From mining to markets: The evolution of bitcoin transaction fees*

David Easley[a,b,c], Maureen O'Hara[c,d,e], and Soumya Basu[a,f]

a. **Cornell University, Department of Economics, United States**
b. **Cornell University, Department of Information Sciences, United States**
c. **University of Technology Sydney, Australia**
d. **Cornell University, Johnson College of Business, United States**
e. **Cornell University, Initiative for Crypto Currencies and Contracts, United States**
f. **Cornell University, Department of Computer Science, United States**

July 2017

Revised May 2018

ABSTRACT

We investigate the role that transaction fees play in the bitcoin blockchain's evolution from a mining-based structure to a market-based ecology. We develop a game-theoretic model to explain the factors leading to the emergence of transactions fees, as well as to explain the strategic behavior of miners and users. Our model highlights the role played by mining rewards, transaction fees, price, and waiting time, discusses welfare issues, and examines how microstructure features such as exogenous structural constraints influence the dynamics of user participation on the blockchain. We provide empirical evidence on the model's predictions and discuss implications for bitcoin's evolution.

JEL classification: G10, G18, D47, D6

Keywords: Bitcoin, Blockchain, Market microstructure, Transaction fees

* We are grateful to the editor (Bill Schwert), two anonymous referees, Levon Barseghyan, Campbell Harvey, Andrew Karolyi, Francesca Molinari, Talis Putnins, Gun Sirer, Clara Vega, David Yermack, Yao Zeng, and seminar participants at the Board of Governors of the Federal Reserve System, Cornell University, HEC Montreal, Johns Hopkins University, Queen's University, the RFS FinTech Conference, and the University of Technology Sydney/ University of Sydney Market Microstructure Conference for helpful comments and discussion. Soumya Basu thanks the Initiative for Crypto-currencies and Contracts (IC3) and National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1650441for support.

Corresponding author telephone: 607 255 3645, fax: 607

E-mail address: mo19@cornell.edu (M. O'Hara)

"In a few decades when the reward gets too small, the transaction fee will become the main compensation for nodes. I'm sure that in 20 years there will either be very large transaction volume or no volume."

Satoshi Nakamoto[1]

1. Introduction

By a variety of metrics, bitcoin is no longer a financial curiosity. From its genesis transaction of 50 bitcoins in January 2009, bitcoins in circulation now number more than 17 million. An estimated 35 million Bitcoin wallets are held worldwide with 100,000 companies accepting payments in bitcoins, some via the newly issued bitcoin debit card. Daily trading volume on November 7, 2017 at major bitcoin exchanges first exceeded a record $5.0 billion, with volume across all cryptocurrencies on Dec. 20, 2017 exceeding $50 billion.[2] While volume has fallen from that record level, it is still substantial, with Nov. 28, 2018 volume of $20.4 billion. A recent study (see Hileman and Rauchs (2017)) estimates that 10 million people now hold a material amount of bitcoins as a financial asset. Approximately 10,467 active global bitcoin nodes exist, each containing a complete copy of the Bitcoin blockchain.[3]

Envisioned in 2008 as a decentralized, trustless digital currency and payment system, the bitcoin blockchain operates on a worldwide basis via a complex set of rules originally proposed in Nakamoto (2008). Fundamental to the bitcoin ecology are miners, who play a crucial role both in creating new bitcoins and in verifying transactions on the blockchain. Mining involves using specialized computer hardware to find a particular mathematical hash function, with the reward for success being payment in new bitcoins. The amount of such payments, as well as a

[1] Cited in https://www.bitcoinmining.com/what-is-the-bitcoin-block-reward/
[2] See https://news.bitcoin.com/markets-update-bitcoins-daily-trade-volume-surpasses-5b/ and http://www.businessinsider.com/daily-cryptocurrency-volumes-vs-stock-market-volumes-2017-12
[3] This is the number of reachable, or active global nodes see https://bitnodes.earn.com

variety of parameters such as the difficulty of the underlying computational problem and even the total amount of bitcoins that can ever be mined are specified exogenously. The bitcoin protocol also exogenously specifies a dynamic adjustment process for these payment and difficulty parameters.

As Fig. 1 illustrates, an additional form of compensation for miners has emerged in the form of transactions fees. These transaction fees are voluntarily appended to bitcoin transactions by buyers and sellers wanting to ensure that their transactions are included in the block of transactions that the miner attaches to the blockchain. While such fees are still not the primary component of miners' total revenues, the endogenous development of transactions fees reflects an important step in the evolution of the bitcoin blockchain from being a mining-based set of rules toward being a market-based system capable of adapting to changing economic conditions. What is less clear is how successful this transition will be.

Insert Fig.1 about here

In this paper, we investigate the evolution of transactions fees in bitcoin. We build a framework for understanding why such fees developed and how they influence the dynamics of the bitcoin blockchain. Examining the bitcoin ecology is complex as it involves strategic behavior on the part of both miners and users, all packaged within a set of exogenous rules. We develop a game-theoretic model to explain the factors leading to the emergence of transaction fees, as well as the interactions between miners and users. Our model also highlights the roles played by the bitcoin price, waiting times, and mining-based revenues known as block rewards. We then provide empirical evidence supporting our theoretical conclusions about the impact of waiting time on fees and we offer estimates of the magnitudes of these effects.

Our research has a number of results on the role and behavior of transaction fees in the bitcoin blockchain. Our model confirms that eventually without transaction fees the blockchain would not be viable as miners' revenues from block rewards are deterministically programmed by the bitcoin protocol to reach zero. Given that this zero block reward level is not projected to occur until the year 2140, our model shows the important role played by the bitcoin price level in sustaining the viability of the bitcoin blockchain. Our model shows why, with increasing bitcoin price levels, transaction fees can long play only a secondary role in explaining the willingness of miners to participate.

Where transaction fees play a bigger role is in affecting the participation of users. We show that, even with transaction fees, there are limits on the size of the blockchain imposed by waiting times confronting users. We show how these waiting times arise in equilibrium and how they are influenced both by endogenous transactions fees and by exogenous dynamic constraints imposed by the bitcoin protocol. More intriguing, we show that waiting times (or for that matter, transaction fees) are not influenced by the block reward; provided mining is viable, our model predicts that the reward level is irrelevant for determining transaction fees in equilibrium.

Our model also demonstrates how equilibrium transaction fees evolve in the bitcoin ecosystem. If the arrival rate of potential transactions is low, transactions without fees attached are written to the blockchain but, as the arrival rate of potential transactions increases, the equilibrium shifts and only transactions with fees attached are posted to the blockchain. These equilibria exhibit strategic complementarity, meaning that while some equilibria are stable, others are not. Moreover, these equilibria can be Pareto-ranked, and we demonstrate that transactions fees are not welfare-improving. We also show that, in these equilibria, transaction costs can induce user nonparticipation. The fees directly induce some users to drop out, while increasing wait times cause other fee-paying users to depart as well. Our model suggests that these welfare and user participation effects are reasons why transaction fees alone are not a panacea for the dynamic challenges facing the evolving bitcoin blockchain.

Our empirical work supports these findings. We show that higher transactions fees are being driven by queuing problems facing users, not by reductions in bitcoin-denominated block rewards. As predicted, in at least most empirical specifications, we find no statistically significant effects for the block reward level. As users battle to get transactions posted on the blockchain, transaction fees rise to levels that discourage bitcoin usage, highlighting an important structural issue confronting the blockchain. While bitcoin can continue to develop as a financial asset, longer waiting times and higher transactions fees could impede its development as a transactional medium. Overall, our results delineate the complex role that transaction fees play in bitcoin's evolution from a mining-based structure to a market-based ecology.

Finally, our general results apply more broadly to the many other systems (for example, Ethereum, Litecoin, and Dogecoin) that use the bitcoin blockchain management protocols. In these and other similar systems, transaction fees also rose dramatically over our sample period, reflecting the interaction of increasing demand and fixed blockchain protocols. For such cryptocurrencies, a fundamental challenge is whether a static rules-based protocol can remain a single entity or whether the disparate needs of users result in fragmentation into multiple coin-based currencies. We argue that such fragmentation can be a natural part of the evolution to a market-based system, underscoring Satoshi's concern about whether volumes will be large or none at all.

This paper is organized as follows. Section 2 provides a literature review. Section 3 gives a brief overview of the bitcoin blockchain, explaining the various exogenous components of the system. Section 4 develops game-theoretic models of the games played by the miners and by the users. We find the Nash equilibria in these games and demonstrate their stability properties and welfare implications. We then characterize the factors influencing these equilibria, with a particular focus on the important role played by waiting times. We examine the equilibrium transaction fee structure and how it influences user waiting times. Section 5 then provides empirical analyses based on predictions from our model. Section 6 concludes by discussing the current structural challenges confronting bitcoin and the role played by transaction fees.

2. Literature Review

Our research joins a growing literature examining bitcoin, digital currencies, and the broader applications of blockchains. A variety of authors (see, for example, Eyal and Sirer (2014); Gans and Halaburda (2015); and Gandel and Halaburda (2016)) analyze design issues of the bitcoin protocol, as well as the dynamic interactions between cryptocurrencies. Other recent research analyzes aspects of the bitcoin ecosystem specifically as they relate to finance and the financial markets (see Boehm. Christin, and Moore (2015); Harvey (2016); Malinova and Park (2016); Raskin and Yermack (2016; 2017); Aune, Krellenstein, O'Hara, and Slama (2017)). Still other research (see Cong and He (2017); FINRA (2017)) examines smart securities and more general uses of the blockchain. A literature also looks at the game played by miners on the bitcoin blockchain. Huoy (2014), for example, analyzes the mining game but his analysis includes neither users nor transactions fees. Kroll, Davey, and Felton (2013) do include transaction fees but argue that they have little importance.

Huberman, Leshno, and Moallemi (2017), written contemporaneously with our paper, analyze a congestion queuing game that includes miners and fees. The basic ingredients of our model and their model are similar: Free entry into mining yields zero expected profit, and waiting time for users lead to fees. But, the focus of our analyses is different. Huberman. Leshno, and Moallemi (2017) are concerned about the viability of mining when mining rewards are zero and fees are necessary to incentivize miners. Their concern is that equilibrium fees could be too low for the blockchain to be viable, and they propose changes to the protocol to address this potential issue. We have the opposite concern, i.e., waiting times and equilibrium fees could be high enough to discourage user participation. Also, our model examines the evolution of these mining rewards and transaction fees in equilibrium, while their analysis looks only at the long-run steady state where mining rewards have disappeared and price effects of bitcoin are assumed irrelevant. This difference is particularly important as whether bitcoin can reach such a steady state depends critically on this evolution from mining to markets. As we demonstrate, the complex interactions in the bitcoin ecology make this evolution far from a sure thing.

3. A brief overview of the bitcoin blockchain protocol

The bitcoin blockchain set out in Nakamoto (2008) involves a decentralized trustless network composed of nodes, with each node containing a complete copy of the blockchain. Miners run nodes, validate transactions, and provide the requisite security for the blockchain. Mining uses computers with dedicated hardware and software to find a specific hash function or string of numbers.[4] The first miner to do so is compensated with a fixed number of newly issued bitcoins known as a block reward. The winning miner is then allowed to post a block (or collection) of pending bitcoin transactions to the blockchain.[5] By doing so, the miner verifies

[4] That is, passing the transactions augmented with a random number through a cryptographic hashing function and cycling through different random numbers to produce a rare result that has many leading zeros.

[5] The block size limit was reduced from 36 megabyte (MB) to 1MB in 2010 to counter risks from spam and denial of service attacks. An improvement in the bitcoin protocol, Segwit2x, activated on August 22, 2017, essentially

that the transactions are valid, and the block is added to the blockchain. If a block is proposed with an invalid transaction, then the block proposal is ignored. The blockchain grows as each validated block is added.

This validation process of posting transactions to the blockchain can also provide another source of revenue to the winning miner via transaction fees embedded in pending bitcoin transactions. The validation process involves miners first gathering up pending transactions into a block and then racing to be first to find the computational solution and, if successful, finally appending the transactions in the block to the blockchain. Transactions specify the number of bitcoins to be taken from one address and the number to be transferred to another address, and any difference in the two numbers can be kept by the miner as a transaction fee. The decision to append a fee is not specified by the bitcoin protocol but is up to the participant in the underlying transaction. We analyze this fee decision in detail in Section 4.

