

The Economics of Cryptocurrencies

– Bitcoin and Beyond*

Jonathan Chiu[†]

Bank of Canada

Victoria University of Wellington

Thorsten V. Koeppl[‡]

Queen's University

First version: March, 2017

This version: September, 2017

Abstract

How well can a cryptocurrency serve as a means of payment? We study the optimal design of cryptocurrencies and assess quantitatively how well such currencies can support bilateral trade. The challenge for cryptocurrencies is to overcome double-spending by relying on competition to update the blockchain (costly mining) and by delaying settlement. We estimate that the current Bitcoin scheme generates a large welfare loss of 1.4% of consumption. This welfare loss can be lowered substantially to 0.08% by adopting an optimal design that reduces mining and relies exclusively on money growth rather than transaction fees to finance mining rewards. We also point out that cryptocurrencies can potentially challenge retail payment systems provided scaling limitations can be addressed.

Keywords: Cryptocurrency, Blockchain, Bitcoin, Double Spending, Payment Systems

JEL Classification: E4, E5, L5

*The views expressed in this paper are not necessarily the views of the Bank of Canada. We thank the audiences at many seminars and conferences for their comments. This research was supported by SSHRC Insight Grant 435-2014-1416. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

[†]Victoria University of Wellington, Wellington 6012, New Zealand and Bank of Canada, Canada (e-mail: jchiu2013@gmail.com).

[‡]Queen's University, Department of Economics, Kingston, K7L 3N6, Canada (e-mail: thor@econ.queensu.ca).

1 Introduction

How well can a cryptocurrency serve as a means of payment? Since the creation of Bitcoin in 2009, many critics have denounced cryptocurrencies as fraud or outright bubbles. More nuanced opinions have argued that such currencies are only there to support payments for illegal activities or simply waste resources. Advocates point out, however, that – based on cryptographic principles to ensure security – these new currencies can support payments without the need to designate a third-party that controls the currency or payment instrument possibly for its own profit.¹

We take up this discussion and develop a general equilibrium model of a cryptocurrency that uses a blockchain as a record-keeping device for payments. Although Bitcoin in its current form has immense welfare costs, an optimally designed cryptocurrency can potentially support payments rather well. First, using Bitcoin transactions data, we show that the welfare cost of a cryptocurrency can be comparable to a cash system with moderate inflation. Second, using summary statistics for US debit card transactions, we find that a cryptocurrency can perform nearly as well as a low-value, retail payment system operating with very low fees.²

Economics research so far has provided little insight into the economic relevance of cryptocurrencies. Most existing models of cryptocurrencies are built by computer scientists who mainly focus on the feasibility and security of these systems. Crucial issues such as the incentives of participants to cheat and the endogenous nature of some key variables such as the real value of a cryptocurrency in exchange have been largely ignored. Such considerations, however, are pivotal for understanding the optimal design and, hence, the economic value of cryptocurrency as a means of payment.

Our focus is primarily on understanding how the design of a cryptocurrency influences the interactions among participants and their incentives to cheat. These incentives arise from a so-called “double-spending” problem. Cryptocurrencies are based on digital records and, thus, can be copied easily and costlessly which means that they can potentially be used several times in transactions

¹Some central banks have recently started to also explore the adoption of cryptocurrency and blockchain technologies for retail and large-value payments. Examples are the People’s Bank of China who aims to develop a nationwide digital currency based on blockchain technology; the Bank of Canada and the Monetary Authority of Singapore who are studying its usage for interbank payment systems; and the Deutsche Bundesbank who has developed a preliminary prototype for blockchain-based settlement of financial assets.

²This raises the issue that many cryptocurrencies currently cannot be scaled sufficiently to function as a true replacement of large retail payment networks. We abstract completely from such scalability issues that mainly arise from technological constraints.

(for a more detailed description see Section 2 below). We formalize this double-spending problem and show how this problem is being addressed by (i) a resource-intensive competition for updating the records of transaction – a process commonly referred to as mining – and (ii) by introducing confirmation lags for settling transactions in cryptocurrency. This implies that a cryptocurrency faces a trade-off between how fast transactions settle and a guarantee (or “finality”) for their settlement. Consequently, cryptocurrencies cannot achieve immediate and final settlement of transactions.

A strength of our analysis is that we take into account the costs of operating a cryptocurrency that prohibits double-spending. This allows us to quantitatively assess how efficient Bitcoin as a medium of exchange can be relative to existing means of payment. Calibrating our model to Bitcoin data, we find that from a social welfare perspective using Bitcoin is close to 500 times more costly than using traditional currency in a low inflation environment.

This is, however, a result of the inefficient design of Bitcoin as a cryptocurrency. Bitcoin uses both currency growth and transaction fees to generate rewards for mining. In its current form, the cryptocurrency reward structure is too generous so that too many resources are being used to rule out double-spending and making it a secure form of payment. We show that the optimal way of providing rewards for mining is exclusively via currency creation at a very low rate rather than by using transaction fees. The optimal design of Bitcoin would generate a welfare cost of only about 0.08% of consumption which is equivalent to a cash system with moderate inflation.

We also evaluate the efficiency of using a cryptocurrency system to support large-value and retail transactions. Using summary data for Fedwire and US Debit cards, we confirm that cryptocurrencies are a much better alternative for low value, high-volume transactions than for large value payments. This is intuitive, as double spending incentives increase with the size of transactions. Hence, more mining and longer confirmation lags (which are both costly) are required when supporting large-value payments in a cryptocurrency. Our exercise shows that cryptocurrency systems can potentially be a valid alternative to retail payment systems that operate at very low fees, as soon as limits on the scale of such systems can be resolved.

The economic literature on cryptocurrencies is very thin. We are not aware of any work that has formalized the design features of a cryptocurrency and that has studied its optimal design under the threat of double spending attacks. We model bilateral exchange based on money, we follow the recently promoted framework of Lagos and Wright (2005) and enrich it by modelling a mining competition to update the blockchain. For formalizing the blockchain itself, we rely on the

theoretical literature of payment systems as a record-keeping device.³

Our work is thus a first attempt to explicitly model the distinctive technological features of a cryptocurrency system which are a blockchain, mining and double-spending incentives within a quantitative economic model. We are also first to theoretically analyze the optimal design of a cryptocurrency and giving a quantitative answer to the efficiency properties of cryptocurrencies.

Some recent contribution have analyzed – from a qualitative perspective – whether Bitcoin can function as a real currency given its security features and lack of usage for making frequent payments (see for example Yermack (2013) and Böhme et al. (2015)). One question in this line of research is to empirically explain the valuation of cryptocurrencies. Gandal and Halaburda (2014) look at network effects associated with cryptocurrencies and investigate how such effects are reflected in their relative valuation. Glaser et al. (2014) look at how media coverage of Bitcoin drives part of the volatility in its valuation.⁴

Another area of research investigates how digital currencies can influence the way monetary policy is conducted. But none of this work can be applied to cryptocurrencies that are based on a blockchain and operate without a designated third-party to issue the currency. Agarwal and Kimball (2015) advocate here that the adoption of digital currencies can facilitate the implementation of a negative interest rate policy, while Rogoff (2016) suggests that phasing out paper currency can undercut undesirable tax evasion and criminal activities. Our findings complements this work as we establish some potential bounds on the costs that can be levied on people through central bank issued digital currency.⁵ Finally, Fernández-Villaverde and Sanches (2016) model cryptocurrencies as privately issued fiat currencies and analyze – in the tradition of the literature on the free banking era – whether competition among different currencies can achieve price stability and efficiency of exchange.

³See for example Koepl et al. (2008) and (2012).

⁴Models have been developed for studying other forms of electronic money technologies. For example, Berentsen (1998) studies digital money such as smart cards; Gans and Halaburda (2013) study platform-specific digital currencies such as Facebook Credits; Chiu and Wong (2015) review e-money technologies including PayPal and Octopus Card.

⁵See the recent discussion of Bordo and Levin (2017) on central bank issued digital currency. Also, Camera (2017) reviews the monetary literature and discusses the challenges of issuing and adopting electronic alternatives to cash.

2 Cryptocurrencies: A Brief Introduction

Our modern economy relies heavily on digital means of payments. Trade in the form of e-commerce for example necessitates the usage of digital tokens. In a digital currency system, the means of payment is simply a string of bits. This poses a problem, as these strings of bits as any other digital record can easily be copied and re-used for payment. Essentially, the digital token can be counterfeited by using it twice which is the so-called double-spending problem.

Traditionally, this problem has been overcome by relying on a trusted third-party who manages for a fee a centralized ledger and transfers balances by crediting and debiting buyers and sellers' accounts. This third-party is often the issuer of the digital currency itself, one prominent example being PayPal, and the value of the currency derives from the fact that users trust the third-party to prohibit double-spending (top of Figure 2.1).

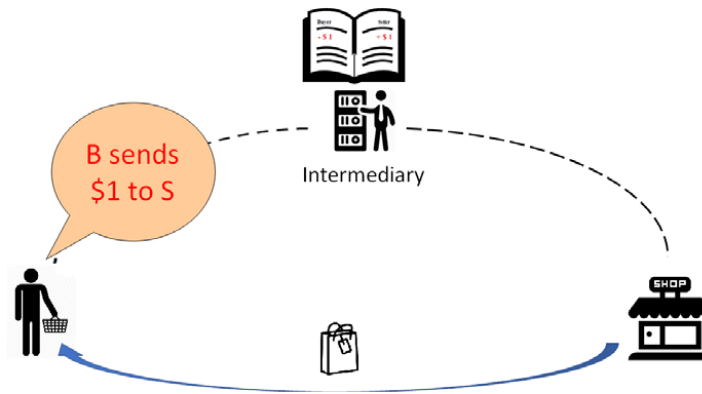
Cryptocurrencies such as Bitcoin go a step further and remove the need for a trusted third-party. Instead, they rely on a decentralized network of (possibly anonymous) validators to maintain and update copies of the ledger (bottom of Figure 2.1). This necessitates that consensus between the validators is maintained about the correct record of transactions so that the users can be sure to receive and keep ownership of balances. But such a consensus ultimately requires that (i) users do not double-spend the currency and (ii) that users can trust the validators to accurately update the ledger.

How do cryptocurrencies such as Bitcoin tackle these challenges? Trust in the currency is based on a *blockchain* which ensures the distributed verification, updating and storage of the record of transaction histories.⁶ This is done by forming a blockchain. A block is a set of transactions that have been conducted between the users of the cryptocurrency. A chain is created from these blocks containing the history of past transactions that allows one to create a ledger where one can publicly verify the amount of balances or currency a user owns. Hence, a blockchain is like a book containing the ledger of all past transactions with a block being a new page recording all the current transactions.

Figure 2.2 illustrates how the blockchain is updated. To ensure consensus, validators compete for

⁶In this sense it is different from traditional money which is merely a partial memory of past transactions as it only records the current distribution of balances and does not record how past transactions generate the current distribution.

Digital tokens with a trusted third party (e.g. PayPal)



Digital tokens without a trusted third party (e.g. Bitcoin)

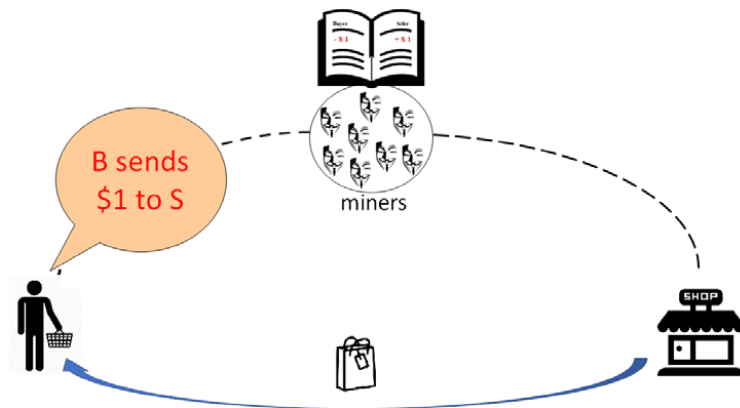


Figure 2.1: Digital Currency vs. Cryptocurrency

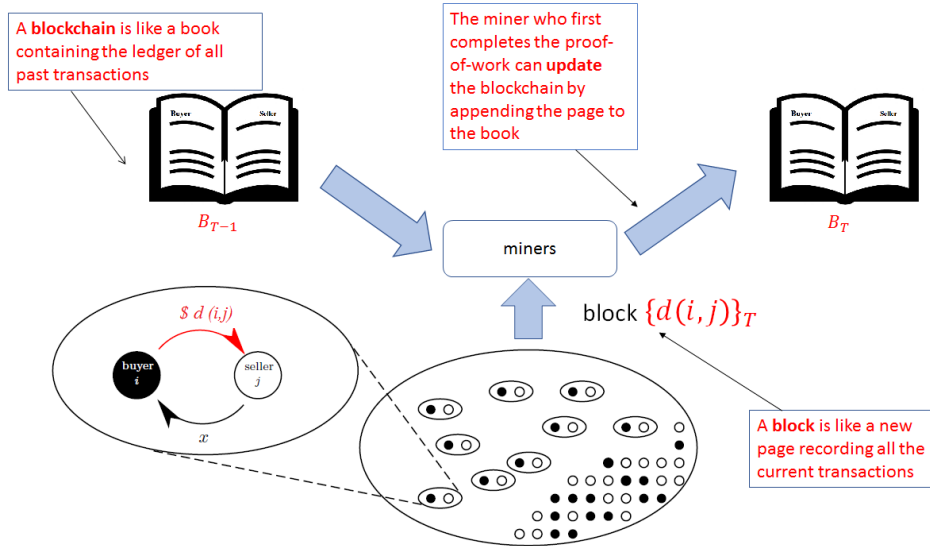


Figure 2.2: Blockchain Based Validation in a Cryptocurrency

the right to update the chain with a new block. This competition can take various forms. In Bitcoin, it happens through a process called *mining*. Miners (i.e. transaction validators) compete to solve a computationally costly problem which is called *proof-of-work* (PoW).⁷ The winner of this mining process has the right to update the chain with a new block. The consensus protocol prescribes then that the “longest” history will be accepted as the trusted public record.⁸ Since transaction validation and mining are costly, a reward structure is needed for mining to take place. In Bitcoin, for example such rewards are currently financed by the creation of new coins and transaction fees.⁹ The main concern for users when trusting a cryptocurrency is the double spending problem: after having conducted a transaction, a user attempts to convince the validators (and, hence, the general public if the blockchain is trusted) to accept an alternative history in which some payment was not conducted.¹⁰ If this attack succeeds, this user will keep both the balances and the product

⁷Other consensus protocols are being explored which we briefly discuss in the Appendix.

