

NBER WORKING PAPER SERIES

THE ICO PARADOX:
TRANSACTIONS COSTS, TOKEN VELOCITY, AND TOKEN VALUE

Richard Holden
Anup Malani

Working Paper 26265
<http://www.nber.org/papers/w26265>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2019

We thank Zhiguo He and Stacy Rosenbaum for helpful comments, and DJ Thornton for exceptional research assistance. Malani is an advisor to Perlin, a firm that offers a DAG-based blockchain ledger for various uses, including trade finance, but he does not presently have a financial stake in the firm. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Richard Holden and Anup Malani. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The ICO Paradox: Transactions Costs, Token Velocity, and Token Value
Richard Holden and Anup Malani
NBER Working Paper No. 26265
September 2019
JEL No. G12,G32,L1,L11

ABSTRACT

Blockchain technology offers firms a novel method of raising capital, via so-called Initial Coin Offerings (ICOs). In the most novel form of an ICO, a firm creates digital assets called “utility tokens” that are tracked on a blockchain-based ledger; requires that its product be purchased only with those tokens; and then raises capital by selling these tokens to investors prior to creating any saleable product. We point out a fundamental paradox with the use of ICOs involving utility tokens. Requiring the use of utility tokens to purchase the firm's product increases the cost of that product by an amount proportional to the cost of running the blockchain that tracks the utility token. In order to increase product revenue—and thus capital raised via an ICO—the firm will want to reduce these blockchain-operating costs. Doing so, however, increases the number of utility-token transactions that take place in any time interval, i.e., increases token velocity and thus the effective supply of tokens. By Fisher's equation, this lowers the dollar value of tokens and thus the amount investors are willing to pay for them. This paradox limits the value of utility-token ICOs. We discuss alternatives to and variations of utility tokens that can mitigate the conundrum.

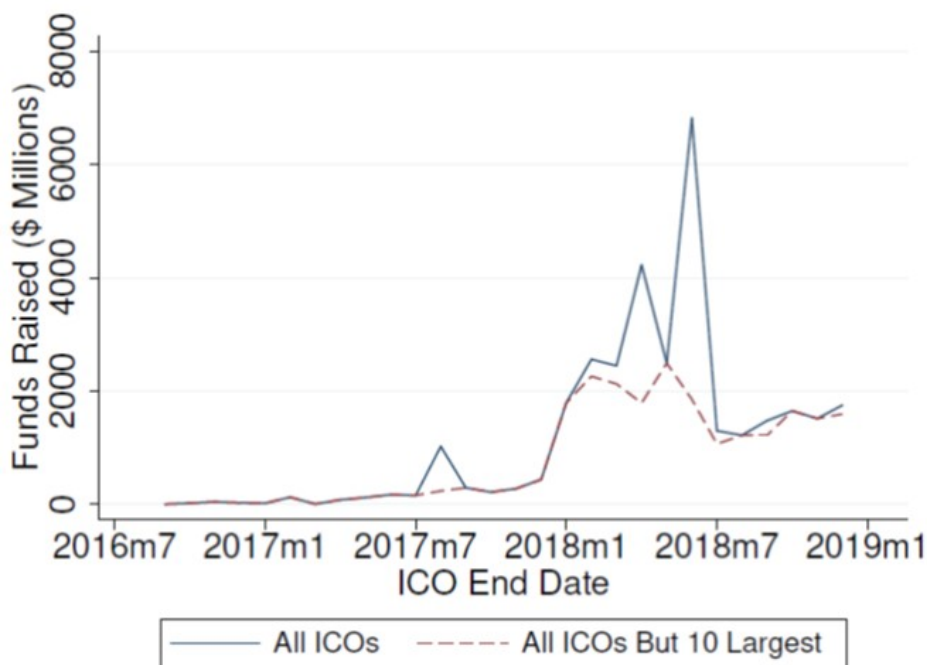
Richard Holden
University of New South Wales
Room 470B
Sydney, NSW, 2052, AUSTRALIA
Australia
richard.holden@unsw.edu.au

Anup Malani
University of Chicago Law School
1111 E. 60th Street
Chicago, IL 60637
and NBER
amalani@uchicago.edu

1 Introduction

In the last three years, there has been a boom in investment in start-up projects that use blockchain, a new technology for creating digital ledgers that do not rely on a centralized authority to manage the ledger. The initial surge of investment came through a new vehicle for raising capital called an Initial Coin Offering (ICO) (see Figure 1). Through December 2018, ICOs raised nearly \$33.4 billion. In 2017 alone, even before the number of ICOs peaked, ICOs raised \$6.5 billion, more than the total amount raised via venture capital investment in all internet projects. While there was a dip in late 2018, that dip was driven by the 10 largest ICOs. When those deals are excluded, the crash is less severe and we see a longer term upward trend.

Figure 1: Capital raised via ICOs.



Source: Davydiuk et al. (2019).

In an ICO, a firm raises capital for a project in three steps. First, it creates a digital asset, called a “token”, that gives owners the right to some value from the firm’s project. This value can either be the dividends of the firm (a “security” token) or rights to the firm’s output (“utility” token).¹ Second, the firm tracks the ownership of the token on a blockchain ledger. Third, the firm sells those tokens to investors. The proceeds of the sale are used to execute the project.

¹ICOs involving utility tokens are sometimes called utility coin offerings (UCOs).

An analogy to an amusement park may clarify the mechanics of an ICO that employs a utility token, the topic of this paper. Amusement parks frequently require that individuals who want to take rides in the park pay with tickets that can be purchased from the park’s sales office. The value of the ticket—and thus the dollar amount a consumer is willing to pay for each ticket—depends on how many tickets are required for a ride and the dollar value of that ride to consumers. Now, an amusement park typically sells tickets to visitors the day of their visit. However, it can, in practice, sell tickets days or months ahead of time.² If a park does so, it would obtain payment for rides before it had to furnish those rides. It could use this capital to invest in expanding or improving rides in the park ahead of delivery, i.e., for financing. Such a strategy is akin to the park selling a utility token that token owners can exchange for rides and using the proceeds of the sale to finance construction at the park.

Understanding the role that ICOs can play as a form of financing is important in determining not only the capital structure of firms, but also the degree to which firms are financially constrained, which affects their business strategy and ability to grow.

In this paper we consider the economics of ICOs that issue utility tokens. To do this, we model two ways in which a firm can generate revenues from a project. One is to charge a fee (or markup over costs) for the output of the project. The other is to sell utility tokens that can later be exchanged for output. In either case, the firm can use revenues to finance the underlying project. In the case of fees, the firm can sell rights to the revenue from fees to investors. With revenue from utility tokens, the firm can directly finance the project.

Setting the optimal fee is akin to the problem of setting a revenue maximizing tax or choosing the optimal markup on costs. By contrast, the optimal utility token policy requires choosing a ledger technology that imposes some friction but not too much. The friction here refers to the cost – in effort and time – of validating transactions on a blockchain ledger, i.e., the cost of what is called “mining” in blockchain parlance. This in turn is determined by the “consensus protocol” used by the ledger. The most popular candidates are proof of work and proof of stake.

Choosing the optimal level of friction involves what we call the ICO paradox. Setting too high a friction discourages sales. But setting to low a friction increases the number of utility token transactions that take place in any time interval, i.e., increases token velocity and thus the effective supply of tokens. This in turn lowers the dollar value of tokens and thus the amount investors are willing to pay for them.

To be more precise, according to Fisher’s equation for valuing currencies (Fisher, 1912), the dollar value of each utility token (p), the currency required to purchase the firm’s product,

²For example, Disney sells tickets to its park ahead of time. See <https://disneyworld.disney.go.com/faq/my-disney-experience/purchasing-tickets-in-advance/>.

is equal to the dollar value of all sales of the firm’s product in a given period of time (qX) divided by the product of token supply (M) and token velocity (V), defined as the number of transactions possible with each token in that period of time: $p = qX/MV$. A decrease in the cost of operating the blockchain increases sales of the firms product in each period (the numerator). But at the same time it also increases the number of possible transactions per unit of time (the denominator).

Our finding is reminiscent of the Grossman-Stiglitz paradox on the impossibility of informationally efficient markets when there are positive costs to information acquisition (Grossman and Stiglitz, 1976, 1980).³ In that environment, if markets are informationally efficient, those who acquire information cannot receive a payoff for so doing. Yet this deters them from paying the cost of information acquisition, thereby making markets informationally inefficient. Thus, an interior amount of informational efficiency arises in equilibrium. So, too, it is with utility tokens, where an interior amount of transaction costs/frictions must arise for the tokens to be worth issuing in the first place. At least, this must hold in any rational asset-pricing equilibrium. This inefficiency requirement puts an upper bound on the value of utility tokens themselves.

We are not the first to raise this paradox, which has previously been called the velocity problem. Vitalik Buterin (2017), the founder of Ethereum, mentioned it in an early blog post.⁴ However, we believe we are the first to connect the paradox to the economics of the underlying project that the firm seeks to finance, i.e., find the optimal behavior of the firm given the paradox.⁵ This is important as it helps trace the paradox back to basic economic phenomena such as the supply and demand functions for the underlying project’s output. That in turn can help promoters determine the precise relationship between miner compensation and velocity, i.e., choose the optimal consensus protocol for the blockchain that tracks the utility token. We believe we are also the first to compare an ICO to alternative methods of financing, e.g., charging a higher price and selling rights to profits. This can help firms determine whether an ICO is even the right method of financing their project.