The bitcoin blockchain mechanism is highly structured, with the total number of bitcoins available for issuance limited to 21 million. Because issuance is tied to mining, the block reward is set to decline as the number of bitcoins in existence grows. The block reward is halved after every 210,000 blocks are mined. In November 2012, the block reward was reduced from its initial level of 50 bitcoins to 25 bitcoins, and then again in July 2016 to its current level of 12.5. The reward will be reduced 32 more times before eventually reaching zero sometime around 2140. While the block reward is declining in numbers of bitcoins, its value to the miner depends upon the bitcoin price. Because the issuance structure is designed to be deflationary, a rising

expanded the total block size to a theoretical upper bound of 4 MB, although actual block sizes are on the order of 1.5 MB -2 MB.

bitcoin price could more than offset any quantity declines, yielding sufficient revenue to sustain the block chain.[6] We discuss these dynamics in Section 4.

The rate at which new bitcoins are issued is also affected by the difficulty of solving the computational problem. In general, the algorithm sets the difficulty level such that on average a block is added to the blockchain every ten minutes, or approximately 144 blocks a day. If new blocks are added faster than this desired level (perhaps because of an increase in the number of miners or an advance in mining technology), the difficulty is increased. It is decreased if the rate is too low. This adjustment is made every 2,016 blocks, or approximately every 14 days.

Transactions to be posted to the blockchain are originally broadcast across the various nodes (designated as either full or partial nodes) in the decentralized bitcoin network. Full nodes contain an exact copy of the bitcoin blockchain, and each full node generally has a holding tank, called a mempool, in which these transactions are held pending their inclusion in a block. Transactions flow into the mempool from these broadcasts, and they leave the mempool when they either are posted to the blockchain or are dropped from the pool (typically after three days or so) if their fee is too low to attract a miner.[7] As transaction volume in the bitcoin network increases, the flows into these mempools also rise. The flows out of the mempool are

[6] The issue of how large to make block rewards is a fundamental issue for many digital currencies. Ethereum, the second largest digital currency, recently reduced its block reward from 5 to 3 ETH. This change was intended, in part, to stem inflation in the number of coins which had reached 14.8% in 2016.

[7] Some nodes have rules over which transactions they will include in their mempool. These restrictions can reflect limitations on the mempool size or minimums on the level of transaction costs that transactions must offer. For details, see https://support.21.co/bitcoin/mempools/what-is-the-mempool-size

circumscribed by the maximum block size and the current difficulty level. The bitcoin mechanism provides no natural means to equilibrate these potentially disparate flows.

Given this protocol, two problems are apparent. First, the deterministic decline in the block reward results in miners' revenues (at least as defined in terms of bitcoins) also falling deterministically over time, raising the potential that miners are unwilling to perform the costly calculations needed to validate the blockchain. Second, the rise in bitcoin transaction volume, coupled with limits on the number of blocks that can be posted to the blockchain, raise the possibility that some, potentially many, transactions are never posted to the blockchain, undermining the willingness of users to transact in bitcoin. With this as a backdrop, the question of interest is whether fees appended to bitcoin transactions represent a potential market solution to these shortcomings.

4. Model

The main participants in the bitcoin blockchain are miners and users. Miners are individuals, often operating in pools, who solve computational problems allowing them to put transactions securely on the blockchain and reap the block reward as well as any fees attached to those transactions. There is free entry into the mining game, so miners play a standard entry-exit game in which any Nash equilibrium will entail zero expected profit.[8] Users submit transactions to the nodes that they want verified and posted on the blockchain. They play a game in which they can chose to pay a fee to move up in the queue and thus reduce their waiting time, or they can not pay a fee and experience a longer waiting time. The Nash equilibria of this game can involve

[8] The zero expected profit claim depends on an assumption that miners are risk-neutral. If they are risk-averse, then equilibrium expected profit will be positive and the entry condition is determined by the expected utility of profit.

none, some, or all of the users paying a fee. Which equilibria occur (there can be multiple equilibria) depends on the parameters of the problem and, most important, on the waiting times for fee-paying and non-fee-paying users.

4.1 Miners

Miners are indexed by m =1,…, M where M is endogenous. We determine the Nash equilibrium value of M after solving the mining problem for a fixed M. We assume that all miners are identical in that they have a common fixed cost of mining $F > 0$, incur a depreciation cost of a common fraction δ of F for each problem they work on, and have a common variable cost per unit of time spent working on a problem of $e > 0$.[9] These costs are denominated in a common currency that, for simplicity, we denote as dollars.[10]

At any moment, the miners are all working on finding a particular hash and having the right to put the next block on the chain. We assume that, independently, for each miner working on the problem, success in solving the problem arrives according to a Poisson distribution with arrival rate λ. The miner who first solves the problem gets to put the next block on the chain. Because

[9] As our focus is on transaction fees, our characterization is simplified to capture the basic features of mining. Mining involves cost differences arising from different computing technology, cost of electricity, and scale, which could all influence specific miner fixed and variable costs. Mining also often feature mining consortia, further complicating entry and exit into the field. See Eyal and Sirer (2014).

[10] Variable costs are primarily electricity costs, and they are substantial. Sirer (2017) notes that the Bitcoin network consumes roughly as much electricity as one-third the annual electricity consumption of Ireland. The resource intensive nature of the blockchain has led to research on alternative green structures for blockchains. See http://hackingdistributed.com/2017/02/23/green-blockchains.

miners are symmetric and assumed to have identical hashing power, the probability that any individual miner is the first to solve the problem is M^{-1}.

The winning miner receives a block reward of S > 0 bitcoins. In addition, the winning miner places the block of up to B transactions from the mempool onto the blockchain, thereby collecting any fees embedded in those transactions.[11] Let f_i be the fee embedded in transaction i, where all fees are also denominated in bitcoins. We assume that in forming blocks miners choose among transactions according to their fees, taking the ones with the greatest fees first and then proceeding down the list of fees until the block is full.[12] If the last transactions added to the block come from a collection of transactions with a common fee (which can be zero) the miner can choose at random among those transactions. Abusing notation slightly, we let B be the block of transactions written to the chain. So, the revenue earned in bitcoins by the successful miner is $S + \sum_{i \in B} f_i$. Denoting the exchange rate between the bitcoin value and the dollar as p, the revenue earned by the successful miner in dollars is thus $p(S + \sum_{i \in B} f_i)$.

Miners are assumed to be risk-neutral. They play a simple game in which they independently chose to participate or not. We assume that there is free entry into mining, thus, in a Nash equilibrium, every miner must be making zero expected profit given the choices of all other potential miners. We ignore integer issues, but with integer constraints, the equilibrium condition is that each miner has a non-negative profit and, given the choices of all other miners, would

[11] We treat all miners as having access to identical mempools and refer to this common pool as the mempool.

[12] We assume that transactions are the same size, although some transactions require more block capacity than others. If there are not enough transactions in the mempool, the winning miner takes all available transactions.

have a non-positive profit from any alternative choice. The number of miners we find is within one of the equilibrium number expressed as an integer.

Any miner's expected revenue from attempting to solve the problem is $p[S+\sum_{i\in B} f_i]/M$. M miners are working on the problem with individual success rates λ, so the expected time until the first success is $1/\lambda M$. Each active miner's expected profit is

$$\frac{p[S+\sum_{i\in B} f_i]}{M} - \frac{e}{\lambda M} - \delta F . \qquad (1)$$

In a Nash equilibrium, the expected profit to mining must be zero as otherwise there would be entry or exit of miners. In an equilibrium with any active miners (M > 0) the number of miners must be such that

$$\frac{p[S+\sum_{i\in B} f_i]}{M} - \frac{e}{\lambda M} - \delta F = 0 \qquad (2)$$

However, if $p[S+\sum_{i\in B} f_i] - e/\lambda < \delta F$, then no miner can make a non-negative profit even if there are no other miners, so in equilibrium there are no active miners. In this case, the value of M that solves Eq. (2) is less than one, and we say that the equilibrium number of miners is zero. For the blockchain to be secure, some number of miners greater than one is necessary. Debate is ongoing about the maximum fraction of mining that can be controlled by one miner, or mining consortium, before security becomes an issue.[13] For our purposes, it is sufficient to set the minimum number of miners for the mining and the blockchain to be viable at $\bar{M} > 1$. If the

[13] See Eyal and Sirer (2014) for analysis of how many miners are needed to keep the bitcoin blockchain viable.

number of miners implied by the zero-profit condition is less than $\bar{M}$ we say that mining is not viable, the blockchain fails, and the number of miners is set to zero.

For any given success rate, λ, Eq. (2) determines the equilibrium number of miners if that number is at least $\bar{M}$. However, the success rate λ is not exogenous. The protocol sets the difficulty of the mining problem so that the arrival rate of new blocks to the chain occurs at approximately some exogenously fixed rate Λ. This rate is part of the security protecting the blockchain as it is designed to allow time for consensus to emerge over the correct ordering of blocks on the blockchain.[14] If there are M miners, each with independent individual success rates of λ, the arrival rate of blocks is λM. In equilibrium, the difficulty of the problem, λ, depends on the number of miners:

$$\lambda M = \Lambda \tag{3}$$

Of course, this requires that there are active miners. If not, then the success rate is undefined.

Result 1: The equilibrium number of miners and arrival rate of success is the simultaneous solution to Eqs. (2) and (3)

[14] The blockchain works by having agreement on the longest chain, and building consensus on what is the longest chain takes time. Shorter block generation times mean more chance of collusion and thus forks, as some miners put blocks on the blockchain that will ultimately be rejected. Because such revisions are possible, most transactions are not viewed as confirmed until at least three or even six blocks have been subsequently added. Shorter block generation times would require even more confirmations for a transaction to be trusted. Bitcoin's confirmation time also works to mitigate slower connections times allowing for broader computer access.

$$M^* = \begin{bmatrix} \dfrac{p[S + \sum_{i \in B} f_i]}{\delta F + e/\Lambda} \;\; if \geq \bar{M} \\ 0 \;\; otherwise \end{bmatrix} \tag{4}$$

and

$$\lambda^* = \begin{bmatrix} \Lambda / M^* \;\; if \; M^* \geq \bar{M} \\ undefined \;\; otherwise \end{bmatrix}. \tag{5}$$

The number of active miners is thus the ratio of the revenue earned by a successful miner to the expected cost of mining. Changes in revenue, $p[S + \sum_{i \in B} f_i]$, that leave $M^* \geq \bar{M}$ affect the equilibrium number of miners but not the rate at which blocks are added to the chain. An increase in the dollar price of bitcoin increases the revenue earned by a successful miner which generates entry of new miners but does not affect the long-run rate at which blocks are added to the chain.

Over time, the block reward, S, declines but this decline can be offset both by fees and by increases in the dollar price of bitcoin. Eventually, S declines to zero so in the limit, absent fees, mining eventually fails because miners would be unable to recover their costs. The timing of when this failure will occur is not straightforward as miners care about total revenue and the declining block reward level can offset (or more) by rising bitcoin prices, at least until the block reward reaches zero. However, as Satoshi predicted, fees have emerged and new blocks continue to be added to the chain.

Remark 1. With no transactions fees and a constant bitcoin price, a decrease in the block reward S to a level at which mining is still viable reduces the number of miners and reduces the difficulty level λ, but it leaves the rate at which new blocks are added to the blockchain, Λ, unchanged.

Remark 2. With a constant block reward level and viable mining, an increase in the bitcoin price or transaction fees increases the number of miners and increases the difficulty level λ but leaves the rate at which new blocks are add to the blockchain, Λ, unchanged.

4.2 Transactions and the mempool

A user who wants to record a transaction moving bitcoins from one address to another address submits a transaction to the node where it is put into the mempool. That transaction succeeds in moving bitcoins only if it is written to the blockchain. We suppose that there are N potential users and the opportunity to engage in bitcoin transactions arises independently for each of them at Poisson rate γ.[15] So, transactions flow into the mempool at rate γN. Initially, we consider a world in which no transactions offer a fee to miners, and miners pick transactions at random from the mempool when they create a block. For expositional convenience, we also treat a block as consisting of one transaction, B = 1.