⁸In general, there is no explicit requirement to follow the consensus protocol in the sense that validators and users can trust an alternative history or blockchain that is not the longest one. Well-designed cryptocurrency, however, try to ensure that there are sufficient incentives to work with the longest history. For example, in Bitcoin the reward paid to a successful miner is contained in the new block itself. Should a different chain be adopted at a later stage, this reward will be obsolete. For a game-theoretic analysis of the incentives to build on the longest chain, see Biais et. al. (2017).

⁹Huberman et al. (2017) explore the reward structure of cryptocurrencies from the perspective of the mining game, but without modelling the double spending problem.

¹⁰While basic cryptography ensures that people cannot spend others’ balances, a miner can exclude from a block

or service he obtained while the counterparty will be left empty handed. Hence, the possibility of such double-spending can undermine the trust in the cryptocurrency.

A blockchain based on a PoW consensus protocol naturally deals with changing transaction history backwards. The blockchain has to be dynamically consistent in the sense that current transactions have to be linked to transactions in all previous blocks.¹¹ Consequently, if a person attempts to revoke a transaction in the past, he has to propose an alternative blockchain (with that particular transaction removed) and perform the PoW for each of the newly proposed block. Therefore, it is very costly to rewrite the history of transactions backwards if the part of the chain that needs to be replaced is long. Hence, the “older” transactions are, the more users can trust them.

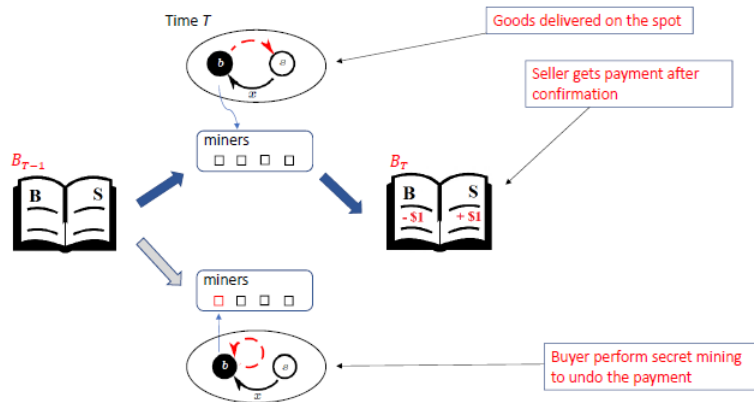
Unfortunately, a blockchain does not automatically protect a cryptocurrency against a double-spending attack that is forward-looking. Figure 2.3 considers a spot trade between a buyer and a seller involving a cryptocurrency. The buyer instructs the miners to transfer a payment to the seller while the seller simultaneously delivers the goods. Notice that the buyer can always secretly mine an alternative history (or submit to some miners a different history) in which the fund is not transferred. The final outcome of the transaction depends on which payment instruction is incorporated into the blockchain first. If the former payment instruction is incorporated, then the double-spending attempt fails. The seller receives the payment and the buyer gets the goods. If the latter is accepted instead, then the double-spending attempt succeeds. In this case, the buyer gets the goods without paying the seller.

Such a double spending attack can be discouraged by introducing a confirmation lag into the transactions. By waiting some blocks before completing the transaction (i.e., the seller delays the delivery of the goods), it becomes harder to alter transactions in a sequence of new blocks. Figure 2.4 illustrates how a confirmation lag of one block confirmation raises the secret mining burden of a double spender. The seller delivers the goods only after the payment is incorporated into the blockchain at least in one new block. Again, the buyer can secretly mine an alternative history in which the payment does not happen. How successful secret mining is depends on the mining competition and the length of the confirmation lag.

some transactions that have been initiated by other people. With positive transaction fees, a miner does not have an incentive to remove other people’s transactions and lose such fees.

¹¹For example, if someone transfers a balance d in block T , it must be the case that the person has received sufficient net flows from block 0 to block $T - 1$ so that the accumulated amount is at least d .

Double spending attempt fails



Double spending attempt succeeds

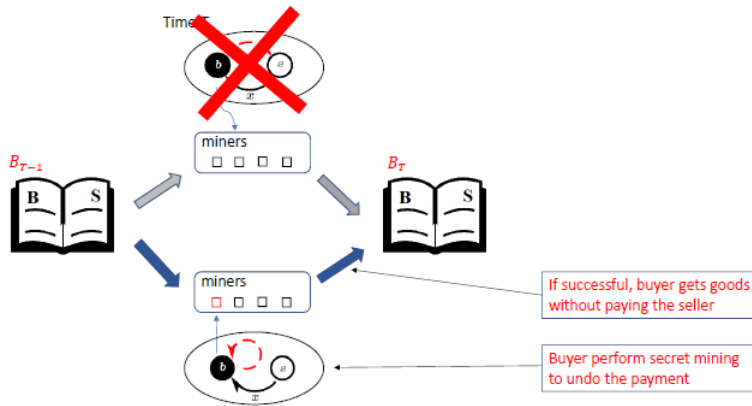


Figure 2.3: Double Spending Attack

Suppose the buyer successfully solves the PoW for the block containing this alternative history. Note that the buyer has an option whether to broadcast the secretly mined block immediately or withhold it for future mining. If he decides to broadcast the block immediately, the seller will not receive the payment, but he will also not deliver the goods as shown on the top of the figure. Hence, the double-spending attack is not successful for the buyer.

Alternatively, the buyer can temporarily withhold the solved block and continue to secretly mine another block (depicted on the bottom of the figure). Specifically, the buyer needs to allow other miners to confirm the original payment to the seller, so as to induce the seller to deliver the goods. At the same time, the buyer needs to secretly mine *two blocks in a row* for which the original transaction is removed.¹² If the buyer is successful in mining two blocks faster than other miners, he can announce an alternative blockchain after the goods are delivered. In this case, the buyer gets the goods without paying the seller. More generally, if the seller delivers the goods only after observing N confirmations of the payment, the buyer needs to solve blocks $N + 1$ consecutive times in order to double spend successfully.

To summarize, trust in a cryptocurrency system involves the interplay of three ideas: the security of the blockchain, the health of the mining ecosystem and the value of the currency.¹³ As shown in Figure 2.5, sufficient mining activities are required for ensuring the security of the blockchain, safeguarding it against attacks and dishonest behaviors. Moreover, only when users trust the security of system will the cryptocurrency be widely accepted and and traded at a high value. Finally, the value of the currency supports the reward scheme to incentivize miners to engage in costly mining activities.

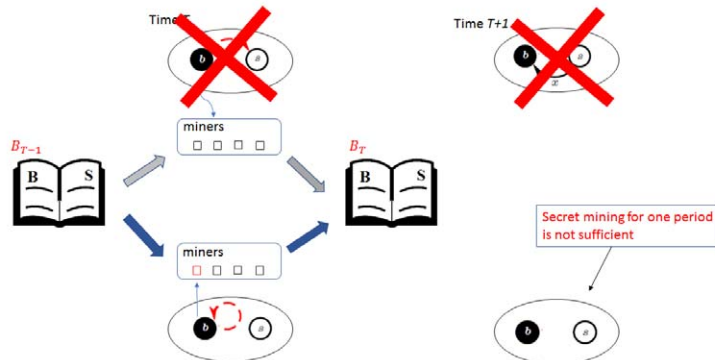
Our model will capture precisely this interdependence by explicitly looking at the joint determination of mining efforts, rewards and cryptocurrency value in general equilibrium. The main features of a cryptocurrency model in Section 3 are therefore given by

- (i) a consensus protocol: miners compete to update a blockchain with the probability of winning being proportional to the fraction of computational power owned by a miner

¹²Since the consensus protocol prescribes that the longest chain is accepted by the miners, the double spender needs to mine two blocks in order to create an alternative, longer blockchain superseding the existing one which has one block solved already.

¹³This has been pointed out already in the computer science literature (see for example Narayanan et al. (2016)), but without making the connection to an equilibrium incentive problem for double spending.

Double spending attempt fails (with one confirmation lag)



Double spending attempt succeeds (with one confirmation lag)

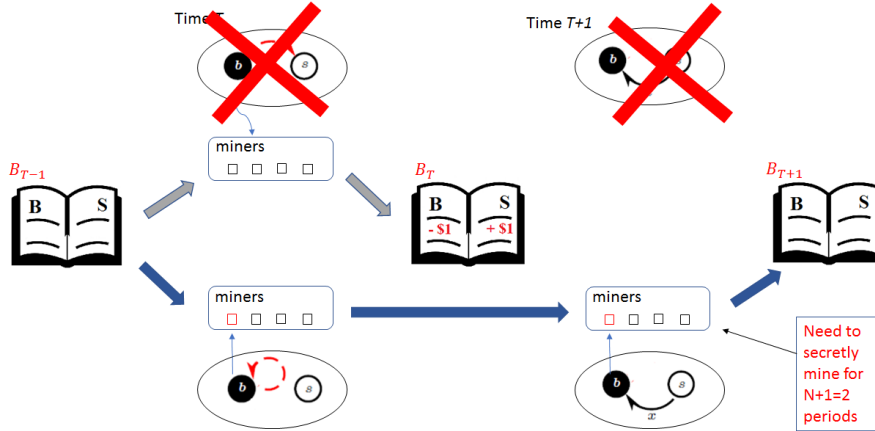


Figure 2.4: Double Spending Attack when Confirmation Lag

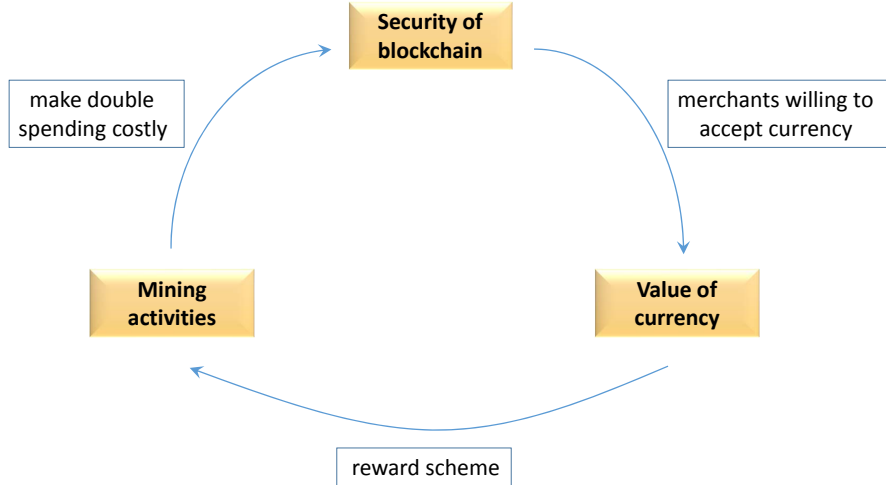


Figure 2.5: Bootstrapping in a Cryptocurrency System

- (ii) settlement lags: double spending is discouraged by sellers waiting for N validations before delivering the goods so that the buyer needs to win the mining game $N + 1$ times in order to revoke a payment
- (iii) a reward scheme: rewards for winning miners are financed by seigniorage (new coins) and transaction fees.

3 The Double Spending Problem

As pointed out in the previous section, due to its digital nature, a cryptocurrency system is subject to the double spending problem. To focus on this problem, this section develops a partial equilibrium model to study the mining and double-spending decision within one payment cycle. Taking as given the price and quantity of balances, the terms of trade and the mining rewards, this basic model determines the mining activities and the buyers' incentives to double spend. In the next section, we will incorporate this basic set-up into a general equilibrium monetary model to perform a full analysis.

3.1 Basic Set-up

We begin our analysis by looking at a single transaction period. As shown in Figure 3.1, there are $\bar{N} + 1$ subperiods within the single period. In subperiod 0, a buyer meets a seller to negotiate a trade. All other subperiods $1, \dots, \bar{N}$ serve as periods for confirming and settling trades that take place in subperiod 0.

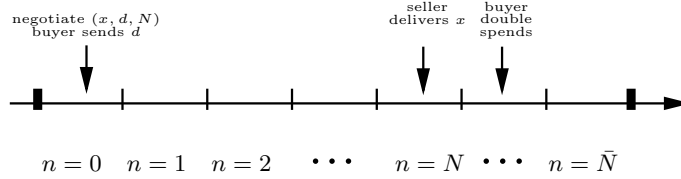


Figure 3.1: Timeline for a single transaction period

The buyer carries a real balance of cryptocurrency equal to z that can be used to buy an amount of goods x from a seller. Upon being matched, the buyer and the seller bargain to determine the terms of trade (x, d, N) which specify that the buyer pays the seller $d \leq z$ units of real balances and that the seller commits to deliver x units of goods after a number of successive payment confirmations $N \in \{0, \dots, \bar{N}\}$ in the Blockchain.¹⁴ We call N the *confirmation lag* of the transaction. For now, the terms of trade are taken as given, but will be determined endogenously in the next section.

The seller produces the good at unit costs, while the buyer's preference for consuming an amount x with confirmation lag N are given by

$$\delta^N u(x) \tag{1}$$

where $\delta \in (0, 1)$ is the discount factor between two subperiods. Hence, discounting across the whole transaction period is given by¹⁵

$$\beta = \delta^{\bar{N}+1}. \tag{2}$$

Finally, both buyers and sellers value real balances linearly and discount all payoffs that arise after the single transaction period at β .