After examining the optimal fee (or markup) and the optimal friction for utility tokens, we compare the revenue raised by each revenue model, and thus the optimal financing model. Surprisingly, a utility token model makes more sense the less technologically efficient is the

³See also Gibbons et al. (2012) for a related rational expectations equilibrium model for goods markets and organizations.

⁴See also Pfeffer (2017); Selkis (2018); Xu (2018). Samani (2017) expands on the use of Fisher’s equation to describe the velocity problem, and Locklin (2018) points out some errors in Samani’s formulas. Locklin’s criticisms do not apply to our model because we assume all consumers and producers are homogeneous.

⁵The paper closest in spirit to ours is Evans (2018), which provides a simulation involving the sale of tokens used to purchase a good. It is not accompanied by a general model about the sale or market in the underlying good.

underlying ledger technology. If the technology is has too little friction, the utility token price will be too low to generate revenues in excess of the simpler fee-based revenue model.⁶

Our final contribution is that we use our economic model of the underlying project and of the price of tokens to help inform technological discussions about how to improve blockchains to mitigate the velocity problem and thus increase the total amount that can be raised via ICOs. Samani (2018) attempts to do this, but, without a model of the value of the underlying project, not all of the recommended methods can avoid the ICO paradox. Moreover, the economic model reveals some new techniques for avoiding the paradox.

The discussion in this paper has applications to revenue models and financing outside the blockchain sector. Although utility token ICOs are typically used to fund underlying projects that themselves use blockchain technology, they can also be used to fund projects unrelated to blockchain.⁷ Indeed, as the amusement park example illustrates, firms can issue the equivalent of utility tokens without a blockchain at all, i.e., not even use blockchain for the financing.

In recent months, a number of companies have attempting what is called a security token offering (STO). This type of offering is driven by two forces. One is that regulatory determination by securities regulators such as the U.S. Securities and Exchange Commission, that blockchain-based tokens, whether technologists call them utility tokens or securities tokens, are what lawyers call securities and will be regulated as such. As a result, there is no advantage to distinguishing between utility and security tokens in marketing. The other force is that many blockchain-based projects that were financed privately seek to sell utility tokens not to raise revenue, but to allow their blockchain projects simply to function. To the extent that promoters want utility tokens sold via STO's to have a positive price, they will also confront the ICO paradox.⁸

While it is true that selling rights to the revenue from either a fee-based or a utility token model can be used to finance the project, this paper does not examine whether it is better to sell rights to these revenue streams or to have an Initial Public Offering (IPO) of equity. Either revenue model can support an IPO to raise capital to begin the project: one could sell net profits after raising revenue either via a transaction fee or via sale of utility tokens.

⁶Nothing stops consumers or the firm from using an existing digital currency blockchain, such as Bitcoin, to send or receive this fee. We ignore that option because it is orthogonal to the core economic question of whether the firm should issue utility tokens to raise revenue and capital.

⁷Examples include, Telegram (<https://bitcoinist.com/telegram-crypto-exchange-liquid-ico-gram/>) and Nagritech (<https://nagricoin.io/>).

⁸The fact that a utility token may not have high price does not imply that the output of the underlying project, even if the project employs blockchain, does not have a high price. The utility token is simply the unit of exchange that a customer must use to purchase the output, not the output itself. We will elaborate on this when we discuss the velocity problem.

The main difference between selling revenue streams and selling profits streams (i.e., equity) is that the latter shares costs with investors. For a direct comparison of ICOs and IPOs, see Catalini and Gans (2018). We do not model the firm’s management of the cost side of its business and therefore do not have anything to say about selling revenue versus profit rights.

Our paper contributes to a growing literature on ICOs. Much of that literature is empirical. It either describes the return to ICOs (Benedetti and Kostovetsky, 2018), as compared to IPOs (Garratt and van Oordt, 2019; Hu et al., 2018) or non-ICO cryptocurrencies (Dittmar and Wu, 2018), or examines correlates of more successful ICOs (Howell et al., 2018; Davydiuk et al., 2019; Momtaz, 2018). One strand of the theoretical literature explains the rationale for ICOs. For example, they solve the first-mover problem that can scuttle platforms that depend on network effects to succeed (Li and Mann, 2018), help internalize the total consumer surplus created by entrepreneurs (Lee and Parlour, 2018), or limit the moral hazard created by sale of equity (Bocks et al., 2019). Another strand examines how utility tokens for a platform are priced (Cong et al., 2018).

A third strand, which includes our paper, examines weaknesses of ICOs. Many of these papers focus on the incentive problems that ICOs create for entrepreneurs (Chod and Lyandres, 2018; Canidio, 2018; Garratt and van Oordt, 2019; Sockin and Xiong, 2018) and possible solutions to such incentive problems (Malinova and Park, 2018). The paper closest in spirit to ours is Catalini and Gans (2018), which provides practical advice, such as the value of entrepreneurs retaining tokens to limit moral hazard and the constraint that a fixed supply monetary policy imposes on future fundraising. Likewise, this paper examines how token velocity can limit the revenue that firms can earn selling utility tokens and what sort of technology can mitigate that effect.

Section 2 provides background on the mechanics of blockchain and ICOs. Section 3 presents an economic model of a market platform that can either raise revenue by charging a transaction fee or by undertaking an ICO. In the main text we specify that the project establishes a market for the trade in some good between consumers and other producers. In the appendix, we extend the model to simpler case where the project itself produces some good. Finally, Section 4 discusses different technological modifications to utility tokens and their potential to mitigate the velocity problem.

2 Background on blockchain, utility tokens and ICOs

2.1 Purpose of and trust in blockchain ledgers

Blockchain is a technology for creating digital ledgers for tracking facts that can be trusted even though there is no central authority maintaining the ledger. The main alternative is ledgers maintained by a central authority, such as a government, a bank, or some other firm. Blockchain is useful when the central authority maintaining the alternative ledger is not trustworthy or the centralized entity charges too high a price for maintaining the ledger. For example, citizens of a country run by kleptocrats may not trust the country's property registry because the kleptocrats may steal their land by manipulating the registry. Another use case is migrant laborers who have to pay very high fees to transfer money back to their home country because of frictions transferring funds across borders or because middle-men can simply steal funds without consequence. In each case, users may want a ledger that eliminates manipulation by (or cuts out completely) the existing intermediaries.

Blockchain utilizes computer algorithms that distribute the task of maintaining a ledger to a large number of entities in a manner that, in theory, gives those entities an incentive not to manipulate the ledger. For this reason, blockchain is sometimes called a distributed ledger technology. Blockchain allows a ledger to be maintained on a digital network with many independent nodes, like the internet. The entities that maintain the ledger are called "miners" on the network. In a well-functioning blockchain-based ledger, parties can trust the miners to honestly maintain the network because blockchain gives miners a reward or compensation for maintaining the network; it makes it costly for miners to manipulate the ledger to benefit themselves; and the rewards and costs are such that the return to maintaining the network is greater than the return to manipulating the network. While this incentive scheme may make blockchain trustworthy, it also drives the major cost to users from blockchain ledger: miner compensation.

A ledger, whether created via blockchain or not, can be used to track any fact about the world. One useful type of fact to track on a ledger is ownership, which can be thought of as an association between the identity of a legal person and the identity of an asset. The ledger tracks change in ownership by recording the fact that the identity of the asset is now associated with a new legal identity. Another useful fact that can be maintained on a ledger is the state of the world. This can facilitate contracting. For example, if A writes an insurance contract with B obligating A to make a payment of L to B if B incurs a medical expense, a ledger that tracks medical expenses facilitates execution of the contract. These sorts of facts can be maintained on many different types of ledgers, but if they are maintained on a blockchain ledger they enjoy the advantages blockchain has over other ledgers.

2.2 Mechanics of miner compensation

There are many ways to build a blockchain. To illustrate their commonalities, we describe the two most widely-used algorithms for ensuring the trustworthiness of a decentralized ledger: “Proof-of-Work” (PoW) and “Proof of Stake” (PoS). PoW is the algorithm at the heart of Bitcoin, the first and largest blockchain ledger ever deployed.⁹ Bitcoin is a ledger that tracks ownership of a digital asset, also called Bitcoin, that can be used as form of payment in the same way a fiat currency like the dollar is. PoS is the algorithm that will soon be used by Ethereum,¹⁰ a ledger that not only tracks ownership of a digital asset called Ether, but also enables initial coin offerings (ICOs), a topic we discuss later. The most important common feature of most blockchains, for the purpose of understanding the economics of ICOs, is that is that miner compensation is inversely related to the velocity with which transactions can be recorded on a blockchain.