Transactions flow into the mempool at rate γN and flow out at rate $\lambda^* M^*$. We are interested in the number of transactions in the mempool and in how long transactions wait on average before being recorded to the blockchain. This can be viewed as a queuing problem with random service, as when a miner builds a block he selects from the mempool at random instead of taking the transaction in the pool that has been waiting the longest as in a standard first-in, first-out

[15] Initially, we take the number of potential users, N, as fixed and focus on parameter ranges in which they all chose to participate in bitcoin. In Section 4.4 we determine N endogenously.

queue.[16] The order in which transactions are removed from the pool does not affect the expected size of the pool, so standard queueing theory results can be applied to determine the pool size.

The dynamics of the pool are straightforward. If the arrival rate of transactions is greater than the rate at which transactions are removed from the pool, then the size of the pool grows without bound. If $\gamma N < \lambda^* M^*$ then there is an equilibrium distribution of pool size with mean $N^* = \rho / (1-\rho)$, where $\rho = \gamma N / \lambda^* M^*$. By Little's Law (Little, 1961) the long-run expected waiting time for a transaction in the mempool to be recorded is $w^* = N^* / \gamma N$.

Result 2. If $\gamma N > \lambda^* M^*$, then the mempool grows at rate $\gamma N - \lambda^* M^*$ and waiting times diverge. If $\gamma N < \lambda^* M^*$, then there is an equilibrium distribution of mempool size with mean $N^* = \rho / (1-\rho)$ and mean waiting time $w^* = [\lambda^* M^* (1-\rho)]^{-1}$ where $\rho = \gamma N / \lambda^* M^*$.

This analysis was done with fixed rates of transactions flowing into and out of the mempool and it focuses on the equilibrium pool size and waiting time. Over time, transaction volume on the bitcoin blockchain has been increasing. Consequently, the rate at which transactions have been flowing into the mempool has increased and the rate at which they are flowing out is not increasing commensurately. As a result, ρ is growing, and, as it approaches one the expected size of the mempool grows and average waiting times diverge. This is a problem for users, as a transaction does not result in bitcoins moving from one address to another until it is written to the blockchain. Or, put another way, potential transactions not eventually written to the blockchain simply cease to exist.

[16] The lack of time priority in the mempool is addressed in Aune, Krellenstein, O'Hara, and Slama (2017).

4.3 Waiting times and transaction fees

Once waiting times become significant, one could expect some users to try to get their transactions recorded ahead of others. One way that a user could do this is to implicitly attach a fee to the transaction by moving more bitcoins out of one address than into another address. The difference can be kept by the miner who writes the transaction to the blockchain. This difference is interpreted as the fee offered to the miner who records the transaction. Provided that mining is viable, an increase in fees does not affect the rate at which transactions are written to the blockchain. Fees increase the number of miners, but the protocol increases the difficulty of the problem to keep constant the rate at which successes occur and blocks can be created.

We view the users as playing a game in which each user decides whether to offer a fee, taking as given the decisions of all other users and the rate at which transactions are written to the blockchain. We assume that any transaction written to the blockchain generates a benefit V to the user, but that this benefit is reduced by the amount of any fee paid and by the waiting time. User expected payoffs are

$$V(f,w)=\begin{bmatrix} V-pf-aw & \text{if } V-pf-aw\geq 0 \\ 0 & \text{otherwise} \end{bmatrix} \tag{6}$$

for a user who pays a fee in bitcoins of f with dollar value pf and whose transaction has an expected delay of *w*, at cost $a>0$ per unit of expected delay, before it is recorded to the blockchain. A user who does not pay a fee still gets the benefit of V if his transaction is recorded to the blockchain, but this benefit is reduced by his longer waiting time. We initially consider only a single fee level $f>0$, and we focus on parameter ranges in which the net benefit of being a user is large enough that all N potential users continue to use bitcoin. Section 4.4

contains an analysis of users' decisions to enter the bitcoin mempool or to complete their transaction without using bitcoin. The fee game can, in principle, have Nash equilibria in which none, some, or all users pay the fee. If all users offer the same fee, then nothing changes other than a transfer from users to miners and an increase in the equilibrium number of miners. If some users offer higher fees than other users offer, then their transactions are selected ahead of those who offer lower, or zero, fees.

It is useful to compute the expected waiting time for a user who pays a fee and for one who does not pay a fee when fraction α of the users pay a fee. Users who offer a fee have their transactions recorded before any transactions without fees, so the fraction of non-fee-paying users does not affect the waiting time for fee-paying users. The expected waiting time for fee-paying users is thus the waiting time in Result 2 with the arrival rate of transactions reduced to $\alpha\gamma N$:

$$w_f(\alpha) = [\lambda^* M^* (1 - \rho_f(\alpha)]^{-1} \tag{7}$$

where $\rho_f(\alpha) = {\alpha N\gamma}/{\lambda^* M^*}$. $\lambda^* M^*$, the equilibrium rate at which transactions are written to the blockchain, is not affected by fees.

The non-fee-paying users face a different situation as their expected waiting time depends on both the number of fee-paying and non-fee-paying users. To derive this waiting time, note that fees do not affect the long-run mean size of the mempool. It remains at $N^* = {\rho}/{1-\rho}$.

Remark 3. Suppose that fees and waiting times are such that all mining is viable and that all N potential users continue to participate. Then, fees do not affect the equilibrium size of the mempool or the equilibrium rate at which transactions are written to the blockchain.

Let $\rho_n(\alpha) = {(1-\alpha)\gamma N}/{\lambda^* M^*}$ and note that for any $\alpha \in [0,1]$, $\rho = \rho_f(\alpha) + \rho_n(\alpha)$. Letting the equilibrium numbers of fee-paying and non-fee paying users be $N_f(\alpha)$ and $N_n(\alpha)$, respectively, we have

$$N^* = N_f(\alpha) + N_n(\alpha) = \frac{\rho_f(\alpha) + \rho_n(\alpha)}{1-(\rho_f(\alpha)+\rho_n(\alpha))} \tag{8}$$

As $N_f(\alpha) = \frac{\rho_f(\alpha)}{1-\rho_f(\alpha)}$ we have

$$N_n(\alpha) = \frac{\rho_n(\alpha)}{(1-\rho_f(\alpha))(1-\rho)} \tag{9}$$

Then, by Little's law, the expected waiting time for non-fee-paying users is

$$w_n(\alpha) = [\lambda^* M^* (1-\rho_f(\alpha))(1-\rho)]^{-1} \ . \tag{10}$$

Calculation shows that for any $\alpha \in (0,1)$, $w_n(\alpha) > w^* > w_f(\alpha)$. So, fee-paying users do not wait as long as non-fee-paying users.

The payoff to a user who pays the fee is $V_f(\alpha) = V - pf - aw_f(\alpha)$ and for a non-fee payer it is $V_n(\alpha) = V - aw_n(\alpha)$. The difference in payoffs is

$$\Delta V(\alpha) = V_f(\alpha) - V_n(\alpha) = a(w_n(\alpha) - w_f(\alpha)) - pf \tag{11}$$

Substituting in waiting times and simplifying yields

$$\Delta V(\alpha) = \frac{a\rho}{\lambda^* M^* (1-\rho)(1-\rho_f(\alpha))} - pf \tag{12}$$

Calculation shows that $\Delta V(\alpha)$ is increasing in α. That is, the benefit to paying the fee is increasing in the fraction of users who pay the fee.

Games in which a player's gain in payoff from taking an action versus not taking it is increasing in the fraction of others taking that action are said to exhibit strategic complementarity. These games have numerous properties that make them easy to analyze. Multiple Nash equilibria can exist but there is an incentive for coordination on one action (fee-paying) or the other (non-fee- paying). Mixed equilibria (a fraction $0 < \alpha < 1$ of users paying a fee) can also exist, but they tend to be unstable. Particularly useful for our purposes, equilibria can be Pareto-ranked and Pareto-superior equilibria tend to be stable.

Result 3. The Bitcoin fee-paying/waiting-time game exhibits strategic complementarity.

The critical parameter in the user's decision problem is

$$z = \frac{a\rho}{pf\lambda^* M^*(1-\rho)} \tag{13}$$

The parameter z summarizes the costs and benefits to the user of paying the fee $\left({pf}/{a} \right)$, the no-fee mean size of the mempool $\left({\rho}/{1-\rho} \right)$ and the expected waiting time for a miner to arrive. Given the rate at which transactions flow out of the mempool, the factor that drives changes in the size of the mempool and thus waiting times is the arrival rate of transactions. If that rate increases over time, then z increases. Restating the payoff differential using z, we have $\Delta V(\alpha) = pf[\frac{z}{1-\alpha\rho} - 1]$ which is negative for all $\alpha \in [0,1]$ if $z < 1-\rho$ and positive for all $\alpha \in [0,1]$ if $z > 1$. If the mempool is not too large ($N^* = {\rho}/{1-\rho}$), and thus mean waiting times are short, not paying a fee is the dominant strategy. If the mempool is large and waiting times are sufficiently long, then paying the fee is dominant.

Figure 2 illustrates the effect of z on the curve describing the payoff difference as a function of the fraction paying the fee. The three curves describe the payoff differential for low, intermediate, and high levels of z. All three curves are upward sloping reflecting strategic complementarity. For intermediate levels of z there are three equilibria, zero, α^* and one. But α^* is unstable as indicated by the arrows in the figure. If users expect an $\alpha < \alpha^*$, then the payoff to paying the fee is less than the payoff to not paying the fee, and if $\alpha > \alpha^*$, the payoff to paying the fee is greater than the payoff to not paying the fee. Although α^* is a Nash equilibrium, an expectation of α^* is delicate and we would not expect to see expectations converge there.

Insert Fig. 2 near here

Result 4. Suppose parameters are such that mining is viable and all N potential users continue to participate. For any parameters in this range, the fee-paying game has at least one Nash equilibrium. For Nash equilibria in the fee-paying game,

1. If $z < 1 - \rho$, not paying the fee is a dominant strategy and the unique Nash equilibrium is one in which no user pays a fee, $\alpha = 0$.
2. If $z > 1$, paying the fee is a dominant strategy and the unique Nash equilibrium is one in which every user pays a fee, $\alpha = 1$.
3. If $1 - \rho \le z \le 1$, there are three equilibria: $\alpha = 0$, $\alpha = 1$ and one in which an interior fraction (α^*) of the users pay a fee, where $\alpha^* = \rho^{-1}[1 - z]$. These equilibria are Pareto ranked (from the point of view of users) with the zero-fee equilibrium yielding the highest user payoffs.

Fig. 3 describes the Nash equilibrium fraction of users paying a fee as a function of the parameter z. Figs. 2 and 3 have interesting implications for the dynamics of fees. Suppose that over time indexed by t the arrival rate of transactions varies exogenously and thus the parameter z_t varies exogenously. If initially, the arrival rate is low, then z_t is low and the only equilibrium is one with a zero fraction of transactions paying a fee, i.e., $\alpha_t = 0$. For an intermediate range of z's, there are multiple equilibria, but as Fig. 2 illustrates the equilibrium with fraction $\alpha^* \in (0,1)$ is unstable. As the zero-fee equilibrium has the greatest payoff, it is stable and as, z_t increases into the intermediate range, it seems likely to persist. Once z_t reaches the critical point at one, the equilibrium switches to everyone paying the fee. As is usual for a game with strategic complementarity, we have persistence of the Pareto optimal equilibrium, instability of the interior equilibrium, and critical behavior of the dynamics.

Insert Fig.3 near here

An important feature of these equilibria is their Pareto ranking. Both the outcomes in which everyone pays a fee, $\alpha = 1$, and no one pays a fee, $\alpha_t = 0$, are stable equilibria. But, from the point of view of the users, they are not equally preferred. Everyone not paying a fee is clearly better than everyone paying a fee. This is because the fee itself does not change the expected waiting time for getting a transaction posted to the blockchain because everyone is paying it.[17] Thus, adding a fee does not change the user experience but results only in paying the miners more which, in turn, simply induces more people to be miners. But, in equilibrium, each miner is still making only zero expected profit. Provided mining is viable, adding fees does not improve miner welfare but only increases the dead-weight loss (from electricity usage, etc.) that arises from more mining. Consequently, transaction fees are not welfare-enhancing.