¹⁴We consider the case in which a seller can commit to deliver the good. If this were not the case, only spot trades can be conducted that – as we show later – will always be subject to double spending.

¹⁵There are two ways to interpret the discount factor δ : (i) buyers prefer earlier consumption or (ii) a buyer's preference can change over time so that a seller's goods will no longer generate utility with probability $1 - \delta$.

3.2 Mining

There are also M miners who compete for updating a blockchain in subperiods $n = 0, \dots, \bar{N}$ with all the transactions from subperiod 0. In each subperiod, miners perform exactly one costly computational task with a random success rate by investing computing power, q , measured in real balances of cryptocurrency.¹⁶ This task is called the Proof-of-Work (PoW). We assume that miners value real balances linearly.

As motivated by the Bitcoin protocol (see Property (i) in Section 2), if the computational power of miner i in a subperiod is $q(i)$, then the probability that a particular miner i will win the mining game is given by

$$\rho(i) = \frac{q(i)}{\sum_{m=1}^M q(m)}. \quad (3)$$

In other words, the probability of winning is proportional to the fraction of computational power owned. We take this feature as given here and provide a micro foundation for this result in the Appendix. By winning the competition in any subperiod, a miner can update the Blockchain (i.e., append the n th block to the Blockchain) and receive a reward R in real balances. We assume that miners receive and consume this reward after the period, discounted by the time preference β .

Note that the mining games are independent across subperiods. Hence, in any subperiod miner i solves

$$\max_{q(i)} \rho_i \beta R - q(i) \quad (4)$$

so that

$$\frac{\sum_{m=1}^M q(m) - q^*(i)}{\left(\sum_{i \neq j} q(m) + q^*(i)\right)^2} \beta R = 1. \quad (5)$$

where he takes as given the choice by all miners $m \neq i$. Imposing symmetry, $q(m) = Q$ for all m , we obtain as the Nash equilibrium of the mining game that

$$Q = \frac{M-1}{M^2} \beta R. \quad (6)$$

Consequently, the total computing cost of mining in any subperiod is

$$MQ = \frac{M-1}{M} \beta R \quad (7)$$

¹⁶One could assume that miners buy computing power q at price α so that z units of real balances yield computing power $q = \alpha z$. We can, however, simply normalize $\alpha = 1$ by changing the unit for q .

The expected profit of a miner in equilibrium across the single transaction period is thus given by

$$\Pi_m = (\bar{N} + 1) \left[\frac{Q}{\sum_{m=1}^M Q} \beta R - Q \right] = \frac{\bar{N} + 1}{M^2} \beta R. \quad (8)$$

To capture the fact that mining tends to be quite competitive and open to new entrants, we let $M \rightarrow \infty^{17}$ to arrive at the following lemma.

Lemma 1. *As $M \rightarrow \infty$, the expected value of miners is zero, and the aggregate computing power of miners dissipates all rewards from mining*

$$MQ = \beta R.$$

3.3 Secret Mining

As discussed in Section 2, an important concern in a cryptocurrency system is a buyer’s double spending attempts (see Property (ii) in Section 2). When trading, the buyer needs to make a payment d to the seller. To do so, he has to send out an instruction to miners to update the Blockchain with the transaction. However, this is insufficient to ensure that the seller receives a payment. A buyer can engage in *secret mining* by attempting to mine a block in which his payment did not occur.¹⁸ A seller can protect himself from not receiving the payment by waiting to deliver the goods until the payment has been incorporated into the blockchain.

Such confirmation of the payment in the Blockchain however may still be not enough. A buyer can secretly mine a different Blockchain which could be released some periods after the seller has delivered the good replacing the original Blockchain.¹⁹ When such secret mining succeeds, the

¹⁷The total number of miners is estimated to be within the range of 5000 to 100,000 (<https://goo.gl/TPFBvA>). In addition, according to blockchain.info, there are altogether 14 mining pools that individually can account for at least 1% of the total hashrate. Finally, it is feasible for miners to use their existing mining capacities to mine different cryptocurrencies. For example, ASICs (Application-specific integrated circuits) manufactured for Bitcoin can be used to mine altcoins that use SHA-256 as the hashing algorithm (e.g., Namecoin and Peercoin).

¹⁸The secret mining can be done either by the buyer himself or by hiring a miner to mine a block with the instruction that the payment did not occur.

¹⁹Notice that, with such secret mining, the buyer cannot spend the balances of any other agent because, to spend other agents’ balances, one would need to obtain the digital signature of other agents. He can only (i) change the payment instructions of his own transaction and (ii) remove other payment instructions from being mined – and, hence, confirmed – in the block. Hence, in reality a buyer trying to double spend has to remove his own payment and all other payment instructions involving his original balance being spent.

buyer keeps his original balances and the goods while the seller will be left empty handed. In response, the seller can choose to postpone the delivery of the goods and wait for N confirmations. This confirmation lag can potentially deter double spending by the buyer. The idea is that, to undo a transaction with a confirmation lag of N subperiods, a dishonest buyer needs to win the mining game $N + 1$ times in a row. As the number of lags increases, the total PoW required to revoke a transfer is increasing, making it more costly for a buyer to double spend. Furthermore, secret mining is deterred by miners' investments in computing power MQ which, according to Lemma 1, is increasing in the reward R . We look next into the incentives to double spend and call an offer (x, d, N) *double spending proof* (DS-proof) if the buyer has no incentive to engage in double spending in subperiod 0 after the acceptance of the offer.

3.4 Double Spending Proof Contracts

Consider then a trade with the terms (x, d, N) . The buyer will receive the goods in subperiod N when exactly N confirmations of the payment d have been observed in the Blockchain. To double spend, a buyer can secretly mine an alternative history to undo his payment after he has received the goods. This implies that the double spender needs to be the first one who solves the PoW problem for $N + 1$ consecutive subperiods; i.e., from subperiod 0 to subperiod N . For each of the first N subperiods, the buyer does not broadcast the new block immediately, so that some other miners will update the Blockchain and confirm his payment to the seller. The buyer broadcasts his secretly mined blockchain only after he receives the goods and solves the $N + 1$ th block. When the double spending attack succeeds, the original payment is cancelled and the $N + 1$ rewards will be given to the buyer (see Figure 3.2).

Note first that the quantity of good x being exchanged does not directly matter for a double spending attack, as it neither changes the payoffs, nor the incentives to engage in an attack. We consider first subperiod N where the seller delivers his goods. Conditional on having been successful N times with secret mining²⁰, a buyer can now double spend if successful with mining a new block in subperiod N . His expected payoff from investing q_N is given by

$$\rho(q_N)\beta[d + (N + 1)R] - q_N. \tag{9}$$

²⁰We implicitly assume here that once a buyer starts a double spending attack, he does not stop after winning a block before subperiod N . As shown in the Appendix, this is without loss of generality because, if a buyer decides to double spend, it is not optimal to quit after winning some consecutive blocks before subperiod N .

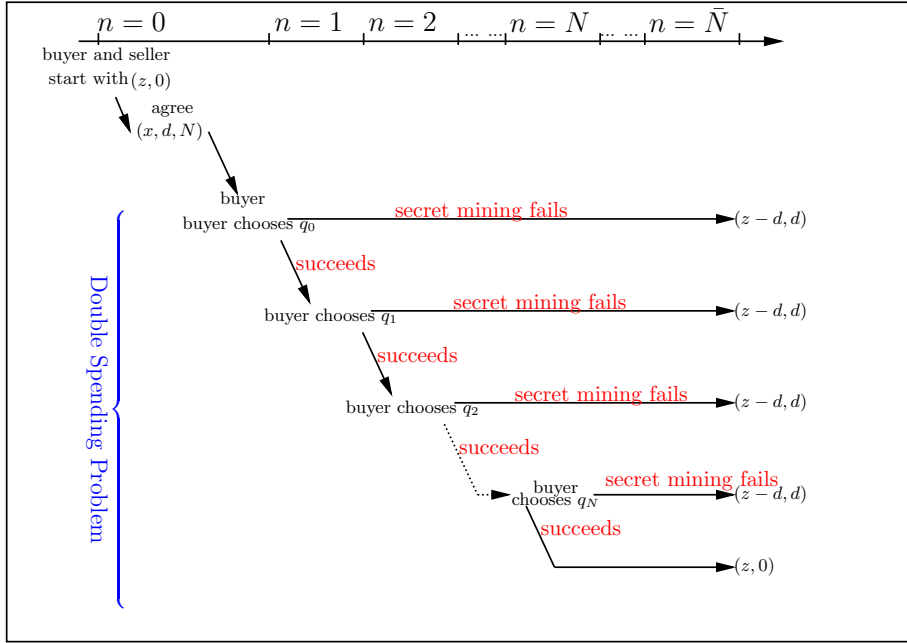


Figure 3.2: The Double Spending Problem

The first term captures the expected revenue from double spending. Conditional on having solved N blocks, with probability

$$\rho(q_N) = \frac{q_N}{QM + q_N} \quad (10)$$

the buyer wins the competition again in the $N + 1$ th round, so that he can double spend. The revenue from double spending is given by the original payment in subperiod 0 which is d . Also, the buyer obtains the rewards from all new blocks in his chain of length $N + 1$. The value in real balances at the end of the period is thus given by $\beta(d + R(N + 1))$.

We work now backwards to derive a condition for offers to preclude double-spending. Define

$$\Delta = \left[\frac{d}{R} + (N + 1) \right]. \quad (11)$$

As $M \rightarrow \infty$, the buyer's optimal choice with respect to investment in computing power in subperiod N is then given by

$$\hat{q}_N(d, N) = \sqrt{QM\beta R\Delta} - QM = \beta R \left[\sqrt{\Delta} - 1 \right] \quad (12)$$

so that the probability of successful double spending in subperiod N is

$$\rho(q_N) = \frac{\hat{q}_N}{QM + \hat{q}_N} = \frac{\sqrt{\Delta} - 1}{\sqrt{\Delta}}. \quad (13)$$

To derive the no-double-spending constraint we work backwards. Given (d, N) and R , the expected payoff for a double spending buyer in subperiod N having been successful N times already is

$$D_N(d, N) = \beta R(\sqrt{\Delta} - 1)^2. \quad (14)$$

Define recursively the expected payoff from double spending in subperiod n for $n \in \{0, \dots, N-1\}$ by

$$D_n(d, N) = \max_{q_n} \rho(q_n) D_{n+1}(d, N) - q_n, \quad (15)$$

which takes into account that the buyer was n times successful, since if he fails once the double spend fails as well. Note that $D_n(d, N)$ can only be positive if $D_{n+1}(d, N)$ is positive and $\hat{q}_n > 0$ only if $D_n(d, N) > 0$. The first-order condition describing the optimal investment is thus given by

$$\hat{q}_n(d, N) = \sqrt{QM \cdot D_{n+1}(d, N)} - QM. \quad (16)$$

By backward induction, we then obtain the following result.

Lemma 2. *As $M \rightarrow \infty$,*

$$D_{N-s}(d, N) = \beta R \left[\sqrt{\Delta} - (s+1) \right]^2, \quad (17)$$

$$\rho_{N-s}(d, N) = \frac{\sqrt{\Delta} - (s+1)}{\sqrt{\Delta} - s}, \quad (18)$$

$$\hat{q}_{N-s}(d, N) = \beta R \left[\sqrt{\Delta} - (s+1) \right]. \quad (19)$$

It follows immediately that $D_n(d, N)$ is strictly increasing in n and, consequently, the investment by the double spending buyer \hat{q}_n is also increasing in n . Hence, if it was optimal to engage in secret mining in subperiod n , it is optimal to continue with secret mining in subperiod $n+1$, if one has been successful in subperiod n . Consequently, double spending is not optimal for the buyer whenever

$$\beta R \left[\sqrt{\Delta} - (N+1) \right] < 0,$$

so that $\hat{q}_0 = 0$ and $D_0(d, N) = 0$. This yields the following *no double spending constraint*.²¹

²¹In a setting where a number n_b of buyers can coordinate in their double-spending attempts, the condition becomes $dn_b < R(N+1)N$. This suggests that a cryptocurrency is more secure in a decentralized environment where it is difficult to coordinate a deviation.

Proposition 3. For $M \rightarrow \infty$, a contract is double spending proof if

$$d < R(N + 1)N. \quad (20)$$

More generally, we have that

$$D_0(d, N) = \beta R \left[\sqrt{\Delta} - (N + 1) \right]^2 \quad (21)$$

whenever double spending yields a strictly positive payoff for the buyer which is decreasing in R and N and increasing in d . The unconditional probability of a successful double spending is

$$P(d, N) = \frac{\sqrt{\Delta} - (N + 1)}{\sqrt{\Delta}}.$$

To study the settlement properties of a cryptocurrency system, we first define the following concepts.

Definition 4. The settlement of a contract (x, d, N) is immediate if $N = 0$ and delayed if $N > 0$. The settlement is final if $P(d, N) = 0$ and probabilistic if $P > 0$.

Inequality (20) provides thus a condition for (full) finality. The reward R helps achieve finality by inducing mining activities which in turn increase the costs of a double spending attempt. Finality can also be supported by either reducing the trade size d or increasing confirmation lag N . Notice that the relationship between d and N defined by (20) is non-linear. This is because, in the double-spending problem, q interacts with d and N in a different fashion. Increasing N raises the mining cost which is linear in q , while increasing d raises the return at the rate P which is a concave function of q .

Given R , there is thus a trade-off between trade size d , settlement lag N and finality as captured by $1 - P(d, N)$. As shown in Figure (3.3), full finality is feasible only for small, sufficiently delayed settlement. Fast settlement of large transactions can only be probabilistic with the probability decreasing in d and increasing in N . This leads us to the following fundamental result.

Theorem 5. For any cryptocurrency based on a PoW protocol, settlement cannot be both immediate and final.

