block confirmations for a particular cryptoasset from a total of 32 cryptoasset exchanges.[15] Different exchanges may view the security levels of the same blockchain differently and thus, could impose different confirmation rules. For example, every exchange has a confirmation rule for BTC transactions, but some exchanges require up to six confirmations for BTC transactions while others require only one or two. To infer an order of security levels *s* for each blockchain in our sample, we take the median confirmation number $k(s)$ for a given blockchain to aggregate across exchanges.

Figure 14 reveals that XLM is the most secure blockchain, as it has a required confirmation number of $k(s) = 1$, followed by BTC with $k(s) = 2$. Blockchains with Nonstandard consensus protocols, such as XRP or NEO, as well as the Hybrid PoW/PoS protocols employed in Decred (DCR) and Dash (DASH) perform comparatively well on security with confirmation numbers below $k(s) = 10$. While this is also true for older PoW blockchains such as Litecoin (LTC), Bitcoin Cash (BCH) or Dogecoin (DOGE), there is much more heterogeneity among the group of PoW protocols, with exchanges requiring far beyond 500 required confirmations for the least secure blockchains in order for transactions to be considered final. With the exception of Binance Coin (BNB), which has a median confirmation

[15] For transparency, Table 16 lists all confirmation rules used in constructing our security metric and Table 17 lists web links to data sources.

number of $k(s) = 5.5$, all PoS and DPoS blockchain transactions require over twelve confirmations and thus place in the lower half of the blockchain sample in terms of security.

Corollary 3.3 directly states that user welfare increases with lower confirmation times, κ, which in turn depend upon confirmation numbers (security), $k(s)$, and the rate at which blocks arrive. We estimate block rates using historical data on block arrivals. For each blockchain, we take the average of daily block counts to estimate its block arrival rates as the time in seconds for one additional block to be added to the blockchain.[16] Table 14 shows block arrival rates as well as confirmation numbers that are used to calculate user confirmation times κ.[17]

The best performing blockchains on this metric employ either Nonstandard protocols or consensus protocols based on delegated Proof-of-Stake. Binance Coin blockchain transactions take an estimated $\kappa = 2.2$ seconds to be considered final by cryptoasset exchanges, making it the blockchain with the shortest confirmation wait time. The two blockchains ranked second and third with respect to this metric are Stellar and EOS, with confirmation times $\kappa = 5.2$ and $\kappa = 6.8$ seconds, respectively, which is more than twice as long as is required for Binance Coin. All of the top five blockchains with respect to confirmation times have short block arrival rates and low confirmation numbers, resulting in overall wait times until final confirmation below one minute, whereas the majority of other blockchains require more than five minutes to complete the confirmation process in an exchange. The set of blockchains performing worst on this metric includes known PoW blockchains such as Bitcoin SV and Ethereum Classic, which exhibit estimated confirmation times required by cryptoasset exchanges of two hours or more. It is worth highlighting that

[16] Block counts are missing in coinmetrics.io for Polkadot (DOT), but we obtain an estimate of its block rate (6 seconds) from the polkascan.io blockchain explorer.
[17] We also consider the mean of confirmation numbers across exchanges as an alternative measure to calculate confirmation times, but the ranking is quantitatively and qualitatively similar to taking the median values.

Ethereum, which captures a considerable amount of the activity in the ecosystem in terms of active users of decentralized applications and payments, is more desirable than Bitcoin in terms of confirmation times, as exchanges involving ETH are settled in a quarter of the time it takes to settle those using BTC (5 versus almost 20 minutes). Thus, although Bitcoin is viewed as relatively secure by exchanges in terms of required confirmation numbers, user utility, ceteris paribus, is higher for other blockchains that have much shorter confirmation times.

5. The Blockchain Frontier

Given Proposition 3.5 and our empirical implementation, we turn to identifying which blockchains are worthy of study as those that are not dominated by any other blockchain with respect to the three attributes. The underlying reasoning for the focus on these blockchains is that, by definition, only such blockchains could provide the maximum utility among all blockchains for some user. To empirically identify whether a blockchain is dominated by another blockchain, we perform pairwise comparisons of confirmation time, scale and adoption metrics to check whether the conditions for dominance in Proposition 3.5 hold. For that purpose, we rank all sample blockchains from 1 (best) to 27 (worst) with respect to each of the empirical metrics and summarize the results in Table 15.

We extract a subset of seven blockchains that are not dominated by any other blockchain, which we refer to as the ``blockchain frontier." The blockchain frontier consists of not only those blockchains that tend to have their relative strength only in adoption, confirmation time (security), or scale, but also blockchains that exhibit a unique combination of attributes, which results in them not being dominated overall.

The financial economics literature has largely focused on the specifics of the Bitcoin blockchain, but our findings suggest that other, less-known blockchains dominate on economically relevant dimensions, such as scale, confirmation time, or adoption within specific user type groups (e.g., *DeFi*). Notably, this frontier does not necessarily list the blockchains with highest market capitalization of its native cryptoasset, but rather the blockchains which bring some unique utility to users. Thus, the frontier translates to the list of blockchains that are optimal for at least *some* users even if that quantity of users is small. We provide an overview of blockchain frontier members below, along with specific examples of users who may prefer the economic attributes of each blockchain.

Bitcoin (BTC): Bitcoin is by far the most adopted blockchain in terms of active users of the *Payment* type. Because it ranks first in this category, Bitcoin cannot be dominated by any other blockchain. Bitcoin benefits from the largest network of other users, facilitating the matching process between these users of the payment type. For example, a user seeking to purchase any goods or services using a cryptocurrency would find it easier to match with a merchant willing to accept BTC, relative to other native cryptoassets. Moreover, Bitcoin's widespread usage renders it as relatively liquid compared to other cryptoassets, and thereby makes it the preferred choice for agents taking a trading position in cryptoasset markets in general. This point can be seen most clearly by the fact that BTC has the highest overall trading volume on cryptoasset exchanges (see, e.g., cryptocompare.com).

Despite its dominance on the adoption by user of the *Payment* type, Bitcoin falls behind on other dimensions such as scale and confirmation time, where it ranks only 11th and 15th, respectively, and it does not support adoption of other user types (e.g., DeFi). Because of these shortcomings, there are multiple blockchains that are not dominated by Bitcoin.

Ethereum (ETH): Ethereum enters the frontier as it leads on *DeFi* user adoption, which exhibits the highest relative importance within the current dApp ecosystem, but it also wins in several other categories such as *Games*. This dominance in terms of *DeFi* user adoption means that users with a preference for more complex but decentralized financial services will find a more suitable counterparty for their needs by using Ethereum over all other blockchains.

An example of such a *DeFi* dApp is Uniswap, which is a decentralized exchange, allowing users to exchange one cryptoasset (e.g., any 'ERC-20' token) for another without an intermediary. For the exchange to function without a centralized market maker, liquidity is provided by the users that are willing to offer specific cryptoassets in exchange for a representative share of a pool of assets. Other users that want to exchange pairs of cryptoassets can do so at a price, which is set by a preset algorithm on the basis of the size and structure of cryptoassets in the liquidity pool. Each trade requires users to pay a flat transaction fee that in turn is distributed among the liquidity providers.

Nevertheless, Ethereum falls behind Bitcoin and TRON in adoption by users of the *Payment* type and it even falls behind TRON in adoption with respect to users interested in gambling and data storage services. The Ethereum blockchain also does not scale well and requires a longer confirmation time for ETH transactions, making it inferior compared to other blockchains when it comes to speed and finality of transactions.

TRON (TRX): The TRON blockchain performs well on all attributes: it is widely adopted by many user types, it has high scale and requires short confirmation times. It does, however, not dominate all other blockchains as there are only a few metrics on which TRON ranks first. The two metrics in which TRON ranks first are adoption by users of the *Gambling* and *Storage* types, i.e., there are more active users of dApps in these categories than on any other blockchain.

An example of a successful dApp for *Storage* type users is the BitTorrent Speed TRON dApp, which enables users to obtain from or send files to other users. The BitTorrent Speed TRON dApp thus facilitates a peer-to-peer marketplace for digital media downloads. Users that offer media files for other users to download are rewarded in the form of platform-specific BitTorrent (BTT) tokens, which in turn can be sold in markets or redeemed within the dApp to download files at higher speeds. This dApp is successful only because it is sufficiently widely adopted by *Storage* type users to generate a large media library for downloads.

In addition to its relative strength in scale, confirmation time, and specific user type adoption, TRON ranks second behind Bitcoin in terms of *Payment* user adoption. However, TRON falls behind Ethereum in the important *DeFi* category.

EOS (EOS): EOS is the blockchain with the highest scale and is therefore included in the frontier. Users who wish to have their transactions quickly recorded on a blockchain at a low fee cost prefer a higher scaling blockchain. For instance, this may include users who wish to use a decentralized application that involves frequent micro-transactions. Apart from scale, EOS has the third shortest confirmation time, which adds to the utility of users that prefer not only fast processing times, but also finality of transactions after only few seconds. Related to this point is the fact that EOS ranks second on adoption by users of the *Games* type, i.e., there are many active users employing the EOS blockchain to interact via gaming dApps for entertainment.

A prime example of a dApp that involves many micro-transactions is Upland, which is a *Games* dApp on the EOS blockchain. In the Upland dApp, users buy, sell, and collect virtual properties that resemble real estate in specific cities such as San Francisco or Manhattan. User actions extend beyond purchase and sale of properties; in fact, users must travel nearby in the

virtual world to purchase properties so that the game involves frequent small transactions associated with traveling. Any such interaction within the game requires a transaction to be recorded on the blockchain and thus Upland players require low fees to make playing the game feasible. As shown by Proposition 3.1, transaction costs decline in scale so that the high scale of EOS makes its particularly appropriate for such dApps.

Note that despite its success on scale and confirmation time, EOS is not a blockchain widely adopted by *Payment* type users and thus, it does not dominate the majority of the blockchains in our sample on all of the metrics considered.

Binance Coin (BNB): Transactions on the Binance Coin (BNB) blockchain have the shortest confirmation times of all blockchains in our sample and thus, the BNB blockchain is part of the frontier. However, it neither ranks highly on *Payment* user adoption nor is particularly fast in terms of scale. This means that BNB user utility is highest for those types of users that require fast confirmation of transactions above all, but do not require the same scale as blockchains like EOS.

A primary application for BNB is in trading on the Binance cryptoasset exchange. In particular, the exchange offers discounts when BNB is used to pay for fees incurred on their platform. Thus, trading possession of BNB itself is important for any trader seeking to trade on the Binance cryptoasset exchange. As BNB possesses the shortest confirmation time, it is particularly well-suited for this purpose.

Stellar (XLM): Stellar is ranked second on confirmation time and scale, respectively. It is not dominated by the most scalable blockchain, EOS, as it performs better on confirmation time, and it is not dominated by the blockchain with the shortest confirmation time, BNB, as it scales better.

The Stellar blockchain is primarily used by financial technology companies, such as digital wallet, trading and payment providers, which transact using the Stellar blockchain on behalf of their customers. For example, TEMPO, a European electronic payment company, offers its customers global remittances denominated in EUR. To do so, they issue a stablecoin, EURT, upon the Stellar blockchain, which is backed by actual EUR fiat currency stored in the company's audited bank accounts. As soon as the digital EUR representation exists, it can be transferred to the counterparty as a transaction upon the Stellar blockchain, accompanied by small a transaction fee paid in XLM.

Given the type of network participants, it is not surprising that Stellar is not among the top five blockchains in terms of *Payment* type user adoption, but rather, it relies on secure and fast transactions between fewer entities (e.g., financial technology companies on behalf of their customers). In typical applications, including the one referenced above, Stellar benefits from shorter overall processing and confirmation times than both traditional and other blockchain-based, alternative payment systems.

Bitcoin Satoshi's Vision (BSV): Bitcoin SV enters the blockchain frontier not because it leads on any particular attribute, but it exhibits a combination of high scale and wide adoption by *Payment* type users. It ranks only third on scale, but the two blockchains that perform better on scale do not dominate Bitcoin SV as they perform comparatively poorly on *Payment* user adoption, where Bitcoin SV ranks high. Given this combination of attribute strengths, Bitcoin SV users prefer a higher scale than Bitcoin, i.e., more transactions can be processed at lower fee costs in the same time period, but also require a certain level of adoption for payment services, which blockchains with even higher scale do not offer. That is, some users value the Bitcoin SV blockchain because of this trade-off between *Payment* user adoption and scale.

6. Conclusion

We provide the first extensive overview of the public blockchain ecosystem and find that although there exist hundreds of blockchains, only a few are of economic relevance. We offer a simple theoretical framework that helps explain and deepen that finding. Our framework suggests that three blockchain characteristics drive user utility - scale, security, and adoption - and that only a few blockchains are not dominated by some other blockchain in terms of utility for some user. We empirically identify an exclusive set of seven such blockchains, and we refer to those blockchains as the blockchain frontier.

An important takeaway from our analysis is that there are several blockchains other than Bitcoin that are of economic relevance. Accordingly, our work highlights the need for more research in financial economics that is focused on alternatives to Bitcoin.

Chapter 3 Cryptocurrency Mining: Asymmetric Response to Price Movement

1. Introduction

Cryptocurrencies, as blockchain technologies, have significant implications for financial markets, innovation, transactional security and more. Since the introduction of Bitcoin (Nakamoto, 2008), the cryptocurrency market has grown substantially, reaching a combined market capitalization of over 390 billion USD for all cryptocurrencies.[18] Cryptocurrency miners play a critical role in this ecosystem, as they verify blockchain

[18] As of November 3, 2020. See https://coinmarketcap.com/charts/.

transactions and provide security to ownership of blockchain-based assets (Easley, O'Hara, and Basu, 2019). In this paper, I investigate the entry and exit decisions of cryptocurrency miners as it relates to the type of computing equipment they use.

Cryptocurrency miners face high fixed-costs associated with the computing equipment required to mine. This investment in computing equipment may be partially reversible if it has value in applications other than mining. The reversibility of these fixed (initial investment) costs has a direct effect on the option value of cryptocurrency mining. Miner reactions to price changes between can be explained using an application of real options theory (Dixit, 1989; Dixit and Pindyck, 1994). BTC miners use Application Specific Integrated Circuit (ASIC) mining equipment, which is built for running a specific algorithm (in the case of Bitcoin this is the SHA-256 algorithm) and has little use outside a mining application. On the other hand, ETH miners typically use Graphics Processing Unit (GPU) equipment, which has a variety of applications including video processing, editing, gaming and machine learning. Therefore, when scrapping computing equipment, GPU miners may enjoy a higher salvage value than ASIC miners. Between entry and exit thresholds, established by non-reversible costs, individual ASIC and GPU miners behave in a similar fashion. In particular, price increases encourage miner production and decreases in prices discourage mining. However, the contrasting salvage value of computing equipment of each set of miners leads to distinctive entry and exit decisions. Consequently, ASIC miners, in aggregate, often have asymmetric reactions to price changes, while GPU miners display more symmetric reactions to price changes.

Confirming the predictions of real-options theory, I use a Threshold Vector Error Correction Model (TVECM) (Balke and Fomby, 1997; Hansen and Seo, 2002), to empirically estimate cryptocurrency miners' adjustment to computing power (hashrate) to disequilibria in

the long-run relationship between price and hashrate. The TVECM specification is particularly useful in this application as it allows for separate error correction terms (mean-reversion coefficients) for positive and negative deviations from the long-run equilibrium, and thus highlights any potential asymmetry in miners' adjustments to prices. I find strong evidence that Bitcoin (BTC) miners have an asymmetric response to price shocks; they increase hashrate following positive price shocks, but do not decrease hashrate following similar negative price shocks. Meanwhile, Ether (ETH) miners have a symmetric response to price shocks, increasing hashrate following positive shocks and decreasing hashrate following negative shocks. Consistent with this notion I show that after the advent of ETH ASIC equipment, ETH miners also have an asymmetric response to disequilibria.

The behavior of cryptocurrency miners is of great interest as the functionality of the blockchain and the value of the cryptocurrency depend on the process of mining. Therefore, these results have various practical and theoretical implications for cryptocurrencies and blockchain applications. First, an asymmetric miner adjustment to disequilibria (born of price shocks) is an important consideration for cryptocurrency pricing and valuation. Many have suggested that cryptocurrency/blockchain specific characteristics are good predictors of returns/growth (Bambhwani, Delikouras, and Korniotis, 2019; Biais, Bisiere, Bouvard, Casamatta, and Menkveld, 2018; Liu and Tsyvinski, 2020).

Second, miner adjustment to disequilibria has implications for the security and functionality of blockchain applications. Yermack (2017) suggests that blockchain technology offers more transparent and lower-cost record keeping in a variety of applications. Concerning financial assets, blockchain applications are no longer a theoretical construct. For example, the Australian Securities Exchange (ASX) is developing a distributed ledger using blockchain technology to replace its current settlement system. According to an ASX press release, this

system will be released in March or April 2021.[19] For financial record keeping, security and transparency are of the utmost importance. As miners provide this security in a blockchain-based ledger, their behavior when faced with price volatility affects the functionality of the application.

The rest of the paper is outlined as follows: Section 2 provides background information about cryptocurrency mining and contains a brief literature review. Section 3 models the activation and abandonment decisions of cryptocurrency miners. Section 4 presents a summary of the data. Section 5 describes the methodology. Section 6 presents the results and econometric interpretations from my main specifications. Section 7 presents a discussion of these results and their implications. Section 8 concludes.

2. Background and Literature Review

2A: Background

Those who engage in the creation of new cryptocurrency on an existing blockchain are called miners. In blockchains that use Proof-of-Work (PoW) consensus protocol, they "mine" by solving a complex mathematical problem (determined by the currency's algorithm) and verifying transactions. The first miner to solve this algorithm proposes the next "block" on the blockchain and is rewarded with an amount of the underlying cryptocurrency (the block reward plus transaction fees). The proposed solution is easily verified by other miners, thus it constitutes "proof" that the miner properly extended the blockchain. This paper is concerned primarily with PoW cryptocurrencies, specifically Bitcoin (BTC) and Ether (ETH). Other consensus algorithms exist, such as Proof-of-Stake (PoS), but the underlying assets have not

[19] See https://www.asx.com.au/documents/about/MediaRelease-ASX-DigitalAsset-and-VMware-join-forces-on-DLT.pdf.

yet reached the scale of PoW cryptocurrencies- see Salah (2020) for an economic model of PoS consensus.

The computing power a miner provides to the cryptocurrency network is called a hashrate (generally measured in terahashes/second). Hashrate, in a broad sense, can be viewed as a function of the resources invested by miners (Bhambhwani, Delikoras, and Korniotis, 2019). For example, Saleh (2021) finds that hashrate is positively associated with the electricity consumption of miners. Hashrate is also a positive function of a miner's expected reward in cryptocurrency. If a miner doubles their hashrate, their expected return in cryptocurrency doubles assuming a constant mining difficulty.

Mining difficulty (commonly called difficulty) refers to the complexity of the task miners must complete to propose a block. The difficulty is adjusted based on the blockchain's underlying algorithm to control the speed at which new blocks are "discovered" by miners. Thus, a hashrate that produces a certain amount of BTC today is not guaranteed to produce the same amount of BTC a year from today.

Engaging in the mining process, a miner faces fixed and flow costs. Dominant among these are the computing equipment required to mine (a fixed cost) and the electricity required to run and cool this computing equipment (a flow cost). It is important to note that a given computing unit has a fixed hashrate at which it can mine. To scale hashrate, more computing equipment must be acquired.

The computing equipment used to mine differs between cryptocurrencies based on the blockchain's underlying algorithm. Bitcoin, using the SHA-256 algorithm, is mined using Application Specific Integrated Circuit (ASIC) miners. Ethash, the algorithm for Ethereum, is

ASIC-resistant, so ETH is mostly mined using consumer graphics processing units (GPUs).[20] As its name suggests, ASIC mining equipment is built specifically for mining a particular algorithm and has little use outside this application. GPUs, on the other hand, are commonly found in in consumer computers and are used for many other purposes. The disparity in outside use or salvage value between ASIC and GPU equipment motivates the hypothesis of this paper.

Given the costs and benefits in engaging in mining cryptocurrencies, there should exist a long-run equilibrium price/hashrate relationship. If a cryptocurrency's price increases, all else equal, mining this cryptocurrency becomes more profitable and the value of increasing hashrate (via purchase of more computing units) increases. If the price decreases, mining becomes less profitable and the value of increasing hashrate decreases. Thus, the aggregate hashrate of cryptocurrency should be sensitive to that cryptocurrency's price. This is consistent with prior literature (Pagnotta and Buraschi, 2018). This relationship is further described in Section 3.

2B: Literature Review

My findings are unique to the literature and suggest a long-run equilibrium price/hashrate relationship dependent on the computing equipment used to mine cryptocurrency. This has implications for cryptocurrency valuation and the use of blockchain technology. Here, I survey the literature surrounding cryptocurrencies and blockchain applications as it relates to this research.

Closely related to my work are papers considering issues surrounding the cryptocurrency mining process. Easley, O'Hara, and Basu (2019) model the strategic behavior of miners, focusing on the impact of Bitcoin transaction fees to miners and users. Various

[20] ASIC miners for the Ethash algorithm do exist, the first of which was developed by Bitmain, the Antminer E3. This hardware was announced in April, 2018 and the first batch was shipped in July, 2018 (see https://cointelegraph.com/news/bitmain-releases-ethash-asic-miners). Comparatively speaking, this hardware has not affected the aggregate hashrate as much as the introduction of SHA-256 ASIC miners; these ASIC miners have not immediately dominated the use of GPU mining equipment. This advent of this technology is considered in my empirical analysis (see Section 6.2).

others use a game theoretic framework to model Bitcoin mining (Dimitri, 2017; Ma, Gans, and Tourky, 2018). Cong, He and Li (2021a) study the impact of decentralized miners using centralized mining pools. Kristoufek (2020) uses a Vector Error Correction Model (VECM) to consider the long run equilibrium dynamics between Bitcoin price and the operational cost of mining.

Similar to my threshold analysis (Section 3), Prat and Walter (2018) study the relationship between BTC price and hashrate in the face of irreversible investment costs. Though related, my work is fundamentally distinguishable from this; Prat and Walter (2018) assume that investment in mining is entirely irreversible and that mining units cannot be switched off to save electricity. Through relaxing these assumptions, I show that miners could have a fundamentally different reaction to positive and negative disequilibria in the price/hashrate relationship based on the salvage value (i.e. reversibility) of their equipment (see Section 3). I also present novel empirical evidence supporting this phenomenon.