In the Bitcoin blockchain, if A wants to transfer ownership of X Bitcoins to B, she announces that to the network in a manner that only she can do.¹¹ Miners on the Bitcoin network receive A’s announcement, along with announcements of a number of other transactions, and attempt to solve a computational problem. Importantly, the exact computation problem the miners must solve depends on the set of transactions to be recorded, i.e., it depends on exactly who sends what to whom. The first miner to solve the puzzle, i.e., provide proof-of-work, earns the right to record those announced transactions (called a block of transactions) onto the ledger (which is a chain of blocks and hence also called a blockchain). Because the PoW algorithm regulates which transactions are recorded on the blockchain, i.e., which transactions have generated consensus that they are true, it is often called a consensus protocol.

Solving the puzzle at the heart of the PoW protocol is costly because it requires electricity to run a problem-solving algorithm on a computer. The miner that solves the puzzle first not only gets to record the next block, but is also compensated with a reward, a combination of a voluntary fee paid by parties to record their transactions and new Bitcoins created (or minted) by the software running the network on which the ledger resides. The amount of the reward must be greater than the cost of attempting to solve the puzzle times the probability of being the first to do so.

A selfish miner may attempt to record transactions that did not take place to enrich

⁹As of February 25, 2019, the market capitalization of Bitcoin was roughly \$68 billion. The next closest blockchain token is Ethereum, valued at around \$14 billion. See <https://coinmarketcap.com/>.

¹⁰There are a number of other blockchains that already use PoS, including Peercoin, Decred, Neo, Navvcoin, Reddcoin, PivX, Tendermint, and Dash.

¹¹Blockchains use cryptographic private keys to create digital signatures that ensure that only A sent the message that she wants to transfer X to B (Narayanan et al., 2016).

themselves. To do so, however, they have to solve a different puzzle, one that contains not just the other announced transactions, but also the fake transactions they wish to have recorded, and to solve that different puzzle before anyone else solves a puzzle that reflects just the other announced transactions. To expect (from a statistical perspective) to be first, it has been shown that a miner would have to have more than half the computational power on the network maintaining the blockchain (Nakamoto, 2008). The cost of acquiring that much computation power on the Bitcoin network, for example, would cost over US\$1.4 billion and consume as much electricity as the country of Morocco does in a typical year (Moos, 2018). This is why it is thought to be hard to manipulate a blockchain ledger. This feature of blockchains, in turn, is why blockchain advocates argue that blockchains can be trusted by parties even though they do not rely on trusted central authority.¹²

PoW has been criticized because the amount of electricity required to power miners' computers is very large. It is estimated that maintaining the Bitcoin blockchain consumed as much electricity in 2018 as did all of Ireland that year. Putting aside the social cost of this energy consumption, this criticism can be treated as a complaint about the transaction costs of Bitcoin and PoW.¹³ Because miners must be compensated for their electricity consumption, high electricity consumption means high miner rewards. To address this problem, computer scientists have proposed alternative consensus protocols to generate trust in a decentralized ledger. The alternative that has received the most attention is PoS.

In this algorithm, instead of solving a puzzle for the right to record new transactions to the blockchain, miners bet or stake digital tokens issued by the network. On the Ethereum blockchain, which is in the midst of moving a PoW to a PoS algorithm, the token is called an Ether.¹⁴ A miner's chance of being selected by the network to record the next set of transactions is proportional to the amount of tokens she stakes. A miner who is selected to record the next block obtains, as a reward, both the transaction fees offered by the parties whose transactions are recorded and any newly tokens minted by the network for the purpose of compensating the winning miner.¹⁵

The way that PoS stops selfish mining is by requiring miners to stake their coins for a period of time and to relinquish their staked coins if they fail to win the right to stake the

¹²Often blockchains are called trustless rather than trustworthy. This is not because they are not trustworthy, but because they do not require trust in a central authority.

¹³There are consensus protocols that have miners solve socially useful computational problems so as to reduce the social costs of PoW-style mining. These include Proof of Exercise (Shoker, 2017) and Proof of Useful Work (Ball et al., 2017). However, these alternatives do not change the private costs of mining.

¹⁴See discussion of Casper at <https://github.com/ethereum/wiki/wiki/Proof-of-Stake-FAQs>.

¹⁵Here we provide a description of a simpler version of PoS than that which Ethereum employs. Ethereum actually employs a Byzantine-Fault-Tolerant version that resembles delegated POS (DPOS). See <https://github.com/ethereum/wiki/wiki/Proof-of-Stake-FAQs> for a more complete description of Ethereum's implementation of PoS.

next block or if they are found to have manipulated the blockchain. If this were not the case, then selfish miners would face no cost to manipulating the blockchain. Specifically, they could just bet a very large amount of tokens (which could be borrowed), almost guarantee themselves the right to record the next set of blocks, and record fake transactions (including amounts required to repay borrowed tokens). Even if they did not win the right to record the next block, they could not be punished because the network could not take away their staked tokens. To reduce this risk, PoS requires miners to forfeit their staked tokens if they do not win the right to record the next set or block of transactions. Moreover, the network requires the staked tokens to remain in an escrow account to which the network code has access and rights to drain if a miner is found to have manipulated the ledger. The larger are the stakes and the longer they are held in escrow, the greater the trust users have in the PoS-based blockchain.

The costs of PoS, like that of PoW, is the amount that miners have to be compensated to maintain the network. These are increasing in the amount that has to be staked to ensure a given probability of winning the rights to record the next block and how long staked tokens are held in escrow. The greater the amount to be staked increases the amount lost if the miner does not win the recording right lottery. Since a miner does not have access to her tokens when in escrow, the longer the escrow period, the larger the lost time value of money due to escrow.¹⁶ The greater is a miner's losses, the higher must be her compensation to induce her to participate as a miner on the blockchain.

2.3 Lower miner compensation associated with higher velocity

In this section we explain the basic empirical relationship that drives the ICO paradox: attempts to reduce miner compensation under either PoW or PoS increases transaction velocity on a blockchain ledger, i.e., the number of transactions that the digital ledger can record in a given interval of time. This relationship is easy to explain for PoW. Because harder computational puzzles require more computing power and thus electricity to solve, they also require greater compensation for miners. Because more difficult puzzles also require more time to solve, all else held constant, there is a positive correlation between the time required to validate and record a transaction and the required compensation for miners. Because greater validation time implies that fewer transactions can be recorded in a given interval of time, there exists a negative relationship between the velocity of transactions and miner compensation.

This negative relationship exists in POS, but it is not mediated by validation time. It is

¹⁶See <https://github.com/ethereum/wiki/wiki/Proof-of-Stake-FAQswhat-about-capital-lockup-costs>.

driven by the fact that PoS ensures trust, in part, by requiring miners to stake their tokens for a period of time. The larger is the amount of time miners must stake tokens, the more miners must be compensated for not being able to use those coins for some other investment. But the larger the amount of time staked tokens are in escrow, the lower is the number of tokens available for transactions and thus the lower is the potential number transactions for the average token.

The discussion above suggests that miner compensation is fixed purely by technology. Depending on the consensus protocol and the rule of the platform, it can also be driven by demand for transactions. If consumers can compete for access to miners to record transactions, then the choice of consensus protocol and the parameters chosen—difficulty level and escrow time—can be thought of as determining the supply curve for mining. For a given demand curve for transactions, the intersection of the supply and demand determines an equilibrium transaction cost. When the demand curve shifts out, however, miner compensation may rise if the protocol and platform permit it. This will be reflected in the transaction fee. When the platform changes the parameters of the consensus protocol, e.g., increasing difficulty level, that will shift the supply curve. If the supply curve shifts in, then again miner compensation will rise.

This equilibrium model of miner compensation does not change the basic economics. Technological parameters in the consensus protocol that shift out supply will tend to reduce miner compensation and thus increase the number of transactions per period conditional on a given demand curve for transactions, increasing velocity.

In the previous section, we offered one reason why a blockchain consensus protocol may not want to reduce miner compensation, namely that it would reduce trust in the network. Holding trust constant, the negative relationship between miner compensation and token velocity provides a second reason why the protocol may not want to reduce compensation: lower velocity undermines the use of ICOs with utility tokens to raise revenue. We next explain this method of raising revenue.

2.4 Utility tokens and Initial Coin Offerings

Setting up a blockchain entails development costs. Most significantly, software engineers must be hired to write the code to create and operate the blockchain. These start up costs can be financed in traditional ways, e.g., loans or sales of equity. To be clear, these costs are not unique to blockchain. Setting up a ledger using any method entails some development costs. In this section we discuss the use of blockchains ledgers to finance start-up costs. It is important to clarify, however, that this use of blockchains can be attached to any

project requiring start-up funds, whether that underlying project also uses blockchains for some other purpose or not. For example, it is possible for Filecoin to raise capital for its blockchain-based marketplace for data storage¹⁷ and for Agenus, a pharmaceutical company, to raise capital to fund its new biotech products,¹⁸ each using blockchain technology.

Recall that a blockchain can be used to track ownership of a digital asset. A firm seeking capital could therefore create a digital asset that represents rights to some value it will subsequently create once it is up and running, and sell those right to investors in return for cash. The purpose of the blockchain would be simply to track the ownership of those rights. Blockchain is not the only way to track ownership of those rights to the firm’s value creation. For example, the firm could list its stock on a stock exchange or sell other investment vehicles via investment banks. But whatever value blockchain might have over centralized ledgers is an advantage it would enjoy when it tracks ownership of rights to the firm’s value creation.