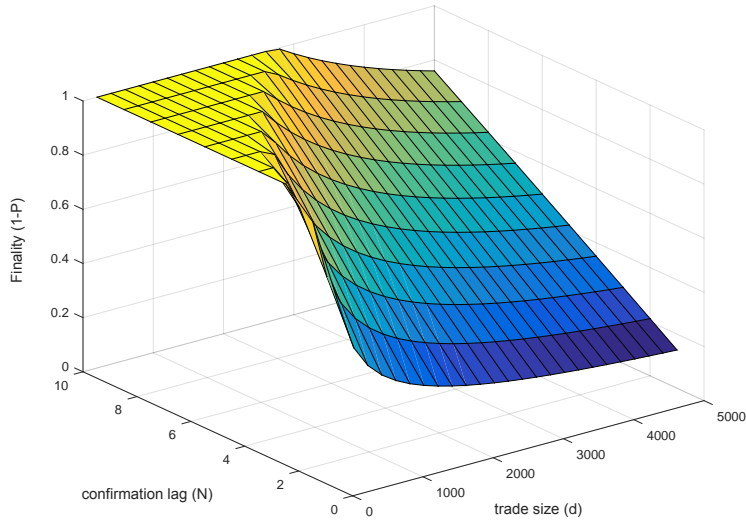


Figure 3.3: Trade Size, Confirmation Lag and Probabilistic Finality

We now incorporate this basic structure of a cryptocurrency into a general equilibrium framework of a monetary economy. This step is essential, as a cryptocurrency is a closed looped system (see introduction) and the cryptocurrency can only have value if it circulates in the economy. The value of the cryptocurrency will determine the rewards that miners receive. These rewards determine the mining effort and, thus, the incentives to double spend. This in turn will feed back into the value of the currency and its usefulness as a medium of exchange. Our framework can then be used to explore the optimal design of a cryptocurrency and to estimate the actual surplus produced by an existing cryptocurrency such as Bitcoin.

4 General Equilibrium Framework

4.1 Dynamic Model

Our model is based on Lagos and Wright (2005).²² Time is discrete and denoted by $t = 0, 1, 2, \dots$. There are a large number B of buyers and a large number $S = B\sigma$ of sellers, where $\sigma \in (0, 1)$. There is also a large number M of miners. While the general setting can also be analyzed, this section focuses on the case of competitive mining ($M \rightarrow \infty$) which, as argued in footnote 17, is

²²This framework is useful because it allows us to study frictions that necessitate the usage of money while still keeping the distribution of balances analytically tractable.

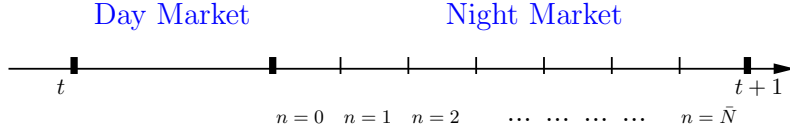


Figure 4.1: Time line

relevant in practice.

In each period, first a day market opens. This is a centralized competitive market for agents to trade a general good h which all agents can produce and consume. Then a night market opens. As shown in Figure 4.1, the night market is further divided into $\bar{N} + 1$ consecutive subperiods with $\bar{N} \geq 1$. In the first subperiod, buyers and sellers are matched with the probability of a buyer being matched to a seller being σ to engage in a bilateral trade. Sellers can produce a good x for buyers at unit cost. To introduce heterogeneity of transactions, buyers' preferences are now given by

$$\delta^N \varepsilon u(x) \tag{22}$$

where $\underline{\varepsilon} \leq \varepsilon \leq \bar{\varepsilon}$ is a preference parameter drawn from a distribution $F(\varepsilon)$ at the beginning of each period and where N is the subperiod in which the good is delivered by the seller. We assume that ε is known to the buyer and seller in a match and that the seller can commit to deliver the good in subperiod N . Finally, buyers discount payoffs across subperiods according to $\delta \in (0, 1)$ and all agents discount payoffs across periods with $\beta = \delta^{\bar{N}+1}$.

4.2 Cryptocurrency

A cryptocurrency is a digital record of ownership of nominal balances m that can be used to pay for transactions. For any transaction, the buyer gives instructions to transfer ownership of a certain amount of his balances to the seller. In the Appendix, we formalize the notion of a blockchain as a transaction-based ledger that records these transfers of cryptocurrency balances throughout time.

In the day market, we assume that there is perfect monitoring in the sense that a person that transfers balances is liable for their authenticity. More precisely, if the balances get lost for the payee, the payer needs to reimburse the payee for the loss. This assumption rules out double spending in the day market.²³ Due to the anonymity in the night market, the exchange of goods

²³This reflects the basic premise that certain parties such as merchants accepting a cryptocurrency and using it could be held legally liable for the losses sustained by other parties. In general, payers in many settings are not fully

necessitates a means of payment which we assume is the cryptocurrency.

The updating of this blockchain follows a PoW protocol where miners compete for the right to add a new block of balance transfers as reported by buyers into the chain for each subperiod. The mining game is identical to the one presented in Section 3. However, the reward R is not exogenous anymore, but depends on the design of the cryptocurrency (see Property (iii) in Section 2). First, the cryptocurrency can create new balances which are paid to miners that win the competition to update the blockchain. We denote the (gross) growth rate of new balances by $\mu \geq 1$. In addition, a fraction $\tau \geq 0$ of balances are paid to the successful miner as a transaction fee for including the transaction in a block. We assume that the $\bar{N} + 1$ block winners of the night market equally share the total reward. Consequently, the block reward is given by

$$R = \frac{Z(\mu - 1) + D\tau}{\bar{N} + 1}. \quad (23)$$

where Z is the aggregate money balances and D is the aggregate spending in the night market. Since rewards are paid in the next period, they are discounted by β/μ taking into account inflation.

4.3 Day Market

In the day market, nominal balances can be exchanged against the general good at price ϕ . We use $z = \phi m$ to denote the real balances associated with m units of balances. All buyers, sellers and miners can produce and consume the general goods with a linear utility function. Miners can also convert general goods 1 – 1 into computing power q at any time. Since they cannot transact in the night market, miners will not hold balances across periods due to discounting. Upon receiving balances as rewards, they will simply trade them for the general good in the day market. Hence, the problem of miners is identical to the one in Section 3 simply adjusted for the discount factor $\beta/\mu \leq \beta < 1$.

To ease exposition, we abuse notation and set $\varepsilon = 0$ for the seller. We use $w(z)$ and $v(z, \varepsilon)$ to denote respectively the day market and the night market value functions. The problem of a buyer

anonymous (e.g. the KYC requirement, online wallets transactions, payments to Bitcoin exchanges, payments made by well-known merchants). Hence, double spenders can be punished either formally (e.g. legal) or informally (e.g. reputational).

that draws ε (or a seller with $\varepsilon = 0$) entering the day market with real balances z is given by

$$w(z) = \max_{z', h} -h + v(z', \varepsilon) \quad (24)$$

subject to

$$z + h \geq z' \geq 0 \quad (25)$$

where $h > 0$ ($h < 0$) denotes production (consumption) of the general good and z' is real balances carried into the night market which have a value of $v(z', \varepsilon)$. The optimal demand for balances is

$$1 \geq v'(z', \varepsilon) \quad (26)$$

with equality when $z' > 0$. Linear preferences imply that

$$w(z) = z + w(0). \quad (27)$$

Hence, the value function before the realization of ε is

$$w(z) = E[w(z, \varepsilon)] = z + W \quad (28)$$

where $W = E[w(0; \varepsilon)]$ is a constant.

4.4 Night Market

If matched with a seller, the buyer holding real balances z makes a take-it-or-leave-it offer (x, d, N) which specifies a payment $d \leq z$ for obtaining x goods to be delivered after confirmations of the payment in N consecutive blocks. The value of a buyer entering the night market with balances z is then given by

$$\begin{aligned} v(z; \varepsilon) &= \sigma([\delta^N \varepsilon u(x) + D_0(d, N)] + \beta w(\frac{z-d}{\mu})) + (1-\sigma)\beta w(\frac{z}{\mu}) \\ &= \frac{\beta}{\mu} z + \sigma([\delta^N \varepsilon u(x) + D_0(d, N)] - \frac{\beta}{\mu} d) + \beta W \end{aligned} \quad (29)$$

where $D_0(d, N)$ is the expected value from double spending. Note that $D_0 = 0$ if the buyer has no incentive to double spend given the offer (x, d, N) . Since the seller has a linear technology to produce x , the value function of a seller is given by

$$v(z; 0) = \frac{\beta}{\mu} d(1-\tau)(1-P(d, N)) - x + \frac{\beta}{\mu} z + \beta W. \quad (30)$$

With probability $P(d, N)$ a double spending attack is successful leaving the seller without balances from the trade. Furthermore, the seller is required to pay the transaction fee which is a fraction τ of the payment in real balances d .

If balances can be used for trading in the night market, we have that in the day market buyers will have a positive demand for balances ($z' > 0$). Furthermore, since $\mu \geq 1 > \beta$, it is costly to carry extra balances into the night market. This implies that sellers do not carry any balances in the night market, while buyers carry only the amount of balances into the market they will use to make an offer; i.e., the cash-in-advance constraint is binding or $z' = d$. Hence, the buyer's optimal take-it-or-leave-it offer is thus a solution to the problem

$$\max_{(x, d, N)} -d + (1 - \sigma) \frac{\beta}{\mu} d + \sigma [\delta^N \varepsilon u(x) + D_0(d, N)] \quad (31)$$

$$\text{subject to} \quad (32)$$

$$x = \frac{\beta}{\mu} d (1 - \tau) (1 - P(d, N)) \quad (33)$$

The values of $D_0(d, N)$ and $P(d, N)$ depend on whether the offer is DS-proof or not. The mining game and the buyer's incentives to engage in double spending are identical to Section 3 except for that discounting of payoffs across periods is given by $\beta/\mu < \beta$. Hence, we have that

$$P(d, N) = D_0(d, N) = 0$$

if $d \leq R(N + 1)N$ – in other words, if the contract is DS proof – or

$$P(d, N) = \frac{\sqrt{\Delta} - (N + 1)}{\sqrt{\Delta}} > 0 \text{ and } D_0(d, N) = \frac{\beta}{\mu} R \left[\sqrt{\Delta} - (N + 1) \right]^2 > 0$$

if the contract involves double spending. Note that double spending gives the buyer an additional payoff $D_0(d, N)$, but also tightens the seller's participation constraint thereby reducing the amount being transacted in the night market.

The buyer's problem in general can have multiple solutions. For example, sometimes a buyer can be indifferent between double spending and a DS-proof contracts. Similarly, the buyer can be indifferent between a contract with a long confirmation lag and large consumption and one with earlier but smaller consumption. For completeness, the seller's value function is given by

$$v(z; 0) = \beta w\left(\frac{z}{\mu}\right) = \frac{\beta}{\mu} z. \quad (34)$$

since they do not receive any surplus.

4.5 DS Proof Equilibrium

In what follows, we focus on an equilibrium in which the cryptocurrency has a positive value and all trades are double-spending proof.²⁴ Given R , define the set of optimal money demand as $\Gamma(\varepsilon; R)$. For a given selection from the solution set $\Gamma(\varepsilon; R)$, the aggregate transfer of balances in the night market and aggregate money demand are described by

$$D = B\sigma E(d) = B\sigma \int_0^\infty z^*(\varepsilon; R) dF_\varepsilon(\varepsilon) \quad (35)$$

$$Z = BE(z) = B \int_0^\infty z^*(\varepsilon; R) dF_\varepsilon(\varepsilon). \quad (36)$$

In equilibrium, nominal balances are growing at rate μ and so do prices. Our definition has only used real balances which stay constant across time.

Definition 6. *A DS-proof cryptocurrency equilibrium with (μ, τ) is given by offers $(x(\varepsilon), d(\varepsilon), N(\varepsilon))$, a money demand $z(\varepsilon) > 0$ and a mining choice q such that*

1. *for all ε the money demand and the contract maximizes the buyer's utility*
2. *the mining choice is a Nash equilibrium of the mining game in every subperiod*
3. *for all ε the contract satisfies condition (20)*
4. *the day market for balances clears.*

Note that when the DS constraint (20) is binding, a buyer may consider offering a DS contract in order to increase the consumption and/or to reduce the confirmation lag. However, the settlement of a DS contract must be probabilistic ($P(d, N) < 1$) and, according to the seller's participation constraint (33), the buyer has to offer a higher balance d for the seller to accept the offer. Consequently, buyers will offer DS proof contracts when the net gain for carrying extra balances is small.

²⁴Bitcoin – like cryptocurrencies in general – is designed to discourage double spending attacks. Under the current system where sellers wait for multiple confirmations, double spending activities do not seem to be a significant concern. However, some successful double-spending were reported when the confirmation lag was insufficient. According to Bicoiwiki, “in November 2013 it was discovered that the GHash.io mining pool appeared to be engaging in repeated payment fraud against BetCoin Dice, a gambling site. Dice sites use one transaction per bet and don't wait for confirmations.”

This is the case when the opportunity cost of holding balances, captured by the nominal interest rate $i = \mu/\beta - 1$, is large and the expected utility from a transaction is sufficiently low (small σ and $\delta\bar{\varepsilon}$). In the appendix, we derive an explicit sufficient condition under which only DS-proof contracts are offered which confirms this reasoning. We next show that under a weak additional condition an equilibrium exists.

Proposition 7. *If B is sufficiently high, then a DS-proof cryptocurrency equilibrium exists.*

To support a positive value of cryptocurrency and to discourage buyers from double spending, the reward R has to be sufficiently high. For any given (μ, τ) , this can be achieved by having a sufficiently large number of buyers B . This result highlights an important feature of a cryptocurrency: the mining reward R is financed by *aggregate* transactions (which scales with B), while the incentive to double spend depends only on the individual transaction size d . Once again, this is also captured by the DS constraint where $d < R(N + 1)N$.