As mentioned, my results have implications for cryptocurrency asset pricing. The literature surrounding this has evolved in recent years. Cheah and Fry (2015) posit that Bitcoin has zero fundamental value. Baur, Hong, and Lee (2018) find that Bitcoin is largely uncorrelated with traditional asset classes. Liu and Tsyvinski (2020) find similar results and suggest that returns can instead be predicted by cryptocurrency specific factors such as proxies for investor attention, price-dividend ratio, realized volatility, and supply conditions. Consistent with this, Bambhwani et al. (2019) find that hashrate and network size are significant pricing factors for cryptocurrencies. Transactional benefits/demand also can predict a fraction of cryptocurrency returns (Biais, Bisiere, Bouvard, Casamatta, and Menkveld, 2018; Cong, Li, and Wang, 2021b).

Of equal interest is research that considers the realized and possible applications of blockchain technology. Abadi and Brunermeier (2019) compare the usefulness of blockchains for record keeping with traditional centralized ledgers. Catalini and Gans (2020) argue that the cost of verifying the digital state of an asset is cheaper on blockchain, as centralized intermediaries retain significant market power and charge fees. They also claim the cost of scaling and operating blockchains is lower, citing the competitive mining landscape and low barriers to entry (Catalini and Gans, 2020). Cong and He (2019) suggest that blockchain smart contracts can reduce information asymmetry and improve consumer surplus, as blockchains have lower barriers to entry and facilitate more competition than traditional brokers. Yermack (2017) studies various potential applications of blockchain technology and discusses their implications for corporate governance. He praises blockchains for offering low-cost, transparent record keeping.

Finally, I consider research that provides economic interpretations of various peculiarities of blockchains and cryptocurrencies. Foley, Karlsen, and Putniņš (2019) study transactional data, and estimate that roughly 46% of Bitcoin transactions involve illegal activity. Makarov and Schoar (2020) identify cross-country arbitrage opportunities across various cryptocurrency exchanges, which are exacerbated in times of large Bitcoin appreciation. They also suggest that trading volume predicts a significant percent of Bitcoin returns (Makarov and Schoar, 2020). Salah (2020) studies possible equilibria arising from PoS consensus protocols. Sockin and Xiong (2020) model utility tokens (such as Ether, Filecoin, and GameCredits) as membership to a digital platform to facilitate transactions and as a speculative investment.

3. Entry and Exit Decisions of ASIC and GPU Miners

It is the hypothesis of this paper that ASIC miners and GPU miners often have differentiated responses to price changes based upon the salvage value of their computing equipment. This section of my paper details the entry and exit decisions of miners subject to stochastic cryptocurrency prices in order to establish a framework for the empirical testing of the paper's main conjecture.[21]

3A: Combined Entry and Exit Strategies

The miner has the choice to be either idle (0) or active (1). I denote $V_0(P)$ as the value of the option to invest (the value of an idle miner) and let $V_1(P)$ denote the value of an active miner. $V_1(P)$ is the sum of two components, the entitlement to profit from operation and the option to abandon should the price fall too low.

If the miner invests, that is enters the market, he obtains a project that produces one unit of output per period.[22] In this illustrative example, the project lasts forever or until his computing equipment is abandoned. For transparency, the cost of operation "*C*" is known and constant. The riskless rate of return is "*r*." The miner must incur a lump sum "*I*" to participate in mining and may or may not face an abandonment cost "*E*" if the project is abandoned. I allow for negative abandonment costs, representing a positive salvage value, in my analysis.

[21] This research makes no claim to the development of real options pricing theory, but rather it is an application of the technique to Bitcoin and Ether mining. The two-threshold model presented below is derived primarily from Dixit (1989), and the four-threshold model is derivative of Dixit and Pindyck (1994). Dixit and Pindyck (1995) forced me to acknowledge the shortcomings of net present value analysis when miners make decisions about expansion and abandonment. This concept is also addressed by Brennen and Schwartz (1985) and Dixit (1989). The value of waiting to invest or abandon is outlined by McDonald and Siegel (1986) and Dixit (1992). Pindyck (1991) explicates the impact of uncertainty and sunk costs upon firm entry and exit thresholds, and I interpret miner behavior with the framework he provides.

[22] For simplicity, I assume that mining difficulty is fixed for our decision horizon. In reality, mining difficulty is regularly adjusted, following procedures based on the blockchain's underlying code. The purpose of this adjustment is, in fact, to regulate the supply of new cryptocurrency so it remains relatively constant. Though this assumption is unrealistic, it allows for a more straightforward presentation of applied real options theory.

The simultaneous consideration of entry and exit means that the active miner enjoys a revenue stream and the option of abandonment, while the idle miner has the option of becoming "live" again. Thus, the value of an active miner and an idle one must be determined simultaneously.

An active miner will remain active as long as $P > P_S$ and an idle miner will remain idle if $P < P_H$, where P_H and P_S are thresholds determined by the following optimization process. The value maximizing strategy between the thresholds P_S and P_H is to continue the status quo.

Over the range of prices $(0, P_H)$ an idle miner retains his option to become active. Non-arbitrage conditions dictate a $V_0(P)$ which satisfies a differential equation over the price interval detailed above. The boundary conditions link the values and derivatives of $V_0(P)$ to those of $V_1(P)$ at P_H. Similarly, from P_S to P_∞, an active miner remains active. $V_1(P)$ satisfies a corresponding differential equation and its boundary conditions link the values and derivatives of $V_1(P)$ to those of $V_0(P)$ at P_S. This system of equations contains just enough information to complete the miner's two threshold system of decision making.

At the threshold P_H, the idle miner pays a lump sum "*I*" to exercise his investment option giving up $V_0(P)$ to get a "live" project which has a value $V_1(P_H)$, so that value matching would suggest that:

$$V_0(P_H) = V_1(P_H) - I \tag{9}$$

and, as a consequence of smooth pasting, we have:

$$V_0'(P_H) = V_1'(P_H). \tag{10}$$

Likewise, at the abandonment threshold, value matching maintains:

$$V_1(P_S) = V_0(P_S) - E$$

(11)

where E is the cost of abandonment, so that when an active miner abandons his project he receives $V_0(P_S)$ and pays an abandonment cost. The smooth pasting condition insists that:

$$V_1'(P_S) = V_0'(P_S).$$

(12)

Using the general solutions for $V_0(P)$ and $V_1(P)$ (see Appendix B for derivation), I express:

$$V_0(P) = A_1 P^{\beta_1}$$

(13)

and

$$V_1(P) = B_2 P^{\beta_2} + \frac{P}{\delta} - \frac{C}{r}$$

(14)

where A_1 and B_2 are option value coefficients, and δ is the convenience yield. Substituting these solutions into Equations (9) through (12) provides a system of four non-linear equations with four unknowns: P_H, P_S, A_1, and B_2. A closed form analytic solution is not possible for this highly non-linear system of equations. However, it can be proven that a solution exists, it is unique, and has intuitive economic properties. For instance, $0 < P_S < P_H < \infty$ with A_1 and B_2 being positive.

3B: Amending the Miner's Two Threshold Model with Mothballing and Reactivation

In the face of deteriorating prices, a miner has options other than permanently abandoning his computing equipment. In particular, the miner can put his equipment into a state of temporary suspension, allowing it to be reactivated in the future at a sunk cost that is

less than the start-up cost of mining. Temporary suspension includes a lump sum cost E_M and a flow cost *M*, while the miner would incur a one-time reactivation cost of *R*.

Clearly, the economics of the situation demand that *M* be less than *C*, the per unit operating costs. For reactivation, the price must rise to P_R where $P_R < P_H$. "Mothballing" or "lay-up" occurs at P_M where $P_M > P_S$. Recall that P_S is the price where a miner abandons his computing equipment in its entirety. For simplicity, I will assume that $E = E_M + E_S$, so that the cost of directly abandoning an active project, *E*, equals the sum of fixed costs of mothballing, E_M, plus the cost of abandoning the "laid-up" equipment, E_S.

The thresholds explicated above (P_H, P_R, P_M, and P_S) are endogenous and, clearly, an artifact of the parameters describing the economic circumstances of the miner. If the miner's maintenance cost *M* is sufficiently high or the reactivation cost *R* is not sufficiently less than the full investment cost, then optimality conditions may determine that it is better for the miner to abandon the computing equipment at P_L without considering a P_M threshold.

Recall that the miner can remain idle over the range of prices $(0, P_H)$ so that his value in this situation is given:

$$V_0(P) = A_1 P^{\beta_1}.$$

(15)

Furthermore, the value of an operating miner is given by:

$$V_1(P) = B_2 P^{\beta_2} + \frac{P}{\delta} - \frac{C}{r}$$

(16)

for prices in the interval (P_M, ∞) when the possibility of mothballing is acknowledged.

The mothballed miner can continue over the range of prices (P_S, P_R) and the value of the "laid-up" computing equipment would be given:

$$V_M(P) = D_1 P^{\beta_1} + D_2 P^{\beta_2} - \frac{M}{r}$$

(17)

where $D_1 P^{\beta_1}$ is the value of the miner's option to reactivate and $D_2 P^{\beta_2}$ is the value of the miner's option to scrap the equipment. Clearly M/r is the cost of maintaining the laid-up equipment forever.

The switching points or the value matching points of the four-threshold model is given below by equations (18)-(21):

$$V_0(P_H) = V_1(P_H) - I$$

(18)

$$V_1(P_M) = V_M(P_M) - E_M$$

(19)

$$V_M(P_R) = V_1(P_R) - R$$

(20)

$$V_M(P_S) = V_0(P_S) - E_s$$

(21)

where Equation (18) recognizes the price at which an idle miner becomes active. Equation (19) acknowledges the price at which an active miner decides to mothball his equipment. Equation (20) details the price at which the mothballed miner reactivates his equipment. Finally, Equation (21) describes the price, P_S, at which the mothballed miner completely abandons his project.

Evaluating equations (18) to (21) at their relevant switching points P_H, P_R, P_M, P_S and using their smooth pasting conditions yields eight non-linear equations that can be used to solve for four threshold prices and the four option value coefficients. (See Appendix B for the

substitution of the general solutions for $V_0(P)$, $V_1(P)$, and $V_M(P)$ into Equations (18) to (21) and their respective smooth pasting conditions).[23]

3C: Comparative Statics

These non-linear equations evade analytical solutions, but numerical results for the endogenous variables are easily obtained. In order to illustrate the comparative static behavior of the four threshold prices, consider a highly stylized Ether miner whose desired level of production is 1,000 ETH per year. To accomplish that objective the requisite level of investment would be, approximately, a quarter of a million dollars, that is $I = \$250{,}000$. Annual operating expenses would be around $\$80{,}000$. Let the flow cost of mothballing be $\$1{,}000$/year and a starting point for the fixed cost of reactivating a "laid-up" project would be $\$40{,}000$. Furthermore, the standard deviation of the geometric brownian motion that represents the behavior of prices is given a base value of $\sigma = 20\%$. The miner enjoys a risk-free rate of $r = 4\%$ and has a convenience yield of $\delta = 10\%$. Finally, abandonment costs for the miner's equipment are given an initial value of $E_S = E_M = 0$.

Examining panel 1 of Figure 15, it is clear that as uncertainty, σ, increases, the zones of the miner's inactivity between entry and exit increases. The area of inactivity between P_H and P_S increases as σ increases because the profitability of becoming active or becoming inactive is more likely to be reversed by changes in price due to the increase in the standard deviation of the price process. In regard to mothballing, if the miner has "laid-up" his computing equipment, then clearly P_R increases with increases in σ. The increased volatility

[23] My analysis above assumes that the miner has a monopoly right to invest in a given project. The miner enjoys the advantage of waiting to make the optimal investment decision. Consequently, the calculations of the optimal thresholds are unencumbered and intuitive. In practice, miners must consider the reactions of other miners. Consider an environment with a large number of risk-neutral single-unit firms, each subject to aggregate price uncertainty. Starting with a price process that all miners respond to then solving for the price that clears the market, it becomes clear that miner entry and exit decisions establish reflecting price barriers P_U (upper bound) and P_L (lower bound) that are identical to P_H and P_S, respectively. Though the value of waiting is eliminated in a competitive industry, the mathematical constraint of zero excess profit enforced by miner behavior provides price barriers that mimic the thresholds of a monopolist miner.

increases the optimal threshold P_R the miner insists upon to offset the elevated risk of a price reversal below the threshold.

Panel 2 documents the impact of increases in the cost of scrapping the project entirely upon two threshold prices. As characterized by the positively sloped P_H, an idle miner needs a higher and higher entry price with increases in E_S because reversing that decision becomes more expensive. As demonstrated by the inverse relationship between P_S and E_S, an active miner becomes more patient with price declines before closing up shop entirely, as it becomes more expensive to do so.

Panel 3 depicts a positive relationship between P_R and E_M. If a miner is idle then reactivation will require a higher P_R as E_M increases because reversing that decision is becoming increasingly expensive. Also displayed in the third panel is the notion that increases in E_M decreases P_M. An active miner continuously reduces the optimal threshold price because "laying up" computing equipment has a larger one-time cost with increases in E_M. Importantly, it should be added there is a magnitude of E_M that is large enough to ensure the intersection of P_M and P_S which would eliminate mothballing as an option for the active miner.

Panel 4 characterizes the impact of *R*, the one-time cost of reactivation, upon P_M and P_R. *R* has a positive impact upon the requisite threshold price, P_R, for an idle miner to become active again. The lump sum price of *R* reduces the price at which a miner mothballs his equipment, P_M. The intuition of both results is clear. An idle miner optimally requires a higher P_R to offset the increased cost of reactivating. An active miner tolerates a lower P_M with an increase in *R* because reversing the decision to mothball becomes more expensive.

Panel 5 displays the impact of the flow cost of mothballing, *M*, upon four endogenous threshold prices. Increases in *M* reduce $(C - M)$ the flow cost savings from mothballing for the miner. Therefore, both P_R and P_M fall; a miner will "lay-up" a live project less readily (P_M

decreases) due to the elevated cost of maintaining the idle equipment and the miner will reactivate computing equipment more readily for the same reason (P_R decreases). In addition, P_H and P_S will both rise; the miner will be more reluctant to invest at all (P_H increases) and will abandon his mothballed computing equipment with less reluctance (P_S increases). Furthermore, a falling P_M and a rising P_S will intersect and reveal a level of *M* such that the miner's option to "lay-up" his equipment is economically eliminated.

Panel 6 depicts a positive relationship between all four threshold prices and the operating expenses, *C*. P_H increases with expenses because increases in *C* reduce the expected flow of profit, and hence the value of the computing equipment, so a higher price is required for the miner to invest "*I*". In addition, the miner will abandon his equipment at a higher threshold price, P_S, because he will lose more money when the active flow cost *C* is higher.

Of particular relevance to my empirics is the effect of abandonment cost, E_S, on the miner's thresholds. Note that in panel 2, negative values of E_S are permitted. A negative E_S implies that the miner enjoys a positive salvage value when scrapping his equipment. As this salvage value increases the thresholds converge. This is the differentiating characteristic driving the empirical results of the paper. GPU miners, benefitting from a higher salvage value for their mining equipment, will have relatively closer P_H and P_S thresholds than ASIC miners, whose mining equipment has a lower salvage value (i.e., a higher abandonment cost).

Thus, given a price somewhere between P_H and P_S, GPU miners will likely be more responsive to both positive and negative price shocks, as this price is relative closer to their relevant thresholds. Conversely, ASIC miners may not respond symmetrically to positive and negative price shocks. For example, a positive price shock could lead an ASIC miner to reactivate his equipment or expand his investment if his P_R or P_S thresholds are reached, but a

negative shock of the same magnitude will not illicit a symmetric response if his P_M or P_S thresholds are not reached.

Though I am primarily concerned with the effect of salvage value on miners' thresholds, I present comparative statics for various other characteristics of miners. These are relevant for two reasons. First, I do not attempt to structurally model the entry and exit decisions of miners. Thus, the starting numbers chosen for the comparative static analysis are mean to be illustrative but not necessarily realistic. Therefore, the sensitivity of miners to these other characteristics may be of interest to an informed reader. Second, there could be heterogeneity among miners; they may face different cost structures and therefore their optimal thresholds for investment, abandonment, reactivation and mothballing differ. As I consider data on aggregate hashrate and the corresponding cryptocurrency price, any price shock is likely to push a certain segment of miners (representing only a portion of the aggregate hashrate) over their respective thresholds. This is an important characterization, as I only seek to differentiate positive and negative price shocks in my empirical model.

4. Data

4A: Data Sources

Daily cryptocurrency price data are collected from Coinmetrics.io for both Bitcoin (BTC) and Ether (ETH). These prices are the "closing prices" at time 00:00 UTC, the (beginning of the) next day. The daily BTC price data cover the period from 07/18/2010 to 10/19/2019 (3,381 days), and the daily ETH price data cover the period from 08/08/2015 to 10/19/2019 (1,534 days).

Average daily BTC hashrate is obtained from data.bitcoinity.org. This measure is the sum of average daily hashrate, measured in terahashes/second (TH/s), provided by each mining

pool and other sources of computing power (miners not using a pool). Average daily ETH hashrate is obtained from etherscan.io. This is the average aggregate hashrate, measured in TH/s, for all miners of the ETH blockchain. Hashrate data cover the same period as price data.

I separate BTC data based on a regime shift in the cointegration relationship between price and hashrate (See Appendix C and Table 24) and focus my analysis on BTC price and hashrate data after 04/14/2014 (2,015 days).

4B: Summary Statistics

Summary statistics are presented in Table 18. Both BTC and ETH have positive mean returns, and positive mean hashrate growth. Consistent with Bhambhwani et al. (2019), the standard deviation of both price and hashrate returns are significantly greater than the respective means for both BTC and ETH. This is symptomatic of the highly volatile cryptocurrency market.

Figure 16 and Figure 17 graph time-series trends in the variables of interest for BTC and ETH, respectively. Panel 1 presents the natural log of the cryptocurrency price in USD over time, while panel 2 shows the natural log of the cryptocurrency hashrate, measured in terahashes/second. Panel 3 graphs the first difference in log- price over time. Panel 4 graphs the first difference in log- hashrate over time. Panel 5 combines Panel 1 and Panel 2 on the same graph and is useful for visualizing the cointegration relationship (see Section 5). For both BTC and ETH, log- price and log- hashrate appear to be non-stationary variables, while their respective first differences appear to be stationary.

5. Methodology

I employ a Threshold Vector Error Correction Model (threshold VECM or TVECM) as my main empirical specification. The process for preparing and implementing this model is

similar to that of the VECM model. First, using the augmented Dickey and Fuller (1979) (ADF) test, I test my variables (price and hashrate) for the existence of a unit root. The null of the ADF test is that there is a unit root present in the time series variable. The existence of a unit root implies that the time-series is $I(1)$ (i.e. non-stationary). The results of the ADF test, presented in Table 19, suggest that both price and hashrate are non-stationary variables.

Given this, I test for bivariate cointegration between the variables. Cointegration between price and hashrate would imply that the variables will tend towards a common equilibrium. Bivariate cointegration implies that a linear combination of two $I(d)$ variables can form a series that is $I(d-b)$ where $b > 0$. In the case of price and hashrate, this can be expressed as:

$$log(hashrate_t) = \beta_1 log(price_t) + \epsilon_t \quad (22)$$

where $log(price_t)$ and $log(hashrate_t)$ are $I(d)$, and ϵ_t is $I(d-b)$. If the series represented by Equation 22 is weakly stationary, i.e. $(d-b) = 0$, then mean reversion towards the common equilibrium between price and hashrate will persist. The Johansen (1991, 1995) tests (both trace and eigenvalue tests) provide strong evidence that price and hashrate are cointegrated for both BTC and ETH (see Table 20).[24]

To reconcile this cointegration relationship (see Equation 22) within a reduced-form model, it can be expressed as a cointegrating vector, A. This vector takes the form of $A' = [1, -\beta_1]$ such that:

[24] The number of lags selected for the Johansen tests are determined using the Bayesian Information Criterion (BIC) (Schwarz, 1978) and the Akaike Information Criterion (AIC) minimizing unrestricted Vector Autoregression (VAR) specification. To account for persistence in the BTC time-series, I err on the side of including more lags in these specifications.

$$A'Y_t = log(hashrate_t) - \beta_1 log(price_t)$$

(23)

where $Y_t' = [log(hashrate_t), log(price_t)]$. For example, if $A'Y_t < 0$ (meaning that $log(hashrate_t) < \beta_1 log(price_t)$) it implies that, based on their long-run equilibrium relationship, hashrate is relatively low compared to price. Likewise, if $A'Y_t > 0$ it implies that hashrate is relatively high compared to price (presuming $\beta_1 > 0$). This allows for the implementation of the VECM and TVECM models. The VECM specification (Johansen, 1991; 1995) can be written generally as:

$$\Delta Y_t = \Phi D_t + \Pi A' Y_{t-1} + \sum_{k=1}^{q} \Gamma_k \Delta Y_{t-k} + \epsilon_t.$$

(24)

Φ represents the constant and trend coefficients, with $D_t = u_0 + u_1 t$ where u_0 is the constant component and $u_1 t$ is the trend component. The long-run impact coefficient, Π, is the mean reversion term representing the speed of adjustment towards the equilibrium. This is referred to as the Error Correction Term (ECT). When $\Pi < 0$, this represents the percent of disequilibrium from time t-1 corrected in time t. If $\Pi > 0$, this is indicative of divergence. $\sum_{k=1}^{q} \Gamma_k$ represents the short-run impact coefficients. These are coefficients for lagged values of ΔY_t and capture short-run deviations from the equilibrium. An intuitive benefit of the VECM specification compared to the Vector Autoregression specification is that the VECM offers interpretations of both the short-run dynamics and the long-run equilibrium of the system.

Drawing on my threshold analysis, a price increase would not necessarily occasion an increase in hashrate for any one miner unless their reactivation or investment threshold is reached. However, given heterogeneity between miners in aggregate, any number of miners would attain either their reactivation or their investment threshold with the price change. This

then will be reflected by an increase in the aggregate hashrate. This reaction to the disequilibrium caused by the price change is what Π captures in Equation 24.