4.4 Extensions

[17] In coordination games, mechanisms that improve coordination could improve user welfare. Transaction fee estimator algorithms, which are now readily available, do just that by helping users estimate better what equilibrium transaction fees will be in the next block.

The model developed thus far is simplified so that we could illustrate our points about bitcoin without distractions that added little to the analysis. Here, we consider a few extensions and show how they modify the details of our analysis but do not change the qualitative points we make.

4.4.1 Multiple fees and heterogeneous users

The analysis has focused on a single fee and users who chose either to pay this fixed fee or to offer no fee. This analysis can be generalized to multiple fee levels. Suppose that users consider fee levels (in bitcoins) of $f_1 > f_2 > \cdots > f_K = 0$ and let the fraction of users who chose fee f_k be α_k. Then, the $\alpha_1 N$ users who pay the highest fee are selected first, those who pay the second highest fee are selected next once the highest fee-paying users are removed, and so on. Their waiting times can be defined recursively as in section 4.3.1. If users are identical, then having positive fractions paying multiple fee levels requires indifference and thus fractions at each level that equalize the costs and benefits of each fee level.

More realistically, users can be heterogeneous with respect to the benefit they receive from a transaction and their utility cost of waiting time. In this case, equilibria in which users are separated by fee levels are possible. This possibility can be illustrated even with only two fee levels, *f* and zero, as in section 4.3.1. Suppose that there are two types of users with values and waiting costs (V_i, a_i) where $V_1 \geq V_2$, $a_1 \geq a_2$ and that fraction α_1 of the users are high types (have high value and high waiting cost). High types pay the fee, f, and low types pay no fee if $V_1 - pf - a_1 w_f \geq 0$, $V_2 - a_2 w_n \geq 0$ and

$$pf/a_2 \geq w_n - w_f \geq pf/a_1 \tag{14}$$

where w_n, w_f are the waiting times for non-fee-paying and fee-paying users respectively.

Calculation shows that

$$w_n - w_f = \frac{N^*}{\lambda^* M^* - \alpha_1 N\gamma} \tag{15}$$

The right-hand side of this equation is the equilibrium size of the mempool divided by the difference in rates between the flow out of the mempool and the flow of high type users into the mempool. Inequality (14) relates the difference in waiting costs to the difference between the flow out of the pool and the flow of high types into the pool. With heterogenous users, the optimality of the non-fee-paying equilibrium may not hold due to the differential impacts of fees across users.

4.3.2 Exit by users

We focused on equilibria in which all potential users chose to submit transactions to the mempool whenever opportunities arise. If waiting times are too long, or if some transactions never get recorded on the blockchain, those users can choose not to submit transactions to the pool. Presumably, they would instead conduct their transactions in an alternative currency. Result 1 implies that the dollar value of the fee necessary to make mining viable is

$$pf \geq p\overline{f} = MAX\{0, \overline{M}\delta F + \frac{\overline{M}e}{\Lambda} - pS\} \ . \tag{16}$$

At this fee per transaction and the waiting time if all users participate, $N^* = \rho / (1-\rho)$, the net benefit of transactions can be so low that not all users choose to participate.

Suppose that users can obtain net benefit $\overline{v} \geq 0$ from conducting their transaction without using bitcoin. Then, the waiting time for a user paying fee *f* must be no greater than $(V - \overline{v} - pf) / a$ as otherwise the user would not submit the transaction to the mempool. As we have an exogenous lower bound on *f*, this implies an upper bound on waiting times of $\overline{w} = {(V - \overline{v} - p\overline{f})}/{a}$. This upper bound on waiting times, in turn, implies an upper bound on the number of users who would be willing to participate in using bitcoins for transactions

$$N \leq \frac{\lambda^* M^* - \overline{w}^{-1}}{\gamma} \tag{17}$$

The equilibrium flow of transactions into the mempool (γN) is bounded by the flow of transactions being written to the blockchain ($\lambda^* M^*$) minus the inverse of the maximum waiting time that users are willing to bear.

In summary, our model provides a variety of insights into the driving forces behind the dynamics of the bitcoin blockchain. It provides a framework for understanding the emergence of fees as the rate of arrival of transactions and thus waiting times have increased, and it illustrates the impact of block reward and the price of bitcoin on the equilibrium number of miners and the absence of any impact on equilibrium fees. It also demonstrates the welfare effects of transaction fees and how they can contribute to the deadweight loss associated with mining. Finally, it demonstrates both why a diversity of fees emerges as heterogeneous users use fees to battle for shorter waiting times and why some user groups can decide to eschew using bitcoin altogether when transactions fees get too high, implying a limit on the potential growth of bitcoin usage.

5. Empirical analyses

Our model suggests a variety of empirical relations affecting the operation and stability of the bitcoin blockchain. The mean waiting time should affect both the fraction of transactions paying a fee and the level of fees. Also, the mean waiting time depends on the flow of transactions into the mempool and the flow of transactions out of the mempool. The flow out is determined by the bitcoin protocol and is exogenous. The flow in is exogenous as long as waiting times and fees do not affect the willingness of users to participate. That is, waiting times can affect the fees that users pay, but waiting times and fees affect only the flow into the mempool once they become long enough or high enough to discourage participation. Once participation is affected by these variables, waiting times are endogenous. Furthermore, the declining block reward should have no impact on fees as it affects only the equilibrium number of miners (as long as mining remains viable) and not the rate at which transactions are added to the blockchain.

In this section, we examine these predictions with respect to the role played by transactions fees. We provide empirical evidence on the predicted relation between waiting times and zero-fee equilibrium, and between waiting times and transaction fee levels. We also test whether the block reward level has any effect on transaction fee usage or levels.

Before turning to the data, it is useful to raise several caveats. One is that volume is very hard to measure accurately in bitcoin because we observe only how many bitcoins go from one address to another.[18] A single person often has many addresses for privacy purposes, so transfers between such addresses add to total reported volume. A related issue is that spam attacks on the mempool can artificially inflate reported mempool levels, introducing noise in this data series. Having noted this concern, it is standard to use reported volumes and mempool levels and we do so in our empirical work.

A second caveat is that the quality and quantity of some data series are limited. While data drawn directly from the blockchain exist from the genesis transaction, the blockchain does not record many variables of interest to our study. Data series such as mempool volume and average

[18] This concern relates to calculating on-chain volume. There is also off-chain volume which does not affect the block chain, but it is essentially impossible to calculate.

waiting times can be found from other sources, but they often feature non-existent or spotty data in earlier time periods and more complete data in recent periods. This results in series having missing data. In our empirical work, we adjust for this when appropriate and restrict our sample periods to only complete data intervals when data quality concerns are too high.

A third limitation is that no data exist on the number of miners, the number of potential users of bitcoin, or the characteristics of these users. Therefore, we cannot directly test the equilibrium prediction from the miner's game. We can ask whether the predicted separation between the mining game and the user experience, waiting times and fees, is borne out in the data. On the user side, we cannot test predictions about who is discouraged from participating. We can test predictions about the relation between waiting times and fees, and we do report some recent data on bitcoin acceptance by retail firms.

A fourth issue arises because data recorded to the blockchain comes from transactions placed in blocks by miners. Transactions in the mempool that do not make it onto the blockchain are not captured in any currently available data set. For many empirical questions, this is not a problem, but, for others, such as analyses of waiting times or transaction fees, this is an issue. Mempool waiting time data, for example, only captures those successful transactions, and omits the ones that miners opt not to select, imparting a downward bias to the waiting time data series. Also, we have data only on median waiting times while our theory provides results for mean waiting times. As empirical mean waiting times can be affected by outliers and are thus less reliable, median waiting times could be more appropriate. Conversely, transaction fee data could be upwardly biased because transactions attaching too low a fee to attract a miner are excluded. More extensive data series may emerge over time, but, for now, these issues preclude testing some predictions of our model.

Finally, our sample period runs from the inception of bitcoin up to April 2017. This interval captures the evolution of bitcoin trading, changing levels of block rewards, and the emergence of transaction fees, but it includes neither the chaotic period in fall 2017 when bitcoin prices reached $19,783 nor the subsequent decline of bitcoin prices in spring 2018. While transaction fees are now lower than during this bubble period, they are comparable to fee levels in our later sample period.[19]

5.1 Data and summary statistics

Our data come from several sources. The bitcoin blockchain stores all transaction information from when the system began, and a copy of this information is stored on every node so that it can validate new transactions as they come up. We started a bitcoin node and allowed it to download and validate all transactions from when bitcoin started until the end of April 2017. Each transaction in the bitcoin node has the following information: time when the transaction settled, the difficulty of the block, the value of the coins transferred, the block reward level, and the mining fees paid. We obtain most of the information in this paper directly from the blockchain or infer it from data on the blockchain. To answer transient queries, we use blockchain.info, which runs many bitcoin nodes and reports on various metrics about the network as a whole. We use blockchain.info for the inflow rate of transactions, the size of the transaction mempool, and the median waiting time for a transaction. Finally, to get the bitcoin to US dollar conversion rate we use the Coindesk Bitcoin Price Index (XBP), which is based on an average of bitcoin prices across leading global exchanges.

5.2 Transactions fees and equilibria

[19] For example, average transaction fees on Nov. 19, 2018 were $1.11 and on April 20, 2017 they were $1.05.

We first look at how the fraction of fee-paying transactions posted to the blockchain has evolved over time. As Fig. 4 shows, prior to approximately the beginning of 2011, very few transactions paid fees, consistent with our no-fee equilibrium. This changed over 2011 and 2012, when transactions with attached fees moved from being a relatively small fraction to predominating trading. By late 2013, an only-fee-paying equilibrium seems to have emerged.[20] Our theory suggests an explanation for why such a dynamic pattern appears.

Insert Fig. 4 near here

Our model generates a negative relation between the arrival rate of transactions to the mempool, which drives growth in the mempool size and thus waiting times, and the propensity of users to pay a fee. Turing first to the percent of users paying a fee, $y = 100\alpha,$ and how it relates to the exogenous parameter, z, we approximate the empirical relation between these two variables as a power law where we expect a negative exponent on z.[21] That is,

$$y = \beta_1 z^{\beta_2}$$

where $\beta_1 > 0, \beta_2 < 0$. This relation is linear in logs,

$$\ln(y) = \ln(\beta_1) + \beta_2 \ln(z)$$

[20] Prior to 2011, most transactions did not offer a fee and miners had little financial incentive to fill a block with transactions (a successful miner earned the block reward even if he did not put any transactions on the blockchain). As result, this period has empty blocks. The percent of zero-fee transactions for these blocks is undefined so we treat them as missing observations. Other blocks had transactions, but these transactions did not pay fees so the other early blocks have 100% zero-fee transactions.

[21] Our model yields a negative relation between these two variables and, as y must lie between zero and one hundred, it cannot be linear.

and in our empirical work we use the approximation $z \approx p^{-1}MWT$. The pairwise correlations coefficient between the logarithm of the percent of zero-fee transactions and the logarithm of the median waiting time over the sample period December 4, 2011 to April 28, 2017 is -0.10, lending support to our hypothesis that β_1 is negative.

Median waiting time data are available from December 3, 2011 to April 28, 2017. Even over this period, we have waiting time data only on alternate days until July 7, 2015. Waiting times were at first highly variable and long, and then they quickly dropped by early 2013 and stayed nearly constant at a level under ten minutes and increased until early in 2016, when they again became highly variable and increased substantially. We split our sample period into two subperiods, and we take waiting time to be exogenous in the early period and endogenous in the late period. The first, early period is December 3, 2011 through April 23, 2016 and the second, late period is April 24, 2016 through April 28, 2017. In addition to the observation about waiting time, there are two other motivations for this choice. First, by April of 2016, almost all transactions written to the blockchain pay a fee and the mean dollar value of transactions fees per byte, which heretofore had been less than \$0.001, increased in early 2016 to about \$0.0026. Second, on April 24, 2016, data on average daily mempool size which seems a reasonable candidate as an instrument for endogenous waiting times first became available on a nearly daily basis. We report summary statistics for each subperiod in Table A1 in the Appendix.

