# Blockchain Economics*

Joseph Abadi and Markus Brunnermeier†

June 18, 2018

**Abstract**

When is record-keeping better arranged through distributed ledger technology (DLT) than through a traditional centralized intermediary? The ideal qualities of any record-keeping system are (i) correctness, (ii) decentralization, and (iii) cost efficiency. We point out a *Blockchain Trilemma*: no ledger can satisfy all three properties simultaneously. A centralized ledger writer extracts rents due to its monopoly on the ledger. Its franchise value dynamically incentivizes honest reporting. Decentralized ledgers provide static incentives for honesty through computationally expensive Proof-of-Work algorithms but eliminate rents through "fork competition." Portability of information between "forks" and competition among miners fosters competition *among decentralized ledgers* that is fiercer than traditional competition. However, fork competition can engender instability and miscoordination. While blockchains can keep track of ownership transfers, enforcement of possession rights is still needed in many blockchain applications.

**Keywords:** DLT, Blockchain, Digital Economics, Platform Economics, Cryptocurrencies, "Fork Competition", Contestable Markets

*We are grateful for helpful comments from Stephen Morris, Zhiguo He, and seminar participants at Princeton, the NYU Intermediation Conference, and the BIS.

†Abadi: Department of Economics, Princeton University, jaabadi@princeton.edu, Brunnermeier: Department of Economics, Princeton University, markus@princeton.edu

# 1 Introduction

Traditionally, records have been maintained by centralized entities. Distributed Ledger Technology (DLT) has provided us with a radical alternative to record information. DLT has the potential to be as groundbreaking as the invention of double-entry bookkeeping in fourteenth-century Italy. It could revolutionize record-keeping of financial transactions and ownership data.

Blockchains are a particular type of distributed ledger written by decentralized, usually anonymous groups of agents rather than known centralized parties. Consensus is attained by making the ledger publicly viewable and verifiable. Ideally, a ledger should (i) record all information correctly and do so (ii) in a cost efficient and (iii) fully decentralized manner to avoid any concentration of power. In this paper we point out a "Blockchain Trilemma": it is impossible for any ledger to fully satisfy the three properties shown in Figure 1 simultaneously.

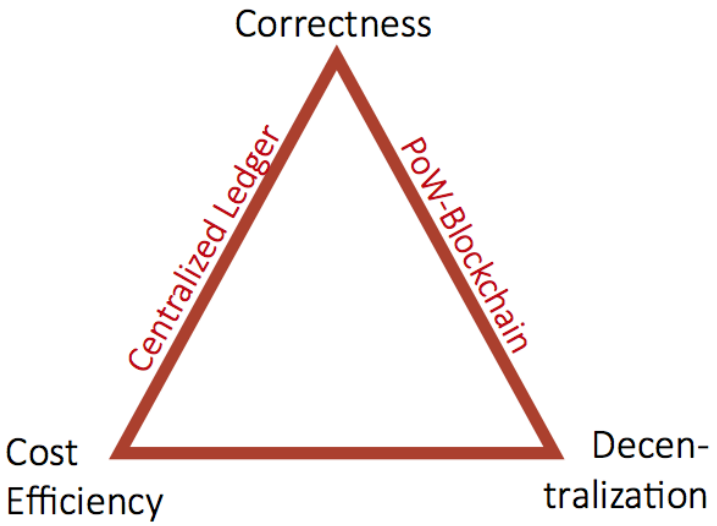


Figure 1: The Blockchain Trilemma.

Traditional ledgers, managed by a single centralized intermediary, forgo the desired feature of decentralization. The correctness of the ledger is maintained by limiting competition. A centralized ledger writer is incentivized to report honestly because he does not wish to jeopardize his future profits and franchise value. That is, a centralized ledger writer is dynamically incentivized. In contrast, decentralized ledgers promote competition but entail real inefficiencies. Competition completely erodes writers' future profits and

franchise values. Consequently, dynamic incentivization of decentralized ledger writers is impossible. The ledger's correctness must rely on a mechanism that provides purely static incentives.

Blockchains allow two forms of competition that lead to two distinct inefficiencies. (i) First, there is free entry of ledger writers. As anybody can become a ledger writer (or miner) on a public blockchain, a consensus mechanism is needed to determine the true history written on the ledger (from possibly conflicting reports). Applying a majority rule is complicated by the fact that individual entities can masquerade as a large number of entities for free, subverting the democratic nature of the distributed ledger. To limit this problem and ensure honest record-keeping, ledger writers must typically perform computationally expensive tasks in order to record information and validate others' reports. The cost of writing on the ledger gives writers static incentives not to report dishonestly. (ii) Second, information on the existing ledger is made portable to possibly competing ledgers via "fork competition". A proposer of a new ledger can "fork off" an existing blockchain by establishing different rules while retaining all the information contained in the original blockchain. Fork competition erodes the rents of a ledger monopolist, but also comes at a cost: too many competing blockchains may coexist. The community of users/readers may be split among too many different ledgers (or cryptocurrencies) and fail to fully exploit positive network externalities. This entails a true efficiency loss, above and beyond the redistributive rent extraction associated with a monopolistic ledger writer or the waste of computational resources resulting from free entry. Finally, current technology limits the scalability of blockchain technology, a third cost.

We emphasize that fork competition eliminates inertia in the adoption of new, competing ledgers. In a traditional setting, ledger users are anchored to an incumbent ledger by the centralized intermediary's monopoly on the information contained in the ledger. Those with high stakes in the existing ledger are reluctant to switch to a competitor. Network externalities amplify this informational anchor, making even those with low stakes in the existing ledger unwilling to switch. When network externalities are strong, the market ceases to be contestable– even with free entry of competing ledgers, the incumbent's advantage is so great that it is able to extract full surplus from users. Fork competition eliminates the anchor on the established ledger due to the portability of information. Network externalities then play no role in amplifying inertia, and the market is *always* contestable: competing forks of the blockchain are at no disadvantage whatsoever against the established ledger.

In addition to the polar cases of completely centralized traditional ledgers and completely decentralized blockchains, there is a third type of ledger called a "permissioned" blockchain that shows promise in many applications. The writers of a permissioned blockchain are known agents rather than anonymous miners, so Proof-of-Work is unnecessary. Permissioned blockchains then seemingly break the Trilemma: they allow for fork competition, like anonymous blockchains, but completely eliminate the waste of resources. We show that the impediments to entry of writers on a permissioned blockchain sub-

stantially weaken fork competition. Permissioned writers have franchise values and therefore can collude to prevent competing forks from surviving, whereas dynamic punishment schemes that sustain collusion are impossible when there is free entry of writers.

Finally, we informally make the important point that while blockchains guarantee transfers of *ownership*, some sort of enforcement is required to ensure transfers of *possession*. For example, in a housing market the owner of the house is the person whose name is on the deed, but the possessor of the house is the person who resides in it. The buyer of the deed needs to be certain that once she holds the deed, her ownership of the house will be enforced. In the stock market, the purchaser of a share has ownership of future dividends but not necessarily possession, since the delivery of dividends needs to be enforced. Broadly, blockchains can record obligations. Punishing those who default on their obligations is another matter. While it is difficult to provide static incentives for blockchain writers to impose discipline on users of the ledger, centralized intermediaries' incentives can be appropriately aligned: if a centralized intermediary fails to guarantee transfers of possession, the ledger's users can abandon the ledger, destroying the intermediary's franchise value.

Blockchains have applications that reach far beyond the realm of cryptocurrencies and tokens. For instance, blockchains could be used in the fintech space to track consumers' transaction and credit histories. Permissioned blockchains have also been suggested as a tool to manage supply chains and track the delivery of items in real time. There are several potential applications of blockchains that, if pursued, will require enforcement by intermediaries or legal entities. Banks could use blockchains to track interbank loans or manage their clients' collateral, both of which require mechanisms to ensure debtors will repay their creditors. Governments may also turn to blockchains to maintain land registries, which could be useful in developing countries where the primary institutional friction is overly bureaucratic record-keeping processes, but seems likely to be unhelpful when the issue is instead that the government enforces ownership selectively.

**Related Literature**. The paper most closely related to ours is Biais et al. (2017), who study the stability of a blockchain-based system. It shows that while the strategy of mining the longest chain proposed by Nakamoto (2008) is in fact an equilibrium, there are other equilibria in which the blockchain forks, as observed empirically. In that model, forks occur for several reasons and are interpreted as causing instability. Writers' payoffs when forking depend exogenously on the number of writers who choose a given branch of the fork. In our model, writers' payoffs are instead determined by readers' preferences, which puts more discipline on exactly how and when a fork may occur. Cong and He (2017) focus mostly on the issue of how ledger transparency leads to a greater scope for collusion between users of the system. In contrast, we consider collusion between writers of the blockchain rather than users and show that collusion can occur only when entry of writers is constrained.

Some of the recent literature on blockchains in economics focuses on the security and the costs of the system. Easley, O'Hara, and Basu (2017) use a game-theoretic framework to analyze the emergence of transaction fees in Bitcoin and the implications of these fees

for mining costs. The R&D race between Bitcoin mining pools is described in Gans, Ma, and Tourky (2018), who argue that regulation of Bitcoin mining would reduce the overall costs of the system and improve welfare. Huberman, Moallemi, and Leshno (2017) study transaction fees in Bitcoin and conclude that the blockchain market structure completely eliminates the rents that a monopolist would extract despite the fact that only one miner processes transactions at a time. We depart from these analyses by endogenizing the mechanism used by the blockchain: in our model, users of the system essentially choose between competing mechanisms on different branches of a blockchain fork. The cost of implementing a given mechanism is pinned down by the free entry condition.

Our framework uses a global game of the type pioneered by Carlsson and van Damme (1993) in order to select a unique equilibrium. Rather than review the massive literature on global games here, we refer the reader to Morris and Shin (2001) for an extensive and general analysis of the global games framework. We use techniques from the more recent literature on global games with non-Gaussian private values pioneered by Sakovics and Steiner (2012) and advanced by Drozd and Serrano-Padial (2017). Our work is also related to the recent literature on the importance of network externalities in blockchain payment systems. Sockin and Xiong (2018) show that strategic complementarities in cryptocurrency holdings lead to fragile equilibria with different cryptocurrency prices. Cong, Li, and Wang (2018) argue that expectations of growth in a blockchain's participation impact the current price of its native token. Our paper differs from these studies in that we analyze the importance of network externalities for arbitrary blockchains rather than just cryptocurrency blockchains and show that these externalities interact with the replicability of information on a blockchain in an important way.

We also relate to the literature on cryptocurrencies. Chiu and Koeppl (2017) develop a macroeconomic model in which the sizes of cryptocurrency transactions are capped by the possibility of a double-spend attack and derive optimal compensation schemes for writers. Schilling and Uhlig (2018) study cryptocurrency pricing in a monetary model and derive necessary conditions for speculation to occur in equilibrium. Pagnotta and Buraschi (2018) derive a pricing framework for cryptocurrencies that explicitly accounts for the interplay between demand for the currency and the cryptographic security provided by miners.

Recent computer science literature has studied blockchain security extensively. Most papers in computer science, such as Gervais et al. (2016), study how to defend against "double-spend" attacks or other types of attacks that could be undertaken by a single individual who holds control over a large portion of the network's computing power. The conclusion of studies in the computer science literature is that a large fraction of the blockchain writers must always play honestly in order for the network to be secure. In such models, writers are prevented from deviating by other writers who discipline them. Writers are implicitly prevented from colluding in any way. In contrast, we study a more general type of attack without explicitly referring to double-spending. Our model shows that the cost of operating a blockchain is intrinsically linked to the cost of preventing attacks, no matter what they may be. Furthermore, our model shows that the implicit assumption of

no collusion is unnecessary. The impossibility of dynamic collusion between writers on a blockchain is a characteristic that emerges naturally from the free entry condition.

Finally, our paper is related to the literature on optimal intermediation structures. Most notably, Diamond (1984) shows that when monitoring is costly, it is most efficient to use a single intermediary. In contrast, in our framework it is optimal to have several intermediaries because competition in writing on the ledger yields outcomes that are more desirable for the blockchain's users. In the computer science literature, Wüst and Gervais (2017) study the applicability of blockchain to several markets from an informal standpoint.

The rest of the paper is structured as follows. Section 2 discusses the basics of blockchain technology. In Section 3, we present the baseline model of a static choice between ledgers. We analyze a specific example where agents choose between two branches of a blockchain fork and another example in which agents choose between traditional ledgers in order to spell out the tradeoffs between decentralization and cost-efficiency. Section 4 extends the static model to a repeated setting and studies permissioned blockchain as well as the security features of traditional ledgers and blockchains. Section 5 discusses practical issues related to blockchain technology including some points that we do not address in our formal model, such as the transfer of physical assets on a blockchain. Section 6 concludes.

# 2 Blockchain Technology

In this section we outline how blockchains work and the distinguishing features of blockchains with anonymous writers.

## 2.1 What is a blockchain?

A blockchain is a ledger in which agents known as writers (or nodes) take turns recording information. This information could consist of payment histories, contracts outlining wagers between anonymous parties, or data on ownership of domain names, among other applications. As discussed later, there are many possible algorithms to select the current writer. The ledger consists of a tree of blocks that contains all the information recorded by writers starting from the first block, which is called the *genesis block*. Each branch of the tree corresponds to a chain leading back to the genesis block (hence the name "blockchain").

A chain of blocks leading back to the genesis summarizes a state. Readers and writers of the ledger must reach a consensus about which state is considered the valid state. Typically, the community coordinates on the longest chain of blocks as the valid state, as suggested in Nakamoto (2008). Each writer is periodically allowed to add a block to the tree. Writers usually extend only the consensus chain, and readers will act only in response to events on that chain. A writer's decision to extend a given chain can be seen as a signal that the writer accepts that chain as valid. Writers are rewarded for achieving consensus through readers' acceptance of the chain they extend. In general, writers accrue rewards

and transaction fees for each block added to the tree, so these rewards are realized only if those fees are on the consensus chain.

However, it is in principle possible for readers and writers to coordinate on a chain other than the longest one or even for different communities to coordinate on separate chains. A "hard fork" occurs when part (or all) of the community decides to change the rules governing the blockchain. To do so, they start their own blockchain that builds off of the old chain, but they ignore any writers who do not follow the new rules. Similarly, writers who use the old rules will ignore all writers who use the new ones, so the blockchain effectively forks and becomes two blockchains. The data contained in the original chain is included in both of the new blockchains, but neither blockchain uses data that was recorded on the other after the fork occurred. Hard forks will feature prominently in our model and will intensify competition between ledgers by allowing information from the original blockchain to be replicated on a competing ledger.

For example, in 2016 the Ethereum community split after a hack that stole $55 million from investors in a contract on that blockchain. Some Ethereum users argued that the currency should be returned to the investors, whereas others believed the blockchain should be immutable. The users who believed the currency should be returned ignored all blocks occurring after the hack and built their own chain on which the hack never occurred. After this point, both sides began ignoring the blocks built by the other side, and each part of the community considered only its own chain to be the valid chain.

On any blockchain, there are some rules that readers and writers tacitly agree to follow. These rules are written into the code distributed by the software developers for that blockchain. For example, cryptocurrency transactions are signed cryptographically by the sender of the transaction. Whenever blockchain writers receive a message to add a given transaction to a block, they can perform a cheap computation to verify that the sender properly signed the transaction. If the verification fails, the transaction is considered fraudulent. Writers who follow the rules will refuse to add any such transaction to a block. In general, blockchain security algorithms work so that it is inexpensive for writers to confirm that the rules are being followed. If a previous writer added fraudulent transactions to a block at the end of the longest chain, the consensus algorithm prescribed by Nakamoto (2008) specifies that all other writers should ignore that particular block and refuse to put other blocks on top of it.

Another example of rules that blockchain users agree to follow are the rules for writers' compensation. For instance, Bitcoin miners are awarded a certain number of coins for finding a block. All other writers must check that the miner who found the last block did not attempt to circumvent the blockchain's policies by minting more coins than what is allowed. In most of our analysis we will suppose that the network is sufficiently secure to ensure that the rules are followed. We focus on which rules for writer compensation emerge in equilibrium when there is scope for competition between ledgers. In an extension of our model, we examine how the rules are enforced in the first place.

An attack on a blockchain involves the addition of blocks that are somehow invalid.

Either the blocks contain outright fraudulent transactions, or they are added somewhere other than the end of the longest valid chain. It is clear that attackers stand to gain by adding fraudulent transactions to their blocks simply because such a strategy allows them to steal from others as long as other readers and writers go along with the attack, but these attacks are usually automatically detected by all users of the system. It is perhaps less obvious why an attacker would want to add valid blocks somewhere other than the end of the longest chain. The key observation is that this type of attack permits dishonest actors to reverse transactions or records written on the longest valid chain. If an attacker or group of attackers controls the majority of the computing power on the network, even if this group's chain of blocks begins behind the longest valid chain written on by others, eventually the length of the attackers' chain will exceed that of the other chain. At this point it becomes the longest valid chain. All writers (both the honest ones and the attackers) then write on the attackers' chain.