However, my threshold analysis highlights a shortcoming of this particular application of the model. The VECM specification assumes that Π is symmetric for both positive and negative disequilibria. This is problematic, as Figure 15 clearly demonstrates that there is "space" between the entry and exit thresholds of miners; that is there is a range of possible prices between the thresholds for which the optimal strategy for the miner will be inaction. Thus, the equilibrium cryptocurrency price (measured by the cointegrating relationship, $A'Y_t$) may reasonably fall anywhere within this range. If this equilibrium price is closer to the activation threshold, then miners will be more responsive to positive price shocks than negative shocks, as these positive shocks are more likely to push miners across the activation price threshold. Likewise, if the equilibrium price is closer to the abandonment threshold, then miners will be more responsive to negative price shocks than positive price shocks. Thus, it would be inappropriate to assume that Π is symmetric for positive and negative disequilibria.

To illustrate the relevance of this point, consider panel 2 on Figure 15 which shows that as salvage value increases, activation and abandonment thresholds begin to converge. As ETH miners enjoy a higher salvage value, their activation and abandonment decisions may be more sensitive to price changes than BTC miners. Meanwhile, the equilibrium BTC price could be closer to either the activation threshold or the abandonment threshold for BTC miners, as the area between these thresholds is greater. If this is indeed the case, BTC miners will react differently to positive and negative disequilibria.

To reconcile my empirical model with this reality, I employ a TVECM specification (Balke and Fomby, 1997; Hansen and Seo, 2002). This allows for separate error correction

terms for positive and negative disequilibria. To accomplish this, I present the dummy variables $d_{1,t}$ and $d_{2,t}$ where:

$$d_{1,t} = 1 \ \ if \ \ A'Y_{t-1} \leq \gamma \tag{25}$$

and

$$d_{2,t} = 1 \ \ if \ \ A'Y_{t-1} > \gamma \tag{26}$$

where γ is the threshold parameter. I set $\gamma = 0$ to differentiate between positive and negative disequilibria. $d_{1,t}$ and $d_{2,t}$ indicate whether the disequilibrium from time t-1 is a negative or positive disequilibrium, respectively.

The introduction of these dummy variables allows for the following TVECM specification:[25]

$$Y_t = \Phi D_t + \Pi_1 d_{1,t} A'Y_{t-1} + \Pi_2 d_{2,t} A'Y_{t-1} + \sum_{k=1}^{q} \Gamma_k \Delta Y_{t-k} + \epsilon_t. \tag{27}$$

The novelty of this specification is the separation of Π into Π_1 and Π_2. Π_1 is the ECT corresponding to a negative disequilibrium (ECT-). In the hashrate equation, this represents the mean reversion (or speed of adjustment) when hashrate is relatively low compared to price. Π_2 is the ECT corresponding to a positive disequilibrium (ECT+), representing the mean reversion when hashrate is relatively high compared to price.

[25] As in the Johansen tests, the number of lags in the TVECM specification is selected using the Bayesian Information Criterion (BIC) (Schwarz, 1978) and the Akaike Information Criterion (AIC) minimizing unrestricted Vector Autoregression (VAR) specification. Though I am less interested in the short-run price-hashrate dynamics, these regressors are necessary due to serial correlation.

6. Equilibrium Results

6A: Single Sample Specification

The results for the TVECM specification are presented in Table 21. There is a long-run cointegrating relationship between price and hashrate for BTC and for ETH, with separable mean reversion coefficients representing miner-responses to positive and negative disequilibria. The ECT- coefficient for the BTC hashrate equation is -0.0111, and is significant at a 99% confidence level, while the corresponding ECT+ coefficient is statistically insignificant. This indicates that BTC miners respond to negative disequilibria (when hashrate is relatively low compared to price) by increasing hashrate, but they do not have a symmetric response to positive disequilibria. Meanwhile, both the ECT- and the ECT+ coefficients in the ETH hashrate equation are significant at 99% confidence levels; ECT- with a coefficient of -0.0123 and ECT+ with a coefficient of -0.0114. This indicates that ETH miners respond to both positive and negative disequilibria.

This result is consistent with my analysis of miners' entry and exit decisions (see Section 3). ETH miners, who enjoy a higher salvage value when scrapping their GPU mining equipment, have closer activation and abandonment thresholds (P_H and P_S respectively). Thus, they are more sensitive to both positive and negative disequilibria, as price shocks push these miners across their respective thresholds. BTC miners, who receive a lower salvage value when scrapping their ASIC computing equipment, have a greater range of prices between their activation and abandonment thresholds. If the equilibrium price (as given in the cointegration relationship) were to bisect these thresholds, we would expect still expect BTC miners to have a relatively symmetric response to disequilibria. However, it is also possible that this equilibrium price is closer to the activation or abandonment threshold, which would lead to miners being more responsive to positive price shocks or negative price shocks, respectively.

The results of the TVECM suggest that the equilibrium BTC price (in the examined sample) is closer to the activation threshold (of the marginal miner or a group of miners representing the largest proportion of aggregate hashrate).

6B: ASIC ETH Miners

On April 3, 2018 Bitmain (a manufacturer of ASIC mining equipment) announced a new Ethash ASIC miner: the Antminer A3.[26] This news was very relevant to the Ethereum blockchain, as Ethereum was designed to be "ASIC-resistant," meaning that ASIC equipment should struggle to run the Ethash algorithm. This is perhaps why this specialized ASIC equipment, first shipped in July 2018, did not seem to affect the aggregate ETH hashrate significantly (see Figure 16, Panel 2). However, the cost-efficiency of this mining equipment cannot be denied as the first batch sold out the day it was announced.[27]

Given this, I create two subsamples: a pre-ASIC subsample from 10/25/2016 to 06/30/2018 and a post-ASIC subsample from 08/01/2018 to 10/19/2019. For each of these, I re-estimate the TVECM specification, following the same process outlined in Section 5. Table 22 presents the results of these specifications. Consistent with the results of the single sample specification, both ETC- and ECT+ are significant at 99% confidence levels for the pre-ASIC subsample: ECT- with a coefficient of -0.0134 and ECT+ with a coefficient of -0.0188. Thus, ETH miners in the pre-ASIC era have a relatively symmetric response to positive and negative disequilibria. For the post-ASIC era, the ECT+ coefficient is significant at the 99% confidence level with a coefficient of -0.0307, while the ECT- coefficient is statistically indistinguishable from 0. This suggests that post-ASIC ETH miners react only to positive disequilibria (when hashrate is relatively high compared to price) by decreasing hashrate.

[26] See https://cointelegraph.com/news/bitmain-releases-ethash-asic-miners.
[27] See https://www.coindesk.com/bitmain-confirms-release-first-ever-ethereum-asic-miners and https://www.businessinsider.com/asic-crypto-mining-ethereum-ether-cryptocurrency-2018-5.

This, at first glance, suggests that the introduction of ETH ASIC mining equipment led to a fundamental change in miner behavior. This would be consistent with my entry and exit analysis of miners; the interpretation of post-ASIC ETH miners would be the opposite of BTC miners, alluding to an equilibrium price closer to the abandonment threshold (of the representative miner or dominant group of miners).

While this conclusion seems appealing, there are a number of other possible explanations. For example, there could be an overlapping trend in this subsample causing the disparity between miners' reactions to positive and negative disequilibria. This could be representative of the price crash in early 2018, when the ETH price fell from an all-time high of $1,448.18 in January, 2018.[28] On August 1, 2018 (the beginning of the post-ASIC sample) the price was $419.02 and reached a low (closing) price of $82.83 on December 15, 2018. It is also important to note that I cannot measure the rate at which these ASIC miners were adopted, as Bitmain did not release the number of Antminer E3 units sold. It is therefore hard to consider the affect this technology had on aggregate hashrate. As the post-ASIC sample is only 438 days, these alternatives must be acknowledged.

6C: BTC Block Reward Halving

For robustness, I separate the BTC sample based on the BTC block reward halving. Bitcoin halving is when the block reward for mining a new Bitcoin block is halved, meaning that miners will receive half the previous reward for mining new blocks. This happens every 210,000 blocks, or roughly once every four years. At the start of my sample, the block reward for mining a new block was 25 BTC per block. This reward "halved" on 07/09/2016, dropping to 12.5 BTC per block. Therefore, I repeat my analysis for the subsample before 07/09/2016 and the subsample after 07/09/2016.

[28] See https://www.coindesk.com/price/ethereum.

The results of this analysis are presented in Table 23. Similar to the single sample results for BTC, the ECT- term in the pre-halving sample is significant at the 99% confidence level with a value of -0.0181 while the ECT+ term is insignificant. Likewise, the ECT- term in the post-halving sample is significant at the 99% confidence level with a value of -0.0135, while the ECT+ term is similarly insignificant.[29] These results suggest that BTC block reward halving does not affect BTC miners' adjustment to disequilibria.

7. Discussion

As mentioned, the importance of miners to blockchain applications is often overlooked by financial economists. It is easy, given the wealth of data offered by public blockchains, to consider the hashrate as a variable for which we should control, but not necessarily study. However, practitioners will point out the pivotal role miners play in PoW blockchains. Miners verify transactions, facilitating liquidity of blockchain based assets (i.e. cryptocurrencies). At scale, this verification process provides security to ownership of these assets. Furthermore, the incentive structure of miners implies that there is indeed a long-run price-hashrate relationship. For empiricists, this alludes to the fact that these variables are cointegrated.

Thus, the implications of an asymmetric miner response to price shocks are relevant to academics and practitioners alike. Foremost among these are pricing implications. As hashrate is an empirically significant factor for predicting cryptocurrency returns (Bambhwani et al., 2019), my results suggest that this may indeed be a non-linear effect on returns in the face of asymmetric miner reactions. This is a promising avenue for future research.

[29] Interestingly, the trend term in the post-halving sample is insignificant. I therefore re-estimate the model excluding the trend term. The optimal number of lags varies based on the inclusion or exclusion of this term. As I do not interpret the short-run impact coefficients (i.e. lagged regressors), I again err on the side of including more lags to account for serially correlated variables.

Moreover, the asymmetric miner response I document has implications for applications of blockchain technology. Miner behavior in the face of volatile prices can affect the security and functionality of blockchain applications. To illustrate this, consider a blockchain based ledger used to record transactions of corporate securities for a particular firm. In this application, miners could be rewarded a cryptocurrency linked to the equity value of this firm. While this application of blockchain technology should increase transparency of ownership and facilitate cheaper auditing, issues unique to blockchain applications, born of the incentive structure of miners, can arise. If a negative externality causes the firm's stock price to plummet, miners might abandon their investment in mining equipment. This abandonment will decrease the security of the blockchain, and by extension, the firm's ledger. In extreme cases the blockchain based ledger would cease to function with no miners to verify transactions. Following a negative shock to price this is a concerning outcome. Therefore, the (a)symmetry of miner reactions to price shocks is an important consideration for blockchain applications in a financial setting.

Finally, the miner responses to price shocks documented in the TVECM results suggest that price-hashrate spirals, as described by Pagnotta and Buraschi (2018), may exist, but that for ASIC mined cryptocurrencies only a positive *or* a negative price-hashrate spiral will exist. This has implications for the scaling of new blockchain applications, as positive price-hashrate spirals may facilitate the growth of blockchains by rewarding early investors (Catalini and Gans, 2020). This is another promising avenue for future research.

8. Conclusion

I model the entry and exit decisions of cryptocurrency miners using a real options framework. I show that the optimal threshold for investment in mining equipment and the

optimal threshold for abandonment of a mining project begin to converge when salvage value increases. This is relevant to the two types of computing equipment used by miners; specialized ASIC equipment, with a lower salvage value, and consumer GPU equipment, with a higher salvage value. Given this, I expect GPU miners to have a more symmetric response to price shocks than ASIC miners.

Using a TVECM specification, I separate the positive and negative mean reversion (error-correction) terms to show miner responses to price-hashrate disequilibria born of price shocks. The TVECM results show that BTC miners, using ASIC equipment, respond only to negative disequilibria (when hashrate is relatively low compared to price) by increasing hashrate, while ETH miners respond to both positive and negative price disequilibria. These results hold for BTC miners both before and after the block reward halving. Furthermore, I show that after the introduction of ETH ASIC mining equipment, the response of ETH miners is asymmetric. This is consistent with my analysis of miners' entry and exit decisions, which are sensitive to the differential salvage value of ASIC and GPU computing equipment.

These results have various theoretical and practical implications. An asymmetric miner response suggests that hashrate, which has been identified as a significant factor in cryptocurrency pricing models, should in fact be a non-linear factor in pricing equations for ASIC mined cryptocurrencies. Second, asymmetric miner response has implications for applications of blockchain technology, as miners facilitate transactions and provide security to ownership of blockchain-based assets. Both the pricing and security implications born of this miner behavior warrant further study.

References

Abadi, J., Brunnermeier, M., 2019. Blockchain economics. *Working paper*.

Angeris, G., Chitra, T., 2020. Improved price oracles: Constant function market makers. *Working paper*.

Angeris, G., Kao, H.-T., Chiang, R., Noyes, C., Chitra, T., 2021. An analysis of uniswap markets. *Cryptoeconomic Systems Journal* 1.

Alsabah, H., Capponi, A., 2020. Pitfalls of bitcoin's proof-of-work: R&d arms race and mining centralization. *Working Paper*.

Aoyagi, J., 2020. Liquidity provision by automated market makers. *Working paper*.

Arnosti, N., Weinberg, S. M., 2018. Bitcoin: A natural oligopoly. *CoRR abs/1811.08572*.

Balke, N., Fomby, T., 1997. Threshold cointegration. *International Economic Review* 38, 627-645.

Barbon, A., Ranaldo, A., 2021. On the quality of cryptocurrency markets: Centralized versus decentralized exchanges. *Working paper,* University of St. Gallen, School of Finance Research Paper Forthcoming, Swiss Finance Institute Research Paper No. 22-38.

Basu, S., Easley, D., O'Hara, M., Sirer, E., 2018. Towards a functional fee market for cryptocurrencies. *Working Paper*.

Baur, D., Hong, K., Lee, A., 2018. Bitcoin: Medium of exchange or speculative assets? *Journal of International Financial Markets, Institutions and Money* 54, 177-189.

Benetton, M., Compiani, G., Morse, A., 2019. Cryptomining: Local evidence from china and the us. *Working Paper*.

Bhambhwani, S., Delikouras, S., Korniotis, G., 2019. Do fundamentals drive cryptocurrency prices? *Working paper*.

Biais, B., Bisiere, C., Bouvard, M., Casamatta, C., 2019. The blockchain folk theorem. *Review of Financial Studies* 32(5), 1662–1715.

Biais, B., Bisiere, C., Bouvard, M., Casamatta, C., Menkveld, A., 2018. Equilibrium bitcoin pricing. *Working paper*.

Brennan, M., Schwartz, E., 1985. Evaluating natural resource investments. *Journal of Business* 58, 135-157.

Budish, E., 2018. The economic limits of bitcoin and the blockchain. *NBER Working Paper*.

Callaway, B., Goodman-Bacon, A., Sant'Anna, P., 2021. Difference-in-differences with a continuous treatment. *Working paper*, Cornell University.

Capponi, A., Jia, R., 2021. The adoption of blockchain-based decentralized exchanges. *Working paper*, Columbia University.

Catalini, C., Gans, J., 2020. Some simple economics of the blockchain. *Communications of the ACM* 63, 80-90.

Chabakauri, G., 2013. Dynamic equilibrium with two stocks, heterogeneous investors, and portfolio constraints. *The Review of Financial Studies* 26, 3104–3141.

Chen, L., Cong, L. W., Xiao, Y., 2019. A brief introduction to blockchain economics. *Working Paper*.

Chiu, J., Koeppl, T. V., 2019. Blockchain-based settlement for asset trading. *The Review of Financial Studies* 32, 1716–1753.

Chiu, J., Koeppl, T. V., 2018. The economics of cryptocurrencies - bitcoin and beyond. *Working Paper*.

Cheah, E.-T., Fry, J., 2015. Speculative bubbles in bitcoin markets? an empirical investigation into the fundamental value of bitcoin. *Economic Letters* 130, 32-36.

Chod, J., Lyandres, E., 2019. A theory of icos: Diversification, agency, and information asymmetry. *Management Science* Forthcoming.

Cong, L. W., He, Z., Li, J., 2021a. Decentralized mining in centralized pools. *Review of Financial Studies* 34(3), 1191–1235.

Cong, L. W., Li, Y., Wang, N., 2021b. Tokenomics: Dynamic adoption and valuation. *Review of Financial Studies* 34(3), 1105–1155.

Cong, L. W., He, Z., 2019. Blockchain disruption and smart contracts. *Review of Financial Studies* 32, 1754-1797.

d'Avernas, A., Maurin, V., Vandeweyer, Q., 2022. Can stablecoins be stable? *Working paper*, University of Chicago, Becker Friedman Institute for Economics Working Paper No. 2022-131.

Davydiuk, T., Gupta, D., Rosen, S., 2019. De-crypto-ing signals in initial coin offerings: Evidence of rational token retention. *Working Paper*.

Dickey, D., Fuller, W., 1979. Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association* 74, 427-431.

Dimitri, N., 2017. Bitcoin mining as a contest. *Ledger* 2, 31-37.

Dixit, A., 1989. Entry and exit decisions under uncertainty. *Journal of Political Economy* 97, 620-638.

Dixit, A., 1992. Investment and hysteresis. *Journal of Economic Perspectives* 6, 107-132.

Dixit, A., Pindyck, R., 1994. Entry, exit, lay-up, and scrapping. In: Investment Under Uncertainty, Princeton University Press, Princeton, NJ, pp. 213-244.

Dixit, A., Pindyck, R., 1995. The options approach to capital investment. *Harvard Business Review* 73, 105-115.

Dwork, C., Naor, M., 1992. Pricing via processing or combatting junk mail. In *12th Annual International Cryptology Conference* pp. 139–147.

Easley, D., O'Hara, M., Basu, S., 2019. From mining to markets: The evolution of bitcoin transaction fees. *Journal of Financial Economics* 134, 91–109.

Foley, S., Karlsen, J., Putnins, T., 2019. Sex, drugs, and bitcoin: How much illegal activity is financed through cryptocurrencies? *Review of Financial Studies* 32, 1798-1853.

Frazzini, A., Pedersen, L. H., 2014. Evaluating natural resource investments. *Journal of Financial Economics* 111, 1–25.

Gan, R. J., Tsoukalas, G., Netessine, S., 2020. Initial coin offerings, speculators, and asset tokenization. *Management Science* Forthcoming.

Gregory, A., Hansen, B., 1996. Residual-based tests for cointegration in models with regime shifts. *Journal of Econometrics* 70, 99-126.

Griffin, J., Shams, A., 2020. Is bitcoin really untethered? *The Journal of Finance* 75, 1913–1964.

Hansen, G., Seo, B., 2002. Testing for two-regime threshold cointegration in vector error-correction models. *Journal of Econometrics* 110, 293-318.

Hardle, W., Harvey, C. R., Reule, R. C., 2020. Understanding cryptocurrencies. *Journal of Financial Econometrics* 18(2), 181–208.

Hardouvelis, G., 1990. Margin requirements, volatility, and the transitory component of stock prices. *The American Economic Review* 80, 736–762.

Hardouvelis, G., Theodossiou, P., 2002. The asymmetric relation between initial margin requirements and stock market volatility across bull and bear markets. *The Review of Financial Studies* 15, 1525–1559.

Harvey, C., 2016. Cryptofinance. *Working Paper*.

Harvey, C., Ramachandran, A., Santoro, J., 2021. DeFi and the Future of Finance. John Wiley & Sons.

Hasbrouck, J., Rivera, T., Saleh, F., 2022. The need for fees at a dex: How increases in fees can increase dex trading volume. *Working paper*.

Hinzen, F., John, K., Saleh, F., 2022. Bitcoin's limited adoption problem. *Journal of Financial Economics* 144, 347–369.

Howell, S. T., Niessner, M., Yermack, D., 2020. Initial coin offerings: Financing growth with cryptocurrency token sales. *Review of Financial Studies* 33(9), 3925–3974.

Hsieh, D., Miller, M., 1990. Margin regulation and stock market volatility. *The Journal of Finance* 45, 3–29.

Huberman, G., Leshno, J., Moallemi, C., 2021. Monopoly without a monopolist: An economic analysis of the bitcoin payment system. *The Review of Economic Studies* 88, 3011–3040.

Irresberger, F., 2019. Coin concentration of proof-of-stake blockchains. *Working Paper*.

Irresberger, F., John, K., Mueller, P., Saleh, F., 2021. The public blockchain ecosystem: An empirical analysis. *Working paper*, NYU Stern School of Business.

Johansen, S., 1991. Estimation and hypothesis testing of cointegration vectors in gaussian vector autoregressive models. *Econometrica* 59, 1551-1580.

Johansen, S., 1995. Likelihood-based inference in cointegrated vector autoregressive models. Oxford University Press.

John, K., Rivera, T., Saleh, F., 2022. Equilibrium staking levels in a proof-of-stake blockchain. *Working paper*.

Jylha, P., 2018. Margin requirements and the security market line. *The Journal of Finance* 73, 1281–1321.

King, S., Nadal, S., 2012. Ppcoin: Peer-to-peer crypto-currency with proof-of-stake. *White paper*: https://peercoin.net/assets/paper/peercoin-paper.pdf.

Kristoufek, L., 2020. Bitcoin and its mining on the equilibrium path. *Energy Economics* 85.

Lehar, A., Parlour, C., 2020. Miner collusion and the bitcoin protocol. *Working paper*.

Lehar, A., Parlour, C., 2021. Decentralized exchanges. *Working paper*.

Lehar, A., Parlour, C., 2022. Systemic fragility in decentralized markets. *Working paper*.

Lee, J., Li, T., Shin, D., 2019. The wisdom of crowds in fintech: Evidence from initial coin offerings. *Working Paper*.

Li, Y., Mayer, S., 2022. Money creation in decentralized finance: A dynamic model of stablecoin and crypto shadow banking. *Working paper*, Fisher College of Business Working Paper No. 2020-03-030.

Li, J., Mann, W., 2020. Digital tokens and platform building. *Working Paper.*

Li, T., Shin, D., Wang, B., 2019. Cryptocurrency pump-and-dump schemes. *Working Paper.*

Liu, Y., Sheng, J., Wang, W., 2020. Do cryptocurrencies have fundamental values? *Working Paper*.

Liu, Y., Tsyvinski, A., 2020. Risks and returns of cryptocurrencies. *Review of Financial Studies* Forthcoming.

Liu, Y., Tsyvinski, A.,Wu, X., 2022. Common risk factors in cryptocurrency. *Journal of Finance* 77, 1133-1177.