For the early period (December 3, 2011 through April 23, 2016), which has 944 observations, we ran the following regression (robust standard errors were used to correct for heteroskedasticity and dependence of the errors)

$$\ln(\%0\,fee_t) = \beta_1 + \beta_2 \ln(MWT_t) + \beta_3 \ln(BTCprice_t) + \beta_4 blockreward_t + \varepsilon_t \quad (20)$$

where MWT is median waiting time, and the variables BTC price and blockreward (number of newly issued bitcoins awarded to the winning miner) capture exogenous effects on fees. Bitcoin price is included because our empirical approximation implies that the coefficient on its log should be -1. Block reward is included to allow for exogenous effects on bitcoin trading. We are also interested in seeing if the prediction of our model that fees are not affected by block reward levels holds in our sample period. Table 1, regression 1, gives these estimation results.

Insert Table 1 near here

The regression results show that median waiting times have a negative and significant effect on the percentage of zero-fee transactions. Thus, as median waiting times increase, the percentage of transactions posted to the blockchain with no transaction fee decreases. A 1% increase in median waiting time results in a 1.8% reduction in the percent of zero-fee transactions. This is consistent with the prediction of our model that as median waiting times increase, the equilibrium shifts away from zero-fee transactions. A 1% increase in BTCprice reduces the percent of zero-fee transactions by 0.8 %, which is significantly different from our prediction of 1%, but it is at least of the correct sign and roughly of the right magnitude.

Regression 2 of Table 1 provides evidence from an alternative specification that adds changes in the log of BTCprice to Regression 1. We include this specification because BTCprice has increased dramatically over our sample period, raising the natural concern that the apparent correlation between the percent of zero-fee transactions and MWT could be driven by a spurious correlation between these variables and BTCprice.[22] This specification does not result in a

[22] A natural approach to addressing this issue would be to look at first differences of these variables. Unfortunately, our data are unevenly spaced over this sample period, making first differences a mix of daily, twice

significant change in the coefficient on MWT. Regression 3 includes a one-week lagged value of MWT. It has a negative and significant effect on the percentage of zero-fee transactions, but no significant effect on the coefficient on MWT.

In the regressions reported in Table 1, we treat median waiting time as exogenous. In the last year of our sample (not included in the regressions reported in Table 1), it is more reasonable to treat median waiting time as endogenous. We ran various two-stage least squares regressions using instrumental variables to capture the influence of exogenous model variables on the endogenous median waiting times. The ideal instrument would be highly correlated with the endogenous variable but uncorrelated with the error term. In our setting, instrument choice is complicated both by the limited number of variables available and by the restricted sample periods of some data series. We use average daily mempool size as an instrument as it should affect median waiting times (pairwise correlation coefficient = 0.65) and does not have an obvious direct effect on the percent of zero-fee transactions (pairwise correlation coefficient = -0.15). However, that variable is available only on a (nearly) daily basis for the period April 24, 2016 – April 28, 2017. For this sample period, with 351 observations, we ran the following instrumental variables regression in first differences, using generalized method of moments (GMM) to correct for potentially heteroskedastic and auto-correlated errors:[23]

$$\Delta \ln(MWT_t) = \alpha_1 + \alpha_2 \Delta Mempool_t + \alpha_3 \Delta \ln(BTCprice_t) + \alpha_4 blockreward_t + \varepsilon_{1t} \qquad (21)$$

and

daily and sometimes many days of observations. We consider first differences for the shorter period in which we have complete data in Table 2.

[23] Regressions in levels are reported in the appendix in Table A3.

$$\Delta \ln(\%0\, fee_t) = \beta_1 + \beta_2 E\Delta \ln(MWT_t) + \beta_3 \Delta \ln(BTCprice_t) + \beta_4 blockreward_t + \varepsilon_{2t} \qquad (22)$$

where Δ ln(MWT) is daily change in log median waiting time and E Δ ln(MWT) is its estimated value from the first-stage regression, and the variables Δ ln(BTCprice), blockreward, and Δ ln(mempool) capture the exogenous effects on median waiting times. With this empirical specification, we explain daily changes in the percent of zero fee transactions using daily changes in median waiting times, mempool size, and BTCprice.[24]

Table 2, regression 1, provides these results. The first stage of regression 1 shows that only the change in ln(mempool) has a significant effect on median waiting time. Specification tests show that the instruments chosen have high correlation with MWT, and the minimum eigenvalue statistic of 51.6 for the first-stage regression (using limited information maximum likelihoos) indicates for all standard levels of acceptable bias that our instrument is not weak (see Craig and Donald,1993, and Stock and Yogo, 2005). Finally, the hypothesis that MWT is exogenous is rejected [the hypothesis that ln(MWT) is exogenous has a C statistic of 4.79 with p-value=0.03] lending support to our instrumental variables approach.

Insert Table 2 near here

Turning to the second stage of Regression 1, we find that changes in the log of estimated median waiting time have a negative and significant effect on changes in the log of the percentage of zero-fee transactions. These results are consistent with our model's prediction that the presence of transaction fees in equilibrium is influenced by user queuing effects. Allowing for the possibility of a time trend of the form $y_t = \beta_1 z_t^{\beta_2} \alpha^t$, the constant in our first differences

[24] Block reward takes on only two values in this sample. We do not consider changes. Removing block reward from the regressions does not significantly change the coefficients on the change in mempool size or the estimated change in median waiting time.

regression provides an estimate of α. The estimate of α is 0.09 with a z-statistic of 0.51, so the presence of the time trend of this form is strongly rejected.

Regressions 2 and 3 of Table 2 provide evidence from specifications that include additional instruments. In Regression 2, we include the one-week lagged value of the change in mempool size as an instrument. The coefficient on estimated change in log MWT in the second stage using this additional instrument in the first stage is not significantly different from the coefficient in Regression 1. None of the coefficients in either the first or second stage regressions changes significantly, but we cannot reject the null hypothesis that the overidentifying restrictions are valid. Hansen's test of validity of the instruments, the J statistic, is Chi(1) = 0.12 with p-value = 0.91, so we cannot reject the null hypothesis that the instruments are valid. Regression 3 includes the log of BTCprice. This specification also does not significantly change the results found in Regression 1. Most important, the block reward is not significant in explaining changes in the percent of zero fee transactions. Taken together, these results provide strong support for the hypothesis arising from our model that the percent of zero-fee transactions is driven by waiting times.

5.2.1 Alternative specifications

A natural concern with the empirical analyses in Tables 1 and 2 is that our log-log empirical model is an approximation of both our theoretical model and the empirical process. So we also consider a linear model. Because the dependent variable, the percent of zero-fee transactions, can take on values only between zero and one hundred, we ran a fractional response model that allows the dependent variable in our regressions (the fraction of zero-fee transactions) to be constrained to take on values (continuously) in (0,1). These results are reported in the Appendix in Table A2, with Panel A providing OLS results for the early period and Panel B providing IV results for the later period. Magnitudes of the coefficients, and the marginal effects of median waiting time on the fraction of zero-fee transactions, are naturally different from those reported in Table 1, but our qualitative conclusion remains unchanged. Median waiting time has a significant, negative effect on the percent of zero-fee transactions.

5.3 *Queueing, mining, and transaction fees*

In the bitcoin blockchain ecology, transaction fees can solve two problems: They can incentivize miners to participate by offsetting declining mining revenues, and they can solve a queuing problem for users. Consequently, both mempool waiting times and the block reward level could influence fee levels.

Fig. 5, Panel A shows that transaction fees per byte were initially zero in 2011, showed substantial volatility in the next two years and increasing again in more recent time periods. Panel B shows transaction fees per byte over the latest year of our sample. This period saw a decrease in the block reward as well as increases in the median daily waiting time. Thus, an interesting question is: How do these variables influence average transaction fee levels? Following our previous analysis, we split our sample period into two pieces. The first period is December 4, 2011 through April 23, 2016 and the second is April 24, 2016 through April 28, 2017. For the early period, we ran the following regression (robust standard errors were used to correct for heteroskedasticity and dependence of the errors):

$$txfee_t = \beta_1 + \beta_2 MWT_t + \beta_3 BTCprice_t + \beta_4 blockreward_t + \varepsilon_t \qquad (23)$$

Insert Fig. 5 near here

The regression results show that median waiting times do not have a significant effect on the daily average transaction fee and that the block reward has a significant positive effect on the daily average transaction fee. Both results are counterintuitive. Given that median waiting times did have a significant and negative effect on the fraction of fee-paying transactions, these results could suggest that over the entire sample period the mean changes in fees are not well captured by either waiting time or block reward effects. Certainly, the volatility of fees in the 2012 - 2013 period suggest that a variety of factors could be at play. Alternative specifications, corresponding to the alternative specifications in Table 1, are provided in Regressions 2 and 3 of Table 3. Although these alternatives change the coefficients, the counterintuitive results remain.

As in our analysis of the percent of zero-fee transactions, we treat the mean transaction fee as endogenous in the later period and use average daily mempool size as an instrument. For the late period, with 351 observations, we ran the following instrumental variables regression in first differences, using GMM to correct for potentially heteroskedastic and auto-correlated errors:

$$\Delta MWT_t = \alpha_1 + \alpha_2 \Delta Mempool_t + \alpha_3 \Delta BTCprice_t + \alpha_4 blockreward_t + \varepsilon_{1t} \tag{24}$$

and

$$\Delta transfee_t = \beta_1 + \beta_2 E\Delta MWT_t + \beta_3 \Delta BTCprice_t + \beta_4 blockreward_t + \varepsilon_{2t} \tag{25}$$

where the variables are as previously defined. Regression 1 of Table 4 provides the results from this regression.

Most important, changes in estimated waiting times are a significant explanatory factor of changes in mean transaction fees and block reward is not a significant factor. An increase of one standard deviation in the first-difference of mean waiting time yields an increase of 0.94 standard deviations in the first-difference of transactions fees. The lack of an impact of the block reward confirms one of the predictions of our model. The coefficients on estimated median waiting times are positive and significant in all three of our alternative specifications, consistent with the predictions of our model.

6. Conclusions and policy implications

Nakamoto (2008) conjectured that, over time, the bitcoin blockchain would have very large volumes or none at all. We show in this paper that his reasoning was mostly correct. In the absence of transaction fees, eventually the blockchain is not viable. But even with transactions fees, an upper bound exists on the size of the blockchain imposed by the waiting time confronting users in the mempool. As this waiting time becomes large, users exit the blockchain in much the way that miners exit the blockchain when their revenues no longer generate profits. Thus, the equilibrium in the bitcoin blockchain is a complex balancing of user and miner participation. Transaction fees play a crucial role in affecting both clienteles and, thus, in influencing the viability of the blockchain.

As the bitcoin ecology migrates to a more market-based system, a variety of interesting issues become apparent. One such issue is the role played by microstructure features such as exogenous structural constraints. While constraints limiting the growth of new bitcoin issuance are in line with the system's original design, the constraints on block size are a relatively recent addition intended to decrease the system's vulnerability to attack. A perhaps unintended consequence is that this constraint exacerbates the imbalances between mempool inflow and outflow, potentially leading to instability in the blockchain. The dramatic increase in bitcoin transaction fees in late 2017 is viewed by some as evidence of just such effects.

Such instability recently led to debate in the bitcoin community over whether to increase the bitcoin block size. Proposals to increase the block size to as much as 20 MB were floated, but little consensus emerged. Caffyn (2017. p) notes that some "maintain that limiting block size in the short-term will create a self-regulating market for transaction fees" and that "will increase miners' incentive to process transactions which will benefit the health of the system." Our results suggest otherwise. Increasing transaction fees will increase the number of miners, but this, in turn, will trigger increases in the difficulty level to control the creation rate of new blocks, thereby raising the costs to miners. In a competitive market, this would not lead to an overall increase in compensation. Moreover, as we have shown here, transaction fees reflect queuing problems facing users. As median waiting times rise, users can choose to forgo transacting in bitcoin, perhaps opting instead for one of many other crypto-currencies or simply remaining with traditional fiat monies. As we argue in this paper, transaction fees alone are unlikely to solve the challenges facing the bitcoin blockchain.