On a blockchain, rights to a firm’s value creation generally takes two forms, each called a token. The first is called a “security” token. This token represents rights to the firm’s profits or dividends, and is akin to equity. The token can be exchanged on a blockchain just as equity is traded on a stock exchange’s servers. The second type of token, which is the subject of this paper, is called a “utility” token. This token obtains value from the requirement—imposed by the firm—that goods or services created by its project can only be purchased using the token. Notable examples of utility tokens are tokens issued by Filecoin, which operates a marketplace for the exchange of hard drive space; Golem, which operates a platform for the exchange of cloud computing resources; and Basic Attention Token, which offers a marketplace for the exchange of access to eyeballs for marketers.¹⁹

This last requirement means that the value of a utility token is proportional to the total revenue—rather than the profit—generated by the project the firm seeks to finance. To understand why, consider a simple example. Suppose the firm’s project produces a widget, but the widget must be purchased using the firm’s utility token, which for fun we will call a “wicket”. Suppose the market price of a widget is \$3, that the firm is able to sell 10 widgets, and the firm creates only 20 wickets. Because all purchases of widgets from the firms must take place in wickets, the value of the 20 wickets must be equal to the dollar value of the 10 widgets people want to buy from the firm. This implies that the value of each wickets is $\$30/20$, i.e., equal to total firm revenues (in dollars) divided by the number of minted tokens. In our model, we will show that this simple characterization of wicket price is not exactly right because it does not indicate the time period within which the firm’s 10 sales

¹⁷See <https://filecoin.io/>.

¹⁸See <http://agenusbio.com/wp-content/uploads/2019/02/BEST-FINAL.pdf>.

¹⁹See <https://filecoin.io/>, <https://golem.network/>, and <https://basicattentiontoken.org/>.

take place. But, to the extent that the example shows that token value is proportional to revenue, it is correct.

Regardless of whether a firm issues a security token or a utility token, it can sell that token to investors even before it completes the underlying project. The capital raised from the sale can then be used to finance that underlying project. This process is called an Initial Coin Offering (ICO).²⁰ It is called an ICO to evoke the idea of an Initial Public Offering for equity security on a stock exchange. However, the similarities are limited. Whereas an IPO involves the sale of stock to the public, an ICO refers also to sales of utility tokens only to small group of private investors. That said, many of the largest ICOs have involved token sales to the public. While the largest IPOs raise billions of dollars, the largest ICOs are also significant. The top 10 highest grossing ICOs, listed in Figure 2 below, have collectively raised \$7.6 billion.

Figure 2: Top 10 ICOs by amount raised.

Promoter	Amount raised	ICO Dates	Project
EOS	\$4.1 billion	6/26/17 - 6/18/18	Smart Contracts
Telegram	\$1.7 billion	01/18 - 02/18	Encrypted Messaging & Blockchain Ecosystem
Dragon	\$320 million	02/15/18 - 03/15/2018	Decentralized Currency for Casinos
Huobi	\$300 million	01/24/18 - 02/28/18	Cryptocurrency Exchange
Hdac	\$258 million	11/27/17 - 12/22/17	IoT Contract & Payment Platform
Filecoin	\$257 million	08/10/17 - 09/10/17	Decentralized Cloud Storage
Tezos	\$232 million	07/01/17 - 07/14/17	Self-Amending Distributed Ledger
Sirin Labs	\$158 million	12/16/17 - 12/26/17	Open-Source Blockchain Smartphone
Bancor	\$153 million	12/6/17	Prediction Markets
The DAO	\$152 million	05/01/17 - 05/28/17	Decentralized VC

Source: <https://www.bitcoinmarketjournal.com/biggest-icos/>.


²⁰J.R. Willet is credited with the idea of an ICO. He conducted the very first ICO for Mastercoin in 2014. See <https://www.forbes.com/sites/laurashin/2017/09/21/heres-the-man-who-created-icos-and-this-is-the-new-token-hes-backing/#563df1731183>.

3 Model of financing a platform with an ICO

We model a firm that has an underlying project that provides a platform for the trade of a good. In the Appendix we offer an alternative model of a firm with a project that supplies a good rather than a platform for trade in that good. We focus on a platform here because it seems there are more examples of platforms producers that use blockchain and ICOs than goods suppliers that use that technology (see Figure 3). So our firm is neither a consumer nor producer of the good; instead it provides a marketplace in which consumers and producers can exchange the good with fewer transactions costs.

Figure 3: Capital raised via ICOs.

Total funds raised by categories, %



Source: <https://cointelegraph.com/news/from-2-9-billion-in-a-month-to-hundreds-dead-trends-of-the-rollercoaster-ico-market-in-18-months>.

We will first model the exchange of the underlying good and associated transaction costs before the firm creates its platform. Then we explore two revenue models for the firm's platform. One is to charge a transaction fee (i.e., a markup over costs) and the other is to issue and sell utility tokens. In doing that we will explain the role of initial coin offerings. In Sections 3.6 and 3.7, we model the price of tokens issued in the ICO and demonstrate the paradox that is the main insight of this paper. Finally, we compare the two revenue models and their implications for equilibrium price and quantity of the underlying good.

3.1 Trade in the underlying good

Consider a marketplace for a good with many consumers and producers. Without loss of generality, we assume that each consumer buys just 1 unit of the good. This will help us

relate quantity sold to token velocity. We assume that the supply side is competitive. (If the supply side for the good is monopolistic, the firm that provides the marketplace likely would be bought by the sole supplier of the good and this model would look like a good supplier that conducts an ICO. We leave that model for the appendix.)

We assume that, before our platform firm enters, trade in the pre-existing marketplace entails some friction $F > 0$. This friction operates like a per-unit sales tax of F would, except that proceeds are burned. This is our model of transactions costs. Let $(q_s(X^*(F, t)), q_d(X^*(F, t)), X^*(F, t))$ be equilibrium price producers receive, the equilibrium price consumers pay, and the amount consumers buy, respectively, in equilibrium in that marketplace at time t . (Subscripts d and s indicate inverse demand and supply functions, respectively.) Because trade in the underlying good is competitive and F operates like a tax, $q_d(X^*(F, t)) = q_s(X^*(F, t)) + F$. Because it is not essential to our basic economic insight, we assume

Assumption 3.1. Supply and demand curves are constant over time.

This allows us to suppress the dependence of prices and quantity on t .

Example 1. Let us explore an example with linear demand and supply to be able to relate equilibrium values more transparently to friction. Suppose that the demand curve is $X_d = a_0 - a_1 q_d$ and the supply curve is $X_s = b_0 + b_1 q_s$ for all t . Equilibrium price and quantity with friction F is found by setting $a_0 - a_1(q + F) = b_0 + b_1 q$. This yields an equilibrium producer price and quantity of

$$q_s(X^*(F)) = \frac{a_0 - b_0 - a_1 F}{a_1 + b_1} \quad X^*(F) = \frac{a_1 b_0 + a_0 b_1 - a_1 b_1 F}{a_1 + b_1}. \quad (1)$$

Consumers face a price of $q_d(X^*(F)) = q_s(X^*(F)) + F$. The elasticity of equilibrium quantity with respect to friction is

$$\eta_{X^*F} = \frac{a_1 b_1 F}{a_1 b_0 + a_0 b_1 - a_1 b_1 F} \quad (2)$$

which is increasing in frictions. ■

3.2 The firm offering a market platform

Assume there is a monopolist firm²¹ that has a technology which can reduce frictions in the exchange of the underlying good. We assume the technology can lower frictions to $f < F$.

To implement the technology, we assume the firm needs to pay a fixed cost of I and that marginal cost of operating the marketplace is 0. Our assumption about marginal cost can be relaxed, again without affecting the qualitative results. This fixed cost will need to be financed, a topic we address a bit later.

Importantly, we have not specified whether the platform technology, i.e., the underlying project that needs to be financed, employs blockchain or something else. It could be either because the revenue model and the technology the model employs can be different than the friction-reducing platform technology. The only substantive assumption we make about the platform technology is that:

Assumption 3.2. The cost of implementing that underlying project is the same regardless of what revenue model the firm chooses.

3.3 A transaction fee revenue model

There are two ways this platform-supplying firm can earn revenue from its innovation. One way is to charge a fee, k , on each unit traded.²² The other is to issue utility tokens and sell them to investors. Here we consider the fee model. We assume the fee is constant over time to simplify the analysis.

A transaction fee must be less than or equal to $F - f$, otherwise consumers and producers will not use the firm's platform. The fee operates like a tax on transactions that increases the price consumers pay and decreases equilibrium quantity to $q_d(X^*(f + k, t))$ and $X^*(f + k, t)$, respectively, where $q_d(X^*(F, t)) \geq q_d(X^*(f + k, t)) > q_d(X^*(f, t))$ and $X^*(F, t) \leq X^*(f + k, t) < X^*(f, t)$. The incidence of the fee and the reduction in trade depends on the elasticities of the demand and supply curve, as per usual tax incidence equations. The present value of total revenue is

$$k \int_0^{\infty} e^{-\delta t} X(f + k, t) dt,$$

²¹If the platform market were competitive, the analysis would look like that for the competitive supplier, which we discuss briefly in the appendix. Applying that discussion to platforms would reveal that, because platforms are competitive, the transaction fee they can charge would be driven by the market price of the platforms rather than the sort of optimization in below in Section 3.3.