Lyandres, E., Palazzo, B., Rabetti, D., 2019. Do tokens behave like securities? An anatomy of initial coin offerings. *Working Paper*.

Ma, J., Gans, J., Tourky, R., 2018. Market structure in bitcoin mining. *Working paper*.

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

Malinova, K., Park, A., 2018. Tokenomics: When tokens beat equity. *Working Paper*.

McDonald, R., Siegel, D., 1986. The value of waiting to invest. *Quarterly Journal of Economics* 101, 707-728.

Mora, C., Rollins, R. L., Taladay, K., Kantar, M. B., Chock, M. K., Shimada, M., Franklin, E. C., 2018. Bitcoin emissions alone could push global warming above 2 c. *Nature Climate Change* 8, 931–933.

Nakamoto, S., 2008. Bitcoin: A peer-to-peer electronic cash system. https://bitcoin.org/bitcoin.pdf.

Pagnotta, E., 2022. Decentralizing money: Bitcoin prices and blockchain security. *The Review of Financial Studies* 35, 866–907.

Pagnotta, E., Buraschi, A., 2018. An equilibrium valuation of bitcoin and decentralized network assets. *Working paper*.

Park, A., 2022. Conceptual flaws of decentralized automated market making. *Working paper*.

Pindyck, R., 1991. Irreversibility, uncertainty, and investment. *Journal of Economic Literature* 29, 1110-1148.

Prat, J., Walter, B., 2018. An equilibrium model of the market for bitcoin mining. *Working paper*.

Rosu, I., Saleh, F., 2021. Evolution of shares in a proof-of-stake cryptocurrency. *Management Science* 67, 661–672.

Rytchkov, O., 2014. Asset pricing with dynamic margin constraints. *The Journal of Finance* 69, 405–452.

Saleh, F., 2019. Volatility and welfare in a crypto economy. *Working Paper*.

Saleh, F., 2021. Blockchain without waste: Proof-of-stake. *The Review of Financial Studies* 34, 1156–1190.

Schwarz, G., 1978. Estimating the dimension of a model. *Annals of Statistics* 6, 461-464.

Shams, A., 2020. The structure of cryptocurrency returns. *Working paper*.

Sockin, M., Xiong, W., 2020. A model of cryptocurrencies. *Working paper*.

Yermack, D., 2015. Is bitcoin a real currency? An economic appraisal. *Handbook of Digital Currency* pp. 31–43.

Yermack, D., 2017. Corporate governance and blockchains. *Review of Finance* 21(1), 7–31.

Appendix A: Proofs for Chapter 2

Lemma A.1. *Stationary Distribution of Memory Pool*

Let $a_f > 0$ *denote the arrival rate of transactions with fees at least as large as* $f \geq 0$ *and assume that arrivals follow a Poisson process and that* $a_f < \Lambda$. *Let* M_f *denote the stationary distribution for number of transactions in the memory pool with a fee of at least f. Then, the distribution of* M_f *is given as follows:*

For any $m \in \{0,1,\dots\}$: $\mathbb{P}(M_f = m) = \left(\frac{a_f}{\Lambda}\right)^m \cdot (1 - \frac{a_f}{\Lambda})$

Proof: Note that the number of transactions in the memory with fees at least as large as *f* can only change in unit increments. In particular, either a new transaction with fee greater than or equal to *f* arrives and the transaction number increases by one or a block is mined and the transaction number reduces by one. Since both arrivals and departures follow a Poisson process, the described process is the classical birth-death process. The steady state distribution of such a process, $\{\pi_f(m)\}_{m\in\mathbb{N}}$, is restricted by the following balance equations:

$$\pi_f(0) \cdot a_f = \pi_f(1) \cdot \Lambda \tag{28}$$

and

$$\pi_f(m) \cdot (a_f + \Lambda) = \pi_f(m-1) \cdot a_f + \pi_f(m+1) \cdot \Lambda \quad for \quad m \in \{1,2,\dots\} \tag{29}$$

which imply $\pi_f(m) = \left(\frac{a_f}{\Lambda}\right)^m \cdot \pi_f(0)$ for all $m \in \{0,1,\dots\}$. Then, since $\{a_f(m)\}_{m\in\mathbb{N}}$ must be a valid probability distribution, we must have that $1 = \sum_{m=0}^{\infty} \pi_f(m) = \sum_{m=0}^{\infty} \pi_f(0) \cdot \left(\frac{a_f}{\Lambda}\right)^m$,

which yields $\pi_f(0) = 1 - \frac{a_f}{\Lambda}$. Then, $P\{ M_f = m\} =: \pi_f(m) = \left(\frac{a_f}{\Lambda}\right)^m \cdot (1 - \frac{a_f}{\Lambda})$ for all $m \in \{1,2,\dots\}$ which completes the proof.

Lemma A.2. *Processing Time*

Let $a_f \geq 0$ denote the arrival rate of transactions with fees at least as large as $f \geq 0$. Assume also that arrivals follow a Poisson process, that $a_f < \Lambda$ and that the arrival rate for transactions with fee exactly f is zero. Then, the expected time for a transaction with fee f to be included in a block (i.e., processing time), $W_P(f)$, is given by:

$$W_P(f) = \frac{\Lambda}{(\Lambda - a_f)^2}$$

Proof. Let X_t denote the number of transactions in the memory pool of equal or higher priority than a transaction with fee *f* at a time *t* units after the transaction with fee *f* is first entered in the memory pool. We define X_0 to include the transaction with fee *f* itself so that the first occurrence of $X_t = 0$ corresponds to the transaction with fee *f* being processed. Thus:

$$W_P(f) = \mathbb{E}[T_f]$$

(30)

with

$$T_f := min\{ t > 0 : X_t = 0\}$$

(31)

Let $\{ Y_t \}_{t\in\mathbb{N}}$ denote the embedded Discrete Time Markov Chain in the Continuous Time Markov Chain, $\{ X_t \}_{t\geq 0}$. Then, the following equation holds by definition:

$$T_f := \sum_{n=1}^{N_f} A_n$$

(32)

with

$$N_f := min\{\, t > 0 : Y_t = 0\}$$

(33)

and $A_n \sim exp(\Lambda + a_f)$ denoting the inter-arrival times for transitions in the $\{ X_t\}_{t\geq 0}$ process. Then, plugging Equation 32 into Equation 30 and applying the Law of Iterated Expectation yields:

$$W_P\,(f) = \mathbb{E}\left[\ \mathbb{E}\left[\sum_{n=1}^{N_f} A_n \ \middle|\ Y_0\ \right]\right]$$

(34)

Since A_n is independent of N_f and X_0, Equation 34 becomes:

$$W_P(f) \ = \ \frac{1}{\Lambda + a_f}\,\mathbb{E}[\ \mathbb{E}\big[N_f \ \big|\ Y_0\big]]$$

(35)

Optional Stopping Theorem then implies:

$$\mathbb{E}\big[N_f \ \big|\ Y_0\big] \ = \ \frac{Y_0}{(\Lambda - a_f)/(\Lambda + a_f)} \ = \ Y_0\ \frac{\Lambda + a_f}{\Lambda - a_f}$$

(36)

Plugging Equation 36 into Equation 35 then yields:

$$W_P(f) \ = \ \frac{\mathbb{E}[Y_0]}{\Lambda - a_f}$$

(37)

Note that $Y_0 \sim M_f + 1$ with M_f defined in Lemma A.1. Thus:

$$\mathbb{E}[Y_0] = \sum_{m=0}^{\infty} (m + 1) \cdot \left(\frac{a_f}{\Lambda}\right)^m \cdot \left(1 - \frac{a_f}{\Lambda}\right) = \frac{\Lambda}{\Lambda - a_f}$$

(38)

Plugging Equation 38 into Equation 37 then completes the proof.

Lemma A.3. *Endogenous Fee Function*

There exists an equilibrium such that each User *i* selects a fee $\phi(c_i)$, with $\phi(c_i)$ being given by:

$$\phi(c_i) = 2 \cdot a \cdot \Lambda \int_0^{c_i} \frac{g(x)}{\left(\Lambda - a \cdot \overline{G}(x)\right)^3} \, dx$$

Proof. Recall that each User *i* selects a fee, f_i, to solve Equation 6. Following the convention in the literature (see, e.g., Huberman et al. 2021 and Hinzen et al. 2022), we solve for a pure strategy equilibrium characterized by a strictly increasing fee function, ϕ, that maps user disutility to optimal fee choices. In particular, we derive a function, ϕ, such that:

$$\phi(c_i) = \arg\max_{f \geq 0} \; u_i \cdot \Psi\left(N_{\tau_i}\right) - c_i \cdot W(f) - f$$

(39)

for all c_i so that $\phi(c_i)$ is the optimal fee choice for each User *i* and thus $f_i = \phi(c_i)$ for each User *i*.

Applying Equation 7 to Equation 39 yields:

$$\phi(c_i) = \arg\max_{f \geq 0} \; u_i \cdot \Psi\left(N_{\tau_i}\right) - c_i \cdot \left(W_P(f) + \kappa\right) - f$$

(40)

In turn, the first order condition implies:

$$- c_i \frac{d}{df}[W_P(f)]|_{f=\phi(c_i)} = 1$$

(41)

Note that, since each User i selects a fee given by $f_i = \phi(c_i)$ then the arrival rate of transactions with fees at least as large as $f \geq 0$ is given by $a_f = a \cdot \overline{G}(\phi^{-1}(f))$ with $\overline{G}(x) := 1 - G(x)$. Moreover, in such a case, the transaction arrivals with fees at least *f* follow a Poisson process and the arrival rate of transactions with exactly fee *f* is zero because total transaction arrivals follow a Poisson process and *G* is strictly increasing. Therefore, the hypothesis of Lemma A.2. is satisfied so that we can apply the result from Lemma A.2. with $a_f = a \cdot \overline{G}(\phi^{-1}(f))$:

$$W_P(f) = \frac{\Lambda}{\left(\Lambda - a \cdot \overline{G}(\phi^{-1}(f))\right)^2}$$

(42)

Then, applying Equation 42 to Equation 41 yields:

$$\frac{2 \cdot c_i \cdot \Lambda}{\left(\Lambda - a \cdot \overline{G}(c_i)\right)^3} \times \left(a \cdot g(c_i)\right) \times \frac{1}{\phi'(c_i)} = 1$$

(43)

with $g(x) := \frac{dG}{dx}$. Solving the differential equation in Equation 43 and imposing $\phi(0) = 0$ (which is implied trivially by Equation 6) then completes the proof:

$$\phi(c_i) = 2 \cdot a \cdot \Lambda \int_0^{c_i} \frac{x \cdot g(x)}{\left(\Lambda - a \cdot \overline{G}(x)\right)^3} \, dx$$

(44)

Proof of Proposition 3.1. Lemma A.3. implies that $f_i = 2 \cdot a \cdot \Lambda \int_0^{c_i} \frac{x \cdot g(x)}{\left(\Lambda - a \cdot \overline{G}(x)\right)^3} dx$. Then, direct verification yields $\frac{d f_i}{d \Lambda} < 0$ thereby establishing that User *i*'s fee decreases in Scale, Λ.

Lemma A.3. also establishes that User *i* pays $\phi(c_i)$ in equilibrium with $\phi(c_i)$ being a strictly increasing function. Accordingly, the set of users paying a higher fee than User *i* in equilibrium is given by $\{j: \phi(c_j) > \phi(c_i)\} = \{j: c_j > c_i\}$ and each user is drawn to be of such type with probability $\overline{G}(c_i)$ so that transactions with fees exceeding that of User *i*'s fee arrive according to a Poisson process with rate $a \cdot \overline{G}(c_i)$. Note that the hypothesis of Lemma A.2. are all satisfied so that Lemma A.2. then implies that User *i*'s expected time for her transaction to be included on a block (i.e., processing time) is given by: $W_P(f_i) = \frac{\Lambda}{\left(\Lambda - a \cdot \overline{G}(c_i)\right)^2}$.

Direct verification yields $\frac{d}{d\Lambda}[W_P(f_i)] < 0$, establishing that User *i*'s processing time decreases in Scale, Λ and thereby completing the proof.

Proof of Proposition 3.2. We first establish that $k(s, \alpha) := min\{k \in \mathbb{N}: p(s,k) \leq \alpha\}$ is well-defined for any $\alpha > 0$ and then prove that $k(s, \alpha)$ decreases in *s* for any $\alpha > 0$. To establish that $k(s, \alpha)$ is well-defined, we only need to show that $\{k \in \mathbb{N}: p(s,k) \leq \alpha\}$ possesses a well-defined minimum for any $\alpha > 0$. Recall that for any *s*, we have $\lim_{k \to \infty} p(s,k) = 0$ and $p(s,k)$ decreases in *k* so that $\{k \in \mathbb{N}: p(s,k) \leq \alpha\}$ is a nonempty set. Let $\tilde{k} \in \mathbb{N}$ be such that $\tilde{k} \in \{k \in \mathbb{N}: p(s,k) \leq \alpha\}$. Then $inf\{k \in \mathbb{N}: p(s,k) \leq \alpha\} = inf\{k \in \mathbb{N}, k \leq \tilde{k}: p(s,k) \leq \alpha\} = min\{k \in \mathbb{N}, k \leq \tilde{k}: p(s,k) \leq \alpha\} = min\{k \in \mathbb{N}: p(s,k) \leq \alpha\}$, with the second equality following from $\{k \in \mathbb{N}, k \leq \tilde{k}: p(s,k) \leq \alpha\}$ being

nonempty and compact. Thus, $k(s,\alpha) := min\{k \in \mathbb{N}: p(s,k) \leq \alpha\}$ is well-defined for any $\alpha > 0$ as desired.

To establish $k(s,\alpha)$ decreases in *s* for any $\alpha > 0$, we need to show that $k(s+\delta,\alpha) \leq k(s,\alpha)$ for $\delta \geq 0$. Note that $\alpha \geq p\big(s,k(s)\big) \geq p(s+\delta,k(s))$ with the first inequality following from the definition of $k(s)$ and the second inequality following from $p(s,k)$ decreasing in *s*. Then $k(s,\alpha) \in \{k \in \mathbb{N}: p(s+\delta,k) \leq \alpha\}$ so that $k(s+\delta,\alpha) := min\{k \in \mathbb{N}: p(s+\delta,k) \leq \alpha\} \leq k(s,\alpha)$, which completes the proof.

Proof of Corollary 3.3. $\kappa = \frac{k(s)}{\Lambda}$ decreases in security, *s*, as an implication of Proposition 3.2 and decreases in Λ via direct verification (i.e., $\frac{d}{d\Lambda}[\frac{k(s)}{\Lambda}] = -\frac{k(s)}{\Lambda^2} < 0$).

Proof of Corollary 3.4. Proposition 3.1, Equation 6 and Equation 7 imply that user utility is given by:

$$u_i \cdot \Psi\big(N_{\tau_i}\big) - c_i \cdot \left(\frac{\Lambda}{\Big(\Lambda - a \cdot \overline{G}(c_i)\Big)^2} + \frac{k(s)}{\Lambda} \right) - 2 \cdot a \cdot \Lambda \int_0^{c_i} \frac{x \cdot g(x)}{\Big(\Lambda - a \cdot \overline{G}(x)\Big)^3}\, dx \tag{45}$$

Then, direct verification reveals that user utility increases in adoption, N_{τ_i}, and scale, Λ. Moreover, Proposition 3.2 implies that $k(s)$ decreases in security, *s*, so that user utility also necessarily decreases in security, *s*. This last implication follows because Equation 45 decreases in $k(s)$.

Proof of Proposition 3.5. $N_{\tau,1} \geq N_{\tau,2}$ for all $\tau \in T$, $\Lambda_1 \geq \Lambda_2$, and $\frac{\Lambda_1}{k(s_1)} \geq \frac{\Lambda_2}{k(s_2)}$ implies that Blockchain 1 is preferred to Blockhain 2 for all users by Proposition 3.1, Corollary 3.3 and the preference structure given by Equation 6 (see the proof of Corollary 3.4 for additional detail).

To establish the converse, we proceed by contradiction. Formally, we wish to show that Blockchain 2 being preferred to Blockchain 1 by some user under some model parameters is inconsistent with the disjunction of the following three propositions:

(i) $N_{\tau,1} < N_{\tau,2}$ for all $\tau \in T$

(ii) $\Lambda_1 < \Lambda_2$

(iii) $\frac{\Lambda_1}{k(s_1)} < \frac{\Lambda_2}{k(s_2)}$

Blockchain 1 being preferred to Blockchain 2 for each user under all circumstances means that users of any type (i.e., for any values of u_i, c_i and τ_i) receive higher utility from Blockchain 1 than Blockchain 2 for any transaction arrival rate (i.e, for any a). More explicitly, letting the equilibrium user utility of User *i* when using blockchain *j* be denoted by $\mathbb{U}_{i,j}$, then Proposition 3.1, Equation 6 and Equation 7 imply:

$$\mathbb{U}_{i,j} = u_i \cdot \Psi\left(N_{\tau_i,j}\right) - c_i \cdot \left(\frac{\Lambda_j}{\left(\Lambda_j - a \cdot \overline{G}(c_i)\right)^2} + \frac{k(s_j)}{\Lambda_j} \right) - 2 \cdot a$$

$$\cdot \Lambda_j \int_0^{c_i} \frac{x \cdot g(x)}{\left(\Lambda_j - a \cdot \overline{G}(x)\right)^3} \, dx$$

so that we must prove that $U_{i,1} \geq U_{i,2}$ for all u_i, c_i, τ_i and a is inconsistent with the disjunction of (i) – (iii). As discussed, we proceed by contradiction, imposing $U_{i,1} \geq U_{i,2}$ for all u_i, c_i, τ_i

and a throughout the following argument and establishing a contradiction with the disjunction of (i) – (iii).

Letting $c_i = 0$ precludes $N_{\tau_i,1} < N_{\tau_i,2}$ so that at least one of (ii) or (iii) must hold. Then, holding c_i fixed but taking $a \to \infty$ and setting $u_i = 0$ precludes $\Lambda_1 < \Lambda_2$ so that (iii) must hold. However, holding c_i fixed but taking $a \to 0^+$ and setting $u_i = 0$ precludes $\frac{\Lambda_1}{k(s_1)} < \frac{\Lambda_2}{k(s_2)}$, thereby delivering the desired contradiction and completing the proof.

Appendix B: Proofs for Chapter 3

Deriving the value of the option to invest, abandon, and mothball mining equipment.

The investment and abandonment decision made by the miner takes the price as given by the following geometric Brownian motion:

$$dP = \alpha\, Pdt + \sigma\, Pdz.$$

(46)

To obtain a differential equation for $V_0(P)$ we recognize that for the value of the option to be arbitrage free, the left-hand side must equal the right-hand side of the following expression:

$$\rho\, V_0(P)dt = E(dV_0)$$

(47)

where the required rate of return, ρ, on the value of the option $V_0(P)$ over dt must equal the expected rate of appreciation in the value of the option to "mine," $E(dV_0)$.

Using a Taylor expansion on the right-hand-side of our non-arbitrage condition, we have:

$$\rho\, V_0(P)dt = E[V_0'(P)dP + \frac{1}{2}V_0''(P)(dP)^2].$$

(48)

Employing Ito's lemma and distributing the expectation operator yields:

$$\rho\, V_0(P)dt = V_0'(P)\alpha\, Pdt + \frac{1}{2}V_0''(P)\sigma^2P^2dt.$$

(49)

Dividing by dt and moving everything to the left-hand side:

$$\frac{1}{2}\sigma^2 P^2 V_0''(P) + \alpha\, P V_0'(P) - \rho\, V_0(P) = 0.$$

(50)

Recognizing that $(\rho - \alpha)$ is often designated as the convenience yield δ:

$$\frac{1}{2}\sigma^2 P^2 V_0''(P) + (\rho - \delta) P V_0'(P) - \rho\, V_0(P) = 0.$$

(51)

The solution for this second order differential equation will take the form of $V_0(P) = AP^{\beta}$, and will eventually provide two solutions for β:

$$\beta_1 = \left[\frac{1}{2} - \frac{(\rho - \delta)}{\sigma^2}\right] + \left(\left[\frac{(\rho - \delta)}{\sigma^2} - \frac{1}{2}\right]^2 + \left[\frac{2\rho}{\sigma^2}\right]\right)^{1/2}$$

(52)

and

$$\beta_2 = \left[\frac{1}{2} - \frac{(\rho - \delta)}{\sigma^2}\right] - \left(\left[\frac{(\rho - \delta)}{\sigma^2} - \frac{1}{2}\right]^2 + \left[\frac{2\rho}{\sigma^2}\right]\right)^{1/2}$$

(53)

Since $V_0 = A_1 P^{\beta_1}$ and $V_0 = A_2 P^{\beta_2}$ each satisfy the second order differential equation for $V(P)$, their sum must also satisfy the second order differential equation. Consequently, for a general solution we have:

$$V_0(P) = A_1 P^{\beta_1} + A_2 P^{\beta_2}$$

(54)

with $\beta_1 > 1$ and $\beta_2 < 0$.

The general solution for $V_0(P)$, $V_0(P) = A_1 P^{\beta_1} + A_2 P^{\beta_2}$, becomes simply $V_0(P) = A_1 P^{\beta_1}$ (with $\beta_2 < 0$) because at very low prices the value of the miner's option to invest gets

very small forcing A_2 to be zero to eliminate the possibility of P^{β_2} with $(\beta_2 < 0)$ approaching $+\infty$ and inappropriately impacting the value of the miner's option. The solution for the value of the option of an idle miner to become active is valid over the interval $(0, P_H)$ of prices.

Next, consider the value of an active miner. The calculation is nearly identical to that of $V_0(P)$ except the "live" part of the composite asset pays a net cash flow of $(P - C)dt$, in which case the differential equation representing the value of the active miner becomes:

$$\frac{1}{2}\sigma^2 P^2 V_1''(P) + (r - \delta)PV_1'(P) - rV_1(P) + (P - C) = 0.$$

(55)

The general solution for $V_1(P)$ would be:

$$V_1(P) = B_1 P^{\beta_1} + B_2 P^{\beta_2} + (\frac{P}{\delta} - \frac{C}{r}).$$

(56)

The last two terms in our solution for $V_1(P)$, $(\frac{P}{\delta} - \frac{C}{r})$, are the value of the "live" project when the miner is required to keep operating over the production horizon despite any losses. The first two terms are the value of the option to abandon mining. As P gets very large, the option to abandon should not be unduly influenced by P^{β_1} where $\beta_1 > 1$. Consequently, B_1 is given the value of zero so that P^{β_1} has no impact on the value of the miner's option to abandon. This means that $V_1(P)$ becomes:

$$V_1(P) = B_2 P^{\beta_2} + \frac{P}{\delta} - \frac{C}{r}$$

(57)

over the range of prices from (P_S, ∞).