Debates also are ensuing on other structural issues. Bitfury, for example, argued for requiring miners to use a "smallest transaction first" rule in selecting transactions for blocks. Eyal, Gencar, Sirer, and Renesse (2016) propose new sequencing in Bitcoin NG, whereby the winning miner is allowed to put all the transactions he can from the mempool onto the blockchain until the next miner is chosen. Huberman, Leshno, and Moallemi (2017) propose shortening the interval in which blocks are put on the blockchain. Perhaps the most successful innovation, SegWit2, changed the format of transactions, allowing for more transactions to fit into a block.

Another, and perhaps more fundamental, problem is that different clienteles have different needs of, uses for, and even philosophies regarding bitcoin and a blockchain.[25] In its current configuration, the bitcoin blockchain can process at most seven transactions per second.

---

[25] The issue is complicated both by technology and by governance. As Caffyn (2017) notes, Chinese miners account for more than half the computing (or hashing) power of the network, and they argued against larger sizes due to bandwidth concerns. Other parties, such as exchanges and wallet companies, have different views on the optimal size. Still others call for dynamic increases in the block size over time tied to volume growth. Because governance issues are settled by computing power votes, disagreement can lead to potential forks in the blockchain.

Increasing the block size affects this processing ability linearly, so even a 20 MB blocksize will still process only approximately 60 - 80 transactions per second. By comparison, Visa, Inc. has the capacity to do up to 50 thousand transactions per second. For users who want bitcoin as a type of digital gold, the lower processing capabilities are more than adequate; for users wanting bitcoin as an alternative to Paypal, they are not. New innovations such as Level II technology (which essentially adds network layers on top of the bitcoin blockchain) show promise in alleviating these processing issues, and proposed Level III solutions could do even better. How successful these approaches will be is unclear.[26]

These bitcoin processing issues, combined with transaction cost issues, raise the prospect that while bitcoin can continue to develop as a financial asset, its development as a transactional medium is more problematic. In 2018, only 3 of the top 500 internet retailers accept bitcoin, a number 40% lower than the year before.[27] Neither Amazon nor Walmart accept the digital currency, and their ranks have been joined in recent months by the likes of Steam and Stripe, companies that used to accept the virtual currency. Perhaps more telling is that 1n 2018 the North American Bitcoin Conference stopped accepting payment in bitcoins.

Designing markets to operate efficiently is always challenging, and it is becoming particularly so in bitcoin as divergent clienteles emerge. The recent forking of the blockchain caused by the introductions of Bitcoin Cash (BCC) and Bitcoin Gold (BTG) reflect exactly such

[26] The Lightning Network is one example of how this works. Here transactions are processed off-chain and net positions are reported on-chain.

[27] See "Bitcoin Acceptance Among Retailers Is Low and Getting Lower", https://bloomberg.com/news/articles/2017-07-13

issues, as some users opt to change the rules to make the blockchain better meet their needs.[28] In other market settings such as equity trading, this divergent clientele problem resulted in fragmentation of trading across both different venues and different platforms. Whether the bitcoin blockchain remains a single entity or fragments into a collection of bitcoin-linked blockchains could well depend upon how these market microstructure issues are resolved. Certainly, these are important issues for future research.

[28] Bitcoin cash began on August 1, 2017 with the production of a 1.9 MB block that was not valid on the bitcoin network. Its larger 8 MB block size is intended to allow for more transactions to be posted to the blockchain. For more discussion, see https://www.bitcoincash.org/ . Bitcoin Gold began on November 12, 2017 with an airdrop of the new currency to existing bitcoin holders. The protocol of this new currency is aimed at blocking the use of specialized chips for mining. See https://www.coindesk.com/bitcoin-gold-goes-live-bumpy-blockchain-launch/.

Panel A: Total transaction fees yearly denominated in millions of US dollars and BTC

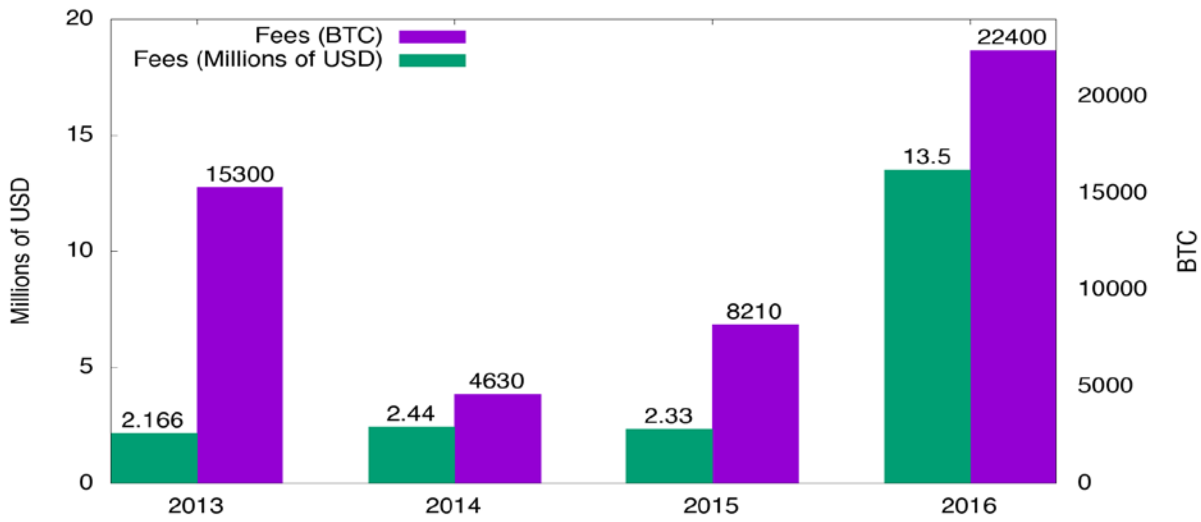

Panel B: Percent of miner revenue derived from transaction fees

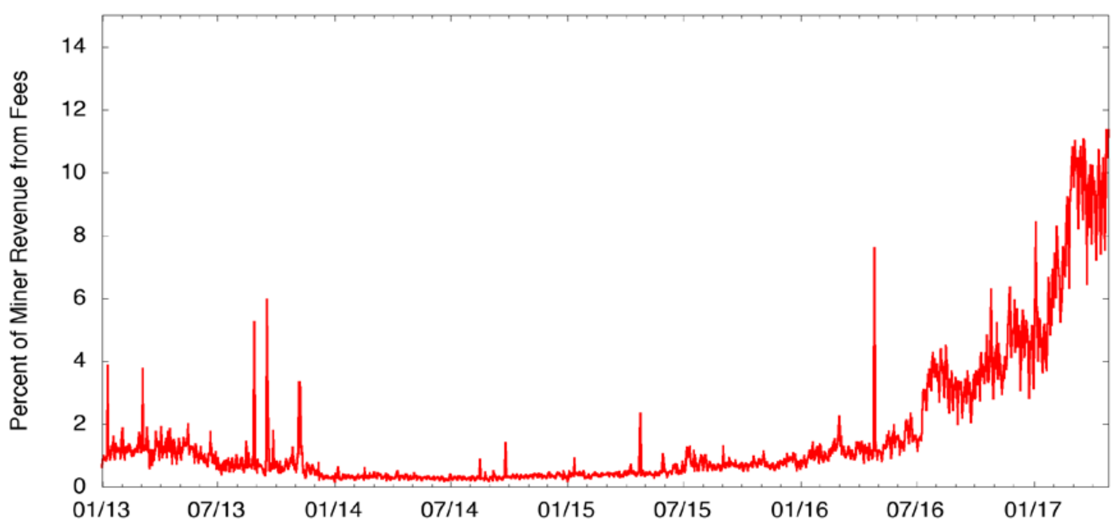

Figure 1. Total bitcoin transaction fees. This figure gives yearly data on the level of transaction fees [measured both in US dollars and in bitcoins (BTC)] and on the relative importance of transaction fees for overall miner revenue.

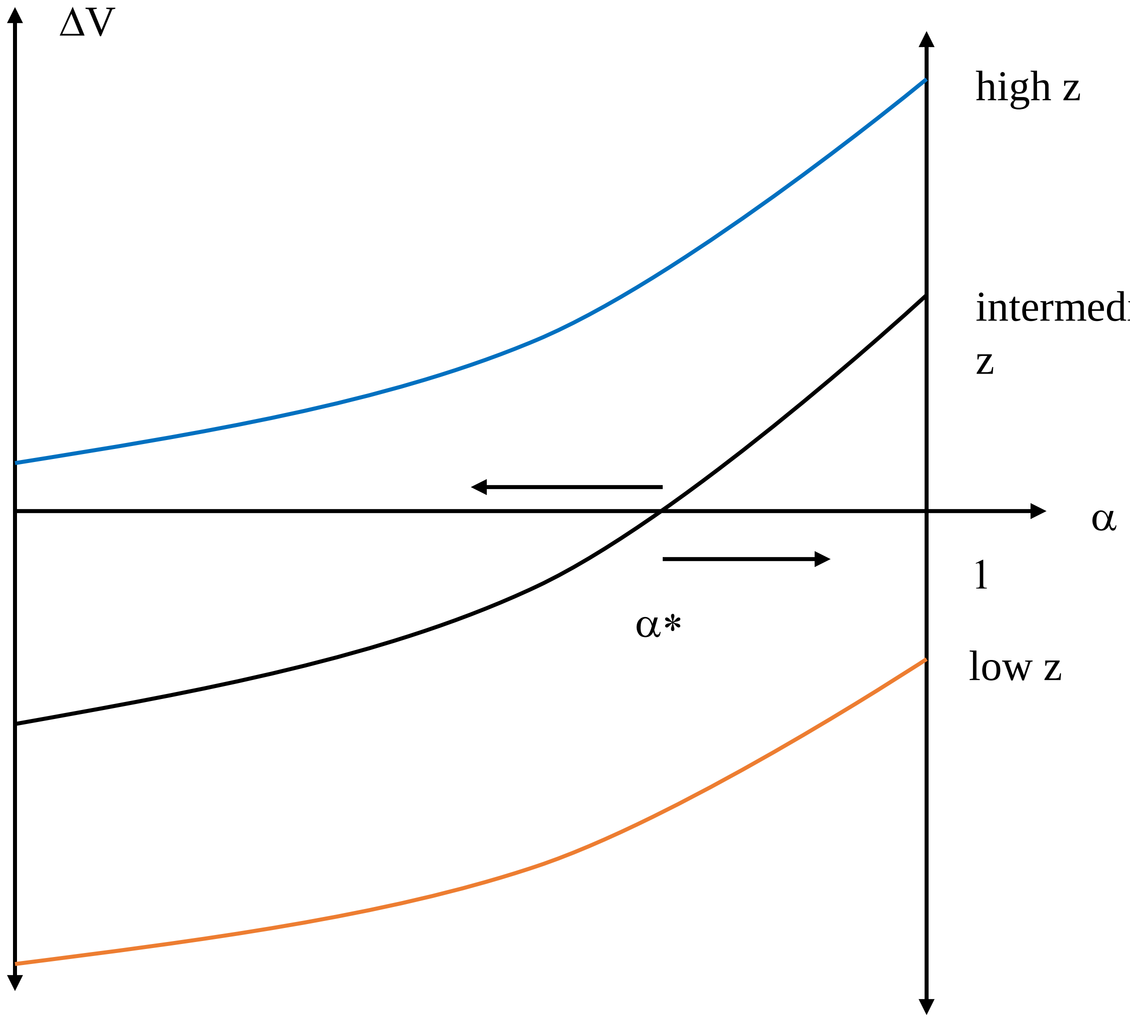

Figure 2. Nash equilibrium and stability. This figure illustrates the effect of z on the curve defining user's payoff difference between paying a fee and not paying a fee as a function of the fraction α of users paying the fee. The three curves describe the payoff differential for low, intermediate, and high levels of z. All three curves are upward-sloping, reflecting strategic complementarity. For intermediate levels of z there are three equilibria, zero, α^* and one. But, α^* is unstable as indicated by the arrows. If users expect an $\alpha < \alpha^*$, then the payoff to paying the fee is less than the payoff to not paying the fee. If $\alpha > \alpha^*$, the payoff to paying the fee is greater than the payoff to not paying the fee.