In cryptocurrency blockchains, this type of attack is commonly referred to as a double-spend attack. An attacker will spend some currency on the longest valid chain, wait to obtain the goods purchased, and then begin building an alternative chain on which the currency was never spent, absconding with both the goods and the money. Double-spends are by far the largest security concern of the cryptocurrency community. This type of attack is also possible when the blockchain in question handles assets other than currency. For example, a financial institution that loses money on a trade may wish to reverse the history of transactions including that trade. Our model extension embeds double spending, but it encompasses a broader class of attacks.

## 2.2 The Types of Blockchains

There are three main types of blockchains. In a private blockchain, a single centralized entity has complete control over what is written on the ledger. That is, there is only one writer. The readers in this situation could be the public, the entity's clients, or a regulator. Different groups may also have different types of read privileges on the ledger: for example, a regulator would likely need to see the entire ledger, whereas a client may be content to see only those transactions that are relevant to her. There is no need for identity management with a private blockchain, since only one entity is permitted to write on the ledger. Therefore, there are no computational costs and the system functions similarly to a privately maintained database that gives read privileges to outsiders. In this system, the writer is disciplined entirely by the readers, who may decide to punish the writer in some way when the writer changes the ledger's rules (or fee structure) or if they detect some sort of fraudulent activity. One way in which this sort of punishment could arise in reality is if an online platform like Amazon decides to raise subscription rates for vendors and vendors respond by switching to a competitor.

A permissioned blockchain is one in which the write privilege is granted not to one entity, but to a consortium of entities. These entities govern the policies of the blockchain and are

the only ones permitted to propagate and verify transactions. The read privilege may be granted to the public or kept private to some extent. The permissioned writers take turns adding blocks to the chain according to a predefined algorithm, so again costly identity management is unnecessary. The writers on a permissioned blockchain are disciplined by readers, just as in a private blockchain, but they are also disciplined by other writers. If one writer deviates and begins validating fraudulent ledger entries by including them in his block, other writers may ignore him and refuse to extend his chain. If a writer proposes a change of the blockchain's policies, other writers may prevent such a change by writing according to the existing policies.

The third and most common type of blockchain is a public blockchain. In a public blockchain, both the read and write privileges are completely unrestricted. Writers are disciplined exactly as in permissioned blockchains. All users of the network are anonymous. However, when writers are allowed to be anonymous, some sort of identity management is necessary. Otherwise, it would be possible for a small entity to pretend to be a large entity, allowing it to add blocks more often than others and hence giving it significant power over which chain of transactions is accepted as valid. This type of attack is known as a "Sybil attack." The typical approach to identity management is to force writers to prove they have accomplished a computationally difficult task before permitting them to write on the ledger. This method is known as Proof-of-Work (PoW) and is used by most major cryptocurrency blockchains, such as Bitcoin, Ethereum, and Litecoin. In order to incentivize writers to perform these expensive computations, they are usually rewarded with seignorage and transaction fees for each block added to the chain. The structure of a blockchain's rewards gives rise to the free entry condition for that particular blockchain. The costs of writers' rewards tend to be economically large. For example, the Bitcoin blockchain currently uses more electricity than Hungary.

# 3 Static Ledger Choice Model

In this section, we present a general model of ledger choice as a coordination game. Our objective is to be able to capture a variety of settings in which readers choose among competing ledgers with different rules or policies. Our leading example applies our model to study competition between two branches of a blockchain fork. We then contrast the model of two competing blockchains with a model in which two traditional ledgers compete. We also examine a hybrid model of competition between a traditional ledger and a blockchain, and in the next section we extend the model to a dynamic setting and analyze the differences between a permissionless blockchain and a permissioned blockchain. The specific examples of competition between different types of ledgers will illustrate the tradeoffs suggested by the Blockchain Trilemma.

We focus on the importance of coordination because many types of ledgers are useful only if they are widely used. For example, consumers will want to hold a fiat currency only

if it is accepted by most vendors. Another situation in which coordination is important is when the ledger contains information about user's creditworthiness (such as Alibaba's Sesame credit score system)– users will not have an incentive to build up their credit score if there are no lenders. Throughout, we will abstract from the specific details of the coordination motive and instead compare different settings by varying a parameter that governs the strength of network externalities.

There are three periods, $t = 0, 1, 2$. There is a set of agents $j \in \mathcal{M}$ known as writers. These agents correspond to those who maintain the ledger. For a cryptocurrency blockchain, these agents would be miners. For a traditional payments ledger, a single centralized intermediary (such as the Federal Reserve or a bank) is usually the sole writer. There is also a continuum of agents $i \in [0, 1]$ known as readers, who are users of the ledger. Finally, there are two agents known as proposers, $P_A$ and $P_B$. These proposers are responsible for choosing the rules under which the ledger operates. Software developers are the "proposers" for a blockchain. When a part of the community wants to fork the blockchain, a developer will write commonly accepted code that implements the desired changes to the rules. On the other hand, for a traditional ledger the proposer is also the writer. That is, the monopolist who runs the ledger also decides on the rules. In what follows, we will allow for the possibility that some writer $j \in M$ is also one of the proposers.

Each ledger $k \in \{A, B\}$ is associated with a fundamental parameter $L_k \in \mathcal{L}_k$ determining the revenues earned by writers. A simple way of thinking about $L_k$ is as an explicit fee charged to readers by the writer(s) of the ledger, but more broadly $L_k$ could be interpreted as an implicit fee. Such implicit fees could arise, for instance, if a monopolist who runs a ledger chooses to sell readers' data to an outside party. The fundamental parameter $L_k$ could also represent a goverment's choice of policy, such as inflation. For example, a government may wish to inflate away its debt, but doing so could be costly for people who hold the currency, who may then collectively decide to abandon the national currency altogether (as in Zimbabwe). Henceforth we will refer to $L_k$ as a fee for ease of exposition.

Readers and writers must both choose ledgers in which to participate. Readers will have homogeneous preferences for coordination on a given ledger as well as heterogeneous fundamental preferences for each ledger, as described below. Writers will choose a ledger $k$ and take an action $a_j \in \mathcal{A}(\pi_k)$ to write on the ledger (where the set of allowable actions may depend on the fraction of readers $\pi_k$ who participate on that ledger). In our applications, this action will generally correspond to the expenditure of computational resources to write on a blockchain, but at times it will also refer to actions taken in order to distort the contents of the ledger.

Readers are heterogeneous in their *fundamental* preferences for ledgers. Each reader is assigned a type

$$s_i = s_{i,A} - s_{i,B}$$

Here $s_{i,k}$ is meant to represent the *stake* that agent $i$ has in ledger $k$ and $\tau_{i,k}$ is a common value for ledger $k$. The stake that a reader has in a given ledger should be interpreted as

the amount of information pertaining to that reader that is encoded in the ledger. For any ledger that keeps track of asset holdings, a reader's stake is simply the set of assets held by that reader, with larger asset holdings being interpreted as a higher stake. However, a reader's stake does not necessarily have to represent the market value of some asset. A reader with a high stake may also be a consumer who has built up a high credit score or a financial institution with a complex set of contracts with other institutions. We denote the population distribution of stakes $s_i = s_{i,A} - s_{i,B}$ by $Q(s)$.

There is also a common value component in readers' preferences, $\tau = \tau_A - \tau_B$. When $\tau > 0$, the common value induces a preference for $A$ among all readers, and when $\tau < 0$, readers prefer $B$. We introduce incomplete information about the common value for equilibrium selection. Formally, we assume that each reader $i$ receives a signal $x_i = \tau + \sigma\eta_i$, where $\eta$ is uniformly distributed on $[-\frac{1}{2}, \frac{1}{2}]$. We typically work in the limit $\sigma \to 0$, so there is an arbitrarily small amount of noise in agents' signals.[1]Incomplete information about this value could be motivated by, for example, uncertainty about the properties of the ledger's technology. With incomplete information about $\tau$, readers' types become two-dimensional. An individual reader's type can be summarized by $\theta_i = (x_i, s_i)$.

Proposers choose the fundamental ledger parameters by choosing $L_k \in \mathcal{L}_k$ and the assignment of stakes to agents by choosing $s_k \in \mathcal{S}_k$. Formally, a mapping $S_k$ of stakes to agents is just a function $S_k : [0,1] \to \mathbb{R}$. Readers are informed about their stakes when they receive their types $s_i$. The proposer's choice of stakes is meant to capture the information encoded in the proposed ledger. When information on ledger $A$ can be replicated on ledger $B$, for example, there would be a set of stakes $s \in \mathcal{S}_A$ that the proposer of ledger $B$ could use as well, so $s \in \mathcal{S}_B$. However, when information on ledger $A$ cannot be replicated, there would be some $s \in \mathcal{S}_A$ such that $s \notin \mathcal{S}_B$. Broadly speaking, information can be replicated across two branches of a blockchain fork, since both branches share the same root blockchain. With a traditional ledger, on the other hand, the centralized intermediary who manages the ledger typically has a monopoly over the information it contains. One of our main results in our applications will be that replicability of information intensifies competition across ledgers– when information can be replicated on a competing ledger, readers no longer face the cost of losing their stakes when switching to a competitor's ledger.

The timing of the game is as follows:
**t=0**: Proposers $P_A$ and $P_B$ choose $(L_A, S_A)$ and $(L_B, S_B)$, respectively.
**t=1**: Readers first observe writers' choices and their own types $\theta_i$. They then choose a ledger $r(i) \in \{A, B\}$ in which to participate.
**t=2**: Writers choose a ledger $k \in F_j \subset \{A, B\}$ and take actions $a_j \in \mathcal{A}(\pi_k)$. Payoffs are realized.

Readers' preferences for each ledger depend on their types, the proportion of other readers who choose that ledger, the revenues (fees) collected by writers, and the actions

[1]When $\sigma \to 0$, agents' priors over $\tau$ become unimportant.

taken by writers at $t = 2$. The actions taken by writers at $t = 2$ may be important to readers for several reasons. If the action at $t = 2$ corresponds to the amount of computational power a writer contributes to a blockchain, readers may prefer ledgers that are more cryptographically secure in the sense that greater computing power is dedicated to it. When the action taken at $t = 2$ corresponds to a distortion of the ledger, readers will prefer ledgers that have not been distorted. Let $\pi_k$ be the proportion of readers who choose ledger $k$, and let $a_k = \{a_j\}_{w(j)=k}$ be the action taken by writers at $t = 2$. A reader who chooses ledger $k$ obtains utility $u(\tau, s_{i,k}, \pi_k, L_k, a_k)$. We assume that $u$ takes the form

$$u(\tau, s_{i,k}, \pi_k, L_k, a_k) = b_\theta(\tau + s_{i,k} - g(L_k) - h(a_k)) + b_\pi \pi_k$$

where $g$ is an increasing function and $b_\theta, b_\pi > 0$. That is, utility is linear in $\tau$, $s_{i,k}$, $g(L_k)$, and $\pi_k$ conditional on the action taken by writers. Linearity in $\tau$, $s_{i,k}$ is natural in this context and is the usual approach taken in the global games literature. Linearity in $\pi_k$ will be useful in deriving the properties of equilibria because it will ease the computation of expected utility across possible realizations of $\pi_k$. We also define

$$\Delta = u(\tau, s_{i,A}, \pi_A, L_A, a_A) - u(\tau, s_{i,B}, \pi_B, L_B, a_B)$$

to be the opportunity cost of choosing ledger $B$. When $a_A = a_B = a$, $\Delta$ takes the form

$$\Delta = b_\theta\big(\tau + s_i - (g(L_A) - g(L_B))\big) + b_\pi(2\pi_A - 1)^2$$

We will define $\hat{\pi}(\tau, s, a, L_a, L_B)$ to be the $\pi_A$ such that $\Delta = 0$ when a reader's type is $s$, the common value is $\tau$, writers take actions $a$, andthe fundamental parameters of the ledgers are $L_A, L_B$. We will usually suppress the dependence on $\tau, a, L_A, L_B$. According to this definition,

$$1 - \hat{\pi}(s) = \frac{1}{2} + \kappa^{-1}\big(\tau + s - (g(L_A) + h(a_A) - g(L_B) - h(a_B))\big)$$

where $\kappa \equiv \frac{2b_\pi}{b_\theta}$. In what follows, it will sometimes be important to impose the following condition.

**Condition SC**: $Q(s)$ and $1 - \hat{\pi}(\theta)$ satisfy a single-crossing property: there exists $\theta^*$ such that $Q(s) > 1 - \hat{\pi}(\theta)$ for all $\theta < \theta^*$ and $Q(s) \leq 1 - \hat{\pi}(\theta)$ for all $\theta \geq \theta^*$.

One way to rephrase Condition SC is to impose monotonicity of the function

$$\zeta(s) = s + \kappa(1 - 2Q(s))$$

in $s$.

Writers' preferences are described by a function $v_w(\pi_{w(j)}, a_{w(j)})$ of participation and actions taken by all writers on the ledger $w_j$ that they choose. In our applications, writers

[2] We use this simple specification of utility to derive sharp analytical results, but our results go through qualitatively as long as readers play a game with strategic complementarities.

will prefer to write on widely used ledgers because their revenues will scale with the number of readers. It is important to allow for dependence on the actions of other writers because when there is competition to write on a given ledger, an individual writer's revenues will depend on the competition faced. Proposer $k$ obtains utility $v_p(\pi_k, a_k)$ at $t = 3$. In our specific examples we elaborate in more detail on how proposers' preferences for participation arise, but one way to motivate these preferences is by thinking of proposers as large stakeholders who benefit when others participate in the ledger through an increase in the value of their stakes. When more readers participate in the proposed ledger, the proposer's stake appreciates by a greater amount.

### 3.1 Characterization of equilibrium with arbitrary competing ledgers

We now prove properties of equilibrium that will hold in all of the settings we consider. First, we show that as noise about the common value vanishes, readers' play is uniquely pinned down in equilibrium. We also characterize the multiplicity of equilibria in a benchmark setting where readers' types are identical. Here we restrict attention to pure-strategy Perfect Bayesian equilibria of the ledger choice game. For a formal definition of Perfect Bayesian equilibrium, we refer the reader to Fudenberg and Tirole (1991).

The main property of equilibria that we can prove at this point is that equilibria will take a "cutoff" form: there will be threshold values $k(s)$ such that all agents with $x_i < k(s_i)$ choose ledger $B$ and all readers with $x_i > k(s_i)$ choose ledger $A$. These cutoffs will be decreasing in $s$, meaning agents with larger stakes in ledger $A$ will be more likely to choose $A$. This is true as long as the actions taken by writers are the same on ledgers $A$ and $B$. That is, readers sort themselves across ledgers according to their preferences. Those whose fundamental preferences for $A$ are above a certain bound will choose $A$ and all other readers will switch to $B$.

**Proposition 1.** *There is an essentially unique equilibrium of the game played by readers at $t = 1$ holding fixed the actions of writers at $t = 2$. There exist weakly monotonically decreasing cutoffs $k(s)$ such that all readers with $x_i > k(s_i)$ choose $r(i) = A$ and all readers with $x_i < k(s_i)$ choose $r(i) = B$. When condition SC holds, the cutoffs are given by*

$$k(s) = -(s - (g(L_A) - g(L_B)) + \kappa(1 - 2Q(s)))$$

The proof of Proposition 1 relies on standard techniques from the global games literature with heterogeneous preferences, as in Sakovics and Steiner (2012) or Drozd and Serrano-Padial (2017). The logic behind the proof is as follows. In this setup, there are certain types $s$ whose fundamental preferences for ledger $A$ are so strong that it is a dominant action to choose $A$ even if all other agents choose $B$. We call this set of types a "dominance region." Then some other types who strongly prefer $A$ will choose $A$ as well, since on top of their fundamental preference for $A$ they know that all types in the dominance region choose $A$. This logic can be iterated to derive a unique equilibrium under certain conditions. The

actions of types with extreme fundamental preferences are "contagious" and induce even types with mild preferences for one ledger over the other to take a given action. It is possible to find the set of types who choose $B$ in exactly the same way.

When Condition SC holds, the cutoffs take a particularly nice form. The reason that the equilibrium is so simple in this case is that readers' types are very dispersed. In fact, their types are so dispersed that even without incomplete information about $\theta$, there would still be a unique equilibrium. The uniqueness of equilibrium comes from the fact that some readers' preferences will be so extreme that they are in the dominance region, meaning there is no need to introduce the relevance of these types through higher-order uncertainty. Hence when preferences satisfy Condition SC, there is effectively no uncertainty about coordination.

In a benchmark case with complete information and identical preferences (captured by stakes), this property does not hold. The introduction of incomplete information or heterogeneous stakes is necessary to select a unique equilibrium. Here we also state a benchmark result that when preferences are identical, there are three equilibria as long as playing $A$ or $B$ is not a dominant action.