Substituting these solutions into Equations (9) to (12) (see Chapter 3, Section 3) yields the following four equations:

$$-A_1 P_H^{\beta_1} + B_2 P_H^{\beta_2} + \frac{P_H}{\delta} - \frac{C}{r} = I \quad (58)$$

$$-\beta_1 A_1 P_H^{\beta_1 - 1} + \beta_2 B_2 P_H^{\beta_2 - 1} + \frac{1}{\delta} = 0 \quad (59)$$

$$-A_1 P_S^{\beta_1} + B_2 P_L^{\beta_2} + \frac{P_S}{\delta} - \frac{C}{r} = -E \quad (60)$$

$$-\beta_1 A_1 P_S^{\beta_1 - 1} + \beta_2 B_2 P_S^{\beta_2 - 1} + \frac{1}{\delta} = 0. \quad (61)$$

These equations may be solved for numerical solutions for the thresholds P_H and P_S, and the corresponding option value coefficients A_1 and B_2.

The derivation of the general solution for the value of the mothballed miner is much like that of the active miner and idle miner. The mothballed miner can continue over the range of prices (P_S, P_R) and the value of the "laid-up" computing equipment would be given:

$$V_M(P) = D_1 P^{\beta_1} + D_2 P^{\beta_2} - \frac{M}{r} \quad (62)$$

where $D_1 P^{\beta_1}$ is the value of the miner's option to reactivate and $D_2 P^{\beta_2}$ is the value of the miner's option to scrap the equipment. Clearly M/r is the cost of maintaining the laid-up equipment forever.

Substituting these solutions into Equations (18) to (21) (see Chapter 3, Section 3) yields the following eight equations:

$$-D_1 P_R^{\beta_1} + (B_2 - D_2) P_R^{\beta_2} - (C - M)\frac{1}{r} + \frac{P_R}{\delta} = R \quad (63)$$

$$-\beta_1 D_1 P_R^{\beta_1 - 1} + \beta_2 (B_2 - D_2) P_R^{\beta_2 - 1} + \frac{1}{\delta} = 0 \quad (64)$$

$$-D_1 P_M^{\beta_1} + (B_2 - D_2) P_M^{\beta_2} + \frac{P_M}{\delta} - (\frac{C}{r} - \frac{M}{r}) = -E_M \quad (65)$$

$$-\beta_1 D_1 P_M^{\beta_1 - 1} + \beta_2 (B_2 - D_2) P_M^{\beta_2 - 1} + \frac{1}{\delta} = 0 \quad (66)$$

$$-A_1 P_H^{\beta_1} + B_2 P_H^{\beta_2} + \frac{P_H}{\delta} - \frac{C}{r} = I \quad (67)$$

$$-A_1 \beta_1 P_H^{\beta_1 - 1} + B_2 \beta_2 P_H^{\beta_2 - 1} + \frac{1}{\delta} = 0 \quad (68)$$

$$(D_1 - A_1) P_S^{\beta_1} + D_2 P_S^{\beta_2} + \frac{M}{r} = -E_S \quad (69)$$

$$\beta_1 (D_1 - A_1) P_S^{\beta_1 - 1} + \beta_2 D_2 P_S^{\beta_2 - 1} = 0. \quad (70)$$

These eight equations may be solved for numerical solutions for the thresholds P_H, P_R, P_M, and P_S, and the corresponding option value coefficients A_1, B_2, D_1, and D_2.

Appendix C: Discussion of Structural Breaks for Chapter 3

The introduction of BTC ASIC mining equipment occurred sometime in 2013. This ASIC equipment was vastly more efficient than the prior mining technology, and its advent represents a paradigm shift in the BTC mining industry. As I seek to distinguish the reactions of ASIC and GPU miners, I consider only BTC data after the widespread adoption of ASIC equipment. Though the first ASIC equipment was sold sometime in 2013, the time-series relationship between price and hashrate indicates that widespread adoption of this equipment did not immediately occur (see Figure 18). This is likely symptomatic of greater demand for the equipment than available supply and imperfect supply chains.

To isolate the "start" date of my sample, I turn to Gregory and Hansen (1996). Gregory and Hansen test for cointegration in the presence of a regime shift at an unknown date. I utilize the search function of their test to isolate the date of this regime shift. The results of the Gregory and Hansen (1996) test are presented in Table 24. Given evidence of a regime shift on 04/13/2014, I start my BTC sample on 04/14/2014. Graphical evidence (see Figure 18) supports this choice. Though the Gregory and Hansen (1996) test also measures cointegration between BTC price and hashrate, I ignore these test statistics as I re-estimate the cointegration relationship using the Johansen (1991, 1995) procedure (see Chapter 3, Section 5).

Figures

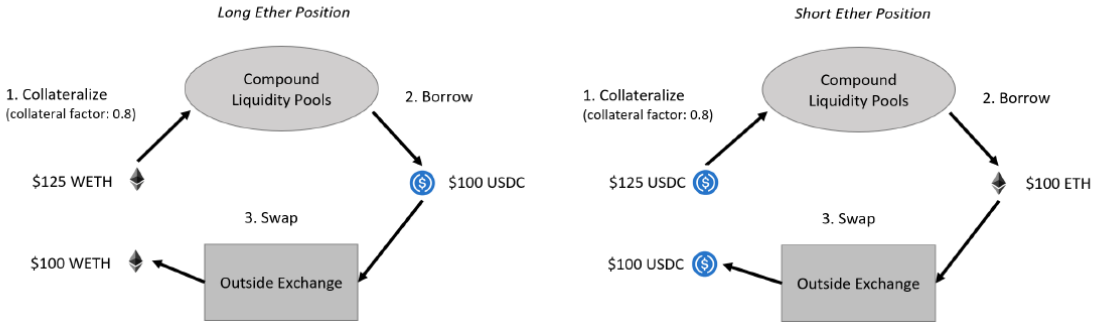

Figure 1: Long and Short Ether Example

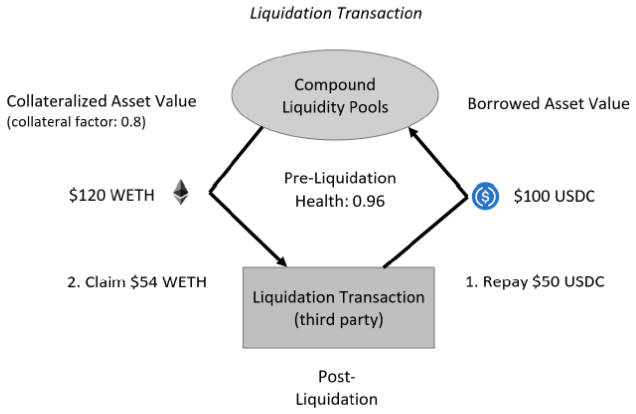

Figure 2: Liquidation Transaction Example

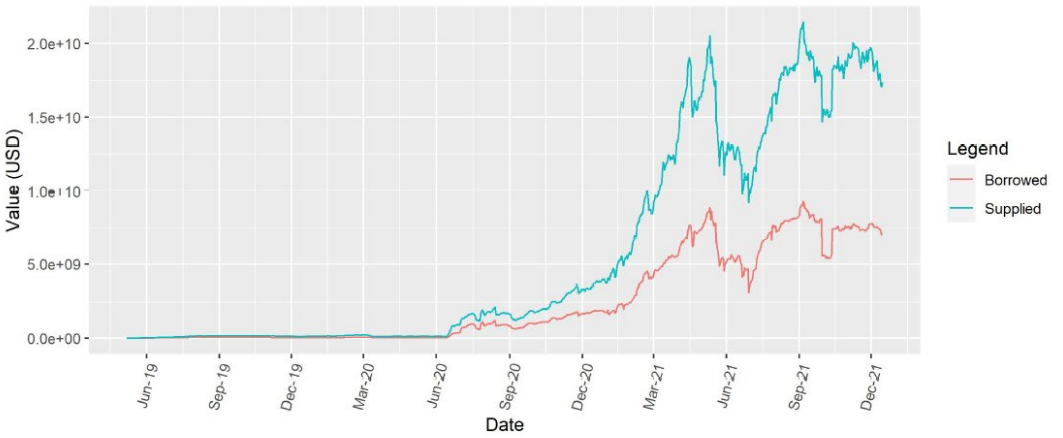

Figure 3: Borrowed and Supplied Asset Values

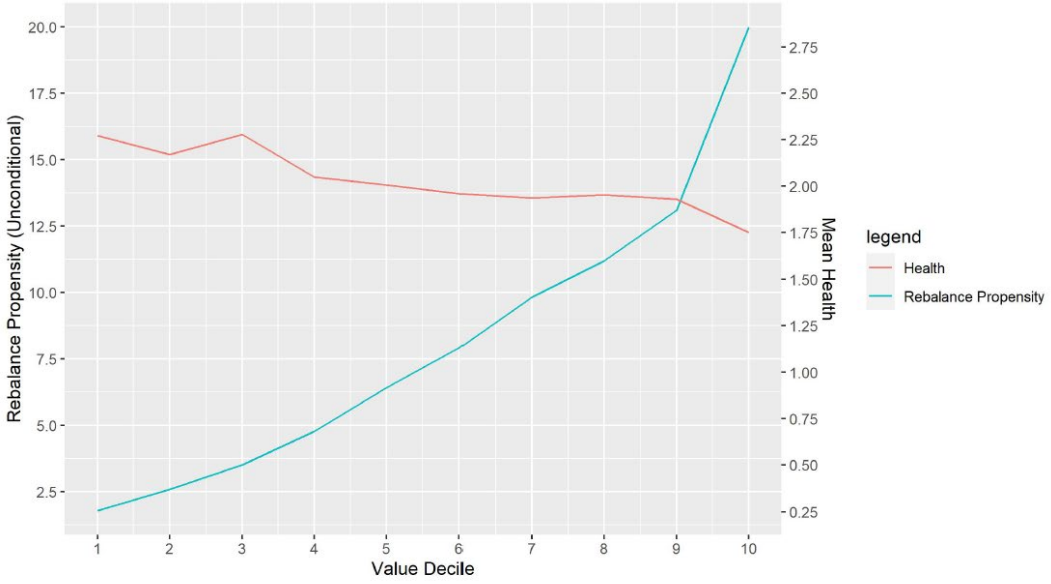

Figure 4: Value Decile Statistics

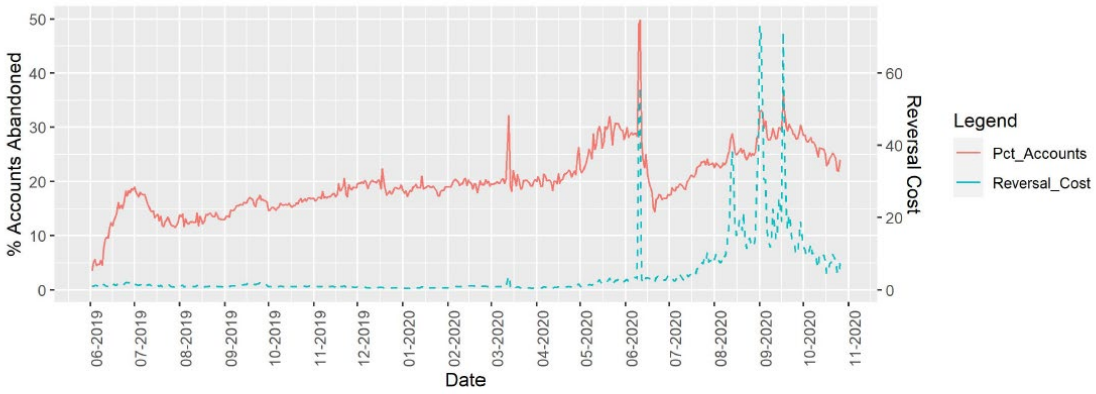

Figure 5: Percent of Accounts with Negative Reversal Benefit

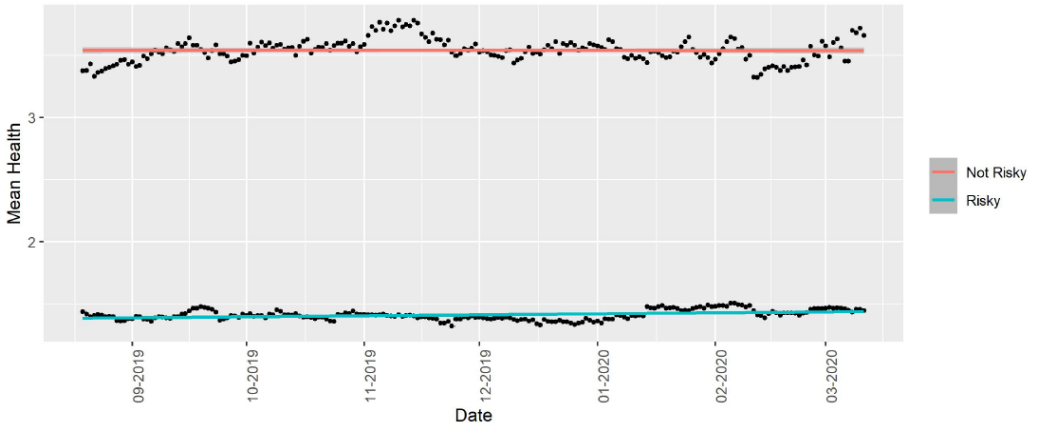

Figure 6: Parallel Pre-Trends

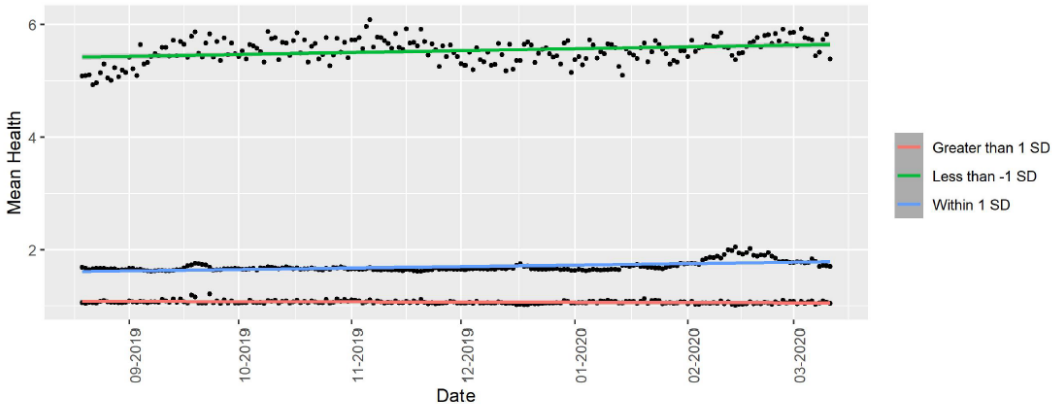

Figure 7: Continious Treatment Parallel Pre-Trends

Figure 8: Blockchain announcement dates and native asset market capital.

This figure shows a scatterplot of blockchains' native asset market capital (log-scale) against blockchain announcement dates. Data are obtained from cryptoslate.com as of February 21, 2021. Whenever the announcement date of a blockchain is not disclosed on cryptoslate.com, we take the timestamped date of the genesis block as the launch date of the blockchain.

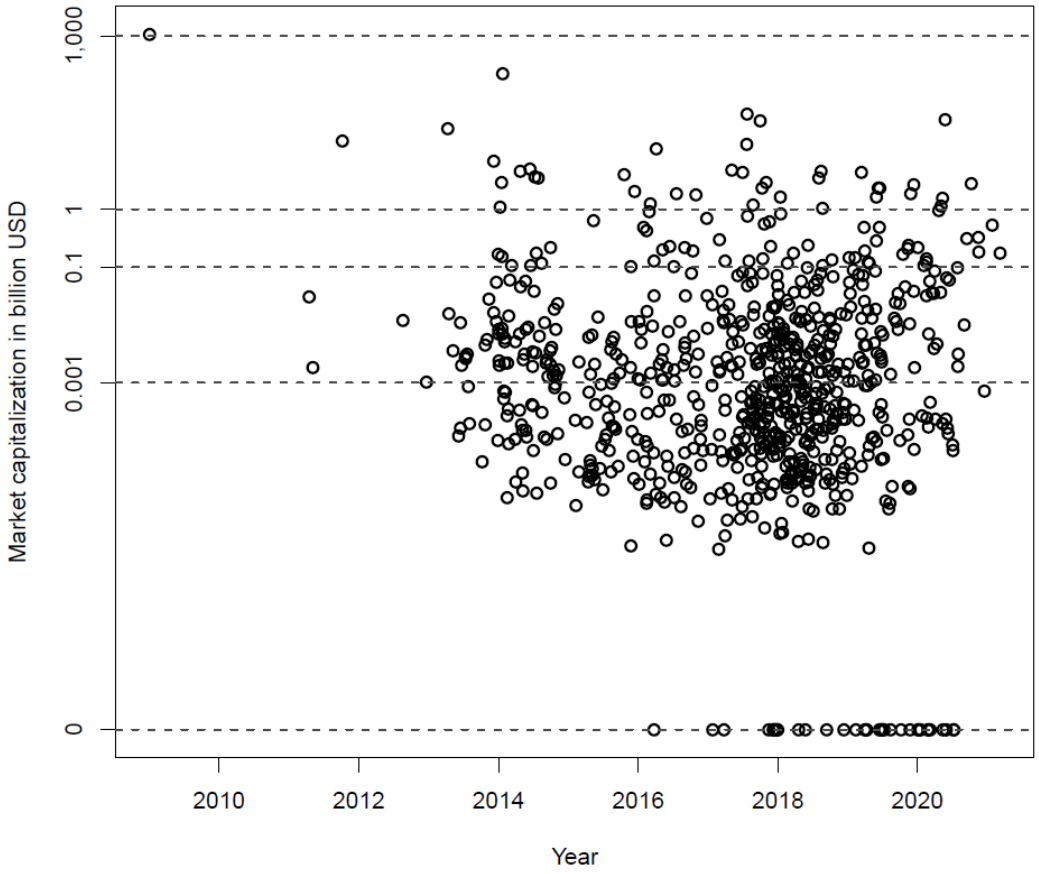

Figure 9: Number of blockchains and market capitalization by consensus protocol.

This figure is a snapshot of the distribution of the raw number of introduced public blockchains by consensus protocol (left) and the native asset market capital share as of February 21, 2021 (right). The full sample consists of 783 blockchains with native cryptoassets listed at cryptoslate.com.

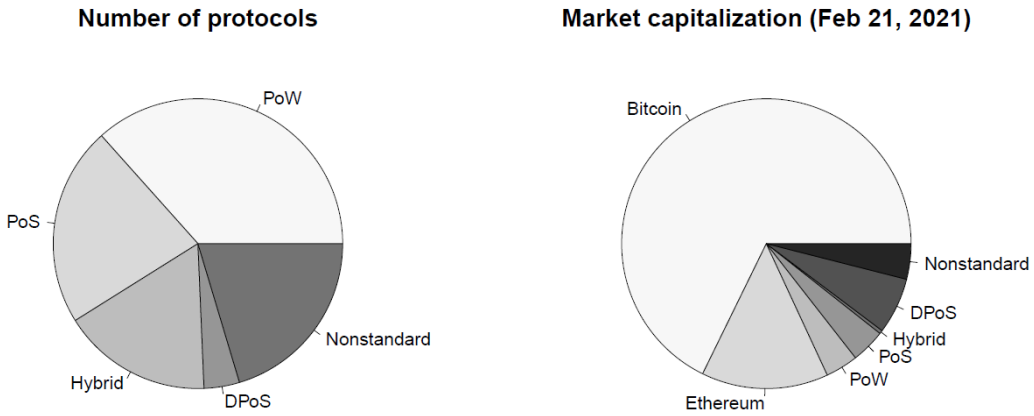

Figure 10: Market capital share within consensus protocol group.

This figure shows native asset market capital share distributions within consensus protocol groups for 783 public blockchains. Data are from cryptoslate.com and are as of February 21, 2021. Major blockchains and their native asset ticker symbol are next to the market capital share. 'Other PoW' excludes Bitcoin and Ethereum.

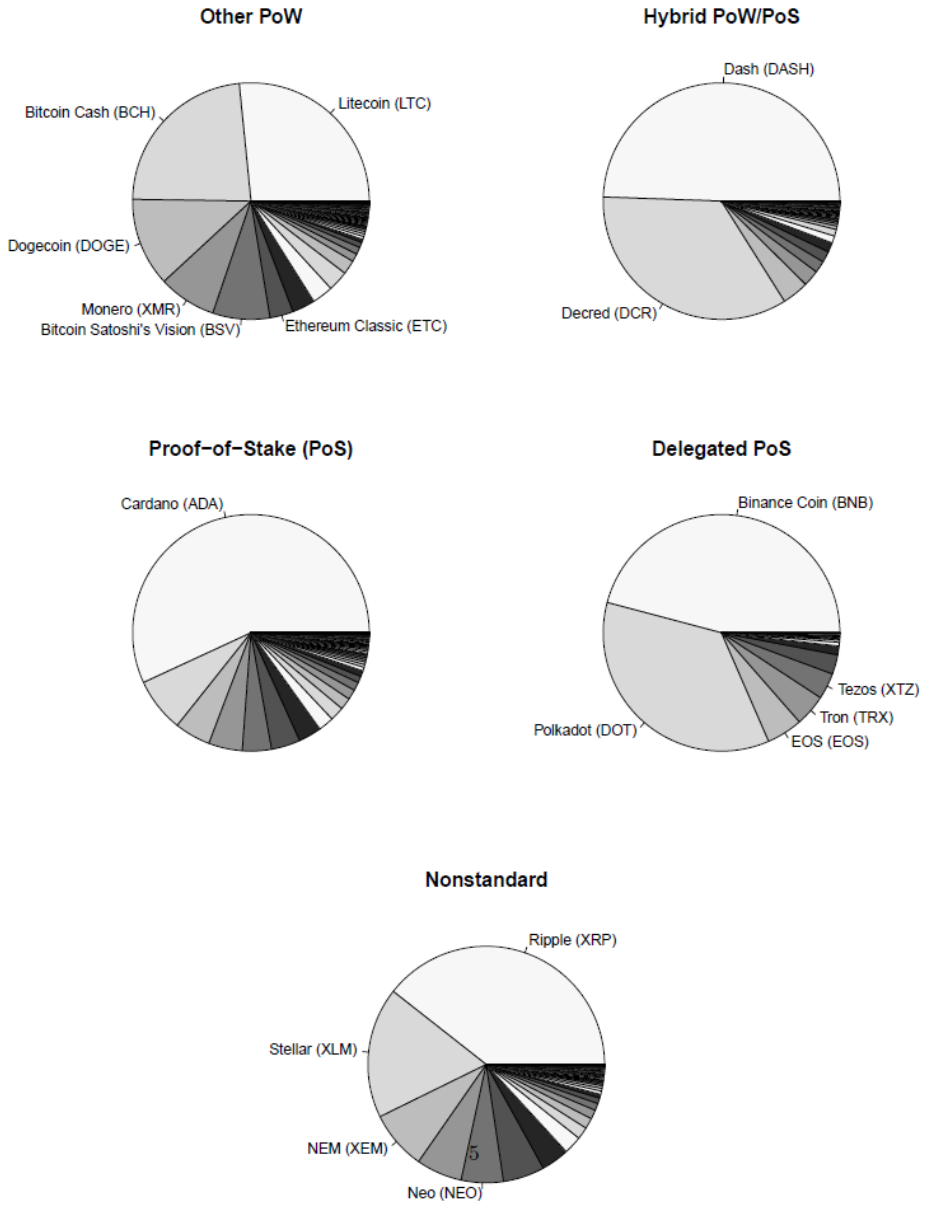

Figure 11: Adoption: Number of active addresses.