4.6 Optimal Cryptocurrency Design

We take now the PoW protocol of a cryptocurrency as given and ask how to optimally structure the reward for miners. We first define the social welfare function as the average utility per period or

$$\mathcal{W} = B\sigma \int_0^\infty [\delta^{N(\varepsilon)} \varepsilon u(x(\varepsilon)) - x(\varepsilon)] dF_\varepsilon(\varepsilon) - C(\bar{N} + 1) \quad (37)$$

The first term is the aggregate surplus generated every period in the night market, while the last term is the cost of mining consisting of the cost C for mining a single block times the total number of blocks $\bar{N} + 1$. Lemma 1 implies that this cost is simply the present value of the total mining rewards. A social planner would design the cryptocurrency by setting the currency growth rate and the transaction fees so as to maximize this social welfare function. The problem for the social

planner is given by

$$\max_{\mu, \tau} B \int \sigma \delta^{N(\varepsilon)} \varepsilon u[x(\varepsilon)] - x(\varepsilon) dF_\varepsilon(\varepsilon) - \frac{\beta}{\mu} R(\bar{N} + 1) \quad (38)$$

subject to

$$(d(\varepsilon), x(\varepsilon), N(\varepsilon)) \text{ is an optimal offer for any } (\mu, \tau) \quad (39)$$

$$d(\varepsilon) = \frac{\mu}{\beta} \frac{x(\varepsilon)}{1 - \tau} \leq R(N(\varepsilon) + 1)N(\varepsilon) \quad (40)$$

$$\frac{\beta}{\mu} R(\bar{N} + 1) = B(\mu - 1 + \sigma\tau) \int \frac{x(\varepsilon)}{1 - \tau} dF_\varepsilon(\varepsilon) \quad (41)$$

This is akin to a Ramsey Problem in the optimal tax literature where the planner takes into account that individuals make optimal decisions taking the policy (μ, τ) as given. Note that this decision is reflected in the first two constraints which can be summarized as the first-order conditions from the buyer's decision problem requiring that there is no DS. The last constraint is the budget constraint for the planner to deliver rewards R to miners for the $\bar{N} + 1$ blocks.²⁵ In the appendix, we show that the optimal design of the system takes a simple form, summarized by the following proposition.

Proposition 8. *The optimal reward structure sets transaction fees to zero and only relies on seignorage; i.e., $\tau = 0$ and $\mu > 1$.*

Positive rewards are required to induce mining activities. Since both transaction fees and inflation distort buyers' incentives to consume in the night market, the optimal designer needs to set (μ, τ) to minimize the total distortion as a result of mining and the loss of trade surplus. Given a level of reward R , lowering transaction fees for trades, but increasing inflation can generate more surplus from trades. The reason is that the inflation tax is shared by all buyers while transaction fees are paid only by the buyers that have a chance to trade. These buyers when facing the fee are precisely the ones who have a high valuation of balances. This implies that levying reward costs upfront in terms of inflation allows distortions to be smoothed out across all buyers upfront.

5 A Numerical Analysis of Bitcoin

Based on our theoretical analysis, we are now seeking to understand the limits of using cryptocurrencies for payments. To do so, we perform several quantitative exercises that calculate the welfare

²⁵Note the \bar{N} could also be a system parameter in our model. We simply impose here that \bar{N} is just sufficiently large to allow for the optimal DS proof contract given $\bar{\varepsilon}$.

Table 5.1: Bitcoin Transaction Characteristics (Source: blockchain.info)

	Per day	Per block
No of transactions	122129.7534	848.1232877
Estimated transaction volume (BTC)	254843.1781	1769.744292
Transaction fees (BTC)	22.45900183	0.15596529

losses relative to other means of payments. We first use Bitcoin trading data to calibrate some parameters for our model. Using this calibration, we compare the current Bitcoin scheme and a fictitious cryptocurrency scheme that operates with an optimal reward structure with an economy that uses traditional cash instead. Then, we conduct a similar exercise by using data from US debit cards and Fedwire to see how well an optimal cryptocurrency can support such transactions.

5.1 Parameterization

We assume that buyers' utility function is

$$u(x) = \log(x + b) - \log b$$

with $b \approx 0$. The length of a period is a day and the length of each trading session is 10 minutes (i.e., average block time). Setting $\beta = 0.9999$ gives an annual discount factor of 0.97. The average Bitcoin supply in 2015 was 14342502.95. Consequently, the money growth rate per day in 2015 was $\mu = (1 + 25/14342502.95)^{6 \times 24} = 1.00025$. This translates into an annual inflation rate of 9.6%.

We use aggregate Bitcoin transactions to calibrate the rest of our parameters (see Table 5.1). We set $\sigma = 0.0178$ to match the average fraction of Bitcoins spent per day, and set $\tau = 0.15596529 / 1769.744292 = 0.000088129$ to match the transaction fees data. The average transaction size is $\tau = 1769.744292 / 848.1232877 = 2.086659237$. Finally, we use $B = 6873428.441$ which is the maximum number of average-sized transactions that the existing stock of Bitcoins can support.²⁶

Table 5.2 summarizes our benchmark parameters. The distribution $F(\varepsilon)$ is set to capture the shape of the empirical distribution of transaction size reported in Ron and Shamir (2013) (see Figure 5.1). Based on this, we obtain an implied density function of the preference shocks ε from our model

²⁶This is close to the number of blockchain wallet users which is 5439181 in 2015 (Source: blockchain.info, year-end number).

Table 5.2: Benchmark parameters

	values	targets
β	0.999916553598325	period length = 1 day
δ	0.999999420487088	$\delta = \beta^{1/(1+\bar{N})}$
μ	1.00025	money growth rate
τ	0.000088	transaction fee
B	6873428	max. no of average-sized transactions
σ	0.0178	velocity per block (block length = 10 mins)
α	1	normalization

and the confirmation lag N as optimally chosen in the transaction in our model given a preference shock (see Figure 5.2).

5.2 Effects of Money Growth

According to Proposition 8, it is optimal to set transactions fees to zero. Before deriving the optimal policy money growth rate, we first study the equilibrium effects of a partial change in the money growth rate around the benchmark equilibrium. Given the benchmark level of τ , Figure 5.2 shows the effects of μ on aggregate trade, average confirmation lags, utility, welfare, rewards and mining costs. By inducing mining activities, a higher μ lowers confirmation lags but increases inflation. The net effect on consumption and utility is positive. Also, a higher μ raises rewards, computational efforts and overall mining costs. The former effect improves welfare while the latter effect reduces welfare. The sum of these two effects results in a hump shape response of welfare to money growth.

5.3 Efficiency of Cryptocurrencies

For the distribution of preference shocks in Figure 5.1, Table 5.3 evaluates the efficiency of Bitcoin as a means of payment relative to a cash system. All computations are for our benchmark model with the same preference parameters, but using different payments systems: cash, Bitcoin, optimal reward structure for Bitcoin. Besides mining costs, we report two measures of the welfare cost.

Figure 5.1: Size Distribution of Bitcoin Transactions (Ron and Shamir, 2013)

Larger or equal to	Smaller than	Number of transactions in the graph of entities	Number of transactions in the graph of addresses
0	0.001	381,846	2,315,582
0.001	0.1	1,647,087	4,127,192
0.1	1	1,553,766	2,930,867
1	10	1,628,485	2,230,077
10	50	1,071,199	1,219,401
50	100	490,392	574,003
100	500	283,152	262,251
500	5,000	70,427	67,338
5,000	20,000	6,309	6,000
20,000	50,000	1,809	1,796
50,000		364	340

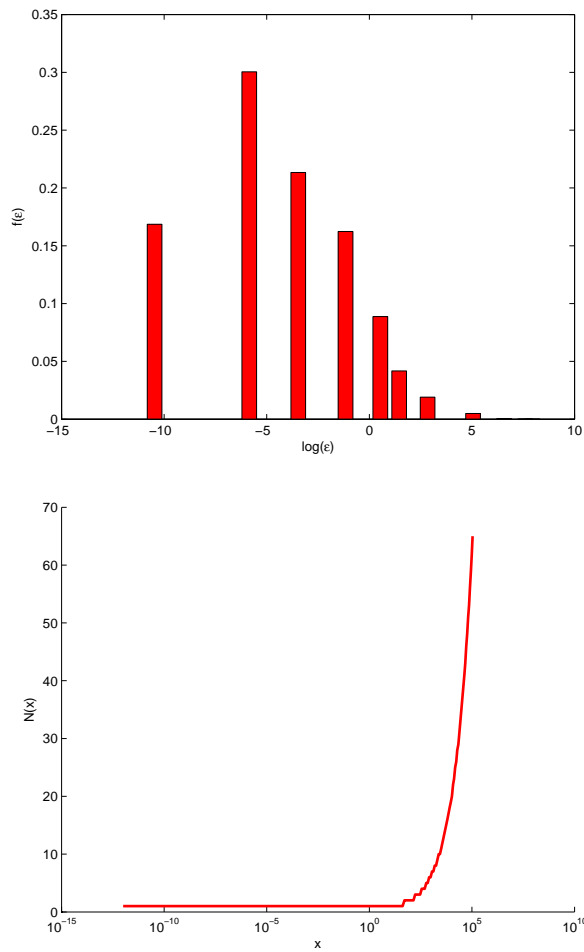


Figure 5.2: Density of Preference Shocks $F(\varepsilon)$ and Confirmation Lag $N(x)$

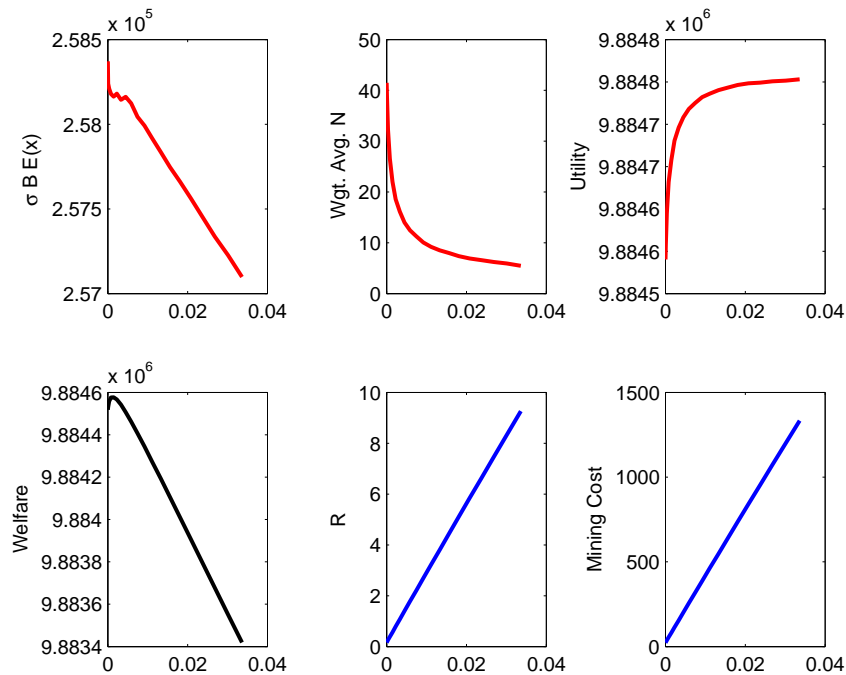


Figure 5.3: Effects of money growth

The first measure gives the fraction of consumption people are willing to sacrifice in order to use cash under the Friedman rule which implies zero welfare costs. The second one computes the inflation rate with traditional cash so that people are indifferent between such system and the cryptocurrency.

The current Bitcoin design is very inefficient, generating a welfare loss of 1.4% relative to an efficient cash system.²⁷ The main source for this inefficiency is the large mining cost, which is estimated to be 360 mn USD per year. This translates into people being willing to accept a cash system with an inflation rate of 230% before being better off using Bitcoin as a means of payment.

However, given the distribution of preference shocks, it is inefficient to set the money growth rate and the transaction fees as high as in the calibrated model for Bitcoin. The optimal policy is to reduce the money growth rate – and to not use transaction fees at all (see Proposition 8) – which will discourage mining substantially. Consequently, an optimally designed reward structure for Bitcoin would reduce its welfare cost to a small fraction of its estimated current cost (0.08%). The corresponding inflation that leaves people indifferent would drop to a more moderate level of

²⁷For comparison, a cash system with a 2% money growth rate generates a relatively small welfare cost of 0.003%.

27.51%.

Still, relative to cash, Bitcoin seems to be a very inefficient payment system for facilitating the observed set of transactions. This result could be driven by the fact that in the data, Bitcoin is being used for both large and small value transactions, and that the total volume of transactions is small. In order to control for this, we examine next the efficiency of a cryptocurrency when it is used to support a large volume of either small or large value transactions.

Table 5.3: Efficiency of Current and Optimal Cryptocurrency Systems

	Bitcoin (benchmark)	Bitcoin (optimal policy)
$\mu - 1$	9.5%	0.17%
τ	0.0088%	0%
welfare loss	1.41%	0.08%
mining cost (per year)	\$359.98 millions	\$6.90 millions
equivalent inflation tax	230.44%	27.51%

5.4 Best Usage of Cryptocurrencies

We now evaluate the efficiency of using cryptocurrencies for retail and large-value settlement systems. In Table 5.4, we present the quantitative results of calibrating our cryptocurrency model to 2014 US retail (debit cards) payments data and US large-value (Fedwire) data. A period is 30 minutes and the block length 1 minute. For the retail data, we pick $B = 30.16$ millions to match the number of debit cards²⁸ and set $\sigma = 0.540853348$ to match the volume of transactions per card per day. For Fedwire, we assume $B = 7866$ to match the number of participants in 2014, and set $\sigma = 0.9795$ which is the average volume of transactions for a participants in 30 minutes. Finally, we chose ε so that the average size of transactions equals the one observed in these payments systems, \$38.29 and \$6.5mn respectively. This is driven by data limitations. Double-spending incentives however increase with transaction size and, hence, we assume that the largest transactions in the debit card and Fedwire systems are 100 times and 5 times the average trade size.