**Proposition 2.** *As long as neither $A$ nor $B$ is a dominant action for any type $s$, generically there are three equilibria taking writers' actions at $t = 2$ as given: one in which all readers play $A$, one in which all play $B$, and a mixed equilibrium.*

In the benchmark case with complete information and identical preferences, there are usually three equilibria. When all agents choose either $A$ or $B$, it is optimal for any individual agents to follow the crowd. However, there is also a mixed equilibrium in which agents are exactly indifferent between the two ledgers: the ledger with a lower value of $L_k$ will have less participation, which induces most agents to choose the ledger on which writers receive larger revenues.

### 3.2 Competition between distributed ledgers

In this section, we present our baseline model of competition between blockchain ledgers. In reality, this competition corresponds to a "hard fork," in which some of the blockchain's writers decide to build their own blockchain with new protocols off of a previously existing (parent) blockchain. Critically, a hard fork preserves all of the data in the parent blockchain. This observation will be crucial for our conclusions: the ability of writers to change the rules of the blockchain but keep readers' stakes in the network intact will allow for perfect competition between ledgers. There will be no inertia in switching ledgers because readers will lose nothing by doing so as long as all other readers switch as well. Blockchains will enhance competition between ledgers, but they will come at the cost of proof-of-work, the first (and most important) cost of decentralization. This example will thus illustrating one aspect of the decentralization-cost efficiency tradeoff postulated in the Trilemma.

The model of blockchain competition falls within the general class of models of ledger competition described earlier. In the game, readers must coordinate on a ledger (branch of a blockchain fork), which corresponds to choosing a ledger $A$ or $B$. We take $A$ to be the branch that keeps the rules of the existing blockchain. This branch has a fundamental parameter $L_A$ and readers have stakes $S$ on that branch. That is, we constrain the proposer $P_A$ to choose $(L_A, S)$. This proposer can be thought of as one of the original developers of the blockchain. The proposer on branch $B$ may choose a new fundamental parameter $L_B$ in a compact set $\mathcal{L} \subset \mathbb{R}_+$ but must choose stakes $S$ as well. Proposer $P_B$ can be thought of as a blockchain software developer who wants to fork the blockchain and therefore chooses new protocols but keeps all users' data intact. If participation on the ledger proposed by $P_B$ is $\pi_B$, $P_B$ receives a payoff $\pi_B(K - g_P(L_B))$, where $g_P$ is an increasing function of $L_B$ and $K$ is a constant. The proposer's payoff is assumed to come from an appreciation of the developer's stake when the proposed ledger is adopted. Function $g_P$ relates the appreciation of the proposer's stake to the fundamental parameter of ledger $B$, so that it is better for the proposer to suggest rules that benefit readers.

In this setting, the set $M$ of writers is a continuum $[0, M]$, where $M$ is taken to be large. We assume there are two branches of the fork, branch $A$ and branch $B$. Writers are responsible for cryptographically securing the ledger, and they are given some surplus for contributing computing power to the blockchain. At $t = 2$, writer $j$ chooses a ledger $w(j) \in \{A, B\}$ and an amount of computational power $c_j \leq 1$ to contribute to that ledger. We assume that writers can observe readers' actions before making a decision because in practice, this is often exactly what happens. Cryptocurrency "mining pools" are set up to automatically mine on whatever blockchain yields the highest profits at that moment. To the extent that the token price on a blockchain proxies for participation on that blockchain, mining pools essentially condition their decisions on users' actions.

Writers pay a linear cost $f(c) = c$ of generating computational power. Let $C_k = \int_{w(j')=k} c_{j'} dj'$ be the total computational power contributed to branch $k$ of the fork, and denote the participation on that fork by $\pi_k$. Then a writer's net profits when contributing computing power $c_j$ to branch $k$ are

$$v_w(\pi_k, c_j, C_k) = \frac{c_j}{C_k} \pi_k L_k - c_j$$

when $C_k > 0$ and $-c_j$ otherwise. The writer's revenues are proportional to participation and the fundamental parameter $L_k$ but are inversely proportional to the computational power contributed by other writers. This revenue function captures two features shared most blockchains. Namely, (1) the total rewards given to writers are fixed, and (2) those rewards tend to be more valuable when the blockchain has been adopted by a larger group of users.

Readers prefer ledgers that are cryptographically secure. Their preferences for cryptographic security are parametrized by a function $h(\frac{C_k}{\pi_k})$ such that $h(\frac{C_k}{\pi_k}) = 0$ whenever

$\frac{C_k}{\pi_k} \geq \underline{C}$ and $h(\frac{C_k}{\pi_k}) = H$, where $H$ is a large constant, otherwise.[3] That is, readers value security in terms of the amount of computational power committed to the blockchain per user, and there is some threshold level $\underline{C}$ of computational power above which readers are completely satisfied with the ledger's security. Below that level, readers are unsatisfied with the ledger's security. For now, we keep the function $h$ exogenous and discuss the benefits of fork competition. In our discussion of attacks on the blockchain we outline how it can be endogenized and discuss the tradeoff between free entry of writers and costly proof-of-work in greater detail.

There is incomplete information about readers' preferences. A common parameter $\tau$ affects readers' preferences for ledger $A$ over $B$ (where $\tau > 0$ pushes readers towards $A$ over $B$). Readers receive signals $x_i = \tau + \sigma\eta_i$, where $\eta_i$ are independently and identically distributed uniformly on the interval $[-\frac{1}{2}, \frac{1}{2}]$ and we take the limit $\sigma \to 0$. The value of $\tau$ is unknown to readers. They may have some prior over its distribution, but in the limit $\sigma \to 0$ this prior will be irrelevant because their signals are extremely precise.[4]This small amount of noise in preferences gives rise to a type distribution $x_i \sim U[\tau - \frac{\sigma}{2}, \tau + \frac{\sigma}{2}]$, since all readers have the same stakes on both ledgers. Adding an arbitrarily small amount of noise to the information structure will ultimately allow us to select a unique equilibrium.

Readers' preferences are summarized by

$$\Delta = E\Big[\big(\frac{1}{2}\kappa\pi_A + \tau - g(L_A) - h(\frac{C_A}{\pi_A})\big) - \big(\frac{1}{2}\kappa\pi_B - g(L_B) - h(\frac{C_B}{\pi_B})\big)|x_i\Big]$$

since each reader's stake is the same on both ledgers. Here $\kappa$ is a coefficient determining preferences for coordination. When $h(\frac{C_A}{\pi_A}) = h(\frac{C_B}{\pi_B}) = h(\underline{C})$, we obtain

$$\Delta(x_i, \pi, L_A, L_B) = E\Big[\tau - (g(L_A) - g(L_B)) + \frac{1}{2}\kappa(2\pi - 1)|x_i\Big]$$

where $\pi$ represents participation on ledger $A$, as before. Critically, here we assume that a proposal $L_B$ induces the same preference among all readers. Later we analyze a case in which readers have heterogeneous preferences for ledger $B$ following a proposal $L_B$.

Finally, we define the publicly information observable to players at each $t$. At $t = 1$, players observe the proposer's action $L_B$. At $t = 2$, all players observe the measure of readers $\pi_k$ who chose ledger $k$ at $t = 1$ for $k \in \{A, B\}$.

Now that we have set up the blockchain game, we may prove our main result.

**Proposition 3.** *Suppose there is $L_B \in \mathcal{L}$ such that $\underline{C} \leq L_B < L_A$. There exists a unique equilibrium when $\tau \leq 0$. In this equilibrium, proposer $P_B$ announces $\tilde{L}_B = \min\{L : L \in \mathcal{L}, L \geq \underline{C}\}$, all readers and writers choose ledger $B$, and writers break even.*

[3]Here we take writers' action set $\mathcal{A}(\pi) = [0, \frac{1}{\pi}]$ to be the computational power produced per blockchain reader. Under this specification, readers' payoffs are of the form assumed in the generic ledger choice model.

[4]We must also assume that the prior on $\tau$ is smooth and has full support to guarantee uniform convergence of the posterior. See Frankel, Morris, and Pauzner (2003).

Proposition 3 is a remarkable result. It states that in a setting in which there is an opportunity to fork a blockchain, readers will always choose the branch of the fork on which writers receive the lowest revenues, and proposers (developers) will propose rules that are beneficial to readers rather than writers.[5] Figure 1 depicts an example of the equilibrium of the blockchain game. Of course, the result that proposers suggest protocols that are beneficial to readers depends partly on the assumption that proposers' incentives are aligned with those of readers, but in a setting with free entry of writers this assumption is not overly restrictive. Writers always make zero profits, so proposing a ledger that increases writers' revenues is pointless. Furthermore, readers choose to switch to ledger $B$ only because they do not stand to lose their stakes when doing so. The replicability of information on ledger $B$ completely removes an obstacle to switching ledgers. We will show that when information cannot be replicated on a competing ledger, readers' stakes impede switching to a ledger where writers earn lower revenue.

Proposition 3 highlights the benefits of a blockchain. When all readers' fundamental preferences for an alternative ledger are identical, the absence of switching costs induces full coordination on the competing ledger. There is perfect competition among ledgers in that as long as it is feasible to set $L_B$ even slightly lower than $L_A$, the competing ledger will win out over the existing one. Remarkably, there is perfect competition between ledgers. Coordination inefficiencies are precluded under these assumptions, but in the next subsection we discuss how coordination can break down when readers have heterogeneous fundamental preferences.

Popular discussion has largely focused on the ways in which blockchains can decrease essentially exogenous costs, such as by inducing faster consensus about a ledger's contents. This result shows that there is an *endogenous* channel through which blockchain reduces the cost of maintaining a ledger: the synergy between *portability of information* and *competition among writers*. When information can be ported to an outside ledger, readers will want to use that ledger if writers are paid lower fees. Individually, writers are better off writing on a ledger with high fees, but competitive forces drive writers to undercut each other by writing on the ledger with lower fees. Writers know that all readers will use the outside ledger when there are enough writers to secure it, so the end result is that all writers must switch to the outside ledger. The downside of a blockchain is that while in a traditional setting writers' fees simply represent a (possibly distortionary) transfer, in the case of blockchain writers' fees are a pure waste of resources. We next examine under what conditions a traditional ledger maintained by a monopolist induces a large distortion due to rent extraction.

[5]Note that the hypothesis $\tau \leq 0$ is not restrictive. It just states that if agents are ex-ante neutral or prefer ledger $B$, there will be a unique equilibrium in which they all switch to ledger $B$. A good benchmark is the case $\tau = 0$.

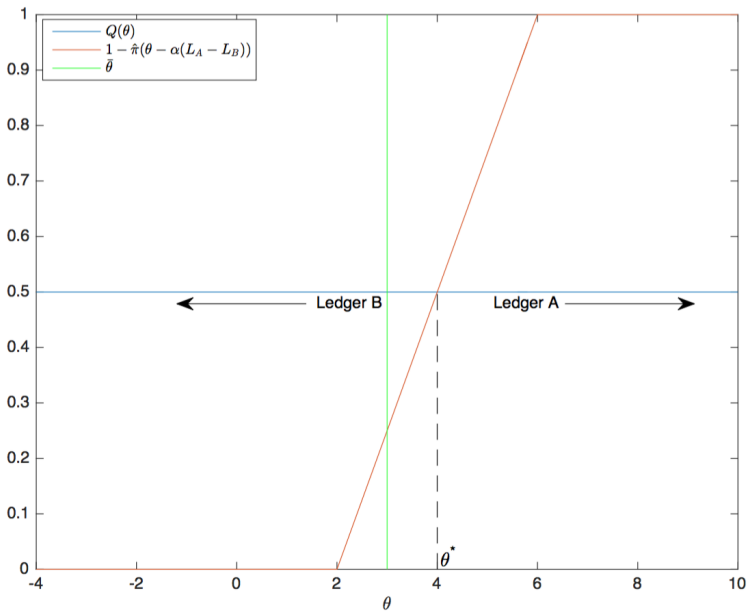


Figure 2: An example of an equilibrium of the blockchain game. Here $\bar{\theta} = 3$, $g(L) = \alpha L$, $L_A = 50$, $L_B = 10$, and $\kappa = 4$. The green line represents the actual CDF of types, which is concentrated in a small interval around $\bar{\theta}$.

### 3.3 A realistic "hard fork"

In this section, we analyze a hard fork that is more realistic than the type highlighted in the preceding analysis where *all* users of the blockchain switch to one branch of the fork and the other is completely abandoned. In reality, hard forks usually lead to a split of the community. For example, the Ethereum community split after hackers stole cryptocurrency from a smart contract. Although the majority of the blockchain's users joined the segment of the community that decided to fork, a significant percentage of users continued to use the original blockchain. The Bitcoin blockchain has also been forked by the (significantly less popular) cryptocurrencies Bitcoin Cash and Bitcoin Gold, both of which changed the rules of Bitcoin in order to benefit users. In these cases, many users of Bitcoin refused to actively use the new cryptocurrencies because they felt that the changes to the rules were actually detrimental or compromised the security of the blockchain. This section will focus on the tradeoff between fork competition and network externality inefficiencies, the second cost of decentralization. Although fork competition can benefit users, we will show that it can also lead to inefficient miscoordination, or "too many ledgers" in equilibrium.

The key mechanism that will underlie realistic hard forks in our model is preference heterogeneity. Although in the benchmark model agents are heterogeneous in their preferences, we take a limit in which this heterogeneity vanishes. We now consider a model identical to the benchmark with the exception of the specification of types. Readers' types are now given by

$$\theta_i = (x_i, f_i)$$

where $f_i \in \{0, f\}$. The type $f_i$ reflects a preference for forking: readers with $f_i = f$ dislike all forks equally, and readers with $f_i = 0$ are not averse to forking the existing blockchain.[6] Types $f_i$ are independently and identically distributed across readers with $\Pr(f_i = \eta) = \mu$. Types $x_i$ are distributed uniformly in the interval $[\tau - \frac{\sigma}{2}, \tau + \frac{\sigma}{2}]$ as before. Readers observe both $x_i$ and $f_i$.

Note that if there exists $L_B \in \mathcal{L}$ such that $L_B \geq \underline{C}$ and $g(L_B) - g(L_A) > f$, we obtain the same result as in Section 3.2. Proposer $P_B$ will propose such an $L_B$ and all readers will switch to branch $B$. In this case, there exists a feasible fundamental parameter $L_B$ that is better than $L_A$ by such a wide margin that all readers, including those who dislike forks, prefer ledger $B$ with parameter $L_B$.

We therefore consider only the case in which all $L_B \in \mathcal{L}$ satisfy $g(L_B) - g(L_A) < f$. In fact, the only situation in which multiple equilibria would arise under complete information is if

$$f + g(L_A) - g(L_B) \geq \kappa(1 - 2\mu) \geq g(L_A) - g(L_B)$$

We derive the unique equilibrium under these conditions. The results are summarized in Proposition 4.

[6] We adopt this specification for simplicity. Allowing for $\eta_i$ to depend on the announced fundamental parameters $L_A$ and $L_B$ would not change the main results. Anecdotal evidence suggests that there are indeed blockchain users who are fundamentally averse to forking.

**Proposition 4.** *Suppose readers face ledgers with fundamental parameters $L_A, L_B$, $W_k \geq L_k$ writers commit to branch $k$ at $t = 1$, and a fraction $\mu$ of readers are of type $f_i = f$. Then if*

$$f + g(L_A) - g(L_B) \geq \kappa(1 - 2\mu) \geq g(L_A) - g(L_B)$$

*the essentially unique equilibrium at $t = 1$ is of one of two types.*

1. *If $f \leq \kappa$, then all readers choose branch $A$ if $g(L_A) - g(L_B) > \mu f$ and branch $B$ if $g(L_A) - g(L_B) < \mu f$.*
2. *If $f > \kappa$, readers of type $f_i = f$ choose branch $A$ iff $f - (g(L_A) - g(L_B)) > (1 - \mu)\kappa$ and readers of type $f_i = 0$ choose branch $B$ iff $g(L_A) - g(L_B) > \mu\kappa$. That is, the miscoordination equilibrium of the complete information game is selected when $f > \kappa$ if such an equilibrium exists.*

This proposition essentially shows that when readers' fundamental aversion to forking is strong relative to the coordination motive, the blockchain is vulnerable to a hard fork that splits the community. Intuitively, when network externalities are weak relative to some readers' dislike of forks, readers who are averse to forks will still prefer not to leave the existing ledger even if all other readers join the new fork. Put another way, coordination motives are a source of strength for a blockchain: when network externalities are weak, coordination among the blockchain community becomes fragile and the community is susceptible to a split.