²²Alternatively, one could model the firm as choosing a fee that is a percentage π of the good price, in which case, e.g., $k = \pi q_s(f)$. We believe our simpler model yields qualitatively the same insight into revenue generation and financing.

where δ is the interest rate.

Going forward, we will focus on the special case where demand is stable over time, i.e., $X(z, t) = X(z)$ for all t . Allowing demand to shift over time complicates the analysis in a manner that obscures the basic economics of raising revenue via transaction fees or sale of utility tokens. In this special case, the present value of total revenues from fees is

$$\left(\frac{k}{\delta}\right)X(f + k).$$

Recall that the firm needs to finance I . If the firm charges a fee, it can sell—in advance—rights to some of the stream of income from that fee to investors in return for capital to cover I . Let ϕ^{TF} be the fraction of transaction fee revenue sold to investors. The financing constraint is

$$\phi^{TF} \left(\frac{k}{\delta}\right)X(f + k) \geq I. \quad (3)$$

This constraint will always just bind as the firm has no reason to raise more capital than is required to launch the platform.

The firm's objective is to maximize profit from the fee subject to the constraint that it raises enough capital to fund implementation of the platform and does not drive away market participants. Assuming (3) just binds, the firm's problem can be written

$$\max_k (1 - \phi^{TF}) \left(\frac{k}{\delta}\right)X(f + k) \quad (4)$$

subject to

$$f + k \leq F. \quad (5)$$

Assuming there is an internal solution, the optimal fee satisfies the condition that the elasticity of equilibrium quantity (not the demand curve) with respect to the fee, η_{x^*k} is equal to one:

$$X^*(f + k^*) = -k^* \frac{\partial X^*(f + k^*)}{\partial k^*} \leftrightarrow \eta_{X^*k}(f + k^*) = 1. \quad (6)$$

Intuitively, raising the fee by 1% increases revenue per transaction by 1%, but also decreases the number of transactions. The optimal fee balances those effects, i.e., sets the elasticity of the transaction effect equal to 1.

If, at the k^* that ensures equilibrium quantity elasticity (with respect to k) is one, the

firm cannot raise enough revenue to cover I , even after selling the rights to all fees (i.e., $\phi^{TF} = 1$), then the transaction fee model is not a viable revenue model because it cannot finance the required start-up investment. If, at the k^* that satisfies the condition above, the market participation constraint (5) binds, we have $k^* = F - f$.

Example 2. Suppose that demand and supply are linear as in Example 1. Using the market clearing conditions from that example, the first-order condition for the firm’s problem, assuming the financing constraint binds but the market participation constraint does not, is

$$k^* = \frac{1}{2} \left(\frac{a_0}{a_1} + \frac{b_0}{b_1} - f \right). \quad (7)$$

The optimal fee decreases as platform frictions rise. ■

3.4 The utility token revenue model

The second way for the platform firm to earn revenue is to mint tokens, require all trade on the platform take place in tokens, and to sell some of those tokens. This type of token is called a “utility token.”²³ Because the tokens are required for trade, they may have value as a medium of exchange, a topic we consider in the next section. If tokens have value, the firm can allocate some of the minted tokens to itself and sell them for dollars. To focus on the nub of the problem, i.e., that choice of technology for the utility token ledger, we assume

Assumption 3.4.1. The firm mints only M tokens and that the firm chooses the technology for maintaining the blockchain ledger for utility tokens once and for all at the start of the project.

Assuming that trade in tokens implies frictions z in the trade of the underlying good,²⁴ the total value of tokens is $p(z, t)M$ at time t , where $p(z, t)$ is the price of a token at time t given trading frictions z .

As with the transaction fee model, we will assume that both frictions due to the requirement that consumers use utility tokens and consumer demand for the underlying good is stable. This will allow us to ignore questions such as when the firm should sell the tokens it has allocated to itself. These questions do not shed light on the core trade-offs between raising revenue with fees versus utility tokens. Going forward, therefore, we shall suppress the dependence of price on time, i.e., we shall write $p(z, t) = p(z)$.

²³This can be done on the firm’s own blockchain or on, e.g., Ethereum blockchain using an ERC-20 token. The technical implementation of the token is not (yet) important for the economics.

²⁴This includes f and frictions added by the trade in tokens.

If the token has positive value, the firm can also sell some of the tokens to investors to finance I . Specifically, it can sell a fraction ϕ^{UT} of the utility tokens at time $t = 0$ to investors such that $p(z)\phi^{UT}M \geq I$.

Note that, even if the firm requires a utility token be used to purchase the underlying good, the price of the good is set in dollars rather than tokens. Define $r(z)$ as the token denominated price of the good. If $p(z)$ is the dollar value of a token, then $r(z)p(z)$ is the dollar price of the good to consumers. If $r(z)p(z) > q_d(X^*(z))$, the price of the good with friction z , producers will profitably lower their prices. If $r(z)p(z) < q_d(X^*(z))$, consumers will bid up the price to $q_d(X^*(z))$. (If there were no frictions in this model, the result would obtain from a simple no-arbitrage condition.) Thus we conclude that

$$r(z) = q_d(X^*(z))/p(z).$$

The issuance of tokens is not without cost. Since utility tokens are issued on the blockchain, they entail a cost that depends on the consensus protocol. For example, with Proof-of-Work (PoW), miners will have to expend energy to validate trades with tokens; consumers and producers will have to compensate miners for that electricity expenditure. With Proof-of-Stake (PoS), miners may have to stake (and thus forego use of) tokens for some time to obtain the right to validate trades; they will have to be compensated for at least the time value of their staked money.

The key choice for a firm hoping to issue utility tokens is what consensus protocol, i.e., technology, to use. This can be parameterized as a choice over the variables s , which is the time for validating a transaction on the blockchain, and θ , which are features of the technology other than s . Choice of time affects miner compensation according to $d(s, \theta)$, where $\partial d(s, \theta)/\partial s \geq 0$ because more time to validate a block means either more energy under PoW or more time tokens are staked under PoS.²⁵ To allow us to invert $d(s, \theta)$ for expositional convenience, we shall assume that d is strictly monotonic in s . Technology features θ can change the compensation required for any given validation time.²⁶

Miner compensation functions like the transaction fee in the previous revenue model, except that the proceeds go to miners rather than the platform firm.²⁷ Miner compensation

²⁵For example, the firm could decrease s by increasing the block size. This allows more transactions per block. That means in any given time period required to validate the block, there are more transactions validated. Since miner compensation is keyed to time, the amount of miner compensation per transaction would also rise.

²⁶For example, if the monetary policy requires voluntary transaction fees and consumers and producers offer transaction fees for token trade validation that are proportional to token price, then an increase in the price of token will increase transaction fees, even though it has no effect on the time to validation.

²⁷It does not qualitatively make a difference if the platform actually used utility tokens to reduce frictions f and not just d . In that case we would replace f with d . Whether the firm uses utility tokens to reduce

therefore raises the friction to trade from f to $f + d(s, \theta)$. Just as greater fees increase consumer price and lowers quantity in equilibrium, so too does greater compensation.

Assumption 3.4.2. Without loss of generality, we assume $d(0, \theta) = 0$.

If miner costs are positive even when validation time is 0, we can simply fold that into frictions f when determining optimal technological choice under the utility token model. These miner costs would affect comparison of the transaction fee and utility token model, but are not relevant to the basic economic difference between the two models, so we ignore them.

Critical to our analysis is the inverse relationship between validation time per transaction s and the velocity of tokens, defined as the number of transactions of a given token in a period of time. Specifically,

$$V(s) = 1/s.$$

Velocity falls in the time required for validation. Moreover, as validation time goes to zero, velocity becomes infinite.²⁸

3.5 Using the equation of exchange to price utility tokens

In order to derive the firm's optimal token policy we need to price tokens. To do that we use Fisher's equation of exchange, which, in this context says that the equilibrium price of a token is that which equates the demand and supply of that token.

Under the utility token revenue model, the underlying good must be purchased with tokens. Thus the demand for tokens at a given point in time is, in dollar terms, equal to total consumer expenditure, $D = q_d(X^*(f + d))X^*(f + d)$.²⁹ (For readability, we have suppressed the dependence of d on s and θ .)

frictions f or not, we will show that the firm must introduce frictions in the trade of utility tokens to ensure a positive price for tokens. When the firm uses utility tokens to reduce frictions, that just means the firm has to raise frictions f above 0 if they use utility tokens to earn revenue.