This figure displays the adoption metric for users of the payment type for each public blockchain in the coinmetrics.io sample. Blockchains are arranged along the x-axis (in no particular order) and are represented by the ticker symbol for their native cryptoassets. The shape of each data point indicates its consensus protocol: Proof-of-Work (PoW), Proof-of-Stake (PoS), Nonstandard, Hybrid PoW/PoS (Hybrid), and Delegated Proof-of-Stake (DPoS). The y-axis represents values of the 30 day average of the number of active addresses as of February 21, 2021 on a log scale. A higher number of active addresses corresponds to a more widely-adopted blockchain.

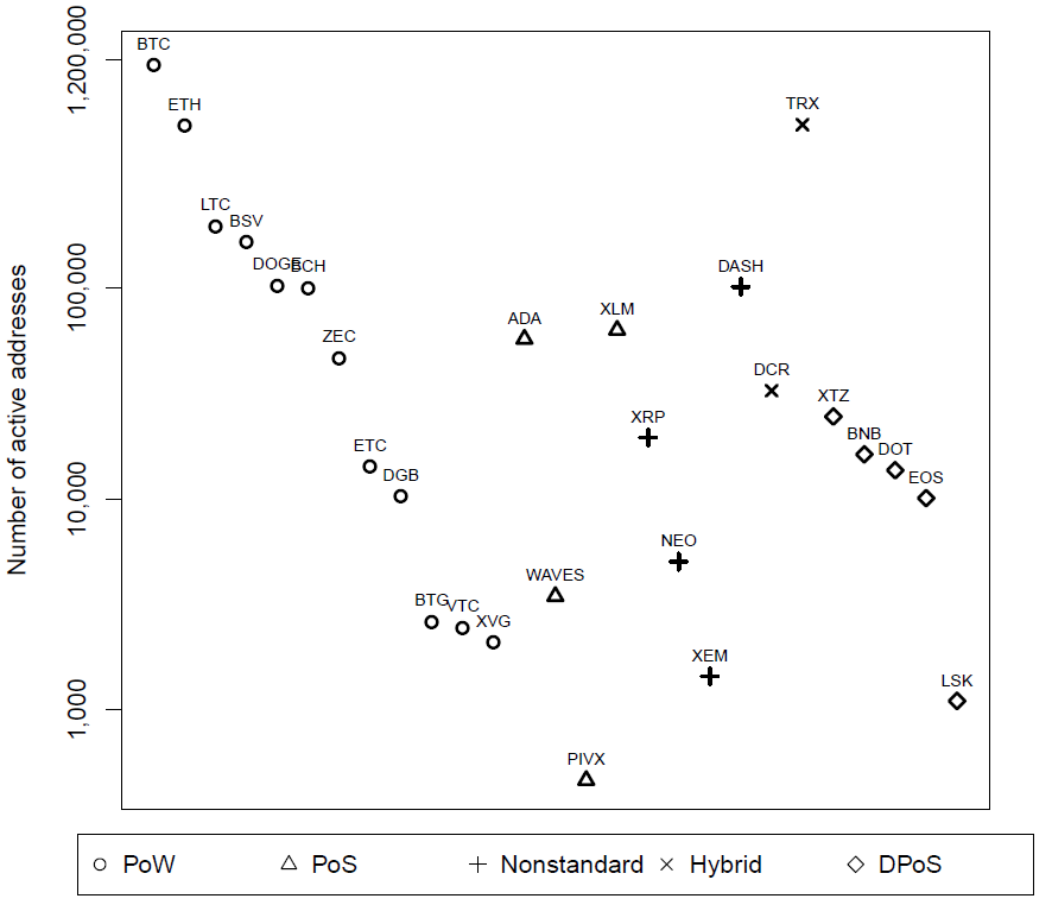

Figure 12: Overview of decentralized applications (dApp)

This figure displays the proportion of dApps and dApp Users in each relevant category. Data for ETH, EOS and TRON are collected from stateofthedapps.com, while data for NEO and WAVES are collected from dappradar.com. Dappradar.com categorization is hand-matched to stateofthedapps.com categorization. MAU is the sum of the unique monthly active users on each dApp over a 30 day period (from 22 January, 2021 to 20 February, 2021) for each category.

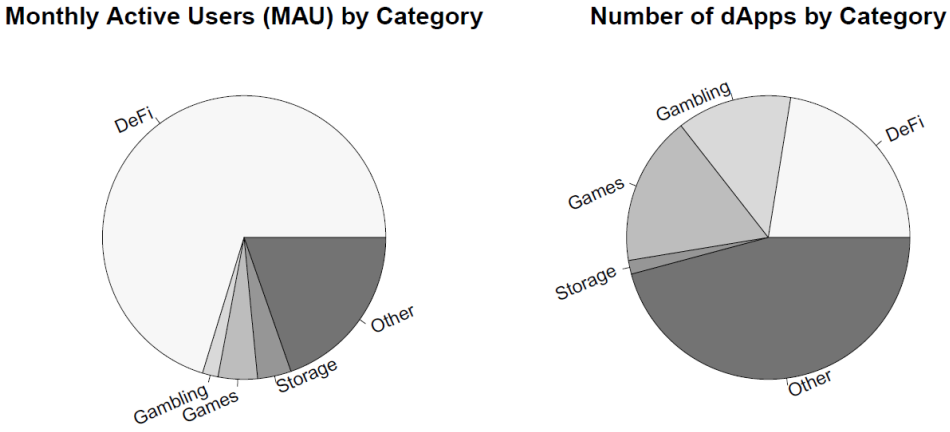

Figure 13: Scale: Transactions per Second (TPS).

This figure displays the scale metric for each public blockchain in the coinmetrics.io sample. Blockchains are arranged along the x-axis (in no particular order) and are represented by the ticker symbol for their native cryptoassets. The shape of each data point indicates its consensus protocol: Proof-of-Work (PoW), Proof-of-Stake (PoS), Nonstandard, Hybrid PoW/PoS (Hybrid), and Delegated Proof-of-Stake (DPoS). Scale is measured as the maximum of daily transaction counts denominated in transactions per second (TPS). The y-axis represents values of TPS on a log scale. A higher value of TPS corresponds to a higher scalability of a blockchain.

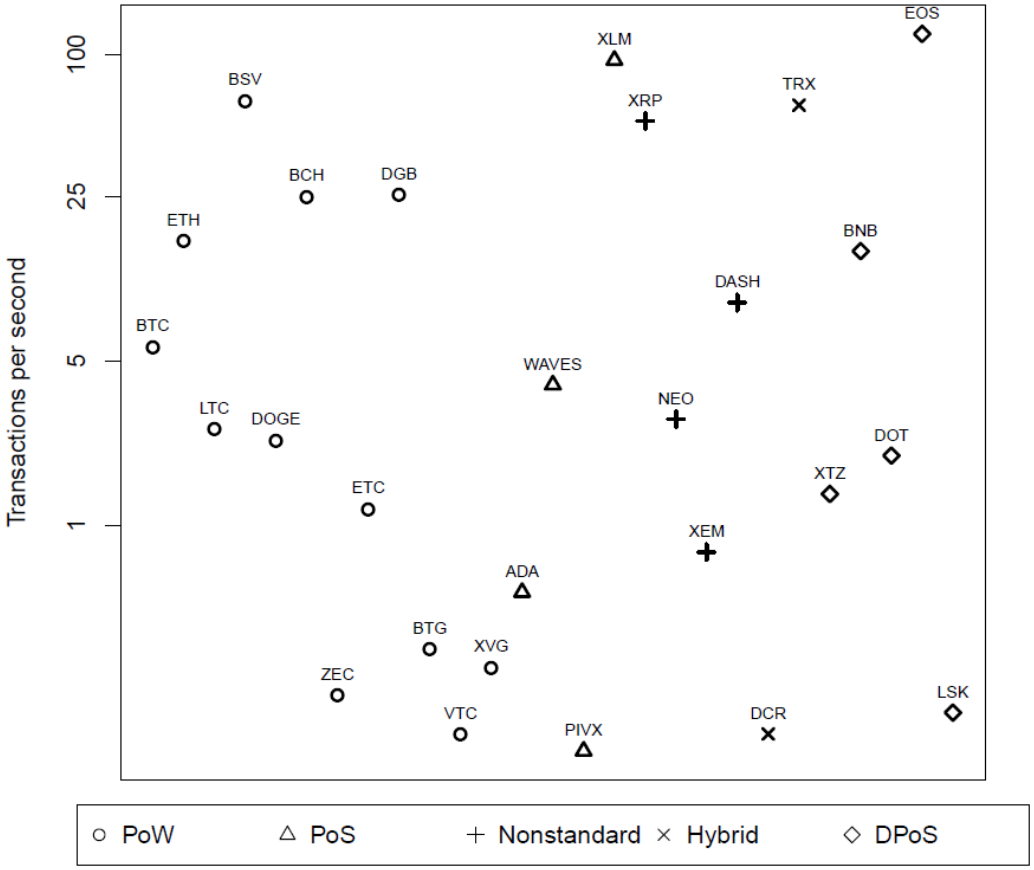

Figure 14: Security: Number of block confirmations.

This figure displays the security metric for each public blockchain in the coinmetrics.io sample. Blockchains are arranged along the x-axis (in no particular order) and are represented by the ticker symbol for their native cryptoassets. Security is measured as the median number of block confirmations required by major cryptoasset exchanges to settle a transaction. Appendix ?? reports data sources and confirmation numbers of individual cryptoassets for all 32 exchanges. The y-axis represents values of block confirmation rules on a reversed log scale. Lower values indicate better security for a blockchain.

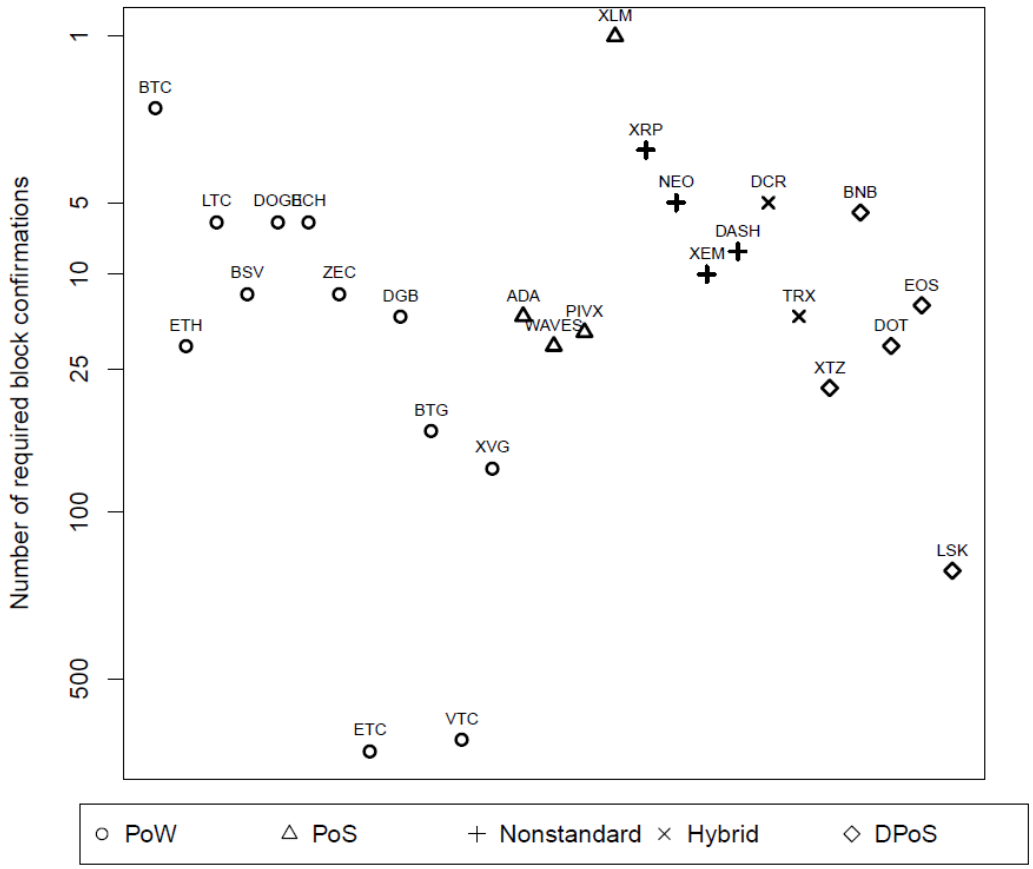

Figure 15: Comparative Statics

This figure presents a comparative static analysis of the effects of various characteristic inputs on the optimal activation, reactivation, mothballing, and abandonment thresholds of cryptocurrency miners.

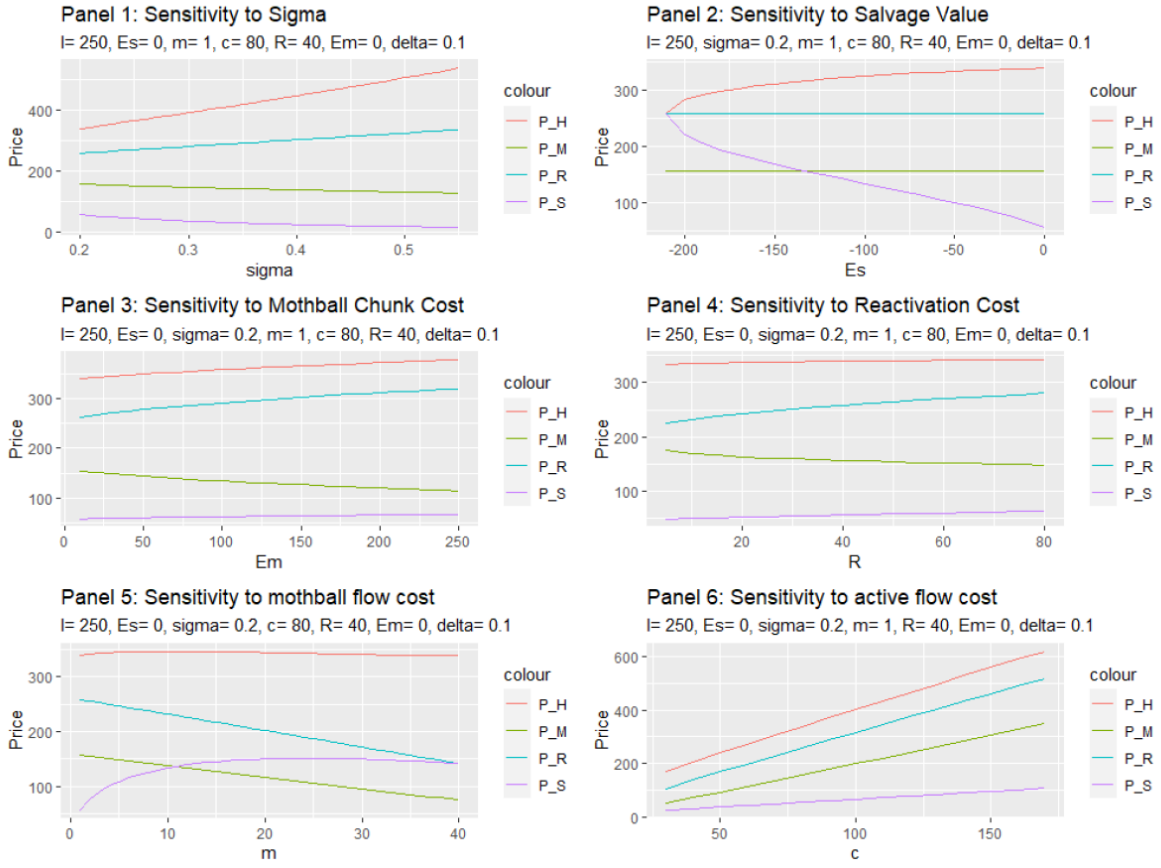

Figure 16: Bitcoin Graphics

This figure plots BTC price, hashrate (measured in terahashes/second), and their respective first differences over time. All variables are log transformed. The sample period is from April 25, 2014 to Oct. 19, 2019.

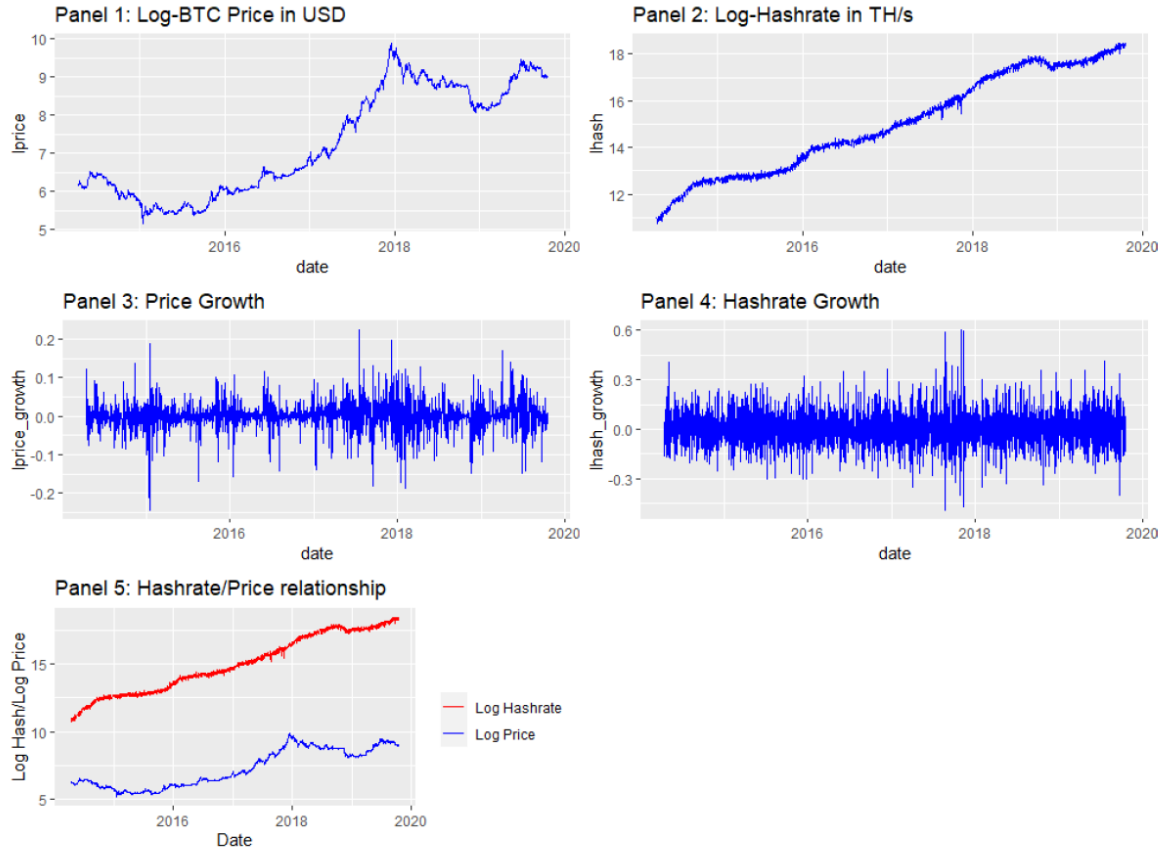

Figure 17: Ether Graphics

This figure plots ETH price, hashrate (measured in terahashes/second), and their respective first differences over time. All variables are log transformed. The sample period is from Oct. 25, 2016 to Oct. 19, 2019.