The possibility of a hard fork that splits the community has important implications for welfare. When no fork is proposed, all readers obtain utility $\mu f - g(L_A) + \kappa$. When a fork is proposed and a community split occurs, on the other hand, readers obtain average utility

$$\mu(f - g(L_A) + \kappa\mu) + (1 - \mu)(-g(L_B) + \kappa(1\mu))$$

Relative to the case with no forking, the welfare gains or losses are

$$(1 - \mu)(g(L_A) - g(L_B)) - 2\kappa\mu(1 - \mu)$$

The first term is the fundamental benefit readers of type $f_i = 0$ obtain by switching to $B$, and the second term is the coordination loss associated with the split. Hence the fork is detrimental to welfare if

$$2\kappa\mu > g(L_A) - g(L_B)$$

The results of the previous section and this one highlight the main tradeoff relevant for determining whether a blockchain is worthwhile. Although a blockchain greatly enhances competition between ledgers and lowers fees, it may also induce an undesirable breakdown of coordination. The possibility of miscoordination is especially strong when network externalities are weak, so a blockchain is likely particularly useful when the coordination motive is important.

### 3.4 Competition between traditional ledgers

In this section, we analyze a competition between a ledger maintained by a monopolist and an outside ledger. The differences between this setting and one with two distributed ledgers will clarify exactly how fork competition differs from standard competition; i.e., what exactly is accomplished by decentralization. We first begin by assuming that the monopolist is the incumbent in the sense that readers have a stake in the monopolist's ledger but not the outside ledger. There are just two writers: the monopolist $\mathcal{M}$ on ledger $A$ and an outside writer (entrant) $\mathcal{O}$ on ledger $B$. In this case, the writers are also the proposers $P_A = \mathcal{M}$ and $P_B = \mathcal{O}$. Each writer may only write on her own ledger. At $t = 1$, the incumbent may choose a fundamental parameter $L_A \geq 0$ and the the entrant chooses $L_B \geq 0$. The incumbent and entrant choose stakes $\hat{S}$ and 0, respectively. The restriction that the entrant must choose zero represents a situation in which readers have no stake in the outside writer's ledger and that writer is unable to replicate the stakes in the monopolist's ledger due to information frictions. Writers do not take actions at $t = 2$.

Readers have preferences summarized by

$$1 - \hat{\pi}(s) = \frac{1}{2} + \kappa^{-1}\big(\tau + s - \alpha(L_A - L_B)\big) \tag{1}$$

Here we use a linear function $\alpha L$ to represent the disutility from paying fees to writers. While less general than equation (1), these preferences will allow us to derive analytical solutions for the monopolist's optimal policy. Readers' types $\theta_i$ are given by their stakes on the monopolist's ledger $s_i$, which has a cross-sectional distribution $Q(s)$ that is uniform on the interval $[S - \frac{d}{2}, S + \frac{d}{2}]$. Here $S$ is the average stake and $d$ is the dispersion of stakes. It is important to distinguish between situations in which Condition SC is satisfied and situations in which it is not. Condition SC is satisfied if and only if $d > \kappa$. We will henceforth assume in this section that the true realization of the common value $\tau$ is zero, but that this is unknown to readers.

The monopolist receives a fee $L_A$ from each reader who participates. The monopolist's objective function is

$$\max_{L_A \geq 0} \pi L_A$$

where $\pi$ denotes participation on ledger $A$. Similarly, the entrant's objective function is

$$\max_{L_B \geq 0} (1 - \pi) L_B$$

In order to proceed, we must determine how the writers' choices of $L_A$ and $L_B$ map to participation $\pi$. Proposition 5 provides an answer.

**Proposition 5.** *When Condition SC holds and $\tau = 0$, all readers for whom*

$$1 - \hat{\pi}(s) = \frac{1}{2} + \kappa^{-1}\bigg(s - \alpha(L_A - L_B)\bigg) > Q(s)$$

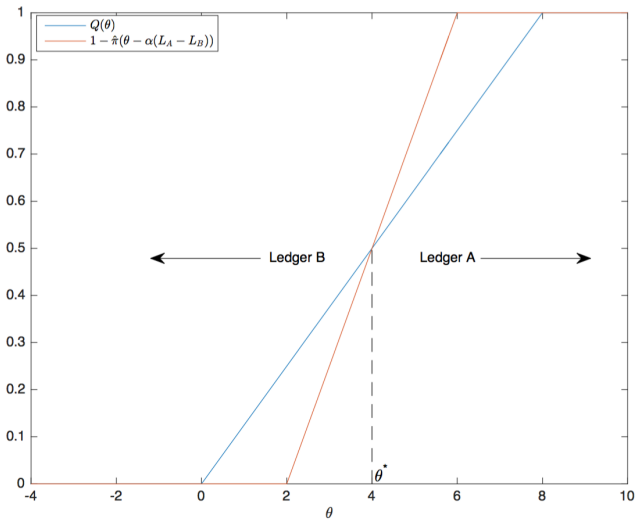


Figure 3: Participation on the monopolist's ledger when stakes are distributed uniformly on $[S-\frac{d}{2}, S+\frac{d}{2}]$ with $\tau=0$, $S=4$, $d=8$, $\kappa=4$, $L_A=40$, $L_B=0$, and $\alpha=0.1$.

*choose to remain on ledger A, and all other readers choose ledger B. When Condition SC does not hold, all readers choose A if* $\alpha(L_A-L_B)<S$ *and B if* $\alpha(L_A-L_B)>S$.

First, we consider the simpler case in which Condition SC does not hold. In this case, network externalities are so strong that all readers will end up choosing the same ledger regardless of how invested they are in ledger $A$. Proposition 5 shows that when network externalities are strong, the incumbent and entrant effectively compete à la Bertrand. Each will try to undercut the other as long as it is possible to do so. However, the incumbent has a competitive advantage corresponding to the average stake $S$ readers have in its ledger. Therefore, in equilibrium the entrant must choose $L_B=0$, and the incumbent monopolist chooses $L_A$ just small enough so that readers do not switch to $B$. By Proposition 5, this yields

$$L_A=\frac{S}{\alpha} \tag{2}$$

In this case, the profits earned by the monopolist depend only on the average stake and $\alpha$, which parametrizes readers' aversion to fees. When the average stake is higher, the monopolist has a larger competitive advantage because there is greater inertia in switching ledgers.

Now we consider the case in which Condition SC holds. By Proposition 5, when the monopolist selects $L_A$, all readers for whom $1-\hat{\pi}(s)>Q(s)=\frac{s-S}{d}+\frac{1}{2}$ choose to remain on ledger $A$. Figure 2 illustrates this situation. To find the cutoff type $\theta^*$ who is indifferent

between remaining on the monopolist's ledger and leaving, we solve

$$\frac{1}{2} + \kappa^{-1}(s - \alpha(L_A - L_B)) = \frac{s - S}{d} + \frac{1}{2}$$

which implies

$$s^* = \frac{d}{d - \kappa}\left(\alpha(L_A - L_B) - \kappa(\frac{S}{d} - 1)\right) \tag{3}$$

so long as the expression on the right-hand side is in the range $[0, S]$. This yields $Q(s^*) = \frac{s^* - S}{d} + \frac{1}{2}$, so we obtain an expression for participation in the monopolist's ledger as a function of $L_A$:

$$\pi(L_A, L_B) = 1 - Q(s^*(L_A, L_B)) = \frac{S + \frac{d}{2} - \frac{\kappa}{2} - \alpha(L_A - L_B)}{d - \kappa}$$

Then the monopolist's problem reduces to

$$\max_{L_A} \left(S + \frac{d}{2} - \frac{\kappa}{2} - \alpha(L_A - L_B)\right) L_A$$

which yields

$$L_A = \frac{S + \frac{d}{2} - \frac{\kappa}{2} + \alpha L_B}{2\alpha} \tag{4}$$

The rents extracted by the monopolist are increasing in the average stake on its ledger because when the average stake is higher, readers must be charged a higher fee before they become indifferent between leaving the ledger and losing their stakes. A high dispersion of stakes also allows the monopolist to extract high fees because when there is a wide distribution of stakes, the sensitivity of the monopolist's revenues to $L_A$ is low. There are fewer marginal readers, so an upwards adjustment of $L_A$ does not result in a large exodus of readers from ledger $A$. Finally, when the parameter $L_B$ is large, readers are reluctant to leave ledger $A$ because they know that they will be charged high fees on the outside ledger regardless, so the monopolist enjoys higher profits.

On the other hand, a strong coordination motive can be detrimental to the monopolist's business. If the coordination motive is strong, when a single marginal reader leaves the ledger it induces many other readers to leave as well. In this case, the sensitivity of participation to $L_A$ is high. Clearly, it will also be the case that when readers' preferences are sensitive to $L_A$, the monopolist must set a lower $L_A$.

Recall that with a blockchain, the fundamental parameter that is chosen in equilibrium is essentially independent of the details of readers' preferences– the ledger that is best for readers is chosen automatically. In the traditional environment, when even partial competition is possible, network externalities work as a disciplining device against the incumbent monopolist. That is, network externalities enhance the importance of ledgers' fundamental parameters when replication of information and perfect, blockchain-style competition between writers is impossible.

Now we analyze the entrant's problem. Participation on the entrant's ledger is

$$1 - \pi(L_A, L_B) = \frac{\frac{d}{2} - \frac{\kappa}{2} - S + \alpha(L_A - L_B)}{d - \kappa}$$

The entrant's problem is then

$$\max_{L_B \geq 0} \left( \frac{d}{2} - \frac{\kappa}{2} - S + \alpha(L_A - L_B) \right) L_B$$

The first-order condition of this problem is

$$L_B = \frac{\frac{d}{2} - \frac{\kappa}{2} - S + \alpha L_A}{2\alpha} \tag{5}$$

The monopolist will choose this value of $L_B$ as long as $S - \alpha L_A \leq \frac{1}{2}(d - \kappa)$. Otherwise, the first-order condition is satisfied only for negative $L_B$, which is impossible, so the entrant sets $L_B = 0$.

Equation (5) shows that the entrant will extract high rents if the dispersion in readers' stakes is large or if the incumbent also extracts large rents. When the dispersion in readers' stakes is large, the sensitivity of the entrant's revenues to $L_B$ is low, as in the case where the monopolist is the incumbent. That is, dispersion in stakes is harmful to readers no matter which ledger they ultimately choose. When the fundamental parameter $L_A$ on the incumbent's ledger is large, readers are more willing to stomach high fees charged by the entrant, so $L_B$ is higher.

The entrant's rents are decreasing in the strength of the coordination motive $\kappa$, the mean stake on the incumbent's ledger $S$, and readers' sensitivity to fundamentals $\alpha$. Network externalities discipline both the incumbent and the entrant– when these externalities are strong, an increase in $L_B$ tends to cause a domino effect that results in a large mass of readers leaving ledger $B$. The fee charged by the entrant is also decreasing in the mean stake $S$ on the incumbent's ledger because that stake gives the incumbent a competitive advantage, so the entrant must charge a lower fee in order to capture a significant segment of the market.

In order to find the equilibrium of the game between the incumbent monopolist and the entrant, we simply combine their first-order conditions. Hence we simultaneously solve equations (4) and (5). This yields

$$L_A = \frac{1}{2\alpha}(d - \kappa) + \frac{1}{3\alpha} S, \; L_B = \frac{1}{2\alpha}(d - \kappa) - \frac{1}{3\alpha} S \tag{6}$$

Then participation on each ledger is

$$\pi_A = \pi = \frac{1}{2} + \frac{1}{3} \frac{S}{d - \kappa}, \; \pi_B = 1 - \pi = \frac{1}{2} - \frac{1}{3} \frac{S}{d - \kappa}$$

We need $0 \leq \pi_A, \pi_B \leq 1$ and $L_k \geq 0$ for $k \in \{A, B\}$. A necessary and sufficient condition is

$$S \leq \frac{3}{2}(d - \kappa) \tag{7}$$

This inequality is a *no-entry bound*. If this inequality does not hold, the incumbent $A$ is in fact able to retain all readers even when $L_B = 0$. That is, the stakes readers have in ledger $A$ endogenously prevent entry by even the most competitive entrant. While network externalities discipline the fees charged by the incumbent, inequality (7) shows that they actually impede entry by competitors as well. When the participation of others is important to readers, it is difficult for a competitor to enter because it cannot attract enough readers to get itself off the ground. On the other hand, when readers' stakes on ledger $A$ are dispersed, it is easier for the entrant to attract the readers with the least to lose by switching, which in turn induces switching by other readers. When the no-entry bound holds,

$$L_A = L^{NE} = \frac{1}{\alpha}\left(S - \frac{1}{2}(d - \kappa)\right) \tag{8}$$

The incumbent sets $L_A$ to be the highest value such that all readers participate in the ledger. We have the following results regarding the case with no entry.

**Proposition 6.** *The no-entry bound on the average stake $S$ is decreasing in the strength of the coordination motive $\kappa$ and increasing in the dispersion of stakes $d$. Readers' welfare under the no-entry bound is decreasing in $S$, increasing in $d$, and decreasing in $\kappa$.*

Now we turn to the case in which there is entry. Equation (6) clarifies that dispersion in stakes and the strength of the coordination motive $\kappa$ affect the fees charged on both ledgers symmetrically. When the coordination motive is powerful, both monopolists are disciplined by the fact that a higher fee will cause a large loss of clientele through spillover effects. When one reader leaves a ledger, other nearly marginal readers follow suit because of the importance of coordination. On the other hand, dispersion in stakes has the opposite effect. When readers' stakes are heterogeneous, only a small mass of readers will be marginal for any given fee, so an increase in the fee does not cause a large loss in a monopolist's client base.

The mean stake $S$ has an asymmetric effect on monopolist's fees. An increase in $S$ increases $L_A$ while decreasing $L_B$. When the mean of readers' stakes on ledger $A$ is high, there is a competitive wedge between ledgers $A$ and $B$. Monopolist $\mathcal{M}$ can extract higher rents than the entrant $\mathcal{O}$ because readers' stake in ledger $A$ acts as an inertial force preventing them from leaving.

We have outlined three types of equilibria: one type in which Condition SC does not hold and all readers choose ledger $A$, an equilibrium in which there is no entry even though Condition SC holds, and an equilibrium with entry. We can now collect our results to determine how the incumbent monopolist's fees vary across the spectrum of equilibria.

**Proposition 7.** *The incumbent monopolist charges fees*

$$L_A = \begin{cases} \frac{1}{\alpha}S & \kappa > d \\ \frac{1}{\alpha}S - \frac{1}{2\alpha}(d-\kappa) & d - \frac{2}{3}S \leq \kappa < d \\ \frac{1}{3\alpha}S + \frac{1}{2\alpha}(d-\kappa) & \kappa < d - \frac{2}{3}S \end{cases}$$

*These fees are decreasing in $\kappa$ for $\kappa < d - \frac{2}{3}S$ and increasing in $\kappa$ for $\kappa > d - \frac{2}{3}S$.*

The main insight of Proposition 7 is that the equilibrium fee charged by the incumbent monopolist is non-monotonic in the strength $\kappa$ of network externalities. When the coordination motive is weak, fees are decreasing in $\kappa$ because a stronger coordination motive leads to a more powerful domino effect causing readers to switch ledgers. Once $\kappa$ reaches a threshold level, network externalities become strong enough to prevent entry, so a larger $\kappa$ actually leads to higher fees because the barrier to entry becomes stronger. When network externalities are more powerful still, the fees charged by the monopolist depend only on its competitive advantage, i.e. the average stake in its ledger, because there is no threat of entry. The market essentially ceases to be contestable, and the monopolist extracts the maximum possible surplus from readers.

Overall, this situation is quite different from the case where two forks of a blockchain compete against one another. When two forks of a blockchain compete, the combination of portability of information and competition between writers drives fees down as far as they can go while still providing sufficient incentives for writers to secure the network. The equilibrium outcome is independent of the distribution of readers' stakes. Welfare losses come mostly from the waste of computational resources and miscoordination due to forking (which tends to occur when network externalities are weak). Under traditional monopolistic competition, even when there is competition both the monopolists may charge high fees. If there is no possibility of entry, strong network externalities protect the incumbent and increase distortionary rents. The incumbent further enjoys high rents because of its monopoly on information, which is detrimental to readers' welfare. Taken together, these results suggest that blockchains should be used as ledgers when coordination motives among users are strong or when switching costs in a traditional setting are high.

### 3.5 Competition between a monopolist and a blockchain

Now we turn to competition between a monopolist and a blockchain. The primary difference from the previous example is that the agent who proposes the fee structure for a blockchain does not care about the fees earned by writers because writers always break even. Rather, the proposer's incentives are aligned with those of readers. As before, the proposer can be thought of as a developer of blockchain software who has a large stake in the network that appreciates when others use the blockchain platform. Formally, there are two ledgers $A$ (monopolist) and $B$ (blockchain) with proposers $P_A = \mathcal{M}$, who is also the writer on ledger $A$, and $P_B = \mathcal{D}$ (for "developer") who is not a blockchain writer.