Table 5.4 confirms that the welfare losses in a retail payment system are much smaller than in a

²⁸Source: <http://www.bis.org/cpmi/publ/d152.pdf>

Table 5.4: Welfare Comparison between Retail and Large-value Systems

	Retail Payments (US Debit cards)	Large Value Payments (Fedwire)
avg transaction size	\$38.29	\$6,552,236
annual volume	59539 millions	135 millions
optimal $\mu - 1$	0.038%	0.53%
optimal τ	0%	0%
confirmation lag	2 mins	12 mins
welfare loss	0.00052 %	0.0060%
mining cost (per year)	\$4.33 millions	\$22.10 billions
equivalent fee (per transfer)	\$0.0002	\$392.56

large-value one. In terms of the consumption equivalent measure, the welfare loss in a larger-value system is 0.006% of consumption, which is about 10 times larger than that in a retail system. A large-value system incurs a huge mining cost of 22 bn USD, which is over 5000 times of that in the retail system. In the last row, we also derive the required transaction fee of a cash system (at 2% inflation) so that people are indifferent between such system and the cryptocurrency. When a cryptocurrency is used for retail transactions, the equivalent transaction fee is a negligible 0.02 cents per transfer. For the large-value system, the corresponding fee becomes a very large \$392.

The basic intuition follows directly from the double spending constraint we have derived in our theoretical model. As the transaction size is smaller in the retail system, the incentives to double spend are also smaller. Furthermore, mining is a public good so that the rewards from money growth can support a large transaction volume. This implies that confirmation lags can be shorter and one needs to induce less mining effort to dwarf double spending. Consequently, money growth can also be smaller in a retail system, making a cryptocurrency system less costly due to inflation. This implies that a cryptocurrency works best when the volume of transactions is larger relative to the individual transaction size. As a result, a cryptocurrency tends to be much more efficient for conducting retail payments.

The transaction fee measures in the last row of the table allow us to also evaluate whether cryptocurrencies can be a viable alternative to current payment systems. The interchange fee in the

current debit card system is about 23 cents per transfer, while the service fee for Fedwire is 82 cents per transfer.²⁹ This suggests that a well-functioning cryptocurrency system can potentially challenge current debit card systems by offering users a competitive transaction fee with large enough transaction volumes.

This comparison – especially for retail payments systems – needs to be interpreted with caution. First, we do not consider certain private costs of running a cryptocurrency system. Examples are costs for data storage, network communication and software such as wallets to operate the system. Second, while the mining cost is a deadweight loss to society, part of the fees collected by retail and large value payment systems are profits earned by the providers so that operating costs tend to be lower than reflected in those fees. Finally, the above comparison does not take into account an important technical limitation of cryptocurrencies. Bitcoin and other implementations of cryptocurrencies face tight limits to their scalability. Unless one can address this issue by changing limits on block size and latency due to network speed, such systems will not be able to handle a large volume of transactions as required by modern retail payment systems.

6 Conclusion

Distributed record-keeping with a blockchain based on consensus through PoW is an intriguing concept. The economics of this technology that underlies most cryptocurrencies are driven by the individual incentives to double-spend and the costs associated with reining in these incentives. These costs are private in the form of settlement delay and social in the form of mining which is a public good. Consequently, as the scale of a cryptocurrency increases, it becomes more efficient. This explains why a double-spending proof equilibrium exists only when the user pool is sufficiently large, and why a cryptocurrency works best when the volume of transactions is large relative to the individual transaction size. This insight seems to be very much ignored in the current debate, but puts scalability of cryptocurrencies front and centre as the main technological challenge to be overcome. Our exercise shows that cryptocurrency systems can potentially be a viable alternative to retail payment systems, as soon as some technological limits can be resolved.³⁰

²⁹For the detail, see <https://www.federalreserve.gov/paymentsystems/regii-average-interchange-fee.htm> and https://www.frbservices.org/servicefees/fedwire_funds_services_2017.html.

³⁰As pointed out by Christine Lagarde, the head of the International Monetary Fund, cryptocurrencies could in time be embraced more broadly: “For now, virtual currencies such as Bitcoin pose little or no challenge to the

For Bitcoin we find that it is not only extremely expensive in terms of its mining costs, but also inefficient in its long-run design. However, the efficiency of the Bitcoin system can be significantly improved by optimizing the rate of coin creation and minimizing transaction fees. Another potential improvement is to eliminate inefficient mining activities by changing the consensus protocol altogether. In the Appendix, we explore the possibility of replacing PoW by a Proof-of-Stake (PoS) protocol. Our analysis finds conditions under which PoS can strictly dominate PoW and even support immediate and final settlement. Notwithstanding, as we point out in this Appendix as well, many fundamental issues of a PoS protocol remain still to be sorted out. There remains much to be learned about the economic potential and the efficient, economic design of blockchain technology.

References

- Agarwal, R. and M. Kimball, (2015). “Breaking through the Zero Lower Bound.” Working Paper WP/15/224, International Monetary Fund.
- Aumann, Robert J. (1965). “Integrals of set-valued functions”, *Journal of Mathematical Analysis and Applications*, 12.1, 1-12.
- Berentsen, A. (197). “Monetary policy implications of digital money”, *Kyklos*, 51(1), 89-117.
- Biais, B., Bisière, C., Bouvard, M. and Cassamatta, C. (2017), “The Blockchain Folk Theorem”, Working paper 817, Toulouse School of Economics.
- Böhme, R., Christin, N., Edelmann, B. and Moore, T. (2015), “Bitcoin: Economics, Technology and Governance”, *Journal of Economic Perspectives*, vol. 29, pp. 213-238.
- Bordo, M. and Levin, A. (2017), “Central Bank Digital Currency and the Future of Monetary Policy”, Economics Working Paper 17104, Hoover Institution.
- Camera, G. (2017). “A Perspective on Electronic Alternatives to Traditional Currencies”, *Sveriges Riksbank Economic Review*, 2017:1, 126-148.

existing order of fiat currencies and central banks ... But many of these [impediments] are technological challenges that could be addressed over time. Not so long ago, some experts argued that personal computers would never be adopted, and that tablets would only be used as expensive coffee trays. So I think it may not be wise to dismiss virtual currencies.” Source: <https://www.imf.org/en/News/Articles/2017/09/28/sp092917-central-banking-and-fintech-a-brave-new-world>

- Chiu, J., and T. Wong. (2015). “On the Essentiality of E-Money”, Working Paper 2015-43, Bank of Canada.
- Fernández-Villaverde, J. and D. Sanches, (2016). “Can currency competition work?”, NBER Working Paper 22157, National Bureau of Economic Research.
- Gandal, N., and H. Halaburda (2014). “Competition in the Cryptocurrency Market”, Working Paper 2014-33, Bank of Canada.
- Gans, J., and H. Halaburda, (2013). “Some Economics of Private Digital Currency”, Working Paper 2013-38, Bank of Canada.
- Glaser, F., Haferkorn, M., Weber, M. and Zimmermann, K. (2014), “How to price a Digital Currency? Empirical Insights in the Influence of Media Coverage on the Bitcoin Bubble”, *Banking and Information Technology*, 15.
- Huberman, G., Leshno, J. and Moallemi, C. (2017), “Monopoly without a Monoplist: An Economic Analysis of the Bitcoin Payment System”, Working Paper 17-92, Columbia Business School.
- Koepl, T., Monnet, C. and Temzelides, T. (2008), “A Dynamic Model of Settlement”, *Journal of Economic Theory*, 142, pp. 233-246.
- Koepl, T., Monnet, C. and Temzelides, T. (2012), “Optimal Clearing Arrangements for Financial Trades”, *Journal of Financial Economics*, 103, pp. 189-203.
- Lagos, L. and Wright, R. (2005). “A unified framework for monetary theory and policy analysis”, *Journal of Political Economy*, 113, pp. 463-484.
- Moore, T. and N. Christin, (2013). “Beware the Middleman: Empirical Analysis of Bitcoin-Exchange Risk”, in: *Financial Cryptography and Data Security*, Springer.
- Rogoff, K.S., (2016). *The Curse of Cash*, Princeton University Press.
- Ron, D. and Shamir, A. (2013). “Quantitative analysis of the full bitcoin transaction graph”, International Conference on Financial Cryptography and Data Security, pp. 6-24.
- Rosenfeld, M. (2014) “Analysis of hashrate-based double spending”, arXiv preprint, arXiv:1402.2009.
- Yermack, D. (2013) “Is Bitcoin a real currency? An economic appraisal”, No. w19747. National Bureau of Economic Research.

A Appendix – Proofs and Derivations

A.1 A micro-foundation for the proof-of-work problem

Miners perform their mining between trading sessions. By investing computing power $q(m)$, the probability that a miner m can solve the computational task within a time interval t is given by an exponential distribution with parameter μ_m

$$F(t) = 1 - e^{-\mu_m t}$$

where $\mu_m = q(m)/D$. The parameter D captures the difficulty of the proof-of-work controlled by the system. The expected time needed to solve the problem is thus given by

$$\frac{D}{q(m)}.$$

Aggregating over all M miners, the first solution among all miners, $\min(\tau_1, \tau_2, \dots, \tau_M)$, is also an exponential random variable with parameter $\sum_{m=1}^M \mu_m$. Hence the expected time needed to complete the proof-of-work by the pool of miners is³¹

$$\frac{D}{\sum_{m=1}^M q(m)}.$$

Furthermore, any particular miner m will be the first one to solve the proof-of-work problem with probability

$$\rho_n = \frac{q(n)}{\sum_{m=1}^M q(m)}.$$

³¹In practice, the parameter D is often adjusted to maintain a constant time for completing the proof-of-work problem given any changes in the total computational power.

A.2 Proof of Lemma 2

Proof. This is true for $s = 0$. Suppose the result holds true for $s = n - 1$. Consider $s = n$,

$$q_{N-n}(d, N) = \sqrt{QM \cdot D_{N-n+1}(d, N)} - QM \quad (\text{A.1})$$

$$= \beta R(\sqrt{\Delta} - n) - \beta R \quad (\text{A.2})$$

$$= \beta R[\sqrt{\Delta} - (n + 1)]. \quad (\text{A.3})$$

$$\rho_{N-n}(d, N) = \frac{q_{N-n}(d, N)}{QM + q_{N-n}(d, N)} \quad (\text{A.4})$$

$$= \frac{\sqrt{\Delta} - (n + 1)}{\sqrt{\Delta} - n}. \quad (\text{A.5})$$

$$D_{N-n}(d, N) = \rho_{N-n}(d, N)D_{N-n+1}(d, N) - q_{N-n}(d, N) \quad (\text{A.6})$$

$$= \frac{\sqrt{\Delta} - (n + 1)}{\sqrt{\Delta} - n} \beta R(\sqrt{\Delta} - n)^2 - \beta R[\sqrt{\Delta} - (n + 1)] \quad (\text{A.7})$$

$$= \beta R[(\sqrt{\Delta} - (n + 1))]^2. \quad (\text{A.8})$$

□

A.3 Suboptimal to quit mining after winning a block

Claim: If a DS plan (q_0, \dots, q_N) generates a positive payoff, then the buyer has no incentives to quit mining after winning a block in $n < N$.

Proof: Consider the case with $N = 2$

$$\begin{aligned} D_0 &= \rho_0 \rho_1 \rho_2 \frac{\beta}{\mu} [d + 3R] - \rho_0 \rho_1 q_2 - \rho_0 q_1 - q_0 \\ &= \rho_0 \left[\underbrace{\rho_1 \left[\underbrace{\rho_2 \frac{\beta}{\mu} [d + 3R] - q_2}_{D_2} \right] - q_1}_{D_1} \right] - q_0 \end{aligned}$$

Define $D_1 = \rho_1 \left[\rho_2 \frac{\beta}{\mu} [d + R(1 + N)] - q_2 \right] - q_1$. Since

$$D_0 = \rho_0 D_1 - q_0 > 0$$

we must have

$$\rho_1 \left[\rho_2 \frac{\beta}{\mu} [d + 3R] - q_2 \right] - q_1 > \frac{\beta}{\mu} R. \quad (\text{A.9})$$

Similarly, define $D_2 = \rho_2 \frac{\beta}{\mu} [d + R(1 + N)] - q_2$. Condition (A.9) becomes

$$\rho_1 D_2 - q_1 > \frac{\beta}{\mu} R,$$

which implies that $D_2 > 2 \frac{\beta}{\mu} R$, or

$$\rho_2 \frac{\beta}{\mu} [d + R(1 + N)] - q_2 > 2 \frac{\beta}{\mu} R. \quad (\text{A.10})$$

Condition (A.10) implies that

$$\rho_2 \frac{\beta}{\mu} [d + 3R] - q_2 + u(x) > 2 \frac{\beta}{\mu} R + \frac{\beta}{\mu} d.$$

Here, the LHS is the expected payoff of continuing DS in session 2, and the RHS is the payoff of quitting DS.

Similarly, condition (A.9) implies that

$$\rho_1 \left[\rho_2 \frac{\beta}{\mu} [d + 3R] - q_2 \right] - q_1 + \delta u(x) > \frac{\beta}{\mu} R + \frac{\beta}{\mu} d.$$

Here, the LHS is the expected payoff of continuing DS in session 1, and the RHS is the payoff of quitting DS.

For a general N , we can show that

$$\begin{aligned}
D_N &= \rho_N \frac{\beta}{\mu} [d + R(1 + N)] - q_N > \frac{\beta}{\mu} NR, \\
D_{N-1} &= \rho_{N-1} D_N - q_{N-1} > \frac{\beta}{\mu} (N-1)R, \\
&\vdots \\
D_1 &= \rho_1 D_2 - q_1 > \frac{\beta}{\mu} R, \\
D_0 &= \rho_0 D_1 - q_0 > 0.
\end{aligned}$$

They imply respectively

$$\begin{aligned}
\rho_N \frac{\beta}{\mu} [d + R(1 + N)] - q_N + u(x) &> \frac{\beta}{\mu} (NR + d), \\
\rho_{N-1} D_N - q_{N-1} + \delta u(x) &> \frac{\beta}{\mu} [(N-1)R + d], \\
&\vdots \\
\rho_1 D_2 - q_1 + \delta^{N-1} u(x) &> \frac{\beta}{\mu} (R + d).
\end{aligned}$$

Hence, there is no incentives to quit after winning a block.