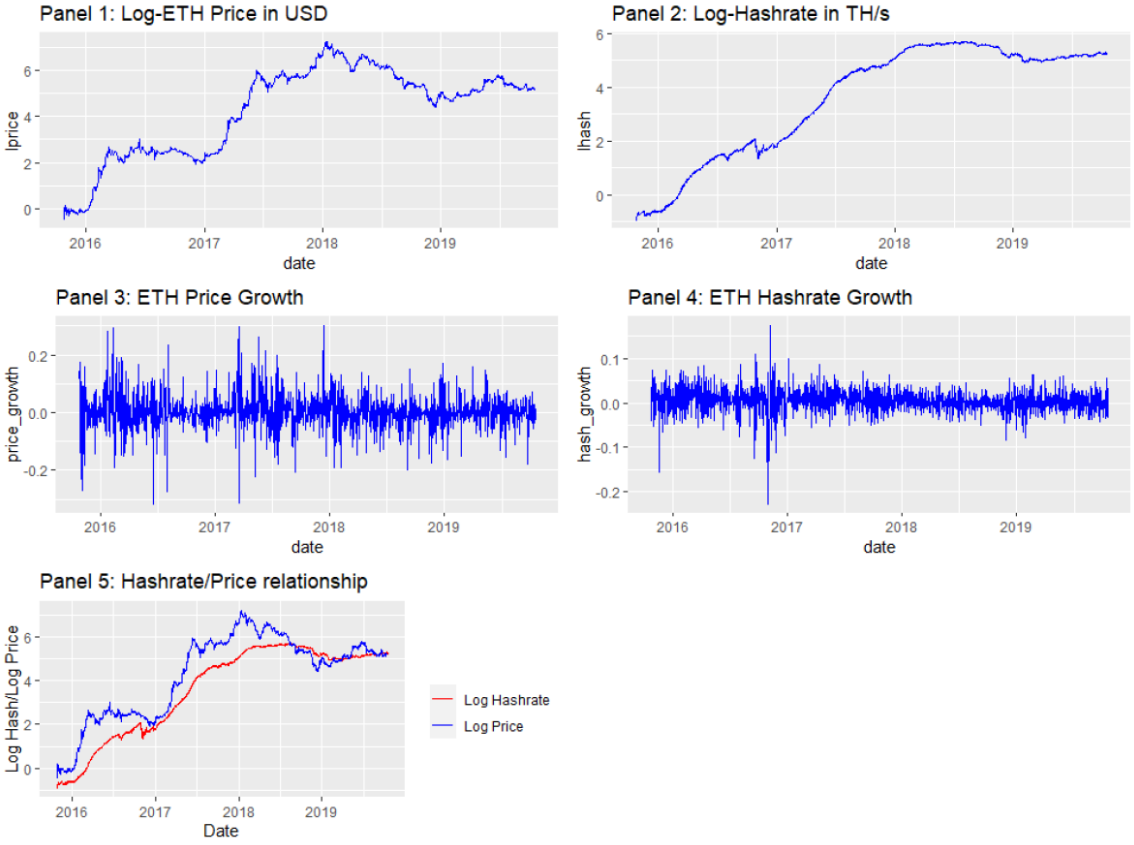

Figure 18: Bitcoin Price/Hashrate

This figure plots BTC price and hashrate (measured in terahashes/second) over time. The variables are log transformed. The sample period is from July 18, 2010 to Oct. 19, 2019.

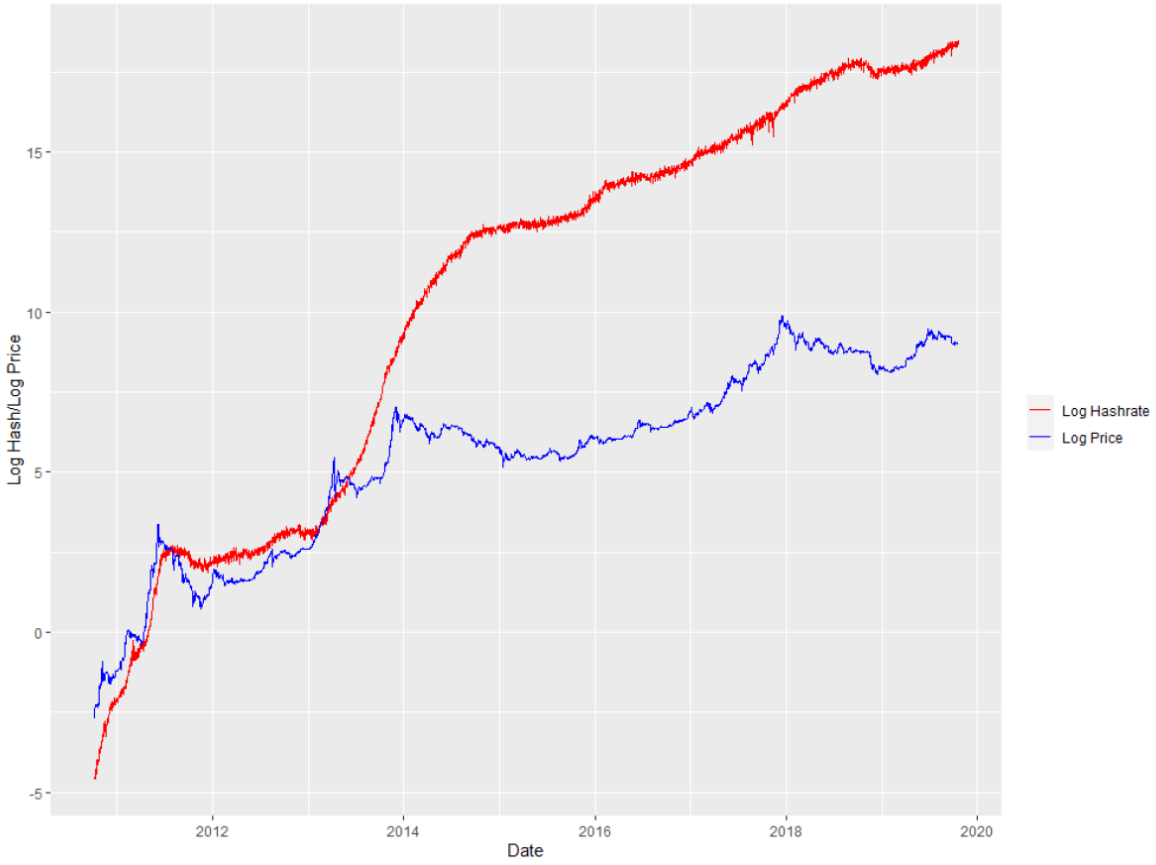

Tables

Table 1: Compound Transaction Gas Estimates

This table presents estimates of different types of transactions calling Compound's smart contracts. Estimates are obtained from the Compound Finance documentation.

Function	Estimated Gas Cost
Lending (mint c-token)	150,000
Redeem Collateral (when borrowing)	250,000
Redeem Collateral (when not borrowing)	90,000
Borrowing	300,000
Repay Borrow	90,000
Liquidation	400,000

Table 2: Asset Summary Statistics

This table presents market-level summary statistics for each asset borrowed on Compound. Data are obtained from the Compound Finance public API.

Asset:	USDC	ETH	WBTC2	DAI	USDT	TUSD
Collateral Factor (Min)	0.75	0.75	0.65	0.75	0	0
Collateral Factor (Max)	0.8	0.8	0.65	0.8	0	0
Borrow Rate (Min)	0.015575	0.003446781	0.03117656	0.00332222	0.004797073	1.1e-23
Borrow Rate (Max)	0.2066547	0.05293333	0.06537511	0.2208212	0.2549495	0.09860798
Supply Rate (Min)	0.000941926	2.06e-05	0.000755311	0.001715456	0.000366504	0
Supply Rate (Max)	0.1721626	0.01204572	0.009805318	0.2091978	0.2198541	0.07571252
Liquidity (Min)	1,235.85	2,708.745	720,956,690	15,607,475	8,514.772	0
Liquidity (Max)	5,175,484,024	7,745,777,020	2,911,268,207	5,015,031,843	900,754,697	127,502,385
Borrows (Min)	100.0001	166.2327	50,762,454	8,443,041	4,000.865	0
Borrows (Max)	4,233,697,120	411,820,480	259,788,019	3,981,437,150	745,026,013	87,958,547
Utilization (Mean)	0.5799558	0.04142269	0.09383332	0.7586446	0.6736732	0.6208167
Stablecoin	Yes	No	No	Yes	Yes	Yes

Asset:	BAT	UNI	COMP	ZRX	MKR	WBTC
Collateral Factor (Min)	0.6	0.6	0.6	0.6	0	0.4
Collateral Factor (Max)	0.65	0.7	0.65	0.65	0.55	0.75
Borrow Rate (Min)	0.02235929	0.003677432	0.0202767	0.02113267	0.02307397	0.0209834
Borrow Rate (Max)	0.3571833	0.2903595	0.740308	0.2701064	0.1330994	0.22614
Supply Rate (Min)	8.23e-05	0.000215608	5.31e-06	5.73e-05	0	0
Supply Rate (Max)	0.2682767	0.1467786	0.4468037	0.1704487	0.03801719	0.1192228
Liquidity (Min)	786.84	3,596,418	243,939.6	220.7863	0	9,752,836
Liquidity (Max)	339,509,963	485,720,460	433,026,963	271,942,246	3,812,393	2,774,558,356
Borrows (Min)	57.08491	428,408.5	80.68284	25.45318	0	25,182.63
Borrows (Max)	301,554,799	170,236,182	76,067,197	44,016,473	667,986.4	202,225,263
Utilization (Mean)	0.1121674	0.1256877	0.215479	0.1563089	0.03777635	0.02953106
Stablecoin	No	No	No	No	No	No

Asset:	LINK	SUSHI	AAVE	SAI	YFI	REP
Collateral Factor (Min)	0	0	0	0	0	0
Collateral Factor (Max)	0.65	0.55	0.6	0.75	0.6	0.5
Borrow Rate (Min)	0.03929899	0.02307397	0.02307397	0.05523447	0.0230741	0.02020078
Borrow Rate (Max)	0.2548372	0.4851858	0.0851879	0.2154964	0.2279085	0.4146844
Supply Rate (Min)	0.001774273	0	3.96e-10	0	8.72e-09	0
Supply Rate (Max)	0.1451377	0.2818102	0.01418486	0.1855215	0.1157746	0.290577
Liquidity (Min)	4,792,347	3,751.862	5,127.473	0	5,129.904	293.786
Liquidity (Max)	225,966,204	2,309,049	14,551,611	39,496,646	2,916,506	37,863,755
Borrows (Min)	3,295,190	0	0.1710201	0	0.1065152	5.21e-15
Borrows (Max)	30,120,560	1,252,444	1,777,145	28,847,709	283,628.7	2,739,959
Utilization (Mean)	0.1413887	0.1937419	0.06520885	0.2559087	0.156291	0.2275341
Stablecoin	No	No	No	Yes	No	No

Table 3: Summary Statistics - Borrower Sample

This table presents the summary statistics concerning the sample of borrowers. Data are obtained from the Compound Finance public API. This represents all accounts with active borrowing positions between 05-15-2019 and 10-26-2020. Accounts with health less than 0.9 or greater than 10 are omitted.

Variable	Mean	Median	Min	Max
Health	2.026712	1.563868	0.9000173	9.999923
Value	382,566.50	869.7692	0.000000004353378	891,779,200
Volatility	0.005423868	0.00431545	0	0.1070818
USD GWEI	0.00002419212	0.00001102024	0.000001210732	0.00021468
Net Interest Rate	1.105438%	0.9506622%	-23.67828%	15.49476%
Update	0.08227875			
Unique Accounts	**12,386**			
Account-Day Observations	**906,917**			
Account Position Updates	**74,620**			

Table 4: Descriptive Analysis - Probability of Update

This table presents pooled results from fixed-effect linear probability model (LPM) and logistic (Logit) specifications. Data are obtained from the Compound Finance public API. This represents all borrowing accounts between 05-15-2019 and 10-26-2020. Accounts with health less than 0.9 or greater than 10 are omitted. All continuous variables are standardized. Standard errors for the LPM are clustered at the account level. The coefficients reported for the logistic specification are average marginal effects, along with their corresponding delta-method standard errors.

Dep. Var: Update_Dummy	**LPM**	**Logit**
USD_GWEI	-0.0096237***	-0.0599463***
	(0.0004546)	(0.0016292)
USD_value	0.0103732**	0.0125491***
	(0.0042141)	(0.0017207)
Health	-0.0117354***	-0.0467384***
	(0.0013057)	(0.0019987)
Volatility	0.0083916***	0.0363983***
	(0.0010493)	(0.0016114)
Net-Interest Rate	0.0000471	-0.0000199
	(0.0010685)	(0.0013849)
Account FE	Yes	Yes
Adj R^2 / Pseudo R^2	0.00315564	0.00884005
N-obs.	906,917	904,691
Unique-Addresses	12,386	11,389

Table 5: Value Decile Analysis - Logit

This table presents the fixed effect logistic model results within value deciles. Data are obtained from the Compound Finance public API and standardized (within decile). This represents all borrowing accounts between 05-15-2019 and 10-26-2020. Accounts with health less than 0.9 or greater than 10 are omitted. The coefficients reported are marginal effects, along with their corresponding delta-method standard errors.

Update_Dummy	**Decile 2**	**Decile 3**	**Decile 4**	**Decile 5**	**Decile 6**	**Decile 7**	**Decile 8**	**Decile 9**	**Decile 10**
USD_GWEI	-0.0595734***	-0.0632715***	-0.1234466***	-0.1153121***	-0.1028746***	-0.0781944***	-0.0370698***	-0.0173035***	0.0024761
	(0.009322)	(0.0091003)	(0.0104266)	(0.0092332)	(0.0073516)	(0.0059497)	(0.0046291)	(0.0037488)	(0.0027653)
USD_value	0.0411652***	0.0427464***	0.0182074**	0.0481679***	0.0398032***	0.0142086**	-0.0132328**	-0.0296596***	-0.0144158**
	(0.0132979)	(0.0068362)	(0.007465)	(0.0071839)	(0.0069517)	(0.0063758)	(0.0062095)	(0.0052924)	(0.0071272)
health	-0.1210949***	-0.1042992***	-0.0693612***	-0.0376495***	-0.0112174*	-0.0052195	-0.0053323	-0.0086001*	0.0090363**
	(0.0118928)	(0.0115995)	(0.0100337)	(0.0075547)	(0.006635)	(0.0056413)	(0.0057998)	(0.0050252)	(0.0036865)
Volatility	0.0150826*	0.0442528***	0.0390146***	0.027455***	0.0207314***	0.0507092***	0.0445223***	0.0599733***	0.0612171***
	(0.0080897)	(0.0069906)	(0.0066802)	(0.0064154)	(0.0057056)	(0.0054684)	(0.0055694)	(0.0047206)	(0.0042105)
Net-Interest Rate	0.0740502***	0.0168154**	-0.0317248***	-0.030107***	-0.0476116***	-0.0364522***	-0.0126131***	0.0029913	0.0323661***
	(0.007549)	(0.0076317)	(0.0059248)	(0.0052284)	(0.0052498)	(0.0047163)	(0.0044982)	(0.0040013)	(0.0030723)
Pseudo R^2	0.02108534	0.017219	0.01628416	0.01298709	0.01155908	0.00908693	0.00380163	0.00515452	0.00535227
N-obs.	75,727	69,999	70,140	75,297	79,574	83,036	85,723	89,121	90,863
Unique-Addresses	817	1,055	1,663	2,114	2,365	2,416	2,436	2,280	1,862
Update Prob.	0.0259507	0.0352072	0.0477661	0.0642371	0.0791865	0.0981976	0.1119602	0.1309858	0.1997537
Mean Health	2.170499	2.278841	2.048316	2.007565	1.958898	1.93553	1.954195	1.930208	1.751461

Table 6: Discrete Exposure to Liquidation Penalty - Difference in Difference Analysis

This table presents results from the discrete difference in difference specification, measuring the effect of accounts' exposure to the increase in Compound's liquidation penalty on their account health (dependant variable) when account rebalancing occurs. Data are obtained from the Compound Finance public API and standardized. This represents all borrowing accounts between 08-19-2019 and 10-26-2020. Accounts with health less than 1 or greater than 10 are omitted, as are low-value accounts that are effectively precluded from transacting due to high fees. The variables of interest are Post, a dummy variable which equals 1 for all periods after the liquidation penalty increased, and Risky, a dummy variable which equals 1 for accounts borrowing greater than 50% of their borrowing capacity. Account-level controls include value, volatility, and net-interest rate of position. I also control for the average daily gas price in USD. Standard errors are clustered at the account level in all specifications.

Dep. Var: Account Health	**(1)**	**(2)**	**(3)**	**Value > $1,000**	**Value > $10,000**
Post Dummy	-0.1989087***	-0.0252013			
	(0.040962)	(0.0546283)			
Risky	-0.5017756***	-0.4087563***	-0.4166223***	-0.5063355***	-0.5178643***
	(0.0298478)	(0.0334502)	(0.0334871)	(0.0408931)	(0.054394)
Post Dummy x Risky	0.2268862***	0.1720284***	0.1728292***	0.2652254***	0.2708656***
	(0.032944)	(0.0369644)	(0.0372716)	(0.0441327)	(0.0593706)
Constant	2.045674***	1.992076***	2.091315***	2.010048***	1.925308***
	(0.0344037)	(0.0424044)	(0.0668231)	(0.0682067)	(0.0662139)
Controls	X	X	X	X	X
Account FE		X	X	X	X
Day FE			X	X	X
Adjusted R^2	0.11828314	0.12178501	0.13223812	0.13709783	0.13699862
N-obs.	64,039	64,039	64,039	51,902	36,209
Unique-Addresses	11,203	11,203	11,203	8,226	5,197

Table 7: Parallel Pre-Trends: Discrete Analysis

This table reports estimates the parallel trends assumption test for the discrete DiD specification outlined in Equation ??. Data are obtained from the Compound Finance public API and standardized. This represents all borrowing accounts between 08-19-2019 and 10-26-2020. Accounts with health less than 1 or greater than 10 are omitted, as are low-value accounts that are effectively precluded from transacting due to high fees. Account-level controls include value, volatility, and net-interest rate of position. I also control for the average daily gas price in USD. Standard errors are clustered at the account level in all specifications.

	Account Health
Risky	-0.3977793***
	(0.062805)
Post Dummy x Risky	0.1542798**
	(0.0654608)
Pre30 x Risky	0.0999255
	(0.0833986)
Pre60 x Risky	-0.0608928
	(0.092142)
Pre90 x Risky	-0.0993356
	(0.1153768)
Pre120 x Risky	-0.0115925
	(0.109846)
Pre150 x Risky	-0.1427329
	(0.1172677)
Pre180 x Risky	-0.0291094
	(0.0811959)
Controls	X
Account FE	X
Day FE	X
Adjusted R^2	0.13256092
N-obs.	64,039
Unique-Addresses	11,203

Table 8: Continuous Exposure to Liquidation Penalty - Difference in Difference Analysis

This table presents results from the continuous difference in difference specification outlined in Equation ??, measuring the effect of accounts' exposure to the increase in Compound's liquidation penalty on their account health (dependant variable) when account rebalancing occurs. Data are obtained from the Compound Finance public API and standardized. This represents all borrowing accounts between 08-19-2019 and 10-26-2020. Accounts with health less than 1 or greater than 10 are omitted, as are low-value accounts that are effectively precluded from transacting due to high fees. The variables of interest are Post, a dummy variable which equals 1 for all periods after the liquidation penalty increased, and liq_harm_frac, which is the fraction of the account's collateral value which would be lost via liquidation penalty. Liq_harm_frac is normalized for ease of interpretation. Account-level controls include value, volatility, and net-interest rate of position. I also control for the average daily gas price in USD. Standard errors are clustered at the account level in all specifications.

Dep. Var: Account Health	**(1)**	**(2)**	**(3)**	**Value > $1,000**	**Value > $10,000**
Post Dummy	-0.1554417***	-0.1075907***			
	(0.0229381)	(0.0297174)			
Liq_harm_frac	-1.582852***	-1.610255***	-1.618525***	1.624498***	-1.595565***
	(0.03371)	(0.0391117)	(0.0386965)	(0.0464867)	(0.0619971)
Post Dummy x Liq_harm_frac	0.3196***	0.2199374***	0.2042506***	0.2595776***	0.2893233***
	(0.037785)	(0.042413)	(0.0417558)	(0.0493557)	(0.0657522)
Constant	2.044748***	2.024134***	1.920932***	1.931897***	1.912928***
	(0.0190847)	(0.0220236)	(0.0483569)	(0.0500261)	(0.0462922)
Controls	X	X	X	X	X
Account FE		X	X	X	X
Day FE			X	X	X
Adjusted R^2	0.52440450	0.52725204	0.53701069	0.54134053	0.53755789
N-obs.	64,039	64,039	64,039	51,902	36,209
Unique-Addresses	11,203	11,203	11,203	8,226	5,197

Table 9: Parallel Pre-Trends: Continuous Analysis

This table reports estimates used to test the parallel trends assumption for the continuous DiD specification outlined in Equation ??. Data are obtained from the Compound Finance public API and standardized. This represents all borrowing accounts between 08-19-2019 and 10-26-2020. Accounts with health less than 1 or greater than 10 are omitted, as are low-value accounts that are effectively precluded from transacting due to high fees. Account-level controls include value, volatility, and net-interest rate of position. I also control for the average daily gas price in USD. Standard errors are clustered at the account level.

	Account Health
Liq_harm_frac	-1.69396***
	(0.0529026)
Post Dummy x Liq_harm_frac	0.2797718***
	(0.055776)
Pre30 x Liq_harm_frac	0.1506234**
	(0.0668681)
Pre60 x Liq_harm_frac	0.0903426
	(0.0747251)
Pre90 x Liq_harm_frac	0.1779247**
	(0.0795194)
Pre120 x Liq_harm_frac	0.077638
	(0.0785497)
Pre150 x Liq_harm_frac	0.0092754
	(0.0762545)
Pre180 x Liq_harm_frac	0.0228001
	(0.0611429)
Controls	X
Account FE	X
Day FE	X
Adjusted R^2	0.5375344
N-obs.	64,039
Unique-Addresses	11,203

Table 10: Number of blockchains by consensus protocol and launch year

This table presents the public blockchain ecosystem by the number of counts of a particular consensus protocol from 2009 to the end of 2020. Data are obtained from cryptoslate.com.

Protocol / Year	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020
PoW	1	0	3	1	17	39	28	26	63	72	26	10
PoS	0	0	0	0	0	17	13	14	26	63	30	11
Hybrid	0	0	0	1	4	16	8	23	32	29	10	8
DPoS	0	0	0	0	0	2	1	8	7	3	6	4
Nonstandard	0	0	0	0	2	7	5	14	35	50	30	16
Total (N=781)	**1**	**0**	**3**	**2**	**23**	**81**	**55**	**85**	**163**	**217**	**102**	**49**

Table 11: Blockchain sample for empirical analysis.

This table lists all blockchain names, their symbol, announcement year, native asset market capitalization in billions USD, and native asset total transaction count since inception. Network data are from coinmetrics.io. Market capitalization is based on quantities obtained as of February 21, 2021. Consensus protocols are divided into Proof-of-Work (PoW), Proof-of-Stake (PoS), Hybrid, Delegated Proof-of-Stake (DPoS), and Nonstandard.