Proposers $P_A$ and $P_B$ choose parameters $L_A, L_B \in \mathcal{L}$ at $t = 0$. Proposer $P_A$ is constrained to choose stakes $\hat{S}_A$, which are uniformly distributed on $[S - \frac{d}{2}, S + \frac{d}{2}]$, and $P_B$ must choose stakes $\hat{S}_B = 0$. When a blockchain competes against a monopolist, there is still perfect competition between blockchain writers, but the blockchain cannot replicate the information contained on the monopolist's ledger.

As in the baseline blockchain model, there is a continuum of writers $j \in [0, M]$. However, there is no longer incomplete information. When readers' stakes on ledger $A$ are distributed in an interval of finite length, an arbitrarily small amount of noise in agent's beliefs will have no effect on the equilibrium. Nevertheless, despite this change to the model, the equilibrium played by writers will be the same as in the baseline model of a blockchain fork.[7] Furthermore, blockchain writers cannot write on the monopolist's ledger, so they all must commit to ledger $B$ at $t = 1$. The equilibrium at $t = 2$ is just like the equilibrium in the case of monopolistic competition so long as Condition SC is satisfied, which again reduces to the inequality $d \geq \kappa$. To see this, note that in this setting the distribution of types is simply the distribution of stakes on ledger $A$ and apply Proposition 1.

We then have equilibrium play along any path for $t \geq 1$, so solving the model reduces to solving the proposers' optimization problems at $t = 0$. The monopolist behaves as if facing a fixed outside ledger with parameter $L_B$, so the optimal $L_A$ is again given by (4). However, $P_B$ has different preferences than an entrant monopolist. As in the baseline blockchain model, $P_B$'s preferences are given by $(1 - \pi)(K - g_P(L_B))$, where $g_P$ is an increasing function. If readers only join ledgers for which the average computing power per user is at least $\underline{C}$, $P_B$ must choose $L_B^* = \min\{L : L \in \mathcal{L}, L \geq \underline{C}\}$. The monopolist then chooses

$$L_A = \frac{\frac{d}{2} - \frac{\kappa}{2} + S + \alpha L_B^*}{2\alpha}$$

as long as

$$S + \alpha L_B^* \leq \frac{3}{2}(d - \kappa)$$

This inequality is the no-entry bound in the presence of a blockchain. Note that the no-entry bound is tighter when $L_B^*$ is larger. This is because when the minimum feasible computational power required to support a blockchain is large, the compensation necessary to attract writers (and thus the minimum blockchain fee) will be higher, thereby dissuading readers from using the blockchain.

The fee charged on the blockchain will be lower than that charged by an entrant monopolist precisely when $L_B^*$ is less than the expression given in (6) for the entrant's fee. Furthermore, in this case the lower fee charged on the blockchain will induce the incumbent monopolist to drop its fee below what it would charge when facing an entrant monopolist.

[7] Indeed, the $t = 3$ part of the proof of Proposition 3 is independent of the information structure so long as all writers observe participation on the ledger.

The condition for a blockchain to lower fees on both ledgers is

$$L_B^* < \frac{1}{2\alpha}(d-\kappa) - \frac{1}{3\alpha}S$$

When $L_B^* = \underline{C}$, this result is particularly stark. A blockchain lowers costs for readers when the computational expenditure required to placate readers' need for cryptographic security is small, when the dispersion of readers' stakes on the monopolist's ledger is high, or when the coordination motive is weak. Surprisingly, a blockchain tends to lower costs when the average stake on a monopolist's ledger is small. This is because when stakes on a monopolist's ledger are large, an entrant monopolist would optimally charge a low fee in order to induce switching by readers. Hence when the incumbent already charges high fees, competition by a traditional intermediary should be enough to lower costs to readers. Blockchain is useful primarily when entrants into the market have incentives to charge high fees. Free entry of blockchain writers implies that there is no incentive for a proposer to choose a policy that gives writers large fees because all writers break even regardless. The feature of the blockchain that allows it to more effectively compete with traditional intermediaries is that it strips writers of their market power.

# 4 Dynamic Ledger Choice

We now consider a repeated version of the static blockchain ledger choice game presented in the previous section. We show that, remarkably, readers and writers must play the static equilibrium of Proposition 1 in every period of the game. In short, this is because the free entry condition guarantees that writers cannot be rewarded or punished by any dynamic scheme. Therefore, writers will not be able to collude with each other on an outcome that is beneficial to them. Importantly, this property of permissionless blockchains with free entry will not carry over to permissioned blockchains where certain known parties write on the ledger. On a permissioned blockchain, it will be possible for collusion between writers to prevent low fees from emerging.

## 4.1 Permissionless blockchain

The repeated game with a permissionless blockchain is played on "days" $T = 1, 2, \ldots$. On each day, proposers, readers, and writers play the static game. Readers are short-lived and die after one period, but writers and proposers $P_A$, $P_B$ live forever and discount payoffs at rate $\delta$. Histories of this game are defined recursively. Let $\mathcal{H}^1 = \{\emptyset\}$. Then define

$$\mathcal{H}^T = \mathcal{H}^{T-1} \times \mathcal{L} \times [0,1] \times [0,M]^2 \times \mathcal{L}$$

The observable quantities are whether the initial writer chose on day $T$ chooses to propose a fork (where 1 indicates that a fork was proposed), which fork $L_B \in \mathcal{L}$ was proposed, how

many readers chose branch $A$, and how much computing power was committed to each branch. The last $\mathcal{L}$ represents the parameter $L_k$ on the ledger $k$ chosen by the majority of readers at $t = 2$, which becomes the reference parameter on branch $A$ in the next period. That is, when readers choose a particular fork of the blockchain, that chain is extended and becomes the default for developers to build off of if they want to fork in the future. The histories $\mathcal{H}^{T,t}$ that are publicly observable within subperiod $t$ of day $T$ are defined in the obvious way. Readers observe their own private signals and writers observe the entire history of their private signals.

We define subgame-perfect equilibrium in the usual way. We now show that in any SPE of the repeated game, writers always make zero profits from contributing computing power to the blockchain. The unique SPE of the repeated game will then be one in which agents play the unique SPE of the static game.

**Proposition 8.** *In any SPE of the repeated game, writers make zero profits. The unique SPE is the equilibrium of Proposition 1 played on every day $T$.*

### 4.2 Permissioned Blockchain

We now consider the case of a permissioned blockchain. One might think that a permissioned blockchain strictly dominates a permissionless blockchain in any application, since it allows the replication of information just like a permissionless blockchain but does not involve any waste of computational resources. If the set of equilibria with permissioned and permissionless blockchains were the same, in a sense permissioned blockchains would break the Trilemma by eliminating the usual waste of resources. However, free entry of writers on a permissionless blockchain actually helps to sustain equilibria that are beneficial to readers because they eliminate the possibility of collusion among writers. The computational costs of a permissionless blockchain can then be seen as the costs of allowing for free entry. On a permissioned blockchain, there is no free entry: the consortium of entities that are allowed to write on the ledger jointly decide whether to admit new members, and then those new members are identified to the blockchain's readers. The lack of free entry represents a failure of true decentralization. In the case of permissioned blockchain, the synergy between replicability of information and competition between writers fails because competition between writers is imperfect, since writers earn rents.

In order to capture this situation, we present a simple model of a permissioned blockchain. The model is similar to the baseline model with the exception that there is a finite number of writers who do not incur computational costs. Play occurs on days $T = 1, 2, \ldots$, and each day consists of subperiods $t = 0, 1, 2$ just as in the benchmark ledger choice model. There are proposers $P_A, P_B$ who choose fixed parameters $L_A > L_B$, respectively, in each period. They both choose stakes $\hat{S}$ (which are irrelevant because information is always replicated across branches of the fork). Here branch $A$ can be seen as the reference ledger. Our main result will be that with a permissioned blockchain, it will be possible for writers to prevent forking to branch $B$.

There are $M \in \mathbb{N}$ writers who discount payoffs at rate $\delta$ and a continuum of short-lived readers $i \in [0,1]$. The timing is as follows. At $t = 0$, proposers announce $L_A$ and $L_B$. At $t = 1$, after learning writers' decisions, readers individually choose a fork $k \in \{A, B\}$ of the blockchain. Writers choose a branch of the fork at $t = 2$. Readers receive a payoff of zero if no writer chooses the branch of the fork that they chose at $t = 1$.

In this setting, there is no question of computational security because there are no computational problems to be solved. Therefore, readers' preferences can be represented by

$$1 - \hat{\pi}(\theta) = \frac{1}{2} + \kappa\big(\theta - (g(L_A) - g(L_B))\big)$$

Writers obtain payoffs $\frac{1}{W_k}\pi_k L_k$ if they write on a branch with participation $\pi_k$, surplus parameter $L_k$, and $W_k$ writers.

Now we show that when $\delta$ is sufficiently large or $M$ is sufficiently small, there is a SPE of this game in which all writers choose ledger $A$ and a new ledger is never proposed. This is in contrast to the permissionless blockchain case, in which readers and writers would always coordinate on ledger $B$ if $L_B < L_A$. Consider the following equilibrium conjecture:

1. After any history in which all writers chose $A$ in all previous periods, all writers choose $A$.

2. After any history in which some writer chose $B$ in some previous period, all writers choose $B$.

Within a given day, writers have an incentive to announce $B$ because then all readers switch to $B$ and they obtain all the revenues on branch $B$. However, afterwards they receive lower payoffs because all writers play $B$, and they cannot deviate to obtain higher payoffs because readers will choose $B$ in every period.

Formally, the incentive constraint that must be satisfied in order for the specified strategy profile to be an equilibrium is

$$L_B + \frac{\delta}{M(1-\delta)} L_B \leq \frac{1}{M(1-\delta)} L_A$$

This inequality can be rearranged to obtain

$$\frac{L_A}{L_B} \geq \delta + (1-\delta)M \tag{9}$$

This inequality holds when $\frac{L_A}{L_B}$ is sufficiently large. Playing $A$ is incentive compatible when $L_A$ is large relative to $L_B$ because when a writer decides to play $B$, she takes an immediate payoff of $B$ but loses future rents proportional to $L_A$. This inequality is also satisfied for large $\delta$ or low $M$. When writers are patient or competition between writers is weak, they have an incentive to conform to equilibrium play.

To restate the main point, there is nothing inherent in the blockchain data structure itself that impedes rent-seeking behavior. Adding a costly identity management system to allow for free entry of writers in fact increases the costs of using the ledger for a *given* set of policies. However, perfect competition among writers combined with the fact that blockchains can be forked *endogenously* decreases the cost of using a ledger because it allows for the selection of rules that are most beneficial to readers. With a permissioned blockchain, there is no computational cost of verification, so it is possible to maintain a decentralized, immutable ledger with no single point of failure without any waste of resources whatsoever. However, when there is no computational expenditure involved in managing a blockchain, writers must earn rents, so collusion via dynamic punishment schemes can reduce incentives for writers to choose non-distortionary policies that are beneficial to readers. Decentralization is critical to competition precisely because it prevents collusion.

### 4.3 Blockchain security

Traditional ledgers have been criticized for being opaque and vulnerable to fraud. One of the principal advantages of blockchain protocols is that the ledger is resilient to fraud by a single bad actor. In this section, we analyze the security of both traditional ledgers maintained by monopolists and blockchains. We outline a simple model of blockchain security and compare the security of a blockchain to that of a ledger written by a monopolist. We show that while centralized intermediaries have *dynamic* incentives not to distort their own ledgers, blockchain writers' incentives are *static*, which makes it expensive to incentivize honest reporting. We discuss the tradeoff between correctness and monopolistic rent extraction on the one hand and that between correctness and decentralization on the other, showing how proof-of-work costs arise naturally.

The model of blockchain security is based off of the dynamic blockchain model. As before, there are two proposers $P_A$ and $P_B$ and a continuum of readers $i \in [0, 1]$. We depart from the earlier model in that we allow for some "large" writers who each command a positive measure of computing power. There is a large writer $J$ with unlimited computing capacity and a continuum $j \in [0, M]$ of infinitesimally small writers with computing power $dj$. This assumption is meant to capture "51% attacks" in which an entity or mining pool able to control a majority of a blockchain's computing power mounts a malicious attack on the network in order to reap financial gains. We will also assume the large writer lives for only one period. We do this in order to abstract away from dynamic punishments for large writers who can attack the network. This assumption is reasonable because (1) large writers would not be able to profitably attack the blockchain on a regular basis given that others would join the attacks and drive their profits to zero, and (2) even if the blockchain completely shut down these writers could simply choose to attack another blockchain.

In subperiod $t = 0$ of each day $T$, proposers $P_A$ and $P_B$ announce a fixed fundamental parameter $L \in \mathcal{L}$. For simplicity, we will assume $L = M$ so that in an equilibrium with no attacks, small writers always expend their entire computing power. The proposers differ in

their announcements of stakes: $P_A$ announces stakes $S_T$ while $P_B$ announces $S_{T-1}$. Here $S_T$ represents the stakes on the longest chain in the blockchain, whereas $S_{T-1}$ represents forking the blockchain back to the state in the previous period. The ability to fork the blockchain backwards will discipline writers who engage in fraudulent activity because their gains will be nullified when such a backwards fork occurs. Subperiod $t = 1$ is the same as in the benchmark model. Readers choose a ledger at $t = 1$ after learning their types.[8]

The main difference from the benchmark model is at $t = 2$. In each period, an attack is possible on ledger $A$ with some small probability $\mu > 0$. We assume an attack is unlikely to ensure that small writers do not play as if the blockchain is constantly under attack, which would imply that they take large losses in periods where attacks succeed and make positive profits when they fail (in contrast to what happens in reality). When an attack is possible, the large writer chooses an action $a_J \in [0, \overline{a}]$ as well as computing power at $t = 2$. The action $a_J$ represents the size of the distortion of the ledger attempted by writer $J$. In order for the attack to have a chance of succeeding, the large writer must choose computing power $c_J > \int c_j dj$, so that the computing power provided by the large writer is sufficient to overwhelm the rest of the network. The type of attack modeled here is one in which the large writer creates an invalid fork of the blockchain on which he distorts the ledger while small writers write on a valid fork. Readers are initially fooled by large writers' reports[9]and transact according to the invalid chain because it has greater proof-of-work.[10]

On each day $T > 0$, a public signal $y_T \in \{0, 1\}$ is revealed. The signal takes value 1 with probability $p(a_{J,T-1})$, where $a_{J,T-1}$ is the action $a_J$ played by the large writer at $T - 1$. We assume that

$$\frac{d}{da}\frac{p(a)}{1-p(a)} \geq 0$$

i.e., the hazard rate is nondecreasing. This signal could correspond to news media revealing that an attack on the blockchain has occurred, large numbers of people realizing that their accounts on the ledger have been compromised and spreading word of the attack, or participants with a vested interest in the blockchain communicating evidence of the attack to the community. In this setting, the assumption $y \in \{0, 1\}$ will be without loss of generality– the equilibrium will be the same regardless of whether readers can perfectly observe $a_{J,T-1}$.

[8] In reality, blockchains have forked after an attack on the network was discovered. Most famously, the Ethereum blockchain forked in 2016 after hackers stole roughly $50 million from a smart contract on the blockchain.

[10] If readers were perfectly able to observe misconduct on the blockchain (as is the case for some blockchains that are not storage-intensive), there would be no possibility of an attack in the first place. In this case, though, a traditional intermediary could arrange the same outcome by being the sole writer on a blockchain of its own with the same protocols, meaning a blockchain would be unnecessary for security.

[10] A 51% attack works because readers look for the longest chain of blocks, so despite the fact that small writers are sending reports as well, these reports are initially ignored by readers.

Readers' preferences are as before. Their fundamental preferences for each branch of the fork are given by $u_{i,T} = \tau - \gamma E[a_{J,T-1}|\{y_s\}_{s=1}^T]$. The term $\tau$ is a (small) preference for the longer chain, reflecting the fact that readers prefer a ledger that does not omit the most recent information. The term $E[\gamma a_{J,T-1}|\{y_s\}_{s=1}^T]$ corresponds to the fact that readers' stakes are impacted by the attack and they can essentially reverse their losses from the *previous* distortion of the ledger by forking from a point in the blockchain before the distortion occurred. Readers receive a noisy signal $x_i = \tau + \sigma\eta_i$ of the common value as before, and their only signals of $a_{J,T-1}$ are the public signals $y_T$. The extent to which $y_T$ is informative about $a_{J,T-1}$ in equilibrium determines what readers learn from the public signal. Readers' types are hence given by $\theta_i = (x_i, y_T)$. Small writers receive revenues $\frac{c_j}{C_k}\pi_k L_k$ when writing on ledger $k$ unless a successful attack occurs, in which case they receive zero. Again, because attacks are infrequent, small writers can neglect the possibility of an attack. Proposers' actions are fixed, so we do not model their preferences. When the large writer attacks ledger $A$ successfully at time $T-1$, he receives revenues $\frac{L+a}{M}\pi_{A,T}$ where $\pi_{A,T}$ denotes the participation on ledger $A$ at time $T$. If readers abandon ledger $A$ on the next day, writers get nothing from their attack.