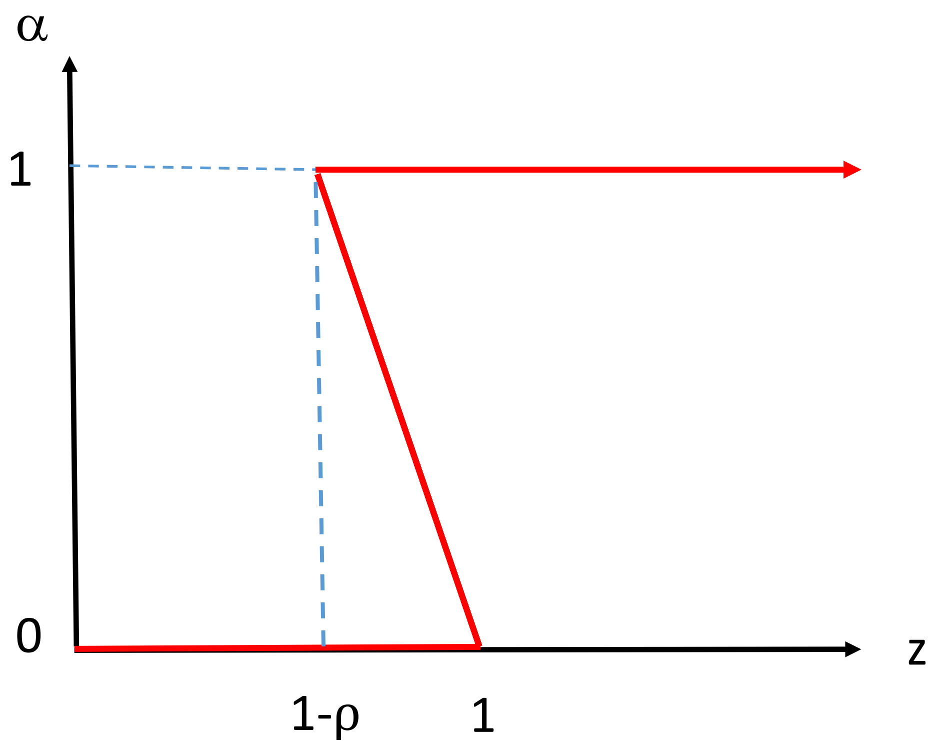

Figure 3. Nash equilibrium and fee-paying users. This figure relates the fraction α of users paying a fee as function of the parameter $z = {a\rho}/{pf(1-\rho)\lambda^* M^*}$, where z summarizes the costs and benefits to the user of paying the fee $({pf}/{a})$, the no-fee mean size of the mempool $({\rho}/{1-\rho})$, and the expected waiting time for a miner to arrive.

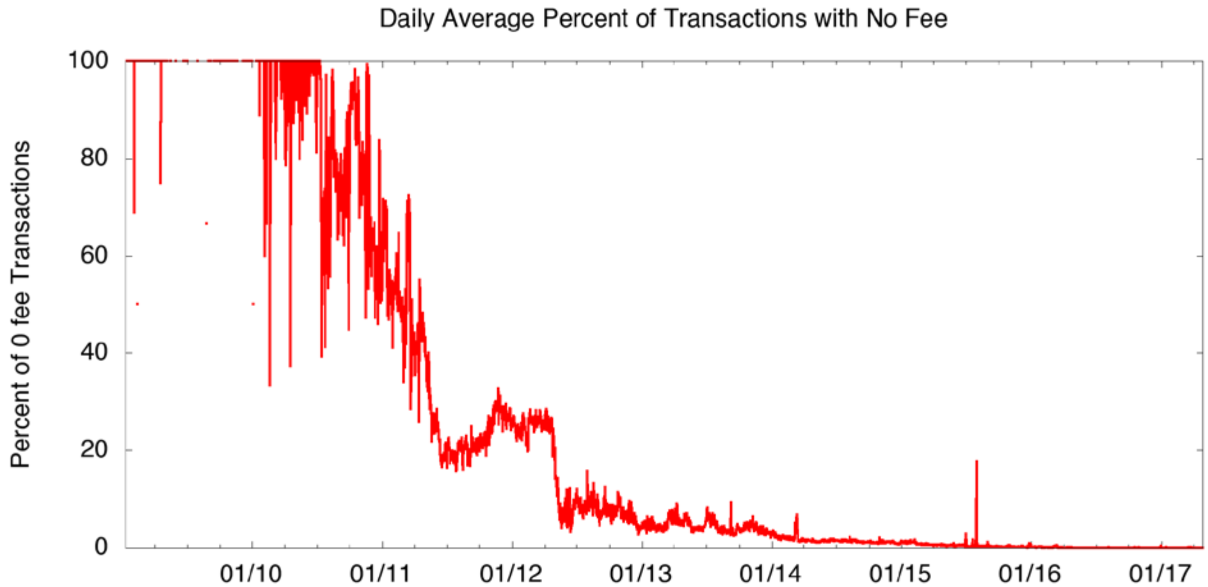

Figure 4. The percentage of transactions written to the blockchain without attached fees. This figure gives the percentage of transactions posted to the blockchain without attached fees. Data are drawn from our bitcoin node.

Panel A: Average daily transaction fee, 2011 - 2017

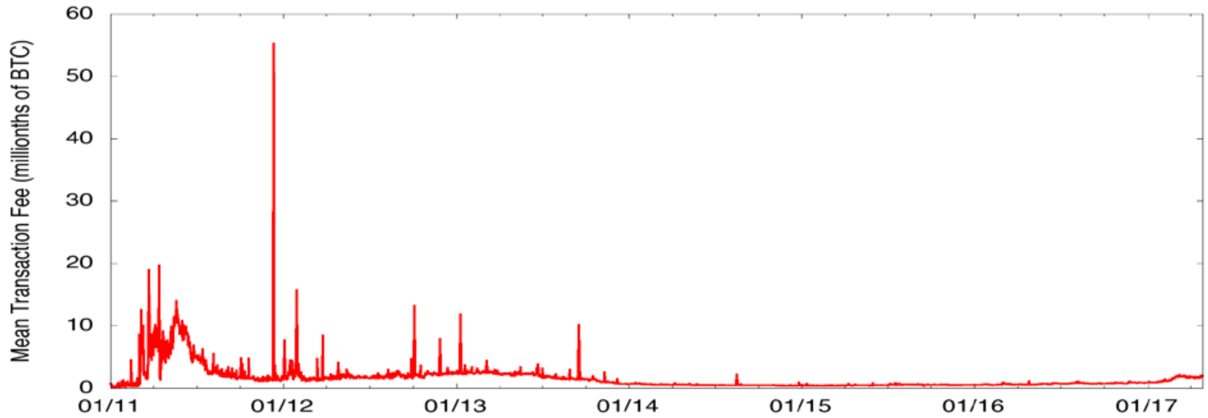

Panel B: Average daily transaction fee, April 2016- April 2017

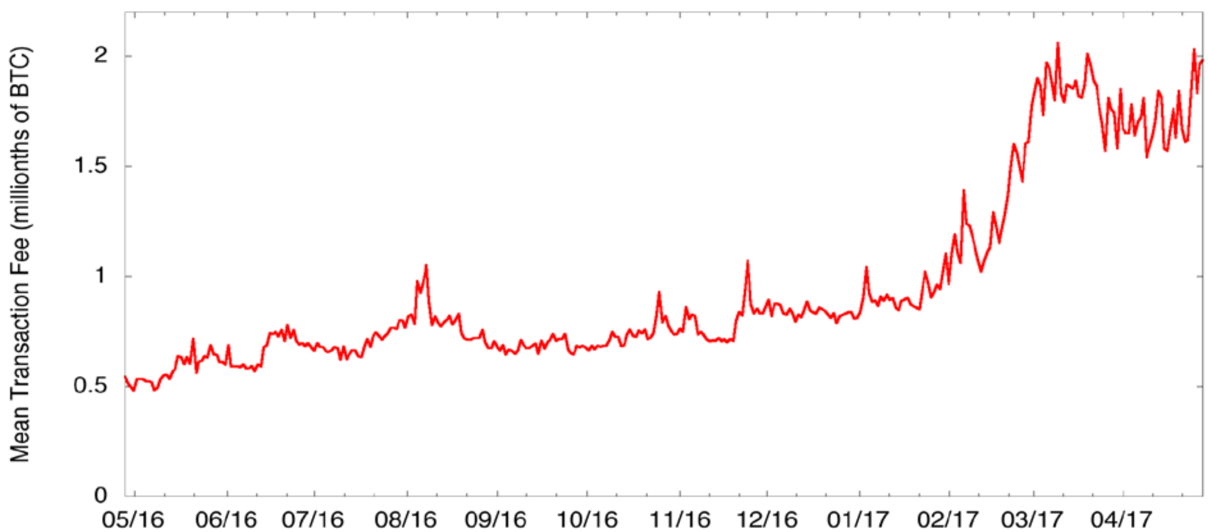

Figure 5. Transaction fees. This figure shows the average daily transaction fee per byte for transactions added to the bitcoin blockchain. Data are taken from our bitcoin blockchain node.

Table 1
Ordinary lest squares regressions: Ln(Percent zero-fee) on Ln(MWT)

This table gives regression results of the log of percent zero-fee transactions on the log of median waiting time [Ln(MWT)] with Ln(BTCprice) and block-reward as exogenous variables in Regression 1, the change in Ln(BTCprice) included in Regression 2 and, one-week lagged Ln(MWT) included in Regression 3. The sample period is December 3, 2011-April 23, 2016. There are 944 observations in Regression 1.

	Ln(Percent zero-fee)					
	(1)		(2)		(3)	
Variable	Coefficient	t-statistic	Coefficient	t-statistic	Coefficient	t-statistic
Ln(BTCprice)	-0.81	-19.1	-0.81	-11.98	-0.85	-19.5
Block-reward	0.03	5.46	0.03	5.51	0.03	6.40
Ln(MWT)	-1.88	-9.34	-1.88	-9.37	-1.47	-6.45
∆Ln(BTCprice)			0.89	0.84		
Lagged Ln(MWT)					-1.03	-4.62
Constant	7.35	13.22	7.32	13.16	8.67	14.6
R^2	0.57		0.57		0.57	

Table 2

Generalized method of moments regression: change in Ln(Percent Zero-Fee) on change in Ln(MWT) short sample

This table gives instrumental variables regression results (using GMM) of the change in the log of percent zero fee transactions on the endogenous variable change in the log of median waiting time [ΔLn(MWT)] with the change in the size of the mempool in hundred thousands as an instrument in Regression 1, lagged change in mempool size in hundred thousands included as an instrument in Regression 2 and, the log of the BTCprice in thousands included as an exogenous variable in Regression 3. The sample period is April 24, 2016-April 28, 2017. There are 351 observations in Regression 1.

			ΔLn(MWT)			
First stage	(1)		(2)		(3)	
Variable	Coefficient	t-statistic	Coefficient	t-statistic	Coefficient	t-statistic
ΔLn(BTCprice)	-0.39	-0.74	-0.38	-0.73	-0.38	-0.73
Blockreward	0.72	0.32	0.42	0.19	0.46	0.18
ΔMempool	1.24	4.49	1.31	4.73	1.24	4.47
Week-lagged ΔMempool			-7.45	-1.69		
Ln(BTCPrice)					-7.69	-0.13
Constant	-0.01	-0.27	-0.01	-0.13	0.04	0.11
Adj. R^2	0.12		0.14		0.12	
			ΔLn(Percent zero-fee)			
Second stage	(1)		(2)		(3)	
Variable	Coefficient	z-statistic	Coefficient	z-statistic	Coefficient	z-statistic
ΔLn(BTCprice)	-4.93	-2.93	-4.95	-2.89	-4.91	-2.85
Blockreward	-5.00	-0.40	-6.17	-0.49	-5.72	-0.37
E ΔLn(MWT)	-0.93	-2.77	-0.95	-2.89	-0.92	-2.77
Ln(BTCprice)					-2.30	-0.14
Constant	0.09	0.51	0.10	0.60	0.25	0.20
Adj. R^2	0		0		0	
C Score [H_0: ΔLn(MWT) exogenous]	4.79(p=0.03)		5.39 (p=0.02)		4.73 (p=0.03)	
Hansen J (H_0: restrictions valid)			0.12 (p=0.91)			

Table 3
Ordinary least squares regressions: mean daily transaction fee on median waiting time

This table gives regression results of the mean daily transaction fees per byte (in e^8 bitcoins) on median waiting time with BTCprice and blockreward as exogenous variables in Regression 1, the change in BTCprice included in Regression 2 and, one-week lagged MWT included in Regression 3. The sample period is December 3, 2011-April 23, 2016. There are 944 observations in Regression 1.