²⁸This can be illustrated with a simple thought experiment. The case where s is zero is akin to the case where any given interval of time t can be divided into infinite sub-intervals and each token is traded once in each sub-interval. Suppose the equilibrium quantity consumed and price (for consumers) at time t is (X, q_d) . As we shall show, $r = (q_d + d(s))/p$ is the price per unit of good denominated in tokens. The total number of token trades required to purchase X units of the good in interval t is then $2 \cdot rX$. In the first $2 \cdot rX/M$ sub-intervals, let market participants use the available M tokens to complete all trades required for purchasing goods be completed. In each remaining sub-interval, let the producers, who hold all tokens after selling goods, exchange each token twice among themselves so that each producer holds the same number of tokens that they received from consumers. Since there are an infinite number of sub-intervals, there are an infinite number of trades per token at t , i.e., velocity is infinite.

²⁹We assume miners are paid in tokens as well, as is often the case.

The dollar-denominated supply of tokens S is a product of the total stock of tokens minted and available for circulation and the velocity of those tokens, $V(s)$. Velocity, i.e., transactions per period, ensures the units of supply are measured per period of time.

The equilibrium price that equates supply and demand for tokens in a period is:

$$D = q_d(X^*(f + d))X^*(f + d) = pMV(s) = S \leftrightarrow p(d) = \frac{q_d(X^*(f + d))X^*(f + d)}{MV(s)}. \quad (8)$$

Token price is increasing in the value of transactions on the platform in a given period, and decreasing in the effective supply of tokens on the market that period.

3.6 Frictionless utility tokens have zero value

Here we demonstrate a basic problem for utility tokens (and thereby hint at a paradox for ICOs in the next section). If the firm chooses a consensus protocol for utility tokens that implicitly sets transaction time s too low and thus velocity too high, the token price will be zero.

To see this, we make the dependency of miner compensation d on validation time s explicit and take the limit of token price p as validation time s goes to 0. From (8), this is

$$\lim_{s \rightarrow 0} p(d(s, \theta)) = \frac{1}{M} \cdot \frac{\lim_{s \rightarrow 0} [q_d(X^*(f + d(s, \theta)))X^*(f + d(s, \theta))]}{\lim_{s \rightarrow 0} V(s)}. \quad (9)$$

We already know the limit of revenue as frictions go to zero: $\lim_{s \rightarrow 0} q_d(X^*(f + d(s, \theta)))X^*(f + d(s, \theta)) = q_d(X^*(f))X^*(f) > 0$. The limit of velocity as friction got zero is infinity: $\lim_{s \rightarrow 0} V(s) = \lim_{s \rightarrow 0} 1/s = \infty$. With no friction, tokens can circulate an infinite number of times in any time interval. Together these results imply that $\lim_{s \rightarrow 0} p = 0$.

A natural question is what happens to the token-denominated price of a unit of the underlying good. Since $r = q_d^*/p$, the token denominated price would be undefined at $p = 0$. However, the limit is well-defined if we replace p using (8):

$$\lim_{s \rightarrow 0} r(d(s, \theta)) = \frac{\lim_{s \rightarrow 0} q_d^*(f + d(s, \theta))}{\lim_{s \rightarrow 0} p(d(s, \theta))} = \frac{\lim_{s \rightarrow 0} MV(s)}{\lim_{s \rightarrow 0} X^*(f + d(s, \theta))} = \infty. \quad (10)$$

That is, as the dollar price of tokens goes to zero, the token price of each unit of the underlying good goes to infinity.

3.7 Optimal friction with utility tokens and the ICO paradox

Assuming the financing constraint binds (i.e., $p(d(s, \theta))\phi^{UT}M = I$), the firm's problem is to choose a quantity of tokens M and validation time s so as to maximize funds raised

$$\max_{M,s} (1 - \phi^{UT})p(d(s, \theta))M \quad (11)$$

subject to the market participation constraint

$$f + d(s, \theta) \leq F \quad (12)$$

and the equation of exchange in (8). We do not worry about whether the firm maximizes token price now or at some future date t because (a) token price is driven by demand for the underlying good and we assumed (A.3.1) demand is fixed over time, and (b) we assumed (A3.4.1) that choice of consensus protocols are fixed at the start. Adding a dynamic choice of technology will not aid our intuition about the choice between revenue models.

Plugging in the equation of exchange for p reveals that M cancels out. The choice of M is arbitrary from an economic perspective: more tokens lower the per token price to exactly compensate. Because $d(s, \theta)$ is assumed strongly (rather than merely weakly) monotonic in s , we can write validation time as a function of compensation: $s(d; \theta)$. Now the objective function is

$$\max_d (1 - \phi^{UT}) \frac{q_d(X^*(f + d))X^*(f + d)}{V(s(d))}. \quad (13)$$

where we have suppressed the dependence of d on θ for expositional ease.

If the financing and participation constraints do not bind, one can take logs of the objective function before maximizing and write the first order condition as

$$\eta_{qdX^*}(X^*(f + d^*))\eta_{X^*d}(f + d^*) + \eta_{Vd}(d^*) = \eta_{X^*d}(f + d^*) \quad (14)$$

Intuitively, the marginal benefits of higher miner compensation are an increase in the price consumers pay and a reduction in velocity, both of which increase token price, both on the left-hand side. The marginal cost is lower equilibrium quantity purchased, on the right-hand side. Optimal miner compensation (equivalently transaction time) balances these two.

The last equation illustrates the basic ICO paradox. If a firm wants to raise revenue via an ICO, it needs to promise investors that their tokens will have positive value. It can generate some value by lowering friction from F to $f + d$. But if it tries to grow the market further by lowering d , i.e., reducing frictions from utility token exchange, then it

faces a headwind from token velocity. *If the elasticity of the revenue effect ($\eta_{qd}\eta_{X^*d} - \eta_{X^*d}$) is smaller than the elasticity of the latter velocity effect (η_{Vd}), then, paradoxically, improving the efficiency of the utility token will decrease token price even as it increases trades on the platform.* Whether this is the case is an empirical question that will vary by market, as it depends on the elasticity of equilibrium price and quantity for the underlying good, and the technology, as it also depends on the relationship between validation time, miner compensation and token velocity.

We can rearrange the expression above to compare it to (6), the first-order condition for the transaction fee model:

$$\eta_{X^*d}(f + d^*) = \frac{\eta_{Vd}(d^*)}{1 - \eta_{qdX^*}(X^*(f + d^*))} \quad (15)$$

where η_{ab} is the elasticity of a with respect to b and q_d is the demand function for the underlying good.

There are two implications. First, the firm will set miner compensation where demand elasticity is less than 1. Because the equilibrium quantity falls in the amount of transactions costs, the elasticity (which is customarily as a positive number) is on the left-hand side. For the right-hand side to correspondingly be positive, it must be that $\eta_{qd} < 1$. Intuitively, if the firm increased miner compensation where demand elasticity was high, the loss from lower quantity would swamp the increased revenue from higher price.

Second, the utility token problem is more complicated than the transaction fee model, which simply sets the elasticity of equilibrium quantity with respect to the transaction fee equal to 1. The firm has to have knowledge about consumer demand and technology.

Example 3. We derive optimal miner compensation assuming linear demand and supply. Equilibrium price and quantity is obtained from the formula's derived in Example 1 by replacing F with $f + d^*$. We assume that miner fees are linear in validation time: $d = \theta s$ where $\theta > 0$. Together this implies, total revenue under the utility-token model is

$$R_{UT} = p^*M = \frac{q_d^*X_d}{\theta/d^*} = \frac{d^*(a_0 - b_0 + b_1(f + d^*))(a_1b_0 + a_0b_1 - a_1b_1^2(f + d^*))}{(a_1 + b_1)^2\theta}.$$

Under our assumption about miner fees, $\eta_{Vd} = 1$. Using the first order condition for d^* under the utility token model, we find that d^* is the solution to the following quadratic equation:

$k_2(d^*)^2 + k_1d^* + k_0 = 0$, where

$$k_2 = a_1b_1^2$$

$$k_1 = 4a_1b_0b_1 - 2a_0a_1b_1 + a_0a_1b_1^2$$

$$k_0 = a_0a_1b_0 + a_0^2b_1 - a_1b_0^2 - a_0b_0b_1 + a_0a_1b_1f + 2a_1b_0b_1f + a_0b_1^2f - a_1b_1f^2.$$

■

3.8 Comparing revenue models

We can compare the two revenue models we model by comparing their objective functions. The ratio of utility token and transaction fee revenues at any given level of transaction cost $f + w$ is

$$\frac{R^{UT}}{R^{TF}}(w) = \frac{q_d(X^*(f+w))X^*(f+w)}{V(s(w))} \frac{1}{wX^*(f+w)} = \frac{q_d(X^*(f+w))}{w} \frac{1}{V(s(w))}. \quad (16)$$

We use two observations to leverage this ratio to discriminate between the revenue models. First, if we evaluate this ratio at the optimal transaction fee $w = k^*$ ($w = d^*$) and find that it is greater (less) than 1, then the utility token (transaction fee) model must generate more revenue than the maximal revenue under the transaction fee (utility token model) model. Second, the ratio of revenue from the utility token to the transaction fee model is decreasing in token velocity.