In a period after no attack has occurred and no attack is possible, the equilibrium is as in Proposition 3. Readers prefer the longer chain slightly, so all readers coordinate on that branch of the fork and writers break even. When no attack occurred in the previous period but an attack is possible at $t = 2$, play at $t = 1$ must be the same as in Proposition 3 because readers and writers are not aware of the possibility of an attack.

In order to understand large writers' incentives at $t = 2$ of day $T-1$, we must analyze the equilibrium after an attack at time $T$. At $t = 1$, the public signal $y_T$ is realized. When $y_T = 0$, the equilibrium must be the one described in Proposition 3. Given that readers slightly prefer the longer chain, the attack is successful and large writers profit. The equilibrium is different when $y_T = 1$, however. Let $a^* = E[a_{J,T-1}|y_T = 1]$ and note that $a^* > 0$. Then when $\tau$ is sufficiently small, we have

$$u_i = \tau - \gamma a^* < 0$$

so were are in the same case as Proposition 3 with $\bar{\tau} = \tau - \gamma a^* < 0$. Hence all readers switch to branch $B$ (the fork of the blockchain in which the attack is rolled back) and the attackers receive zero.

We may now analyze the large writer's choices at $t = 2$ when an attack is possible. Of course, the only interesting case is the case in which the large writer chooses $c_J = M$ and $a_J > 0$. We look for conditions under which he never does so in equilibrium. We have argued that whenever $y_T = 1$, the attack is not successful. When the attack is successful, the large writer gets fees $L$ plus the revenue $a$ from the distortion. Therefore, the large writer must solve

$$\max_a (1 - p(a))(L + a)$$

The first-order condition implies

$$1 = \frac{p'(a^*)}{1 - p(a^*)}(L + a^*) \tag{10}$$

This equation implies that since $L = M$, an equilibrium in which writers attempt to steal may exist only when the hazard rate $H(a) \equiv \frac{p'(a)}{1-p(a)}$ is uniformly low, i.e. $H(a) < \frac{1}{L}$ for all $a \in [0, \overline{a}]$. By our earlier assumption that $H(a)$ is nondecreasing, a sufficient condition to ensure $a^* = 0$ is

$$H(0) \geq \frac{1}{L}$$

When the hazard rate is large, the probability of detection is high enough to completely dissuade the large writer from even attempting an attack. Note that this condition is satisfied for sufficiently large $L$, meaning that when the fee earned by blockchain writers is high, even agents with the ability to subvert the network prefer not to attack it because they stand to lose the fee that they would earn through honest writing. The second force that prevents cheating by writers is that even if $a^* > 0$, it may be that profits earned through ledger distortion are negative. This occurs when

$$(1 - p(a^*))(L + a^*) - L = (1 - p(a^*))a^* - p(a^*)L < 0$$

This second condition reflects the fact that even if the large writer's fee is not large enough to dissuade him from distorting the ledger, the cost $L$ of mounting a 51% attack is enough to render the attack unprofitable. In equilibrium, the cost of the attack $L$ is equal to the fee earned through honest writing, but conceptually they are two distinct objects. Proposition 9 summarizes these findings:

**Proposition 9.** *The large writer chooses not to attack the blockchain if and only if*

$$\max_a (1 - p(a))a - p(a)L \leq 0$$

*A sufficient condition that guarantees this inequality will hold is*

$$H(0) \equiv \frac{p'(0)}{1 - p(0)} \geq \frac{1}{L}$$

This bound on $L$ (the first cost of blockchain) characterizes the tradeoff between decentralization and cost efficiency required to maintain correctness.

Proposition 9 has a striking implication. When the probability of detection is sufficiently large, it is unnecessary to set up an expensive fee structure for writers that leads to a large waste of computational resources. Writers will abstain from distorting the ledger regardless because each marginal unit of computational power spent on an attack earns less on average than one spent on writing honestly. The cost of conducting an attack, which is

exactly equal to the fee earned in equilibrium, acts as further protection against attacks. Crucially, in this framework the equilibrium is unique– readers always abandon the ledger after detecting an attack. The uniqueness of equilibrium is a direct consequence of fork competition. If an attack makes all readers worse off they will coordinate on an alternative ledger on which the attack never happened but the rest of the information on the ledger is intact. Competition among writers will cause writers to coordinate on that ledger as well, and the attacker will get nothing. As we will show, this mechanism is quite different from the one that secures a traditional ledger.

### 4.4 Monopolistic ledger security

Now we analyze the case where a monopolist is able to distort its own ledger while facing competition from a fixed outside ledger. The structure of the game is similar to the dynamic blockchain game where the ledger can be attacked by a group of writers. There is a monopolist who discounts payoffs at rate $\delta$, a manager of the outside ledger, and a continuum $i \in [0,1]$ of readers who live for one period. On each day $T$ at $t = 0$, the monopolist proposes a fixed pair $(L_A, \hat{S}_T) \in \mathcal{L} \times \mathcal{S}$ and the outside proposer announces a fixed $L_B \in \mathcal{L}$ and stakes equal to zero. The stakes announced by the monopolist depend on the history up until period $T$ because the actions taken by the monopolist to distort the ledger may also distort the stakes. Here the stake announcement should be interpreted as a set of private signals received by readers corresponding to their stakes in the ledger. At $t = 1$, each writer chooses its own ledger.

As in the blockchain model of security, the monopolist is able to distort the ledger at $t = 2$ of each period. The monopolist chooses an action $a \in [0, \bar{a}]$ at $t = 2$ and immediately receives a payoff of $\pi_{A,T} a$ (in addition to the fees it usually receives). The structure of public signals is also the same as in the blockchain model. On each day $T$, a public signal $y_T \in \{0,1\}$ is observed at $t = 2$ with $\Pr(y_T = 1|a) = p(a_{T-1})$. When the monopolist's distortion is severe, it both affects more agents directly and is more likely to be revealed to the public. Readers' fundamental preferences for ledger $A$ are given by $u_{i,T} = \tau + s_i - \gamma E[a_T | \{y_s\}_{s=1}^T]$, where $s_i$ is reader $i$'s stake on ledger $A$ and readers receive signals $x_i = \tau + \sigma\eta_i$ as usual. As in the example with a blockchain, a reader's utility is decreased when the monopolist distorts the ledger and plays $a_T > 0$. These preferences differ from those in the blockchain security example in an important way. Reader's stakes in the ledger at $T-1$ are not relevant. This is because readers do not have the option to fork to a ledger on which the distortion that occurred at $T-1$ never happened. The monopolist's action to distort the ledger is *final*. Whereas in the blockchain model readers' play was affected by public signals because it was informative about the utility gains from switching to the alternative ledger, in this model public signals matter only because they affect readers' expectations about the continuation play. Expectations of future attacks can affect readers' actions because the monopolist is able to distort the ledger in all periods.

There will be multiple equilibria because we have no mechanism to pin down readers'

expectations of future play. However, we can establish a lower bound on the fee required by the monopolist to ensure that $a = 0$ is played in all periods, which is a proxy for the cost of maintaining a ledger under a centralized intermediary above and beyond the rents extracted due to its competitive advantage. We will assume that readers punish the monopolist in the harshest way possible– they play on ledger $B$ in all future periods after the public signal $y_T = 1$ is realized. In order to ensure this is an equilibrium for readers, it suffices to assume that there is an action $\tilde{a}$ the monopolist can take so that $\max_i s_i - \gamma\tilde{a} - \alpha(L_A - L_B) < 0$, meaning even the type who is most anchored to ledger $A$ by a personal stake in the system prefers to leave the ledger when readers expect $\tilde{a}$ to be played going forward. The expectations that justify this equilibrium, then, are

$$E[a_T|\{y_s\}_{s=1}^T] = \begin{cases} 0 & y_s = 0 \ \forall \ s \leq T \\ \tilde{a} & \exists \ s \leq T, \ y_s = 1 \end{cases}$$

If we wish to derive a lower bound on $L_A$, we may also assume that participation on the monopolist's ledger is $\pi = 1$ whenever $y_s = 0$ for all $s \leq T$. The monopolist's problem is stationary: as long as $y_t = 1$ has not been realized, the monopolist can achieve some value $V$ in expectation, and after $y_t = 1$ is realized the monopolist gets zero. Hence the monopolist solves

$$\max_a \ a + \delta(1 - p(a))V$$

The first-order condition is

$$1 = p'(a^*)\delta V$$

When the monopolist plays $a^*$, we have $V = \frac{L_A + a^*}{1-\delta(1-p(a^*))}$. A sufficient condition for a unique optimum $a^* \in [0, \bar{a}]$ to exist is then just $\frac{d}{da}\frac{p'(a)(L_A+a)}{1-\delta(1-p(a))} > 0$. This condition is similar to the increasing hazard rate assumption made in the previous section. To ensure the monopolist plays $a^* = 0$, we need

$$\frac{\delta p'(0)}{1 - \delta(1 - p(0))} \geq \frac{1}{L_A}$$

For small $\delta$, this condition is usually significantly weaker than the one derived in the previous section for blockchain security, so when the monopolist is punished as harshly as possible, it is not necessary to pay the monopolist as much in fees as blockchain writers. The intuition for this result is simple: while a blockchain writer is punished for misbehavior only through nullification of the profits obtained by attacking the blockchain, a monopolist is punished via the destruction of its franchise value, which consists of all future fees earned through honest play. This result is restated in Proposition 10.

**Proposition 10.** *There is a threshold value of $L_A$ such that the monopolist never distorts the ledger:*

$$L_A = \frac{1 - \delta(1 - p(0))}{\delta p'(0)}$$

Proposition 10 says that the monopolist's ability to distort the ledger imposes an endogenous lower bound on its fees above and beyond the bound due to the barriers to entry resulting from readers' stakes on the ledger. The less likely the monopolist is to be detected in its deviations, the higher this bound must be. It is worth noting that while the fee required for correctness provides a sharp bound on the cost of a blockchain (due to fork competition), the fee charged by a monopolist may be far from this bound. If the rents earned by a monopolist are large, there is no force to push the fee it charges down to the level derived in Proposition 10.

There are several drawbacks that make the security of a traditional ledger less robust than that of a blockchain, however. First, in a setting with a traditional ledger equilibrium is not unique, so while it may be the case that under the harshest possible punishment scheme it is not necessary to pay an intermediary large fees to obtain ledger security, the equilibrium fee required to ensure good behavior may be much higher. Second, the signal structure $p(a)$ may well be more revealing for a blockchain than for a traditional ledger, since blockchains are designed specifically to provide transparency about attacks on the ledger. Finally, in this case the assumption that signals $y_T$ are either zero or one is not without loss of generality. A richer signal structure would lead to even greater multiplicity that would allow the monopolist to "nickel and dime" readers by proving to them that although she is distorting the ledger, she is not doing so to the extent that readers would prefer to switch to a competitor and lose their stakes in the established ledger. Fork competition is thus important in securing a ledger as well as forcing competition among writer compensation schemes.

Another interesting difference between securing a blockchain and securing a traditional ledger is that the equilibrium in the blockchain game is unique and independent of the nature of public signals while in the traditional setting there are multiple equilibria, and the set of equilibria depends on the information structure. This dichotomy stems from the fact that past actions can be "rewound" by a fork on a blockchain but not on a traditional ledger. The equilibrium in the blockchain game is backwards-looking: readers decide whether they want to switch to a different ledger on which an attack never occurred, meaning their actions are determined by their expectations of malevolent writers' *past* actions. The equilibrium in the game with a traditional ledger is forward-looking: the public signal acts as a coordinating device that determines readers' expectations of the intermediary's *future* actions, but there is no possibility of undoing past events. The uniqueness of equilibrium in the blockchain game can be seen as a security feature. When any attack is revealed to the public, it will always be undone via a blockchain fork. Multiplicity of equilibrium in the game with a centralized intermediary means there are no such guarantees in the traditional setting.

# 5 Discussion

In this section, we informally discuss some practical matters related to the application of blockchain and distributed ledger technology that we do not address formally in our model. The first (and most important) issue is that while distributed ledgers are useful for transferring *ownership* of assets, they do not necessarily guarantee transfers of *possession*. Consider a simple example in which a buyer wishes to purchase a car from a seller on a blockchain. In this case, ownership of the car would be represented by a token in the seller's account on the blockchain. The blockchain's writers would be able to transfer ownership of the token to the buyer, but they would not be able to verify that the buyer was physically in possession of the car after the transaction. To ensure transfers of possession, it is necessary to have some entity that enforces contracts on the blockchain when those contracts involve the transaction of physical assets. This type of enforcement would likely be the role of the government, which would then have to explicitly make reference to the cases in which it would enforce blockchain contracts.

The need for an enforcer alongside a distributed ledger raises two issues. First, while several commentators claim that distributed ledger technology will benefit those in developing countries without strong property rights, one needs to identify why property rights are weak in the first place before concluding that a distributed ledger is the solution. If the government is overly bureaucratic and incapable of setting up good institutions to track property rights, then a distributed ledger is an effective alternative. However, if the government is corrupt to the point that it would outright refuse to enforce some contracts in a publicly available database, a distributed ledger will be useless. Again, the readers of the ledger are the ultimate source of discipline, so a distributed ledger is useful only insofar as it helps them to discipline a corrupt government (through greater disclosure of information, most likely). If the enforcer is itself a private firm, such as a bank that enforces debt obligations, it may be optimal for the enforcer to maintain the ledger as well. The enforcer will have an incentive to fulfill its obligation for fear of losing the rents it earns by maintaining the ledger.

The second issue is the incorporation of blockchains into the legal code. A government cannot simply commit to enforce all contracts on a blockchain because the blockchain may fork. The government could say it will enforce all contracts so long as certain policies are followed, which prevents hard forks that change blockchain's rules. Of course, this enforcement policy would be detrimental because it would essentially destroy the potential for competition between ledgers. Furthermore, if an attack on the blockchain were to occur, such as the one on the Ethereum blockchain in 2016, the government would have enormous power to resolve the issue in its own favor.

# 6 Conclusion

We present a general model of ledger competition and apply it to understand when a blockchain is more economically beneficial than a traditional ledger managed by a centralized intermediary. Our analysis of the tradeoffs between centralized and decentralized record-keeping is guided by the Blockchain Trilemma. We focus the analysis of our static model on the tradeoff between decentralization and cost efficiency. We find that with a blockchain, the rules that are most beneficial to readers of the ledger always emerge in equilibrium via hard forks. This surprising result arises due to the combination of portability of information and competition between writers that are possible with a blockchain. Readers are not reluctant to abandon an older version of a blockchain because all the information contained in the old blockchain is contained in the new one with updated policies, so writers compete to write on the blockchain preferred by readers. A centralized intermediary that maintains a traditional ledger, on the other hand, is able to extract rents from readers by exploiting their desire to keep their stakes in the established ledger. When the coordination motive is sufficiently strong, entry by a competing traditional ledger is ruled out altogether, which suggests that blockchains may help lower intermediaries' rents in situations where the coordination motive is strong. This result suggests that, for example, retail platforms like Amazon's might be better suited to a blockchain, since the coordination motive among buyers and sellers is powerful. Decentralized ledgers do have costs, however. In addition to the waste of resources required by proof-of-work, there is a second cost of blockchains: miscoordination inefficiencies. Blockchains forks can lead to a split of the community and too many competing ledgers in equilibrium.

We also present an extension of our static model to a repeated setting. This extension allows us to show that there is no possibility of collusion among writers of a permissionless blockchain in the repeated game. Free entry of writers rules out any sort of dynamic reward and punishment scheme, so writers must play myopically in every period. Thus the optimal outcome for readers emerges with a permissionless blockchain even in the repeated game. By contrast, collusion is possible among writers of a permissioned blockchain because they earn rents in equilibrium. With a permissioned blockchain, it is not always the case that writers' rents are competed down by hard forks. Permissioned blockchains, then, do not break the Trilemma because they fail to fully meet the decentralization criterion due to lack of free entry.

We also explicitly examine the costs of incentivizing writers to report honestly (correctness). On the one hand, centralized intermediaries are incentivized dynamically by ensuring that the future profits they will earn are high enough to guarantee they do not want to risk losing them. On the other hand, blockchain writers must be incentivized statically by raising the proof-of-work to the point that attacks become unprofitable.

We highlight the important distinction between ownership and possession. Blockchains can only effect transfers of ownership, but the discipline imposed by the security of ownership on a blockchain can also prevent bad actors from defaulting on delivery of possession.

In this paper, we have outlined the incentive mechanisms of two particularly important types of ledgers. What we have not developed so far is a general theory of the interactions between writers and readers on an arbitrary ledger. An investigation of the optimal technological restrictions on the communication between writers and readers is a fruitful avenue for future research.