	Mean daily transaction fee					
	(1)		(2)		(3)	
Variable	Coefficient	t-statistic	Coefficient	t-statistic	Coefficient	t-statistic
BTCprice	-0.23	-11.3	-0.23	-11.5	-0.02	-11.2
Blockreward	2.68	2.32	2.68	2.32	2.53	2.32
MWT	1.17	0.61	1.11	0.58	0.79	0.48
∆BTCprice			0.20	1.31		
Lagged MWT					1.34	0.67
Constant	85.4	2.22	86.5	2.25	81.4	1.96
R^2	0.12		0.12		0.12	

Table 4

Generalized method of moments regression: change in mean transaction fee on change in median waiting time short sample

This table gives instrumental variables regression results (using GMM) of the change in the mean transaction fee in e^{-7} bitcoins on the endogenous variable change in the median waiting time with the change in the size of the mempool in ten thousands as an instrument and the change in the BTCprice and blockreward in thousands as exogenous variables in Regression 1, lagged change in mempool size in ten thousands included as an instrument in Regression 2 and, the BTCprice in thousands included in Regression 3. The sample period is April 24, 2016-April 28, 2017. There are 351 observations in Regression 1.

	∆MWT					
First stage	(1)		(2)		(3)	
Variable	Coefficient	t-statistic	Coefficient	t-statistic	Coefficient	t-statistic
∆BTCprice	-2.52	-0.74	-2.14	-0.27	-2.21	-0.28
Block reward	2.02	-0.32	-2.36	-0.07	-5.32	-0.23
∆Mempool	1.74	3.97	1.69	3.72	1.74	3.95
Lagged ∆Mempool			-0.26	-0.73		
BTCprice					-0.33	-0.34
Constant	-0.05	-0.11	0.04	0.08	0.32	0.34
Adj. R^2	0.14		0.13		0.14	

	∆TransFee					
Second stage	(1)		(2)		(3)	
Variable	Coefficient	z-statistic	Coefficient	z-statistic	Coefficient	z-statistic
∆BTCprice	3.01	1.33	2.89	1.27	2.79	1.25
Blockreward	-1.22	-0.11	-1.46	-0.12	3.94	0.37
E(∆MWT)	0.21	3.36	0.21	3.24	0.20	3.33
BTCprice					0.24	0.84
Constant	0.05	0.30	0.05	0.29	-0.02	-0.77
Adj. R^2	0		0		0	
C Score [H_0: Ln(MWT) exogenous]	17.44(p=0.00)		16.92 (p=0.00)		16.91 (p=0.00)	
Hansen J (H_0: restrictions valid)			0.22(p=0.64)			

Appendix.

Table A1

Summary statistics for data used in regressions

Variable	Mean	Standard deviation	Minimum	Maximum
Panel A: Sample period December 3, 2011-April 23, 2016				
Percent zero-fee	5.05	7.16	0.01	29.49
MWT	8.85	2.72	4.77	47.73
BTCprice	255.8	233.4	2.79	1147.3
ΔBTCprice	0.28	19.80	-208	198.09
Blockreward	30.63	10.45	25	50
Txfeemean	1.24	2.04	0.29	55.2
Panel B: Sample period April 24, 2016-April 28, 2017				
Percent zero-fee	0.021	0.024	0	0.364
ΔPercent zero-fee	-0.00025	0.0312	-0.3120	0.3450
MWT	10.91	3.77	6	29.25
ΔMWT	0.0031	3.819	-20.83	20.22
Avg mempool	10,825.2	13,950.1	693.25	9,2429.7
ΔAvg mempool	151.88	8,548.36	-37,885.8	3,8338.1
BTCprice	775.58	235.08	436.73	1,329.19
ΔBTCprice	2.51	26.02	-128.81	99.05
Blockreward	15.069	5.058	12.5	25
Txfeemean	0.94	0.40	0.48	2.06
ΔTxfeemean	0.004	0.08	-0.57	0.58

This table gives summary statistics for variables used in our analysis for the early period (2011-2016) in Panel A and the late period (2016-2017) in Panel B. Waiting times are expressed in minutes, BTCprice is expressed in dollars per bitcoin, blockreward is the number of bitcoins awarded to the winning miner, mempool sizes are in number of transactions waiting to be confirmed, and transaction fees are given in millionths of a bitcoin per byte. Transaction data are taken from our bitcoin blockchain node and represent on-chain transactions. Other data are from blockchain.info. Source: Bitcoin blockchain node data and blockchain.info

Table A2
Fractional logit regressions and FRACIV regressions: fraction of zero-fee transactions on MWT

This table gives regression results (using fractional logit) in Panel A and second-stage instrumental variables (IVs) regression results (using FRACIV) in Panel B of the fraction of zero-fee transactions on median waiting time (MWT) with BTCprice in thousands and blockreward in hundreds as exogenous variables and mempool size as an instrument in Panel B, BTCprice replaced by the change in BTCprice in Regression 2 and, lagged values of MWT included as an instrument in Regression 3. The sample period is December 3, 2011-April 23, 2016 in Panel A and April 24, 2016-April 28, 2017 in Panel B.

Panel A: Early Period						
OLS regressions	(1)		(2)		(3)	
Variable	Coefficient	z-statistic	Coefficient	z-statistic	Coefficient	z-statistic
BTCprice	-0.003	-7.81	-0.003	-7.94	-0.003	7.81
Blockreward	0.08	17.7	0.08	17.8	0.08	18.0
MWT	-0.06	-4.42	-0.06	-4.42	-0.45	-3.31
∆BTCprice			0.004	0.63		
Lagged MWT					-0.39	-3.19
Constant	-4.68	-25.1	-4.68	-25.5	-4.56	-23.9
Pseudo R^2	0.16		0.16		0.16	
Panel B: Late Period						
Second stage IV	(1)		(2)		(3)	
Variable	Coefficient	z-statistic	Coefficient	z-statistic	Coefficient	z-statistic
BTCprice	0.18	0.85			2.04	0.89
Blockreward	-1.15	-4.54	1.48	-4.39	-1.24-	-4.73
EMWT	-0.04	-2.21	-0.03	-4.16	-4.43	-2.14
∆BTCprice			0.27	0.48		
Constant	-3.03	-32.3	-2.94	-20.0	-3.02	-30.8
Wald test of exogeneity (H_0: MWT exogenous)	3.25 (p=0.07)		5.73(p=0.02)		3.128 (p=0.07)	

Table A3
Second-stage results of instrumental variables regressions on levels of Ln(percent zero-fee) and mean fee

This table give the second-stage results from an instrumental variable regression (using generalized method of moments) in levels for regressions with dependent variables: the logarithm of the percent of zero-fee transactions in Panel A and the mean fee in Panel B. The variables are as reported in Table 2 for Panel A and Table 4 for Panel B.

Panel A: Ln(Percent zero-fee)

Second stage	(1)		(2)	
Variable	Coefficient	t-statistic	Coefficient	t-statistic
Ln(BTCprice)	-0.74	-2.35	-0.71	-2.16
Blockreward	-0.09	-5.30	-0.09	-5.42
E(ln(MWT))	-0.58	-1.73	-0.62	-1.70
Constant	3.32	2.09	3.26	2.01
Wald test of exogeneity [H_0: ln(MWT) exogenous]	3.19 (p=0.07)		3.00 (p=0.08)	
Hansen's J (H_0: restrictions valid)			0.15 (p=0.70)	

Panel B: Mean fee

Second stage	(1)		(2)	
Variable	Coefficient	z-statistic	Coefficient	z-statistic
BTC price	1.36	13.5	1.34	12.4
Block reward	2.89	-4.54	2.81	1.81
EMWT	2.31	-2.21	2.55	2.80
Constant	-4.04	-7.50	-4.13	-7.49
Wald test of exogeneity (H_0: MWT exogenous)	10.5 (p=0)		11.0(p=0)	
Hansen's J (H_0: restrictions valid)			0.21 (p=0.64)	

References

Aune, R., Krellenstein, A., O'Hara, M., Slama, O. 2017. Footprints on a blockchain: trading and information leakage in distributed ledgers. Journal of Trading, 12 (2), page 5-13.

Bohme, R., Christin, N., Edelman, B., Moore T., 2015. Bitcoin: economics, technology, and governance. Journal of Economic Perspectives 29, 213-238.

Caffyn, G., 2017. What is the bitcoin block size debate and why does it matter? http://www.coindesk.com/what-is-the-bitcoin-block-size-debate-and-why-does-it-matter/.

Cong, L. W., He, Z. 2017., Blockchain disruption and smart contracts. Social Science Research Network. https://ssrn.com/abstract=2985764.

Cragg, J. G., Donald, S. G., 1993. Testing identifiability and specification in instrumental variable models. Econometric Theory 9 Social Science Research Network. 222–240.

Eyal, I., Gencer, A., Sirer, E. G., Renesse, R., 2016. Bitcoin-NG: a scalable blockchain protocol. Proceedings of the 13th USENIX Symposium on Networked Systems Design and Implementation (NSDI '16), Santa Clara, CA, March 16-18, 45-69.

Eyal, I., Sirer, E. G., 2014. Majority is not enough: bitcoin mining is vulnerable. In Financial Cryptography and Data Security. Christin, N., Safavi-Naini, R. (Eds.), Springer Heidelberg, Germany, pp. 436-454.

FINRA, 2017. Distributed ledger technology: implications of blockchain for the securities industry. Financial Industry Regulations Authorities, Washington, DC. http://finra.org/sites/defaultfiles/FINRA_Blockchain__Report.pdf.

Houy, N., 2014. The bitcoin mining game. Social Science Research Network. http://ssrn.com/abstract=2407834.

Gans, J. S., Halaburda, H. 2015. Some economics of private digital currency. In Goldfarb, A., Greenstein, S.M., Tucker, C. E. (Eds.), Economic Analysis of the Digital Economy. University of Chicago Press, Chicago, IL, Chapter 9.

Gandel, N., Halaburda, H., 2016. Can we predict the winner in a market with network effects? Competition in cryptocurrency market. Games 7 (3), 1-21.

Harvey, C., 2016, Cryptofinance. Social Science Research Network http://ssrn.com/abstract=2438299.

Hileman, G., Rausch, M., 2017. Global cryptocurrency benchmarking study. Cambridge University, Cambridge Center for Alternative Finance, Cambridge, UK.

Huberman, G., Leshno, J. D., Moallemi, C. 2017. Monopoly without a monopolist: an economic analysis of the bitcoin payment system. unpublished working paper. Columbia Business School, New York.

Kroll, J. A., Davey, I., Felton, E., 2013. The economics of bitcoin mining, or bitcoin in the presence of adversaries. In: Friedman, A. (Chair), The 12th Workshop on the Economy of Information, Security, Georgetown University, Washington, DC, June 11-12 Proceedings of WEIS, pp1-21.

Little, J. D. C., 1961. A proof for the queuing formula: L = λW. Operations Research 9(3), 383-387.

Malinova, K., Park, A., 2016. Market design with blockchain technology. Social Science Research Network. https://papers.ssrn.com/sol3/Delivery.cfm?abstractid=2785626.

Nakomoto, S., 2008, A peer-to-peer electronic cash system. White paper. https://bitcoin.org/bitcoin.

Raskin, M., Yermack, D., 2018. Digital currencies, decentralized ledgers and the future of central banking. In: Conti-Brown, P., Lastra, R., (Eds.), Research Handbook on Central Banking. Edward Elgar, Publishing, 474-486.

Raskin, M., Yermack, D., 2017. Corporate governance and blockchains. Review of Finance 21, 7-31.

Stock, J. H., Yogo, M., 2005. Testing for weak instruments in linear IV regression. In: Andrews, D.W. K., Stock, J. H. (Eds.), Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg. Cambridge University Press, New York, pp. 80–108.