These observations lead to the following proposition, which suggests that the velocity of transactions determines which revenue model is optimal:

Proposition 1 (a) *The utility token model generates more revenue if $V(s(k^*)) < q_d(X^*(f + k^*))/k^*$.* (b) *The transaction fee model generates more revenue if $V(s(d^*)) > q_d(X^*(f + d^*))/d^*$.*

This implies that if velocity is sufficiently low, then the utility token model is preferable, and vice versa.

Unsurprisingly, whether the firm chooses to pursue the utility token model depends critically on the protocol consensus the firm chooses. The ICO firm will want to search for an available protocol that lowers η_{sd} , the elasticity of validation time on miner compensation, as that reduces the ICO paradox. To see why, think of the ICO firm's problem as maximizing token price via choice of miner compensation. The ICO paradox is that lowering compensation to increase equilibrium revenue also increases velocity, lowering price. The

latter effect is driven by the fact that lower compensation means less validation time. This drag on token price is mitigated, however, if the firm chooses a technology where validation time is not responsive to miner compensation, i.e., low η_{sd} .

Example 4. Let us examine whether a firm facing linear supply and demand is better off with a fee or an ICO. To answer this questions we again assume that miner fees are linear in validation time: $d = \theta s$ where $\theta > 0$. Plugging in revenue from each model into (16) yields

$$\frac{R^{UT}}{R^{TF}}(w) = \frac{a_0 - b_0 + b_1(f + w)}{\theta(a_1 + b_1)}.$$

Using the formula for k^* from Example 2 and the proposition above, we see that the utility token model will be superior if $X^*(f) < 2(a_0 - \theta a_1)$, i.e., if quantity when there is no fee is sufficiently small. In that case, the return to the transaction fee model, where revenue is the fee times quantity, is low because quantity is low with no fee and a fee only lowers quantity further.

We can use the formula for d^* from Example 3 and the proposition to determine when the transaction fee model is superior. The formulas are complicated and offer no insight about d^* levels. However, they do suggest that if

$$\theta > \frac{a_0 - b_0 + b_1(f + d^*(a_0, a_1, b_0, b_1, f))}{a_1 + b_1}$$

the transaction fee model is superior. We get this condition because d^* does not depend on θ . When θ is sufficiently large, while the price of tokens is higher, the quantity of sales is sufficiently suppressed that quantity of trade in the underlying good is very low, lowering the market cap of tokens. ■

4 Moving out the technological frontier

In this section we examine ways to reduce the cost of using utility tokens to raise revenues. Specifically, we explore techniques to reduce velocity conditional on miner compensation. We do this assuming either that the firm has decided to finance its platform using an ICO rather than transaction fees or that the firm might choose an ICO over those fees if it can increase the price of tokens. Of the three approaches we explore, only one—*burn and mint*—holds promise as a way to mitigate the paradox and increase the total amount raised via an ICO.

Before we proceed, it should be noted that the firm that does not utilize the same blockchain to reduce frictions to f on its platform as it does for its ICO could always

raise ICO token value by lowering f further. In other words, the ICO paradox pertains to attempts to raise revenue by lowering the frictions associated with utility tokens, not necessarily the platform generally.

4.1 Tokenize transaction fees

One possible approach is to pursue the transaction fee model, but to tokenize the fee. Specifically, the firm can require that the fee be paid not in dollars, but in newly created tokens.

Before we explain why this will not work, one should ask why the firm would want to tokenize the fee rather than collecting it directly in dollars. It is not obvious that the issues of trust or costs that typically motivate the use of blockchain based ledgers over centralized ones apply in the context of collecting fees. Moreover, tokenizing fees will actually increase overall fees because validating the payment of the fee via a blockchain would require compensating miners. Finally, tokenizing the transaction fee means the firm is selling a right to the fee rather than a right to the full price of the underlying good. This lowers the overall amount of funds that can be raised by selling tokens.

Even ignoring these issues, tokenizing fees may exacerbate the negative relationship between lower miner compensation and higher velocity. The price of a token that must be used to pay the transaction fee is:

$$p = \frac{fX^*(f + k + d(k))}{MV(d(k))} \quad (17)$$

where $d(k)$ is the compensation required to induce miners to validate fee payment k made in tokens. As before, lower miner compensation increases velocity. Before, the compensating benefit was that the numerator would rise because lower miner payment would increase both equilibrium good price and quantity. Now, however, the benefit is smaller since the numerator only includes equilibrium quantity, not price.

The benefit of lowering miner compensation would be greater if the transaction fee were a percentage of the equilibrium good price. But even then, this change would simply move the firm closer to the case where it tokenized payment for the good. It would not obviously improve the firm's position relative to that case.

4.2 Work token: requiring miners to hold utility tokens

A second approach is that the firm uses PoS and requires miners to stake tokens that are used for goods payments to earn the right to write the next block (Samani, 2018). A number of

existing projects, including Augur (a prediction market)³⁰, Filecoin, Keep (off-chain private data storage)³¹, Livepeer (video services marketplace)³², Truebit (off-chain computation)³³ and Gems (decentralized mechanical Turk)³⁴ use this work token tactic.

To a large extent, the existing model already captures this approach to lowering the ratio of transactions costs to velocity. As we pointed out in Section 2.2, PoS lowers transactions costs relative to PoW. Whether it offers a better ratio of miner costs to velocity is uncertain. The ratio depends on the duration of time that miners must stake tokens, which is an open variable. In any case, PoS does not escape the negative relationship between miner compensation and velocity because miner compensation and velocity are both driven (though in opposite directions) by the amount of time miners have to stake their tokens. The longer a miner must stake tokens, the larger the compensation she will demand due to the time value of money. Fanti et al. (2019) examine how this requirement to hold tokens affects token price using the equation of exchange.³⁵

A variant of standard work tokens is to create tokens, one to pay for the underlying good and the other that miners must stake to obtain the right to record the next block. The reward that the winning miner gets is paid in the first token, the one used to buy the underlying good. The ICO would focus on sale of the miner token. This modification looks a lot like tokenization of the transaction fees and does not obviously improve matters.

4.3 Burn and mint: platform holds tokens between transactions

A final and more promising approach to improving the value of utility tokens is to destroy utility tokens that are used to pay for the underlying good and re-mint them after some delay (Samani, 2018; Lau, 2018). Projects such as Factom³⁶ and Gnosis/Spunkchain (platform for adult entertainment videos)³⁷ use this burn and mint tactic.

The key value of burn and mint is that, by introducing a delay between the use of a token to purchase a good and the time it is reintroduced to the market functionally increases validation time, but does not require an increase in miner compensation because the network

³⁰See <https://www.augur.net/>.

³¹See <https://keep.network/>.

³²See <https://livepeer.org/>

³³See <https://truebit.io/>.

³⁴See <https://expand.org/>.

³⁵In theory, POW could use a similar tactic as work tokens in POS. The firm would pay miners in tokens but require that those tokens be held in escrow for some period of time, taking them out of circulation. However, the longer is this period, the more miners will have to be compensated due to the time value of money.

³⁶Factom pioneered burn and mint. See <https://www.factom.com/factom-blockchain/>.

³⁷See <https://www.spunkstream.com/>. See Lau (2018) for differences between Factom and the Gnosis/Spunkchain implementations of burn and mint.

holds the token from the market *after* the miner has been paid. In other words, it lowers the denominator of (8) without affecting the numerator.

5 Conclusion

This paper has examined the economic choice between a transaction-fee revenue model and a utility-token token model for a marketplace (which is equivalent to a markup versus utility token for a product supplier.) We characterized the optimal fee and revenue from the transaction fee model and showed that the value of the utility-token model depends on the friction in the underlying blockchain technology used to track tokens. In the process of doing so, we illustrated the ICO paradox or velocity problem: lowering the costs of operating the blockchain ledger too far actually reduces the price of individual tokens and thus the value of the utility-token model.

We compared the two models and demonstrated that the choice of revenue model depends on the the extent of frictions and the amount of trade on the marketplace. Specifically, too high a level of frictions makes the utility-token model inferior. This suggests that for the utility-token model to work, blockchain frictions can be neither too low nor too high. Moreover, if the price of products sold on the marketplace is too high, and thus the quantity of trade on the marketplace is too low, the transaction fee model is inferior.

Finally, we examined different types of blockchain consensus protocols that would allow a firm to better control frictions and thus optimize the market capitalization of all tokens. We showed that tokenizing transaction fees exacerbates the negative relationship between lower miner compensation and higher velocity, while the work token model is just a variant of PoS and does not solve the velocity problem. A more promising option is burn-and-mint, which allows the firm to directly control velocity without affecting miner compensation. This decoupling allows the fee consumer to stay low so as to support the quantity of trade, but the velocity to be low to support the token price.

There are limitations to our model. Notably we assumed that demand is constant, which allowed us to sidestep issues such as the timing of sale. We leave it to future work to relax these assumptions. In practice, firms that conduct ICO's with utility tokens face difficult decisions on whether and when to hold tokens or liquidate them. Since token prices can vary, these choices can dramatically affect the purchasing power of capital raised.