# Appendix A: Proofs

**Proof of Proposition 1**:

*Proof.* This proposition is an immediate consequence of Theorem B.4 and Proposition B.5 in Appendix B. □

**Proof of Proposition 2**:

*Proof.* If neither action is dominant for type $\theta$, then clearly it must be that $1-\hat{\pi}(\theta) \in [0,1]$. When $1-\hat{\pi}(\theta) \in \{0,1\}$, then there are just two equilibria: one in which all agents play $A$ and one in which all play $B$. For all other values of $\theta$, there will be three equilibria. Since $1-\hat{\pi}(\theta) \in (0,1)$, there are equilibria in which all agents play $A$ or $B$. There is also an equilibrium in which $1-\hat{\pi}(\theta)$ agents play $B$ and $\hat{\pi}(\theta)$ agents play $A$ (by the definition of $\hat{\pi}(\theta)$, which is the point at which type $\theta$ agents are indifferent between $A$ and $B$). □

**Proof of Proposition 3**:

*Proof.* We prove the proposition by backwards induction.

**t=2**: At $t=2$, writers know the value of $\pi_k$, $k \in \{A,B\}$. We show that $C_k = \pi_k L_k$ in equilibrium. Suppose first that $C_k < \pi_k L_k$. Then there exists a writer $j$ such that $c_j < 1$, but writer $j$ could make profits by setting $c_j = 1$ because

$$\frac{1}{C_k}\pi_k L_k - 1 > 0$$

Now suppose $C_k > \pi_k L_k$. This means that any writer $j$ for whom $c_j > 0$ would benefit by setting $c_j = 0$, since

$$\frac{c_j}{C_k}\pi_k L_k - c_j = (\frac{1}{C_k}\pi_k L_k - 1)c_j < 0$$

Hence $C_k = \pi_k L_k$.

**t=1**: We will guess and verify that in any equilibrium, $L_B < L_A$. Writers' optimal play at $t=2$ implies that $\frac{C_k}{\pi_k} = \underline{C}$ for each branch of the fork. Then it must be that

$$1-\hat{\pi}(s) \leq \frac{1}{2} + \kappa^{-1}\big(s - (g(L_A) - g(L_B))\big)$$

in equilibrium. Furthermore, all readers have identical stakes, so $s_i = 0$ for all $i$. According to Theorem B.4, type $s_i$'s cutoff signal $k(s_i)$ is $x_i = g(L_A) - g(L_B)$, so as long as $\tau \leq 0$, all readers will have such signals when $\sigma$ is sufficiently close to zero. Therefore all readers play $B$.

**t=0**: Now we confirm our guess that $L_B < L_A$. The equilibrium derived above shows that whenever $L_B < L_A$, the proposer obtains a payoff of $K - g(L_B)$. It is never possible for the proposer to obtain a higher payoff by choosing $L_B \geq L_A$. Furthermore, the proposer can never choose $L_B < \underline{C}$, since in that case readers would know that the disparity in utility between branches $A$ and $B$ is at least $H$. When $H$ is sufficiently large, it is dominant to play $A$. Then it must be that the proposer chooses the lowest possible $L_B$ in order to maximize payoffs, so $L_B = \min\{L : L \in \mathcal{L},\ L \geq \underline{C}\}$.

□

**Proof of Proposition 4**:

*Proof.* These statements follow from Proposition B.7. □

**Proof of Proposition 5**:

*Proof.* The equilibrium follows from Propositions B.5 and B.6. □

**Proof of Proposition 6**:

*Proof.* These properties are a result of equation (8). □

**Proof of Proposition 7**

*Proof.* This equation follows from equations (6), (7), and (8). □

**Proof of Proposition 8**

*Proof.* First we show that at any history $h^{T,2}$, on either branch $k$ of the fork, the total computing power contributed by writers must be $\pi_k L_k$. Suppose that $C_k < \pi_k L_k$. Then there must be some writer $j$ who contributes $c_j < 1$. By deviating to $c_j = 1$ on branch $k$, this writer can achieve positive profits in the current period. Furthermore, this writer's deviation does not affect any publicly observable signal in the future history, since the writer is of measure zero. An analogous argument shows that $C_k$ cannot be greater than $\pi_k L_k$, so $C_k = \pi_k L_k$ at any history.

Second, we must check that proposers play static best responses. Given that both readers and writers play the same strategies that they do in the static game, a proposer can maximize her flow of payoffs by playing $L_B = \min\{L : L \in \mathcal{L}, L \geq \underline{C}\}$. □

**Proofs of Propositions 9 and 10**:

*Proof.* (Under construction) □

## Appendix B: Global Games with Heterogeneous Preferences

We begin by describing the model considered in this Appendix. There is a continuum of players $i \in [0,1]$ who play a one-shot coordination game in which they choose between two options, $A$ and $B$. Players' fundamental preferences consist of heterogeneous private values $\theta_{i,k} \in \mathbb{R}$ for choice $k \in \{A, B\}$ and common values $\tau_k$ for each choice. Players also obtain utility $\kappa\pi_k$ by playing $k$ if $\pi_k$ other players make the same choice. Let $\theta = \theta_A - \theta_B$, $\tau = \tau_A - \tau_B$, and $\pi = \pi_B$. We assume that $\theta$ is iid across players with distribution $F(\theta)$. For now, we assume $F$ is a discrete distribution with finite support but later take the limit of a continuous distribution $F$. Players' preferences can be described by the function

$$v(\theta, \tau, \pi) = \theta + \tau + \kappa(1 - 2\pi)$$

When $v(\theta, \tau, \pi) > 0$, it is a best response for a player of type $\theta$ to choose $A$. Conversely, a player of type $\theta$ should choose $B$ if $v(\tau, \theta, \pi) < 0$.

Henceforth we will assume that players have incomplete information about the common value $\tau$. We assume players have an improper uniform prior over $\tau$[11]and receive signals $s_i = \tau + \sigma\eta_i$ ($\sigma > 0$), where $\eta_i$ is iid across players and independent of $\tau$. The noise term $\eta_i$ is distributed with CDF $H(\eta)$ with support on the interval $[-\frac{1}{2}, \frac{1}{2}]$. In what follows, we will frequently consider the limit $\sigma \to 0$.

By Theorem 5 in Milgrom and Roberts (1990), this is a supermodular game. Therefore, if the signal profile is $\mathbf{s}$, there are largest and smallest rationalizable strategy profiles $\underline{\mathbf{k}}(\mathbf{s})$ and $\overline{\mathbf{k}}(\mathbf{s})$. Furthermore, every equilibrium strategy profile $\mathbf{k}(\mathbf{s})$ satisfies $\underline{\mathbf{k}}(\mathbf{s}) \leq \mathbf{k}(\mathbf{s}) \leq \overline{\mathbf{k}}(\mathbf{s})$. Given that agents observe only their own signals, it must be that all agents play cutoff strategies: for each type $\theta$, there is a signal $k(\theta)$ such that $\theta$ plays $A$ if $s_i > k(\theta)$ and plays $B$ if $s_i < k(\theta)$. When agents play a cutoff equilibrium $\mathbf{k}$, we will denote the expected utility derived from playing $A$ for the cutoff type $k(\theta)$ by $E[v|\mathbf{k}, k(\theta)]$. The equilibrium condition is just

$$E[v|\mathbf{k}, k(\theta)] = 0 \tag{11}$$

for all $\theta$. The following lemma establishes that there is a unique equilibrium in cutoff strategies. The proof is essentially the same as that in Drozd and Serrano-Padial (2017).

**Lemma B.1.** *If $\boldsymbol{k}$ is a cutoff strategy equilibrium and $\Delta > 0$, then $E[v|\boldsymbol{k}, k(\theta)] < E[v|\boldsymbol{k} + \Delta, k(\theta) + \Delta]$.*

[11]The results do not change if we instead assume $\tau$ is uniformly distributed on an interval of finite length as long as that interval is sufficiently large.

*Proof.*

$$
\begin{aligned}
E[v|\mathbf{k}, k(\theta)] &= E[\theta + \tau + \kappa(1-2\pi)|\mathbf{k}, k(\theta)] \\
&= \int_{k(\theta)-\frac{\sigma}{2}}^{k(\theta)+\frac{\sigma}{2}} \left(\tau + \theta + \kappa(\int_{\theta'} (2H(\frac{k(\theta')-\tau}{\sigma}) - 1)dF(\theta'))\right) h(\frac{k(\theta)-\tau}{\sigma})d\tau \\
&< \int_{k(\theta)-\frac{\sigma}{2}}^{k(\theta)+\frac{\sigma}{2}} \left(\tau + \Delta + \theta + \kappa(\int_{\theta'} (2H(\frac{k(\theta')-\tau}{\sigma}) - 1)dF(\theta'))\right) h(\frac{k(\theta)-\tau}{\sigma})d\tau \\
&= \int_{k(\theta)+\Delta-\frac{\sigma}{2}}^{k(\theta)+\Delta+\frac{\sigma}{2}} \left(\tau + \theta + \kappa(\int_{\theta'} (2H(\frac{k(\theta')+\Delta-\tau}{\sigma}) - 1)dF(\theta'))\right) h(\frac{k(\theta)+\Delta-\tau}{\sigma})d\tau \\
&= E[v|\mathbf{k}+\Delta, k(\theta)+\Delta]
\end{aligned}
$$

□

From Lemma 1 it is immediate to see that there is a unique equilibrium. Suppose that $\underline{\mathbf{k}} < \overline{\mathbf{k}}$. Let $\hat{\Delta} = \max_{\theta} \overline{k}(\theta) - \underline{k}(\theta)$, and let $\hat{\theta}$ be the value of $\theta$ that achieves this maximum. Then

$$E[v|\underline{\mathbf{k}}, \underline{k}(\hat{\theta})] < E[v|\underline{\mathbf{k}} + \hat{\Delta}, \underline{k}(\hat{\theta}) + \hat{\Delta}] \leq E[v|\overline{\mathbf{k}}, \overline{k}(\hat{\theta})]$$

where the last inequality comes from the fact that $\overline{\mathbf{k}} \leq \underline{\mathbf{k}} + \hat{\Delta}$.

In what follows, it will be useful to define the following object: for all $\theta \in \Theta$, where $\Theta$ is some set contained in the support of $F$, set

$$\psi(\tau, \Theta) = \frac{1}{\sum_{\Theta} f(\theta)} \sum_{\Theta} H\left(\frac{k(\theta)-\tau}{\sigma}\right) f(\theta)$$

This expression is the expectation of the number of agents in $\Theta$ who play $B$ given the common value $\tau$. We now prove an important lemma (called the "Belief Constraint") about the function $\psi$ due to Sakovics and Steiner (2012) and Drozd and Serrano-Padial (2017):

**Lemma B.2.** *For any subset $\Theta \subset supp(F)$ and any $z \in [0,1]$,*

$$\frac{1}{\sum_{\Theta} f(\theta)} \sum_{\Theta} \Pr(\psi(\tau, \Theta) < z | s = k(\theta)) f(\theta) = z$$

*Proof.* Begin by defining "virtual types" $\delta(s, \theta) = s - k(\theta)$. This reduces the two-dimensional type space to a one-dimensional one. Agents play $A$ whenever $\delta(s, \theta) > 0$ and $B$ when

$\delta(s,\theta) < 0$. With this definition,

$$\psi(\tau,\Theta) = \Pr(\delta(s,\theta) < 0|\tau,\Theta)$$

First we show that $\Pr(\psi(\tau,\Theta) \leq z|\delta(s_i,\theta_i) = 0) = z$. This property is due to Morris and Shin (2003). For brevity, we will denote $\delta(s_i,\theta_i)$ by $\delta_i$.

Define $\tilde{\eta}_i = \frac{\delta_i - \tau}{\sigma}$, and denote the distribution of $\tilde{\eta}$ conditional on $\theta \in \Theta$ by $\tilde{H}$. This variable is iid across players. We have

$$\begin{aligned}
\Pr(\psi(\tau,\Theta) < z|\delta_i = 0) &= \Pr(\Pr(\delta_j > 0|\tau) < z|\delta_i = 0) \\
&= \Pr\left(\Pr\left(\tilde{\eta_j} < -\frac{\tau}{\sigma}\right) < z|\delta_i = 0\right) \\
&= \Pr\left(1 - \tilde{H}\left(-\frac{\tau}{\sigma}\right) < z|\delta_i = 0\right) \\
&= \Pr(1 - \tilde{H}(\tilde{\eta}_i) < z) \\
&= \Pr(\tilde{\eta}_i > \tilde{H}^{-1}(1-z)) \\
&= 1 - \tilde{H}(\tilde{H}^{-1}(1-z)) = z
\end{aligned}$$

Now to complete the proof, observe that

$$\Pr(\psi(\tau,\Theta) < z|\delta = 0) = \sum_{\Theta} \Pr(\psi(\tau,\Theta) < z|s = k(\theta)) \Pr(\theta|\delta = 0,\Theta)$$

Given the uniform prior over $\tau$, the information environment is translation-invariant, so

$$\Pr(\theta|\delta = 0,\Theta) = \frac{f(\theta)}{\sum\limits_{\Theta} f(\theta')}$$

That is, knowing $\delta = 0$ yields no additional information about $\theta$, since each type is equally likely to observe $\delta = 0$. Hence

$$\frac{1}{\sum\limits_{\Theta} f(\theta)} \sum_{\Theta} \Pr(\psi(\tau,\Theta) < z|s = k(\theta)) f(\theta) = z$$

as desired. □

Up until this point, none of the results have depended on taking the limit $\sigma \to 0$. Now we specialize to the case considered in the text where $\sigma$ becomes arbitrarily small and define $\mathbf{k}^\sigma$ to be the threshold equilibrium played for variance parameter $\sigma$. Correspondingly, we denote a specific type $\theta$'s cutoff by $k^\sigma(\theta)$. We then define

$$A_\theta(z|\mathbf{k}^\sigma,\Theta) = \Pr(\psi(\tau,\Theta) < z|s = k^\sigma(\theta))$$

to be the *strategic belief* of type $\theta$– that is, it is the probability that type $\theta$ assigns to the event that a proportion less than $z$ of agents in $\Theta$ play action $B$. Now we prove the final lemma we will need before proving the main result (due to Drozd and Serrano-Padial (2017)).

**Lemma B.3.** *There exist a unique partition $\Theta_1, \ldots, \Theta_S$ and thresholds $k_1 > \cdots > k_S$ such that, as $\sigma \to 0$, $k^\sigma(\theta) \to k_i$ uniformly for all $\theta \in \Theta_i$ and all $i \in \{1, \ldots, S\}$. Furthermore, the cutoffs $k_i$ satisfy the limit conditions*

$$\int_0^1 \left( k_i + \theta + \kappa(1 - 2 \sum_{\Theta_j, j<i} f(\theta') - 2z \sum_{\Theta_i} f(\theta')) \right) dA_\theta(z|\boldsymbol{k}, \Theta_i) = 0$$

*where $\boldsymbol{k}$ denotes the set of limit cutoffs.*

*Proof.* Fix $\tilde{\sigma} > 0$ and define a partition of types $\Theta_1, \ldots, \Theta_S$ by placing two types $\theta, \theta'$ in the same equivalence class whenever $|k^{\tilde{\sigma}}(\theta) - k^{\tilde{\sigma}}(\theta')| < \tilde{\sigma}$. Define $Q_\theta^{\tilde{\sigma}}(\chi|\mathbf{k}^{\tilde{\sigma}}, z) = \Pr(\tau \leq \chi | s = k^{\tilde{\sigma}}(\theta), \psi(\tau, \Theta_i) = z)$ (for $\theta \in \Theta_i$) to be type $k^{\tilde{\sigma}}(\theta)$'s belief about $\tau$ conditional on the event that a proportion $z$ of players in the same equivalence class of the partition play $B$. We have

$$E[v|\mathbf{k}^{\tilde{\sigma}}, k^{\tilde{\sigma}}(\theta)] = \int_0^1 \int_{k^{\tilde{\sigma}} - \frac{\tilde{\sigma}}{2}}^{k^{\tilde{\sigma}} + \frac{\tilde{\sigma}}{2}} \left( \chi + \theta + \kappa(1 - 2 \sum_{\Theta_j, j<i} f(\theta) - 2z \sum_{\Theta_i} f(\theta)) \right) dQ_\theta^{\tilde{\sigma}}(\chi|\mathbf{k}^{\tilde{\sigma}}, z) dA_\theta(z|\mathbf{k}^{\tilde{\sigma}}, \Theta_i)$$

The term $\chi$ in the integrand is bounded by $k^{\tilde{\sigma}} \pm \frac{\tilde{\sigma}}{2}$, so

$$E[v|\mathbf{k}^{\tilde{\sigma}}, k^{\tilde{\sigma}}(\theta)] \leq \int_0^1 (k^{\tilde{\sigma}} + \frac{\tilde{\sigma}}{2} + \theta + \kappa(1 - 2 \sum_{\Theta_j, j<i} f(\theta) - 2z \sum_{\Theta_i} f(\theta))) dA_\theta(z|\mathbf{k}^{\tilde{\sigma}}, \Theta_i) \quad (12)$$

and

$$E[v|\mathbf{k}^{\tilde{\sigma}}, k^{\tilde{\sigma}}(\theta)] \geq \int_0^1 (k^{\tilde{\sigma}} - \frac{\tilde{\sigma}}{2} + \theta + \kappa(1 - 2 \sum_{\Theta_j, j<i} f(\theta) - 2z \sum_{\Theta_i} f(\theta))) dA_\theta(z|\mathbf{k}^{\tilde{\sigma}}, \Theta_i) \quad (13)$$

Note that as $\tilde{\sigma} \to 0$, the right-hand side of (2) converges to the right-hand side of (3) as long as $dA_\theta$ is bounded, which is shown in Lemma 8 of Drozd and Serrano-Padial (2017).