Blockchain	**Symbol**	**Market cap** (in USD billions)	**Transaction count** (in millions)	(in %)	**Launch year**
PoW					
Bitcoin	BTC	1080.00	617.55	5.53	2009
Ethereum	ETH	225.78	1019.50	9.13	2014
Litecoin	LTC	15.45	59.13	0.53	2011
Bitcoin Cash	BCH	13.46	61.30	0.55	2017
Dogecoin	DOGE	7.04	67.68	0.61	2013
Bitcoin SV	BSV	4.54	320.55	2.87	2018
Ethereum Classic	ETC	1.86	58.18	0.52	2016
ZCash	ZEC	1.82	6.84	0.06	2016
Digibyte	DGB	1.12	20.73	0.19	2014
Bitcoin Gold	BTG	0.57	1.75	0.02	2017
Verge	XVG	0.42	4.66	0.04	2016
Vertcoin	VTC	0.05	2.66	0.02	2017
Hybrid					
Dash	DASH	3.01	96.52	0.86	2014
Decred	DCR	2.10	8.60	0.08	2015
PoS					
Cardano	ADA	34.50	4.24	0.04	2017
WAVES	WAVES	1.29	47.49	0.43	2016
PIVX	PIVX	0.10	1.20	0.01	2015
DPoS					
Binance Coin	BNB	46.00	225.07	2.02	2017
Polkadot	DOT	35.75	11.01	0.10	2020
EOS	EOS	4.93	4184.65	37.47	2017
TRON	TRX	4.34	1554.38	13.92	2017
Tezos	XTZ	3.59	33.38	0.30	2014
Lisk	LSK	0.49	3.53	0.03	2016
Nonstandard					
Ripple	XRP	24.86	1749.45	15.67	2013
Stellar	XLM	11.28	943.16	8.45	2016
NEM	XEM	5.08	9.40	0.08	2014
NEO	NEO	3.63	54.75	0.49	2014

Table 12: dApp statistics by blockchain

This table presents dApp statistics for ETH, EOS, TRON, NEO and WAVES blockchains. Data for ETH, EOS and TRON are collected from stateofthedapps.com, while data for NEO and WAVES are collected from dappradar.com. dApp count is the total number of dApps on each blockchain. dApp count (>100 MAU) is the total number of dApps with more than 100 unique monthly active users on each blockchain. MAU is the sum of the unique monthly active users on each dApp over a 30 day period for each blockchain. The 30 day period (considered in dApp transactions and MAU) is from 22 January, 2021 to 20 February, 2021.

Blockchain	ETH	EOS	TRON	NEO	WAVES
dApp count	1,965	233	66	16	26
dApp count (>100 MAU)	163	31	16	4	2
MAU	1,766,762	79,532	109,566	4,887	7,166

Table 13: Monthly Active Users (MAU) by dApp category and blockchain

This table shows the number (and percentage) of Monthly Active Users (MAU) within the dApp categories *Decentralized Finance* (*DeFi*), *Gambling*, *Games*, *Storage*, and *Other*, hosted across five blockchains: Ethereum (ETH), EOS (EOS), TRON (TRX), NEO (NEO) and WAVES (WAVES). Data for ETH, EOS and TRON are collected from stateofthedapps.com, while data for NEO and WAVES are collected from dappradar.com. Dappradar.com categorization is hand-matched to stateofthedapps.com categorization. MAU is the sum of the unique monthly active users on each dApp over a 30 day period (from 22 January, 2021 to 20 February, 2021) for each category.

dApp category	ETH	EOS	TRX	NEO	WAVES
DeFi	1,333,306	19,809	17,888	4,581	6,988
(MAU in %)	(96.44)	(1.43)	(1.29)	(0.33)	(0.51)
Gambling	8,012	8,658	18,546	0	35
(MAU in %)	(22.73)	(24.56)	(52.61)	(0.00)	(0.10)
Games	45,780	41,694	667	48	1
(MAU in %)	(51.91)	(47.28)	(0.76)	(0.05)	(0.00)
Storage	6,934	421	68,248	0	0
(MAU in %)	(9.17)	(0.56)	(90.27)	(0.00)	(0.00)
Other	372,730	8,950	4,217	258	142
(MAU in %)	(96.49)	(2.32)	(1.09)	(0.07)	(0.04)

Table 14: Block arrival rates and confirmation wait times.

This table shows block arrival times, confirmation numbers $k(s)$, and confirmation times κ for all blockchains in the coinmetrics.io sample. Block arrival times are denominated in seconds and are estimated as 14,400 seconds (one day) divided by the average daily block count. Confirmation numbers are given as median values of block confirmations required by 32 major cryptoasset exchanges for a given blockchain (cf. Appendix ?? for details and data sources). Confirmation time is the product of block arrival time and confirmation number. Blockchains are displayed according to their rank on confirmation time, where 1 corresponds to lowest (best) and 27 to highest (worst) confirmation times.

Blockchain	**Symbol**	**Consensus**	**Block arrival rate** (in seconds)	$k(s)$	**Confirmation time** κ (in seconds)	**Rank**
Binance Coin	BNB	DPoS	0.4	5.5	2.2	1
Stellar	XLM	Nonstandard	5.0	1	5.0	2
EOS	EOS	DPoS	0.5	13.5	6.8	3
Ripple	XRP	Nonstandard	4.2	3	12.5	4
TRON	TRX	DPoS	3.0	15	45.3	5
NEO	NEO	Nonstandard	20.9	5	104.6	6
Polkadot	DOT	DPoS	6.0	20	120.0	7
Digibyte	DGB	PoW	18.0	15	270.3	8
Ethereum	ETH	PoW	14.8	20	295.3	9
Cardano	ADA	PoS	20.1	15	301.0	10
Dogecoin	DOGE	PoW	62.9	6	377.3	11
NEM	XEM	Nonstandard	60.5	10	604.8	12
Litecoin	LTC	PoW	147.6	6	885.8	13
PIVX	PIVX	PoS	58.5	17.5	1023.1	14
Bitcoin	BTC	PoW	570.3	2	1140.5	15
WAVES	WAVES	PoS	60.2	20	1203.0	16
Dash	DASH	Hybrid	157.0	8	1255.8	17
Zcash	ZEC	PoW	118.0	12	1416.1	18
Decred	DCR	Hybrid	299.6	5	1498.1	19
Lisk	LSK	DPoS	10.2	176	1792.4	20
Tezos	XTZ	DPoS	61.7	30	1851.7	21
Verge	XVG	PoW	42.2	65	2745.5	22
Bitcoin Cash	BCH	PoW	569.6	6	3417.4	23
Bitcoin SV	BSV	PoW	603.6	12	7243.1	24
Ethereum Classic	ETC	PoW	14.0	1000	14049.7	25
Bitcoin Gold	BTG	PoW	574.8	45.5	26153.1	26
Vertcoin	VTC	PoW	148.4	900	133579.2	27

Table 15: Blockchain Ranking by Attributes

This table shows rankings of blockchains with respect to our attributes: adoption by different user types, scale and the wait time it takes to confirm transactions entered on a particular blockchain. The 27 blockchains in the coinmetrics.io sample are ranked from 1 (best) up to 27 (worst) to indicate their performance on each individual attribute. Blockchains highlighted in the final column are members of the blockchain frontier, i.e., there exists no other than the given blockchain that performs better in each of the categories.

		Adoption								
Blockchain	**Consensus**	**Payment**	**DeFi**	**Gambling**	**Games**	**Storage**	**Other**	**Scale**	**Confirmation**	**Frontier?**
BCH	PoW	8	-	-	-	-	-	7	23	
BSV	PoW	5	-	-	-	-	-	3	24	x
BTC	PoW	1	-	-	-	-	-	11	15	x
BTG	PoW	22	-	-	-	-	-	21	26	
DGB	PoW	18	-	-	-	-	-	6	8	
DOGE	PoW	6	-	-	-	-	-	15	11	
ETC	PoW	16	-	-	-	-	-	18	25	
ETH	PoW	3	1	3	1	2	1	8	9	x
LTC	PoW	4	-	-	-	-	-	14	13	
VTC	PoW	23	-	-	-	-	-	26	27	
XVG	PoW	24	-	-	-	-	-	22	22	
ZEC	PoW	11	-	-	-	-	-	23	18	
ADA	PoS	10	-	-	-	-	-	20	10	
PIVX	PoS	27	-	-	-	-	-	27	14	
WAVES	PoS	21	4	4	5	-	4	12	16	
NEO	Nonstandard	20	5	-	4	-	-	13	6	
XEM	Nonstandard	25	-	-	-	-	-	19	12	
XLM	Nonstandard	9	-	-	-	-	-	2	2	x
XRP	Nonstandard	14	-	-	-	-	-	5	4	
DASH	Hybrid	7	-	-	-	-	-	10	17	
DCR	Hybrid	12	-	-	-	-	-	25	19	
BNB	DPoS	15	-	-	-	-	-	9	1	x
DOT	DPoS	17	-	-	-	-	-	16	7	
EOS	DPoS	19	3	2	2	3	2	1	3	x
LSK	DPoS	26	-	-	-	-	-	24	20	
TRX	DPoS	2	2	1	3	1	3	4	5	x
XTZ	DPoS	13	-	-	-	-	-	17	21	

Table 16: Required block confirmations by cryptoasset exchanges

This table reports the number of block confirmations required by 32 cryptoasset exchanges for a deposit/transaction to be considered as settled. The list of exchanges is obtained from cryptocompare.com on March 25, 2021 and consists of all exchanges with a daily trading volume above $1 million (n=122). We search the FAQ and support web pages of each cryptoasset exchange for disclosed number of block confirmations for each blockchain from the coinmetrics.io sample.

Exchange/Blockchain	ADA	BCH	BNB	BSV	BTC	BTG	DASH	DCR	DGB	DOGE	DOT	EOS	ETC
Bilaxy					3								
Binance					1								
Bitbank		20			2								
Bitbay		10		50	3	120	50						
Bitbuy		6			3							1	
Bitcoin.com	20	2			1		10					25	60000
Bitmax		3			3							15	3
Bitso		6			4								
Bittrex	20	12		12	2		6	6	20	6	20	1	400
BTCBOX		18			2								
CEX.IO	21	10			3	41	6						
Coinbase	10	12			3		2					1	10000
Coineal					1								
CoinEx	12	1	10	12	1		12	1	10	3	5	12	1000
CoinField		2			2		10						
Coinfloor					3								
Coins Pro		12			3								
Currency.com		6			2								
EXMO	15	6			6	50	6	12		4		30	500
FTX		2	1		2							1	
Gemini		15			3								
GOPAX		2			2							45	
Independent Reserve		5		5	2						11	300	300
itBit		15			3								
Korbit	20	3	12	12	2						30	360	1000
Kraken	15	15			4		6			40	25	1	80640
Liquid		4			2		1				60		
Nominex					3								
OKEX	12	2			1	2	10		10			1	100
Polionex		11	1		1		50	4	40	6	1	350	200000
PrimeXBT					3								
VALR					2								
Median	15	6	5.5	12	2	45.5	8	5	15	6	20	13.5	1000

Exchange/Blockchain	ETH	LSK	LTC	NEO	PIVX	TRX	VTC	WAVES	XEM	XLM	XRP	XTZ	XVG	ZEC
Bilaxy	32			3							3			
Binance	12													
Bitbank	24		6							1	6			
Bitbay	12	150	6			10				10	6			20
Bitbuy	12		6							1	1			
Bitcoin.com	20		5			5		50		10	20			6
Bitmax	12		3								1			
Bitso	30		6								1			
Bittrex	36	202	6	10	30	20	600	20	10	2	10	20	100	30
BTCBOX	12		3											
CEX.IO	25		3	5		21				1	1	30		
Coinbase	35		12							1		30		24
Coineal	30													
CoinEx	12	12	3	1	5	3		10	1	1	3	1	30	1
CoinField	30		4								5			12
Coinfloor														
Coins Pro	20										6			
Currency.com	12		4											
EXMO	6	15	4	5		20		20	15	1	1			
FTX	10		2								1			
Gemini	12		12											24
GOPAX	45		6							1				12
Independent Reserve	20		10							6	6			
itBit	12		12											
Korbit	24		6			20				3	1			
Kraken	20	303	12			20		10		1	1	20		24
Liquid	21		6							10	20	30		
Nominex	12													
OKEX	12		1	1		1			1	1	1			6
Polionex	12	303	4	18		1	1200		18	2	2	30		8
PrimeXBT	10													
VALR	30										10			
Median	20	176	6	5	18	15	900	20	10	1	3	30	65	12

Table 17: Sources of block confirmation rules.

Name	Country	URL
Bilaxy	Republic of Seychelles	https://bilaxy.zendesk.com/hc/en-us/articles/360020196932-Deposit-Instructions
Binance	Malta	https://www.binance.com/en/support/articles/360030775291
Bitbank	Japan	https://bitbank.cc/docs/deposit
BitBay	Estonia	https://bitbay.net/en/helpdesk/payments-and-withdrawals/when-funds-will-be-added-to-my-account
Bitbuy	Canada	https://support.bitbuy.ca/hc/en-us/articles/360032483832-Cryptocurrency-deposit-processing-times
Bitcoin.com	Saint Kitts and Nevis	https://support.bitcoin.com/en/articles/3851763-cryptocurrency-deposit-processing-times
BitMax	Singapore	https://bitmaxhelp.zendesk.com/hc/en-us/articles/360012169694-Why-Haven-t-I-Received-My-Deposits
Bitso	Mexico	https://bitso.com/fees/transactions
Bittrex	U.S.A	https://bittrex.com/api/v1.1/public/getcurrencies
BTCBOX	Japan	https://blog.btcbox.jp/en/fees-introduction-2
Cex.io	United Kingdom	https://support.cex.io/en/articles/4383405-processing-your-crypto-deposits-and-withdrawals
Coinbase	U.S.A	https://help.coinbase.com/en/coinbase/trading-and-funding/sending-or-receiving-cryptocurrency/why-is-my-transaction-pending
Coineal	Republic of Seychelles	https://support.coineal.com/hc/en-001/articles/360022989414-Deposit-Missing
CoinEx	Unknown	https://support.coinex.com/hc/en-us/articles/900004316823-FAQ-about-Deposit-
CoinField	Estonia	https://coinfield.freshdesk.com/support/solutions/articles/36000104043-how-many-blockchain-confirmations-are-needed-in-order-to-see-my-cryptocurrency-deposits-
Coinfloor	United Kingdom	https://coinfloor.zendesk.com/hc/en-us?mobile_site=false
Coins Pro	Philippines	https://support.coins.ph/hc/en-us/articles/203165130-When-will-my-balance-update-after-making-a-blockchain-transfer-to-from-my-wallet-
Currency.com	Belarus	https://currency.com/cryptocurrency-faq
Exmo	United Kingdom	https://info.exmo.com/en/wallet/how-quickly-will-my-cryptocurrency-deposit-be-added-to-my-balance/
FTX	Antigua and Barbuda	https://help.ftx.com/hc/en-us/articles/360034865571-Blockchain-Deposits-and-Withdrawals
Gemini	U.S.A	https://support.gemini.com/hc/en-us/articles/205424836-How-long-until-my-digital-asset-deposit-reaches-my-account-
GOPAX	South Korea	https://help.gopax.com/support/solutions/articles/42000020261-how-to-check-the-transaction-status-of-funds-deposited-to-gopax-com
Independent Reserve	Australia	https://www.independentreserve.com/help/faq
itBit	U.S.A	https://help.paxos.com/hc/en-us/articles/360042320931-Daily-Deposits-Withdrawal-Schedule
Korbit	South Korea	https://www.korbit.co.kr/faq/articles/?id=5fFcnrIQ6ILjsSelbxYqjr
Kraken	U.S.A	https://support.kraken.com/hc/en-us/articles/203325283-Cryptocurrency-deposit-processing-times
Liquid	Japan	https://help.liquid.com/en/articles/4713100-crypto-deposit-processing-time
Nominex	Seychelles	https://support.nominex.io/article/7-deposit-withdrawal-does-not-arrive
OKEX	Malta	https://www.okex.com/support/hc/en-us/articles/360000205532-Token-Confirmation-Requirements-Blockchain-Explorers
Poloniex	U.S.A	https://poloniex.com/public?command=returnCurrencies
PrimeXBT	Republic of Seychelles	https://help.primexbt.com/deposits-withdrawals/how-to-deposit#direct-crypto-deposits
VALR	South Africa	https://support.valr.com/hc/en-us/articles/360029561452-How-do-I-deposit-cryptocurrencies-on-VALR-

Table 18: Sample Description

This table reports summary statistics for price growth ($\Delta Price$) and hashrate growth ($\Delta Hashrate$) for Bitcoin (BTC) and Ether (ETH), measured as the first difference of $log(Price)$ and $log(Hashrate)$ respectively. The sample period for BTC is from April 25, 2014 to Oct. 19, 2019 while the sample period for ETH is from Oct. 25, 2016 to Oct. 19, 2019.

Variable	Mean	Median	St. Dev	Obs.
$\Delta Price$(BTC)	0.001464	0.001684	0.03853339	2015
$\Delta Price$(ETH)	0.003937	0.000106	0.06348898	1456
$\Delta Hashrate$(BTC)	0.003666	0.002364	0.1189652	2015
$\Delta Hashrate$(ETH)	0.004243	0.004176	0.02781021	1456

Table 19: Sample Augmented Dickey Fuller Test

This table reports the Augmented Dickey Fuller (ADF) test results (Dickey and Fuller, 1979). Number of lags is chosen according to the Bayesian Information Criterion and Akaike Information Criterion minimizing Vector Autoregression (VAR) specification. I allow for drift and trend terms. The sample period for BTC is from April 25, 2014 to Oct. 19, 2019 while the sample period for ETH is from Oct. 25, 2016 to Oct. 19, 2019.

Variable	ADF	p-value	Lags
$log(Price)$(BTC)	-1.86	0.635	12
$log(Price)$(ETH)	-1.12243	0.91898	2
$log(Hashrate)$(BTC)	-1.86	0.6367	12
$log(Hashrate)$(ETH)	-0.05282	0.99	2

Table 20: Johansen Test

This table reports the Johansen (1991, 1995) trace test results for cointegration between price and hashrate for both BTC and ETH, allowing for a linear trend in cointegration. Number of lags is chosen according to the Bayesian Information Criterion and Akaike Information Criterion minimizing Vector Autoregression (VAR) specification. The sample period for BTC is from April 25, 2014 to Oct. 19, 2019 while the sample period for ETH is from Oct. 25, 2016 to Oct. 19, 2019.

Test	Test Statistic	1% Critical Value	5% Critical Value	10% Critical Value
BTC	46.27	30.45	25.32	22.76
ETH	78.52	30.45	25.32	22.76

Table 21: Full Sample TVECM Specification

This table reports the results from the full sample Threshold Vector Error Correction Model (TVECM) specification for both Bitcoin (BTC) and Ether (ETH). Lagged first-differences are not reported to conserve space. Number of lags is chosen according to the Bayesian Information Criterion and Akaike Information Criterion minimizing Vector Autoregression (VAR) specification. Threshold value is set at 0. Percent of observations below the threshold value for BTC and ETH is 43.5% and 56.7% respectively. A constant and trend are included in both specifications. The sample period for BTC is from April 25, 2014 to Oct. 19, 2019 while the sample period for ETH is from Oct. 25, 2016 to Oct. 19, 2019.

	BTC		ETH	
equation	**Δ lhashrate**	**Δ lprice**	**Δ lhashrate**	**Δ lprice**
ECT-	-0.0111***	0.0002	-0.0123***	-0.0005
	(0.0021)	(0.8949)	(4.4e-08)	(0.9322)
ECT+	-0.0063	0.0028	-0.0114***	-0.0023
	(0.1803)	(0.1427)	(4.9e-05)	(0.7349)
Const	0.0294***	-0.0035	-0.0024	0.0100
	(6.2e-05)	(0.2360)	(0.3371)	(0.0879)
Lags	12	12	2	2
N	2015	2015	1456	1456

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 22: Pre- vs. Post- ETH ASIC Results

This table reports the Threshold Vector Error Correction Model (TVECM) results using two Ethereum (ETH) subsamples. ETH Pre-ASIC represents the period from Oct. 25, 2016 to June 30, 2018 and ETH Post-ASIC represents the period from Aug. 1, 2018 to Oct 19, 2019. Lagged first-differences are not reported to conserve space. Threshold value is set at 0. Percent of observations below the threshold value for the Pre- and Post- periods is 53.8% and 49.4% respectively. A constant and trend are included in both specifications.

	ETH Pre-ASIC		ETH Post-ASIC	
equation	**Δ lhashrate**	**Δ lprice**	**Δ lhashrate**	**Δ lprice**
ECT-	-0.0134***	-6.5e-05	-0.0012	0.0142
	(2.2e-05)	(0.9930)	(0.8377)	(0.3301)
ECT+	-0.0188***	-0.0025	-0.0307***	0.0119
	(0.0004)	(0.8407)	(9.3e-06)	(0.4867)
Const	-0.0039	0.0114	0.0026	-0.0092
	(0.2310)	(0.1471)	(0.3009)	(0.1349)
Lags	2	2	2	2
N	980	980	446	446

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 23: Pre- vs. Post- BTC Halving Results

This table reports the Threshold Vector Error Correction Model (TVECM) results using two Bitcoin (BTC) subsamples. BTC Pre-Halving represents the period from April 14, 2014 to July 8, 2016 and BTC Post-Halving represents the period from July 10, 2016 to Oct. 19, 2019. Lagged first-differences are not reported to conserve space. Threshold value is set at 0. Percent of observations below the threshold value for the Pre- and Post- periods is 32.8% and 48.3% respectively. A constant and trend are included in the Pre-Halving subsample, while only a constant is included in the Post-Halving subsample (see footnote 29).

	BTC Pre-Halving		BTC Post-Halving	
equation	**Δ lhashrate**	**Δ lprice**	**Δ lhashrate**	**Δ lprice**
ECT-	-0.0181***	-0.0037	-0.0135**	0.0016
	(0.0021)	(0.0952)	(0.0057)	(0.4318)
ECT+	-0.0145	0.0107	-0.0021	0.0009
	(0.1619)	(0.0071)	(0.7204)	(0.7098)
Const	0.0002	-0.0097	0.0155***	0.0017
	(0.9830)	(0.0181)	(0.0028)	(0.4213)
Lags	9	9	14	14
N	980	980	446	446

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 24: Gregory and Hansen (1996) Test

This table reports the Gregory and Hansen (1996) test results for cointegration between BTC price and hashrate, allowing for one unknown regime shift in the cointegration relationship. For my purposes, this test is used solely to estimate the structural break date in the cointegration relationship corresponding to the full adoption of ASIC mining; the cointegrating relationship is later re-estimated. The sample period is from July 18, 2010 to Oct. 19, 2019.

Test	**Test Statistic**	**Break Date**	**5% Critical Value**	**10% Critical Value**
ADF	-4.38	4/13/14	-4.61	-4.34
Zt	-3.57	4/13/14	-4.61	-4.34
Za	-39.37	4/13/14	-40.48	-36.19