We also gloss over the connection between a revenue model and financing. As we mentioned, there are many ways to finance a revenue model. While utility-token models are naturally connected to ICOs, it is theoretically possible to issue such tokens, put them in a vehicle, and then sell equity in that vehicle. Likewise, one could take IOU's from a

transaction-fee model, create an asset-backed token, and then conduct an ICO for that token. There may also be important tax implications connected to this decision. We leave it to future economic and legal research to consider these topics.

The conclusions we draw, as well as this future research, have important implications for management strategy. They directly affect decisions about revenue generation. They might indirectly affect choices over blockchain technology, as those can affect revenue. Finally, the amount of revenue that firms receive will affect a litany of other decisions the firm makes, including financing, marketing and technology adoption.

References

- Ball, M., Rosen, A., Sabin, M., and Vasudevan, P. N. (2017). Proofs of useful work. *IACR Cryptology ePrint Archive*, 2017:203.
- Benedetti, H. and Kostovetsky, L. (2018). Digital tulips? returns to investors in initial coin offerings. *Returns to Investors in Initial Coin Offerings (May 20, 2018)*.
- Bocks, K., Haas, C., and Heyden, T. (2019). A theory on pre-ico venture capital involvement. *Available at SSRN 3402623*.
- Buterin, V. (2017). On medium-of-exchange token valuations.
- Canidio, A. (2018). Financial incentives for open source development: the case of blockchain.
- Catalini, C. and Gans, J. S. (2018). Initial coin offerings and the value of crypto tokens. Report, National Bureau of Economic Research.
- Chod, J. and Lyandres, E. (2018). A theory of icos: Diversification, agency, and information asymmetry. *Agency, and Information Asymmetry (July 18, 2018)*.
- Cong, L. W., He, Z., and Li, J. (2018). Decentralized mining in centralized pools.
- Davydiuk, T., Gupta, D., and Rosen, S. (2019). De-crypto-ing signals in initial coin offerings: Evidence of rational token retention.
- Dittmar, R. F. and Wu, D. A. (2018). Returns to initial coin offerings: An empirical examination. *Available at SSRN 3259182*.
- Evans, A. (2018). On value, velocity and monetary theory: A new approach to cryptoasset valuations. *Medium*.

- Fanti, G., Kogan, L., and Viswanath, P. (2019). Economics of proof-of-stake payment systems.
- Fisher, I. (1912). *The purchasing power of money: its' determination and relation to credit interest and crises*. The MacMillan Company.
- Garratt, R. and van Oordt, M. R. (2019). Entrepreneurial incentives and the role of initial coin offerings. *Available at SSRN*.
- Gibbons, R., Holden, R., and Powell, M. (2012). Organization and information: Firms governance choices in rational-expectations equilibrium. *The Quarterly Journal of Economics*, 127(4):1813–1841.
- Grossman, S. J. and Stiglitz, J. E. (1976). Information and competitive price systems. *The American Economic Review*, 66(2):246–253.
- Grossman, S. J. and Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *The American Economic Review*, 70(3):393–408.
- Howell, S., Niessner, M., and Yermack, D. (2018). Initial coin offerings: Financing growth with cryptocurrency token sales. Report, National Bureau of Economic Research.
- Hu, A., Parlour, C. A., and Rajan, U. (2018). Cryptocurrencies: Stylized facts on a new investible instrument. *Available at SSRN 3182113*.
- Lau, W. (2018). On the velocity problem for cryptoasset value. *Medium*, (February 23).
- Lee, J. and Parlour, C. A. (2018). Crowdfunding, initial coin offerings, and consumer surplus. *Available at SSRN 3300297*.
- Li, J. and Mann, W. (2018). Initial coin offering and platform building. *SSRN Electronic Journal*.
- Locklin, S. (2018). Token economics: Considering “token velocity”.
- Malinova, K. and Park, A. (2018). Tokenomics: when tokens beat equity. *Available at SSRN*.
- Momtaz, P. P. (2018). Initial coin offerings, asymmetric information, and loyal ceos. *Asymmetric Information, and Loyal CEOs (July 12, 2018)*.
- Moos, M. (2018). Analysis: Bitcoin costs \$1.4 billion to 51% attack, consumes as much electricity as morocco. *Cryptoslate*, (November 29).

- Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system.
- Narayanan, A., Bonneau, J., Felten, E., Miller, A., and Goldfeder, S. (2016). *Bitcoin and cryptocurrency technologies: a comprehensive introduction*. Princeton University Press.
- Pfeffer, J. (2017). An (institutional) investors take on cryptoassets.
- Samani, K. (2017). Understanding token velocity.
- Samani, K. (2018). New models for utility tokens. *Medium*, (February 13).
- Selkis, R. (2018). 95 crypto theses for 2018. *Medium*, (Jan. 2).
- Shoker, A. (2017). Sustainable blockchain through proof of exercise. In *2017 IEEE 16th International Symposium on Network Computing and Applications (NCA)*, pages 1–9. IEEE.
- Sockin, M. and Xiong, W. (2018). A model of cryptocurrencies. *Unpublished manuscript, Princeton University*.
- Xu, K. (2018). Not a velocity problem: My perspective on payment tokens. *Medium*, (April 8).

Appendix: Modeling the supplier of a good

In the main text we model a firm that is a monopoly provider of a platform for the sale of a good. In this appendix, we show that model is nearly isomorphic with a model of a monopolistic firm that directly supplies a good (rather than a platform for the sale of that good). We also examine the case of a competitive supplier, which also maps onto a model of a competitive platform.

A monopolistic supplier

We first consider a firm that is a monopolistic supplier of a good. We assume that there is some cost F to selling the good for cash or credit card payments. For example, cash may be hard to transport and keep secure and credit cards may charge a fee. Moreover, we assume that selling the good via a new technology, such as blockchain, reduces costs to $f < F$. Let $(q_s(X^*(F)), q_d(X^*(F)), X^*(F))$ now be the equilibrium price producers receive, the equilibrium price consumers pay, and the amount consumers buy, respectively, in a monopolistic market if the firm were not to use, e.g., blockchain, to sell its good.

If the firm were to use the new technology but use a transaction fee revenue model to make revenue and sell a portion of those fees to raise capital for investment, then its maximization problem would look just like that in (4) subject to the financing constraint (3) and the market participation constraint (5). Assuming no corner solution, the optimal fee k^* is given by (6). If the firm were to use utility tokens to raise revenue, its maximization problem would be the same as in (13) subject to the financing constraint $\phi^{UT}pS \geq I$. Assuming there is no corner solution, (15) defines the optimal miner compensation s^* . Given this, we can use the analysis in Section 3.8 to compare the merits of the two revenue models.

A competitive supplier

Now consider a firm that supplies a good in a competitive market. Define F and f as in the last section. Let $(q_s(X^*(F)), q_d(X^*(F)), X^*(F))$ now be the equilibrium price producers receive, the equilibrium price consumers pay, and the amount consumers buy, respectively, in a competitive market if the firm were not to use, e.g., blockchain, to sell its good.

We first examine the optimal fee in a transaction fee revenue model. We can think of $F - f$ as the firm's cost advantage relative to other competing firms. Because this market is competitive, the firm can charge a fee equal to $k^* = F - f$ without losing market share. In other words, we would get the same fee as the corner solution to the maximization problem in (4) subject to the constraints in (3) and (5) if the market participation constraint (5) was binding. Although a monopolistic supplier such as that in the previous section may sell the good at a higher price, it absorbs some of the burden of the selling fee, whether it is F or f , depending on the firm's technology. By contrast, the competitive firm passes that fee on entirely.

If the firm employs a utility token revenue model, its maximization problem would be the same as in (13) subject to the financing constraint $\phi^{UT}pS \geq I$. Assuming there is no corner solution, (15) defines the optimal miner compensation d^* . Recall that, whereas the transaction fee model imposes a fee $k^* = F - f$ on top of f so that the consumer faces the competitive price, the utility token model imposes a miner compensation cost d on top of f . If d is less than $F - f$, the firm can raise d without losing any sales. Moreover, because higher d is achieved by increasing the validation time s and because increasing s lowers $V(s)$ and thereby increases the price of tokens under the equation of exchange (8), the firm's optimal strategy is to set $d = d^*$ such that $d^* = F - f$. This implies a token price of $q_s(X^*(F))X^*(F)/V(s^*)$.

To compare the two revenue models in the competitive supplier case, we now need to compare revenue from the transaction fee model, $(F - f)X^*(F)$, to revenue from the utility

token sale model, $q_d(X^*(F))X^*(F)/V(s^*)$. The transaction fee model raises more revenue if $(F - f)/q_d(X^*(F)) > s^*$, i.e., the cost advantage or markup as a percentage of price is greater than the time required for each transaction (one over velocity). Let \tilde{s} be the validation time that ensures $F - f = q_d(X^*(F))/V(\tilde{s})$. If the optimal s^* in the utility token model is greater than \tilde{s} , then the utility token model is better because $V'(s) < 0$, so

$$q_d(X^*(F))/V(s^*) > q_d(X^*(F))/V(\tilde{s}) = F - f.$$

That is, if validation technology is sufficiently good that, to increase miner compensation to eat up the firm's cost advantage decreases velocity a lot, then the utility model is worthwhile.