A.4 A condition for no double spending

We now derive a sufficient condition under which double-spending contracts are dominated in equilibrium. Define the nominal interest rate as $i = \mu/\beta - 1$.

Lemma A.1. *Only DS-proof contracts are offered if*

$$\delta \varepsilon_{\max} u'(\bar{x}_1)(1 - \tau) \frac{3}{4} - 1 < \frac{i}{\sigma} \quad (\text{A.11})$$

where $\bar{x}_1 = (1 - \tau)(\beta/\mu)2R$ and $i = \mu/\beta - 1$.

Proof: We examine the optimal DS contract and the optimal DS-proof contract.

DS-proof contract

Fix N . The optimal DS-proof contract is a solution to

$$\max_{d,x} -d + (1 - \sigma) \frac{\beta}{\mu} d + \sigma \delta^N \varepsilon u(x)$$

subject to

$$\begin{aligned} \frac{x}{1 - \tau} \frac{\mu}{\beta} &= d \\ d &\leq R(N + 1)N \end{aligned}$$

Note that the participation constraint of the seller is always binding.

If the incentive constraint is not binding, then the FOC is then given by

$$1 = \frac{\sigma}{i} [\delta^N \varepsilon u'(x)(1 - \tau) - 1],$$

where $i = \mu/\beta - 1$ is the nominal interest rate. We denote this solution by (x^*, d^*) where it is understood that the solution depends on N .

If the incentive constraint is binding, then we have that

$$\begin{aligned} \bar{d} &= R(N + 1)N \\ \bar{x}(N) &= \frac{\beta}{\mu} (1 - \tau) R(N + 1)N \end{aligned}$$

and

$$1 < \frac{\sigma}{i} [\delta^N \varepsilon u'(\bar{x}(N))(1 - \tau) - 1]$$

We denote this solution by (\bar{x}, \bar{d}) where it is understood that the solution depends on N .

DS contract

Fix N . The optimal DS contract is a solution of the problem

$$\max_{d,x} -d + (1 - \sigma) \frac{\beta}{\mu} d + \sigma (\delta^N \varepsilon u(x) + D_0(d, N))$$

subject to

$$\begin{aligned} \frac{x}{1 - \tau} \frac{\mu}{\beta} &= d(1 - P(d, N)) \\ 1 - P(d, N) &= \frac{(N + 1)}{\sqrt{\Delta}} \\ d &\geq R(N + 1)N \end{aligned}$$

Note that the participation constraint for the seller is binding again.

Some preliminaries first.

$$\begin{aligned} \frac{\partial D_0(d, N)}{\partial d} &= \frac{\beta}{\mu} P(d, N) \\ \frac{\partial P(d, N)}{\partial d} &= \frac{1}{2R\Delta} (1 - P(d, N)) \\ \frac{\partial d(1 - P(d, N))}{\partial d} &= \left(1 - \frac{d}{2R\Delta}\right) (1 - P(d, N)) \end{aligned}$$

We look next at the function $d(1 - P(d, N))$. First, note that this expression is only valid when $d \geq \bar{d}$. Its minimum is achieved at \bar{d} . It is strictly increasing in d and strictly concave.

Differentiating the objective function w.r.t. to d , we obtain up to a factor of $\frac{\beta}{\mu}$

$$-i + \sigma(1 - P(d, N)) \left(\delta^N \varepsilon u'(x) (1 - \tau) \left(1 - \frac{d}{2R\Delta}\right) - 1 \right)$$

Case 1:

Suppose now that for a given N we have $x^* < \bar{x}$, that is at the best DS-proof contract the constraint is not binding.

Since $(1 - P(d, N)) < 1$ and $(1 - d/(2R\Delta)) < 1$, we have immediately that the objective function is decreasing in d . Hence, the best DS contract has $d = \bar{d}$. But this is worse than (x^*, d^*) .

Case 2:

Suppose now that for a given N we have $\bar{x} < x^*$ so that the constraint is binding for the optimal DS-proof contract.

Note first that the objective function is strictly concave. This implies that there is a unique maximizer and – by the previous argument – the solution needs to satisfy $\hat{x} \in [\bar{x}, x^*)$.

A sufficient condition is thus that the objective function is decreasing at \bar{x} . The first-order condition at \bar{x} is given by

$$-i + \sigma \left(\delta^N \varepsilon u'(\bar{x})(1 - \tau) \left(1 - \frac{1}{2} \frac{N}{N+1} \right) - 1 \right) \quad (\text{A.12})$$

Finally, note that this equation is strictly decreasing in N as $u'(\bar{x})$ is decreasing in N . Hence, a sufficient condition is that the objective function is decreasing at $N = 1$ and \bar{d} or, equivalently, that

$$\delta \varepsilon u'[(1 - \tau)(\beta/\mu)2R](1 - \tau) \frac{3}{4} \leq \frac{i + \sigma}{\sigma}.$$

So we can conclude that DS is never optimal when

$$\delta \varepsilon_{\max} u'(\bar{x}(1))(1 - \tau) \frac{3}{4} - 1 < \frac{i}{\sigma}$$

where $\bar{x}(1) = (1 - \tau)(\beta/\mu)2R$. ■

How to interpret condition (A.11)? Note that in general, the marginal value of an additional unit of money balances is (proportional to)

$$-i + (1 - P)\sigma[\delta^N \varepsilon u'(x)(1 - \tau)\mathcal{E}(x) - 1]$$

where

$$\begin{aligned} \mathcal{E}(x) &= \frac{\partial x}{\partial d} \frac{d}{x} \\ &= \frac{\partial}{\partial d} [d(1 - P(d, N))] \frac{1}{d[1 - P(d, N)]} \\ &= \begin{cases} 1 & , \text{ if } x < \bar{x} \\ 1 - \frac{d}{2R\Delta} & , \text{ if } x \geq \bar{x} \end{cases} \end{aligned}$$

is the elasticity of consumption with respect to money balances. When the incentive constraint is not binding, $\mathcal{E} = 1$. When it is binding, $\mathcal{E} < 1$. Define \bar{x}_N as the maximum DS-proof quantity given N . Evaluating at \bar{x}_N ,

$$\mathcal{E}(\bar{x}_N) = 1 - \frac{N}{2(1 + N)}$$

which is a decreasing function with its maximum equals to $3/4$ when $N = 1$. The idea is that when the incentive constraint is binding, a further increase of the trade size will raise the buyer's incentive to double spend after trade, hence lowering the effective value of a marginal dollar.

This condition tends to be satisfied when (i) the cost of bringing the extra money balances is high (i is high), (ii) the probability of spending that extra balances is low (σ is low), and (iii) the utility gain from spending that balances is low ($\delta, \varepsilon_{\max}$ are low or \bar{x} is high).

A.5 Proof of Proposition 7

We aim to show that a DS-proof cryptocurrency equilibrium exists for a sufficiently large B .

We will proceed in two steps. First, we show the existence under the assumption that only DS-proofs contracts are offered. Second, we will show that condition (A.11) is satisfied so that the derived equilibrium is DS-proof.

Step 1

Suppose condition (A.11) is satisfied. Define the correspondence

$$\varphi(R) = B \left(\frac{\mu - 1 + \sigma\tau}{\bar{N} + 1} \right) \int z(\varepsilon, R) dF(\varepsilon)$$

where $\int z(\varepsilon, R) dF(\varepsilon)$ is aggregate money demand. The correspondence φ expresses total rewards for mining as a function of aggregate money demand which itself is a function of R through the double-spending constraint.

We need to show that the correspondence φ has a fixed point for some R . Given this R , the money demand will determine the optimal offer $(x(\varepsilon), d(\varepsilon), N(\varepsilon))$ for all ε , where $d(\varepsilon) = z(\varepsilon)$, so that we have an equilibrium.

We will use Kakutani's fixed point theorem. This implies that

(i) we need to restrict the mapping φ to be from a non-empty, convex and compact set

and

(ii) φ is non-empty, upper hemicontinuous, convex-valued and closed-valued.

To do so, it is useful to define $z_n(\varepsilon; R)$ which is the optimal money demand given R conditional on choosing confirmation lag $n \in \{1, \dots, \bar{N}\}$. The functions $z_n(\varepsilon; R)$ are continuous by the Theorem of the Maximum and weakly increasing in R .

Lemma A.2. *Fix R_{\min} arbitrarily close to 0. There exists a sufficiently large $B(R_{\min})$ such that $\varphi(R) \geq R_{\min}$ for all $R \geq R_{\min}$.*

Proof. Choose R_{\min} small. For any $R > 0$, we have that $z^*(R; \varepsilon) > 0$ for any ε given our assumptions on the utility function. Since $z_n(\varepsilon; R)$ is weakly increasing in R , we can find a B sufficiently large

for any R_{\min} close to 0 such that

$$B \left(\frac{\mu - 1 + \sigma\tau}{\bar{N} + 1} \right) \int z(\varepsilon; R) dF(\varepsilon) \geq B \left(\frac{\mu - 1 + \sigma\tau}{\bar{N} + 1} \right) \int \min_n z_n(\varepsilon; R) dF(\varepsilon) > R_{\min} \quad (\text{A.13})$$

for all $R \geq R_{\min}$. □

Lemma A.3. *Given $B(R_{\min})$, there exists R_{\max} such that $\varphi(R_{\max}) \leq R_{\max}$.*

Proof. The surplus from bringing z balances in the market is bounded by

$$\beta\sigma(\varepsilon u(x^*) - x^*) - (1 - \beta/\mu)z \quad (\text{A.14})$$

where x^* satisfies $\varepsilon u'(x^*) = 1$. Hence, there is an upper bound $\bar{z}(\varepsilon)$ on money demand that is independent of R . Thus, for all R , we have that

$$\varphi(R) \leq R_{\max} = B(R_{\min}) \left(\frac{\mu - 1 + \sigma\tau}{\bar{N} + 1} \right) \int \bar{z}(\varepsilon) dF(\varepsilon) \quad (\text{A.15})$$

□

These two lemmata allow us to restrict the correspondence φ to a compact interval $[R_{\min}, R_{\max}]$ where $\varphi(R) : [R_{\min}, R_{\max}] \rightarrow [R_{\min}, R_{\max}]$. We turn next to the properties of the correspondence φ .

Lemma A.4. *The money demand correspondence $z(R; \varepsilon)$ is non-empty, compact-valued and u.h.c.*

Proof. Take the optimal money demand functions conditional on n and form the constraint correspondence

$$\mathcal{Z} = \bigcup_{n=1}^{\bar{N}} z_n(R; \varepsilon). \quad (\text{A.16})$$

Note that this correspondence is both u.h.c. and l.h.c, hence continuous. Furthermore, it is compact-valued as it is comprised of a finite number of values.

Let $V_n(R)$ be the value associated utility with $z_n(R; \varepsilon)$. Then, the optimal money demand is given by

$$\max_n V_n(R) \quad (\text{A.17})$$

subject to

$$z_n(R) \in \mathcal{Z} \quad (\text{A.18})$$

By the Theorem of the Maximum, we have that the correspondence of the maximizer $z(R; \varepsilon)$ is non-empty and u.h.c. Since it comprises only a finite number of elements, it is also compact-valued. \square

It follows immediately that the correspondence $\varphi(R)$ is also non-empty, compact-valued and convex-valued (see Aumann (1965), Theorem 1,2 and 4). Furthermore, $\varphi(R)$ is u.h.c (see Aumann (1965), Corollary 5.2). By Kakutani's Fixed Point theorem, there exists $R^* \in [R_{\min}, R_{\max}]$ such that

$$R^* \in \varphi(R^*)$$

which completes the proof.

Step 2

To ensure that only DS-proof contracts are offered, notice that condition (A.11)

$$\delta \varepsilon_{\max} u' \left[\frac{2\beta(1-\tau)}{\mu} R \right] (1-\tau) \frac{3}{4} - 1 < \frac{\mu - \beta}{\sigma\beta}$$

defines a threshold \bar{R}_{\min} such that the condition is satisfied for all $R > \bar{R}_{\min}$. By setting $R_{\min} = \bar{R}_{\min}$ in step 1, we can ensure that the equilibrium determined by the fixed point is DS-proof.

A.6 Proof of Proposition 8

To establish the result, we will look at policy changes that leave the budget constraint and the DS constraints unchanged. We first define $x^*(n)$ as the optimal quantity chosen by the buyer for any given n , that is

$$x^*(n) = \arg \max_x -\frac{x}{1-\tau} \frac{\mu}{\beta} + (1-\sigma) \frac{\beta}{\mu} \frac{x}{1-\tau} \frac{\mu}{\beta} + \sigma \delta^n \varepsilon u(x)$$

subject to

$$\frac{x}{1-\tau} \frac{\mu}{\beta} \leq R(n+1)n. \quad (\text{A.19})$$

The FOC implies that

$$\sigma \delta^n \varepsilon u'(x) \geq \frac{1}{1-\tau} \left(\frac{\mu}{\beta} - (1-\sigma) \right) \quad (\text{A.20})$$

with equality when the DS constraint (A.19) is non-binding. We first need to establish two preliminary results.

Lemma A.5. *Let N^* be the optimal choice of the settlement lag. Then, the DS constraint must be binding for any $n < N^*$.*

Proof. Suppose to the contrary that the DS constraint is not binding for some $n < N^*$. Denoting the optimal choice conditional on n by $x^*(n)$, we then have that

$$\sigma \delta^{N^*} \varepsilon u'(x^*(N^*)) \geq \sigma \delta^n \varepsilon u'(x^*(n)) = \left(\frac{\mu}{\beta} - (1-\sigma) \right) \frac{1}{1-\tau} \quad (\text{A.21})$$

with strict inequality if the DS constraint is binding at N^* .

Since $n < N^*$, it must be the case that

$$x^*(n) > x^*(N^*). \quad (\text{A.22})$$

But this is a contradiction, since consuming earlier and a greater amount is preferred by the buyer. \square

We want to show that, as long as $\tau > 0$, one can construct a welfare-improving policy change by reducing τ . To do that, we can slightly rewrite the Ramsey problem by eliminating the reward function to obtain for the DS constraint

$$\frac{x}{\beta} \leq B \left(\frac{(N+1)N}{\bar{N}+1} \right) (\mu - 1 + \sigma\tau) \int x(\varepsilon) dF_\varepsilon(\varepsilon) \quad (\text{A.23})$$

and for the objective function

$$B \int \sigma \left(\delta^{N(\varepsilon)} \varepsilon u(x(\varepsilon)) - x(\varepsilon) \right) dF_\varepsilon(\varepsilon) - B \left(\frac{\mu - 1 + \sigma\tau}{1 - \tau} \right) \int x(\varepsilon) dF_\varepsilon(\varepsilon) \quad (\text{A.24})$$

Note that a policy change that leaves $B(\mu - 1 + \sigma\tau) \int x(\varepsilon) dF_\varepsilon(\varepsilon)$ constant will not affect the DS constraints for any ε and N . We will now examine such a policy.

Lemma A.6. *Consider any policy change that leaves the DS constraints unchanged and satisfies $d\mu/d\tau \geq -\sigma$ with $d\tau < 0$. A buyer will either leave his original choice $x^*(N^*)$ unchanged or choose $x^*(n) > x^*(N^*)$ with $n \geq N^*$.*

Proof. For any given n , this policy change will induce the buyer to increase consumption if the DS constraint is not binding. To see that, differentiate the right-hand side of the first-order condition (A.20) of the buyer's problem to obtain

$$\begin{aligned} \frac{d \left(\frac{\mu/\beta - (1-\sigma)}{1-\tau} \right)}{d\tau} &= \frac{1}{\beta} \frac{1}{(1-\tau)^2} \left[\frac{\partial \mu}{\partial \tau} (1-\tau) + \mu - \beta(1-\sigma) \right] \\ &= \frac{1}{\beta} \frac{1}{(1-\tau)^2} [\mu - 1 + \sigma\tau + (1-\sigma)(1-\beta)] > 0 \end{aligned}$$

since $\mu - 1 + \sigma\tau > 0$. Therefore, for any given n , the choice of $x^*(n)$ is weakly increasing in the policy change, with the choice stays unchanged when the DS is binding.

Furthermore, by the previous lemma, any choice with $n < N^*$ or $x(N^*) \leq x^*(N^*)$ is dominated by the original choice since the DS constraints remain binding.

Finally, if the buyer chooses any $n > N^*$, it must be the case that $x^*(n) \geq x^*(N^*)$. To see this, suppose to the contrary that $x^*(n) < x^*(N^*)$. The choice $x^*(n)$ is then a feasible choice with N^* as well. Since

$$\delta^N u(x^*(n)) - x^*(n) \geq \delta^n u(x^*(n)) - x^*(n) \geq \delta^N u(x^*(N^*)) - x^*(N^*), \quad (\text{A.25})$$

$(x^*(n), n)$ is not optimal. A contradiction. This completes the proof. \square

Consider any policy with $\tau > 0$. Change this policy according to $d\tau < 0$ and $d\mu/d\tau \geq -\sigma$ such that the DS constraints do not change for any ε and N . Note that this needs to take into account any changes in total consumption $\int x(\varepsilon) dF_\varepsilon(\varepsilon)$ and is feasible, since by the lemmata above any

buyer will only increase his consumption $x(\varepsilon)$, but never lower it. In particular, if $\int x(\varepsilon)dF_\varepsilon(\varepsilon)$ stays unchanged, then $d\mu = -\sigma d\tau$. Otherwise, $d\mu < -\sigma d\tau$ since $d\tau < 0$.

Neglecting the reward costs, from the first term of the objective function the planner prefers a consumption allocation $(x(\varepsilon), N(\varepsilon))$ whenever the buyer does. In addition, the reward costs, which is the second term of the objective function, must decline since $B(\mu - 1 + \sigma\tau) \int x(\varepsilon)dF_\varepsilon(\varepsilon)$ remains constant, but $d\tau < 0$. Hence, the policy change increases the value of the objective function of the planner. This completes the proof, since one can always change policies by lowering τ leaving the DS constraints unchanged until $\tau = 0$.

B Appendix [FOR ONLINE PUBLICATION ONLY] – A formal description of the blockchain

Aggregate State

Let $m_t^D(i) \geq 0$ denote the cryptocurrency balance held by agent i in the period t day market. We then use $\mathcal{S}_t^D = \{m_t^D(i)\}$ to denote the entire public record of the holdings of these balances, called the (*aggregate*) *state*. Similarly, $m_{t,n}^N(i) \geq 0$ and $\mathcal{S}_{t,n}^N$ denote the balances and the state at the beginning of the n th trading session of the period t night market. The economy starts with a given initial state \mathcal{S}_0^D .

Payments

We use $\Delta_t^D(i, j)$ and $\Delta_{t,n}^N(i, j)$ to denote respectively day and night transfers of balances from agent i to agent j and call these transfers *payments*. A day payment is *feasible* if

$$\begin{aligned}\Delta_t^D(i, j) &\geq 0, \\ m_t^D(i) &\geq \sum_j \Delta_t^D(i, j).\end{aligned}$$

Similarly, a night payment is *feasible* if

$$\begin{aligned}\Delta_{t,n}^N(i, j) &\geq 0, \\ m_{t,n}^N(i) &\geq \sum_j \Delta_{t,n}^N(i, j).\end{aligned}$$

An agent can pay positive amounts to others and the total payments are bounded by the balances one has accumulated.³² An agent's balances accumulate over time as a result of net payment inflows. Therefore, given any payments, the state is updated in the two markets according to

$$m_{t,0}^N(i) = m_t^D(i) + \sum_j \Delta_t^D(j, i) - \Delta_t^D(i, j) + T_t(i), \quad (\text{B.1})$$

$$m_{t,n}^N(i) = m_{t,n-1}^N(i) + \sum_j \Delta_{t,n-1}^N(j, i) - \Delta_{t,n-1}^N(i, j), \text{ for } n = 1, \dots, \bar{N} \quad (\text{B.2})$$

$$m_{t+1}^D(i) = m_{t,\bar{N}}^N(i) + \sum_j \Delta_{t,\bar{N}}^N(j, i) - \Delta_{t,\bar{N}}^N(i, j) \quad (\text{B.3})$$

³²Through cryptography, the authenticity of payments is protected by digital signatures corresponding to the sending addresses. As a result, only the owner of the digital signature can transfer balances to another address. This gives rise to the non-negativity and the cash-in-advance constraints.

where $T(i)$ denotes the potential transfers of new balances to agent i in the day market by the cryptocurrency system.

Blockchain

Due to public monitoring, feasible payments during the day automatically update the aggregate state according to the rule (B.1). The new state at the start of the night market is thus given by

$$\mathcal{S}_{t,0}^N = \Psi_0^N(\mathcal{S}_t^D, \mathcal{B}_t^D) \quad (\text{B.4})$$

where $\mathcal{B}_t^D = \{\Delta_t^D(i, j)\}$ is the entire set of day transfers and is called a *block*. The update function Ψ_0^N is defined according to (B.1).

Payments in the night market, however, enter the state through a process we call *mining*. When agent i makes a payment to agent j in the night, he needs to send out an instruction for a feasible payment $\Delta_{t,n}^N(i, j)$ to a pool of miners who compete to update the state with a new block of feasible payments in session n of the night market. The set of feasible payment instructions $\mathcal{B}_{t,n}^N = \{\Delta_{t,n}^N(i, j)\}$ is the n th block of period t payments in the night market.

A sequence of blocks $\{\mathcal{B}_t^D, \{\mathcal{B}_{t,n}^N\}_{n=0}^{\bar{N}}\}_{t=0}^T$ iteratively generates a sequence of states $\{\mathcal{S}_t^D, \{\mathcal{S}_{t,n}^N\}_{n=0}^{\bar{N}}\}_{t=0}^{T+1}$ according to

$$\begin{aligned} \mathcal{S}_{t,0}^N &= \Psi_0^N(\mathcal{S}_t^D, \mathcal{B}_t^D) \\ \mathcal{S}_{t,n}^N &= \Psi_n^N(\mathcal{S}_{t,n-1}^N, \mathcal{B}_{t,n-1}^N), \text{ for } n = 1, \dots, \bar{N} \\ \mathcal{S}_{t+1}^D &= \Psi^D(\mathcal{S}_{t,\bar{N}}^N, \mathcal{B}_{t,\bar{N}}^N), \end{aligned}$$

where the update functions Ψ_n^N, Ψ^D are defined by (B.1)-(B.3). We call the sequence of blocks $\mathcal{B}_T = \{\mathcal{B}_t^D, \{\mathcal{B}_{t,n}^N\}_{n=0}^{\bar{N}}\}_{t=0}^T$ a *blockchain*. Determined by the process of mining, one specific blockchain is used to construct the public state and can be observed by everyone in the economy at all times.

C Appendix [FOR ONLINE PUBLICATION ONLY] – Proof of Stake

We assume that the right for updating of the blockchain is allocated randomly across people. The probability that one can update the chain is proportional to the unused balances in the night market.

This is identical to pledging balances that are not used in transactions in order to gain a right to update the chain. For simplicity, we assume that people can pledge balances after they learn whether they have a match or not.

Denote the balances chosen in the day market for transaction purposes as d and the additional amount of balances chosen for having a stake in the voting procedure as z' .

Whenever, $z' > 0$, we have that there is a strictly positive probability that there is a double spend, since the buyer in a transaction will propose a block that undoes the payment whenever he is allowed doing so.

Fix N . The problem is given by

$$\max_{d, z'} - (z' + d) + \sigma \left[\delta^N \epsilon u(x) + \frac{\beta}{\mu} \left(z' + \left(\frac{z'}{(1-\sigma)Z} \right)^{N+1} d + \left(\frac{z'}{(1-\sigma)Z} \right) R(\bar{N} + 1) \right) \right] + (1-\sigma) \frac{\beta}{\mu} \left[z' + d + \left(\frac{z' + d}{(1-\sigma)Z} \right) R(\bar{N} + 1) \right]$$

When having a transaction, a buyer is only allowed to update the chain if $z' > 0$. Given N , he needs to win the right to update $N + 1$ in order to double spend. If a buyer has no transaction, he will pledge all his balances as a stake. Note that participation of the seller requires

$$x = \frac{\beta}{\mu} d (1 - \tau) \left(1 - \left(\frac{\mathbb{E}(z')}{Z} \right)^{N+1} \right).$$

Notice that the buyer's choice of z' is not observed by the seller, so x depends on seller's expectation $\mathbb{E}(z')$ instead.

Lemma C.1. *The objective function is strictly convex in z' for all $N \geq 1$ and linear in z' for $N = 0$.*

Proof. Differentiating the objective function, we obtain

$$-i + \left(\frac{1}{1-\sigma} \right) (\sigma\tau + (\mu - 1)) + \left(\frac{\sigma}{1-\sigma} \right) (N + 1) \left(\frac{z'}{(1-\sigma)Z} \right)^N \left(\frac{d}{Z} \right) \quad (\text{C.1})$$

Differentiating again, we obtain

$$\left(\frac{\sigma}{1-\sigma}\right)(N+1)N\left(\frac{z'}{(1-\sigma)Z}\right)^{N-1}\left(\frac{d}{Z}\right) > 0 \quad (\text{C.2})$$

or 0 if $N = 0$. □

We are only looking for a sufficient condition so that we can support the policy $N = 0$, $\mu = 1$ and $\tau = 0$. Interestingly, for $N = 0$, we only need to check the slope at $z' = 0$. This case implies settling of the trade on the spot without delay. The idea is that the buyer has zero probability of updating the chain and, hence, can never double spend.³³

Importantly, setting $\mu = 1$ and $\tau = 0$ implies that people have nothing at stake when updating the chain. They do not receive a reward and there are no penalties. In turn, setting positive rewards increase the incentives for traders to acquire more balances.

Proposition C.2. *Set $N = 0$, $\mu = 1$ and $\tau = 0$. Then the PoS protocol avoids double-spending for B sufficiently large and σ sufficiently small.*

Proof. For the policy parameters, (C.1) implies that the objective function is decreasing in z' if

$$\left(\frac{\sigma}{1-\sigma}\right)\left(\frac{d}{Z}\right) \leq i = \frac{1}{\beta} - 1$$

which is satisfied for B sufficiently large and σ sufficiently small since d is bounded. □

Note that, even though a system based on PoS can avoid the double spending attacks studied above, it may still be subject to “long range attacks”. For example, if all coins are owned initially by one issuer, then he can easily propose a long fork starting from the initial block to claim all existing balances in any future dates. By assuming that the blockchain cannot be rewritten after the end of the night period, our model has ruled out these types of long-range attacks. This captures the feature of checkpoints introduced in Bitcoin.

Interestingly, these are exactly the concerns and proposals raised by practitioners. For example, BitFury Group (2015) pointed out that “In PoW-type systems, this [long range] attack is prevented

³³We might need to look at $N = 1$ since the buyer can still procast a different transaction so that the seller at least needs to see one transaction in the chain to hand over the good. This will, however, not influence the analysis, since the condition for $z' = 0$ being optimal is identical except for that we do not need to check the slope, but the value of the function at $z' = (1 - \sigma)Z$. Or we could require a negative derivative at $z' = (1 - \sigma)Z$.

by the enormous amount of computational power needed to build the blockchain from scratch; however, this task is within the realm of possibility with proof of stake. As the attacker is able to move coins freely in the blockchain he is building, he has a much higher dimensionality of the search space ... To prevent a long range attack, the protocol can specify the maximum allowed depth of a branching point.”

Finally, note that there are issues with PoS other than double spending. Consensus could be a problem especially if there is a nothing-at-stake problem or if we need to ensure that new parties can decide unambiguously between competing chains. Our model is not developed to address these issues.