Now, for each $i$, take some arbitrary $\theta_i \in \Theta_i$ and set $k_i = k^{\tilde{\sigma}}(\theta_i)$. As $\sigma$ is taken to zero from $\tilde{\sigma}$, set cutoffs $\hat{\mathbf{k}}_\sigma^{\tilde{\sigma}}$ so that $\Delta_{\theta_i, \theta_i'} = \frac{k_i - \hat{k}_\sigma^{\tilde{\sigma}}(\theta_i')}{\sigma} = \frac{k_i - k^{\tilde{\sigma}}(\theta_i')}{\tilde{\sigma}}$ for all $\theta_i' \in \Theta_i$. Note that

$A_\theta(z|\hat{\mathbf{k}}_\sigma^{\tilde{\sigma}}, \Theta_i)$ is constant as $\sigma \to 0$ under these transformed cutoffs. Then as $\sigma \to 0$,

$$E[v|\mathbf{k}_\sigma^{\tilde{\sigma}}, k^{\tilde{\sigma}}(\theta)] \to \int_0^1 \left( k_i + \theta + \kappa(1 - 2\sum_{\Theta_j, j<i} f(\theta) - 2z\sum_{\Theta_i} f(\theta)) \right) dA_\theta(z|\mathbf{k}_\sigma^{\tilde{\sigma}}, \Theta_i)$$

$$= k_i + \theta + \kappa(1 - 2\sum_{\Theta_j, j<i} f(\theta)) - 2\kappa\sum_{\Theta_i} f(\theta) \int_0^1 z dA_\theta(z|\mathbf{k}^{\tilde{\sigma}}, \Theta_i)$$

Fix $\epsilon > 0$. If we pick $\tilde{\sigma}$ close to zero, we can ensure that

$$\left| E[v|\mathbf{k}_\sigma^{\tilde{\sigma}}, k^{\tilde{\sigma}}(\theta)] - E[v|\mathbf{k}^{\tilde{\sigma}}, k^{\tilde{\sigma}}(\theta)] \right| < \epsilon$$

for all $\sigma < \tilde{\sigma}$. This is because the solution of the system of equations $E[v|\mathbf{k}^{\tilde{\sigma}}, k^{\tilde{\sigma}}(\theta)] = 0$ can be seen as the correct choice of $k^{\tilde{\sigma}}(\theta_i)$ and $\Delta_{\theta_i, \theta_i'}$ for each $i$ holding $W_i$ fixed (which is possible as long as $\tilde{\sigma}$ is sufficiently small). The solution to this system of equations lies in a compact set, so for small $\tilde{\sigma}$ the limit condition will not differ from the equilibrium condition $E[v|\mathbf{k}^{\tilde{\sigma}}, k^{\tilde{\sigma}}(\theta)] = 0$ by more than $\epsilon$. Therefore the limit condition holds as $\tilde{\sigma} \to 0$. □

We now prove the main theorem.

**Theorem B.4.** *In the limit $\sigma \to 0$, the equilibrium strategies are given by a monotone partition $\Theta_1, \ldots, \Theta_S$ of $\Theta$ and cutoffs $k_1 > \cdots > k_S$ such that*

*(i) For all $\theta \in \Theta_i$, $k(\theta) = k_i$;*

*(ii)* $-\underline{\theta}_i - \kappa(1 - 2F(\underline{\theta}_i)^-) \leq k_i \leq -\overline{\theta}_i - \kappa(1 - 2F(\overline{\theta}_i))$

*(iii)* $k_i + \kappa\left(1 - 2\sum_{\Theta_j, j<i} f(\theta) - \sum_{\Theta_i} f(\theta)\right) = -E[\theta|\theta \in \Theta_i]$ *for all $i$.*

*where $\underline{\theta}_i = \min \Theta_i$, $\overline{\theta}_i = \max \Theta_i$.*

*Proof.* Point (i) is a consequence of Lemma 3. We now show the partition is monotone. Suppose that $\theta_1 > \theta_2$ but $\theta_2 \in \Theta_j$, $\theta_1 \in \Theta_m$ with $j > m$. Then

$$-\theta_1 \geq k_m + \kappa(1 - 2\sum_{\Theta_n, n \leq m} f(\theta)) \geq k_j + \kappa(1 - 2\sum_{\Theta_i, i<j} f(\theta)) \geq -\theta_2$$

a contradiction. From this it immediately follows that

$$-\underline{\theta}_i - \kappa(1 - 2F(\underline{\theta}_i)^-) \leq k_i \leq -\overline{\theta}_i - \kappa(1 - 2F(\overline{\theta}_i))$$

which is point (ii).

To see (iii), note that Lemma 3 implies that for all $\theta \in \Theta_i$,

$$0 = k_i + \theta + \kappa\left(1 - 2\sum_{\Theta_j, j<i} f(\theta) - 2\sum_{\Theta_i} f(\theta) \int_0^1 z dA_\theta(z|\mathbf{k}, \Theta_i)\right)$$

Multiplying by $\frac{f(\theta)}{\sum_{\Theta_i} f(\theta')}$ on both sides and moving the $\theta$ term to the left-hand side,

$$\begin{aligned} -E[\theta|\theta \in \Theta_i] &= k_i + \kappa\left(1 - 2\sum_{\Theta_j, j<i} f(\theta) - 2\sum_{\Theta_i} f(\theta) \int_0^1 z d\left(\frac{1}{\sum_{\Theta_i} f(\theta)} \sum_{\Theta_i} f(\theta) dA_\theta(z|\mathbf{k}, \Theta_i)\right)\right) \\ &= k_i + \kappa\left(1 - 2\sum_{\Theta_j, j<i} f(\theta) - \sum_{\Theta_i} f(\theta)\right) \end{aligned}$$

where the second line follows from Lemma 2, the belief constraint. This is precisely the desired result. □

Equipped with Theorem 4, we may now prove some properties of equilibria when the distribution $F$ satisfies certain conditions. We consider three scenarios:

1. $F$ is continuous and $\theta + \kappa(1 - 2F(\theta))$ is monotonically increasing;
2. $F$ has a symmetric, single-peaked density $f$ and $\theta + \kappa(1 - 2F(\theta))$ is non-monotonic;
3. $F$ is a two-point discrete distribution.

The next three propositions characterize the equilibrium in these three cases. Henceforth we assume $H$ is the uniform distribution on $[-\frac{1}{2}, \frac{1}{2}]$.

**Proposition B.5.** *When $F$ is continuous and $\theta+\kappa(1-2F(\theta))$ is a monotonically increasing function, the cutoffs $k(\theta)$ satisfy $k(\theta) = -\theta - \kappa(1 - 2F(\theta))$.*

*Proof.* We show that the partition described in Theorem 4 must consist of singletons in this case. Suppose that $\theta_1 < \theta_2$ are the boundaries of equivalence class $i$ of the partition. By property (ii) of Theorem 4, we have

$$-\theta_1 - \kappa(1 - 2F(\theta_1)) \leq k_i \leq -\theta_2 - \kappa(1 - 2F(\theta_2))$$

By assumption, $-\theta_1 - \kappa(1 - 2F(\theta_1)) > -\theta_2 - \kappa(1 - 2F(\theta_2))$, so this is impossible. Hence the partition is indeed a collection of singletons, and $k(\theta) = -\theta - \kappa(1 - 2F(\theta))$ (again by property (ii)). □

**Proposition B.6.** *Suppose $F$ has a symmetric, single-peaked density $f$ and $\theta+\kappa(1-2F(\theta))$ is non-monotonic. Let $\hat{\theta} = \arg\max_\theta f(\theta)$. The equilibrium is characterized by a parameter $\Delta$ such that*

- $k(\theta) = -\theta - \kappa(1 - 2F(\theta))$ *for* $\theta \notin [\hat{\theta} - \Delta, \hat{\theta} + \Delta]$,
- $k(\theta) = -\hat{\theta} - \kappa(1 - F(\hat{\theta} - \Delta) - F(\hat{\theta} + \Delta))$ *for* $\theta \in [\hat{\theta} - \Delta, \hat{\theta} + \Delta]$,
- *The parameter* $\Delta$ *is the unique nonzero solution to*

$$\Delta = \kappa(F(\hat{\theta} + \Delta) - F(\hat{\theta} - \Delta))$$

*Proof.* Observe that under the assumptions on $F$, there must be only one interval $[\underline{\theta}, \overline{\theta}]$ where $\theta + \kappa(1 - 2F(\theta))$ is decreasing. All $\theta$ in this interval must belong to the same equivalence class of the partition described in Theorem 4. We show this by contradiction. If $\theta \in [\underline{\theta}, \overline{\theta}]$ is at the upper boundary of an equivalence class $\Theta_i$, then by point (ii) of Theorem 4 we have

$$k_i \leq -\theta - \kappa(1 - 2F(\theta)) < k_{i+1}$$

which is impossible because the cutoffs are monotonically decreasing in $i$.

Hence the entire increasing region $[\underline{\theta}, \overline{\theta}]$ belongs to a single equivalence class of the partition. At the boundaries of the equivalence class containing that interval, $k(\theta)$ must be continuous (which follows by again applying the argument showing that there cannot be two equivalence classes containing points in the increasing region). By the argument in Proposition B.5, there cannot be an equivalence class of the partition containing only points in the decreasing region, so it must be that the partition consists of a single equivalence class $[\underline{\theta}, \overline{\theta}]$ containing all values of $\theta$ in the increasing region and singletons for all $\theta$ outside that interval.

Let $k(\theta) = k$ for $\theta \in [\underline{\theta}, \overline{\theta}]$. Point (iii) of Theorem 4 implies that

$$k = -\Big(E[\theta|\underline{\theta} \leq \theta \leq \overline{\theta}] + \kappa(1 - F(\underline{\theta}) - F(\overline{\theta}))\Big) \tag{14}$$

Continuity of the cutoff at the boundaries of the interval implies

$$-\big(\underline{\theta} + \kappa(1 - 2F(\underline{\theta}))\big) = k = -\big(\overline{\theta} + \kappa(1 - 2F(\overline{\theta}))\big)$$

Rearranging these expressions, we find

$$\frac{\overline{\theta} + \underline{\theta}}{2} = E[\theta|\underline{\theta} \leq \theta \leq \overline{\theta}] \tag{15}$$

$$\frac{\overline{\theta} - \underline{\theta}}{2} = \kappa(F(\overline{\theta}) - F(\underline{\theta})) \tag{16}$$

The symmetry of the density $f$ and (5) imply that $E[\theta|\underline{\theta} \leq \theta \leq \overline{\theta}] = \hat{\theta}$ and there exists $\Delta$ such that $\underline{\theta} = \hat{\theta} - \Delta$, $\overline{\theta} = \hat{\theta} + \Delta$. Then (4) reduces to

$$k(\theta) = -\hat{\theta} - \kappa(1 - F(\hat{\theta} - \Delta) - F(\hat{\theta} + \Delta))$$

for $\theta \in [\hat{\theta} - \Delta, \hat{\theta} + \Delta]$ and (5) reduces to

$$\Delta = \kappa(F(\hat{\theta} + \Delta) - F(\hat{\theta} - \Delta))$$

as desired. Finally, we must show that there is a unique nonzero solution $\Delta$ to the above equation. The derivative of the left-hand side with respect to $\Delta$ is 1, and the derivative of the right-hand side is $2\kappa f(\hat{\theta} + \Delta)$ by the symmetry of $f$. The derivative of the right-hand side is greater than 1 for $\Delta = 0$ (since $\theta + \kappa(1 - 2F(\theta))$ is increasing at $\hat{\theta}$) and monotonically decreasing towards zero, so there is a unique crossing point. □

**Proposition B.7.** *When $F$ is a two-point distribution with support $\{\theta_L, \theta_H\}$ (and $\theta_L < \theta_H$) such that $\Pr(\theta = \theta_L) = \mu$, $\Pr(\theta = \theta_H) = 1 - \mu$, the equilibrium cutoffs are*

- $k(\theta) = -(\mu\theta_L + (1 - \mu)\theta_H)$ *for all $\theta$ if $\theta_H - \theta_L \leq \kappa$,*
- $k(\theta_L) = -(\theta_L + (1 - \mu)\kappa)$ *and $k(\theta_H) = -(\theta_H - \mu\kappa)$ if $\theta_H - \theta_L > \kappa$.*

*Proof.* There are two possible cases when the support of $F$ consists of two points: either $k(\theta_L) = k(\theta_H)$ or $k(\theta_L) > k(\theta_H)$. We first suppose that the cutoffs are equal and derive the restriction $\theta_H - \theta_L = \kappa$ in that case. Recall from Lemma 3 that when $k(\theta_H) = k(\theta_L) = k$,

$$0 = k + \theta_H + \kappa\left(1 - 2\int_0^1 z dA_{\theta_H}(z|k)\right) = k + \theta_L + \kappa\left(1 - 2\int_0^1 z dA_{\theta_L}(z|k)\right)$$

We will derive an expression that allows us to evaluate the integrals on the right-hand side in terms of the cutoffs for small $\sigma$.

Consider the equilibrium with finite, nonzero $\sigma$. We have

$$\begin{aligned}
E[v|\mathbf{k}^\sigma, k^\sigma(\theta_H)] = &\int_{k^\sigma(\theta_H)-\frac{\sigma}{2}}^{k^\sigma(\theta_H)+\frac{\sigma}{2}} \left(\tau + \theta_H + \kappa\right) h\left(\frac{k(\theta_H) - \tau}{\sigma}\right) d\tau \\
&- 2\kappa \int_{k^\sigma(\theta_H)-\frac{\sigma}{2}}^{k^\sigma(\theta_H)+\frac{\sigma}{2}} \left(\mu(1 - H(\frac{k(\theta_L) - \tau}{\sigma})) + (1 - \mu)(1 - H(\frac{k(\theta_H) - \tau}{\sigma}))\right) h\left(\frac{k^\sigma(\theta_H) - \tau}{\sigma}\right) d\tau \\
&= \kappa^\sigma(\theta_H) + \theta_H + \kappa(1 - \mu(1 + \Delta_{H,L}^2) - (1 - \mu)) \\
&= \kappa^\sigma(\theta_H) + \theta_H - \kappa\mu\Delta_{H,L}^2
\end{aligned}$$

where the third line uses the fact that $H$ is the uniform distribution on $[-\frac{1}{2}, \frac{1}{2}]$. Similarly, we find

$$E[v|\mathbf{k}^\sigma, k^\sigma(\theta_L)] = \kappa^\sigma(\theta_L) + \theta_L + \kappa(1 - \mu)\Delta_{H,L}^2$$

Suppose that as $\sigma \to 0$, $\Delta_{H,L} \equiv \frac{k(\theta_H)-k(\theta_L)}{\sigma} \to \xi$. Then these equations imply

$$k + \theta_H - \kappa\mu\xi^2 = k + \theta_L + \kappa(1-\mu)\xi^2$$

so

$$\theta_H - \theta_L = \kappa\xi^2$$

Clearly, $\xi^2 \in [0,1]$, so we obtain

$$\theta_H - \theta_L \leq \kappa$$

when the cutoffs are equal.

Now consider the case in which the cutoffs are not equal. Then when $\sigma \to 0$, the cutoff type $k(\theta_H)$ is certain that all type $\theta_L$ players received signals below $k(\theta_L)$, and type $k(\theta_L)$ is certain that all type $\theta_H$ players received signals above $k(\theta_H)$. The equilibrium conditions are then

$$0 = k(\theta_H) + \theta_H + \kappa(1 - 2\mu - (1-\mu)) = k(\theta_L) + \theta_L + \kappa(1-\mu)$$

by part (iii) of Theorem 4. Rearranging, we get

$$k_L - k_H = (\theta_H - \kappa\mu) - (\theta_L + \kappa(1-\mu))$$

Given that $k_L > k_H$, we must have

$$\theta_H - \theta_L > \kappa$$

which completes the proof. □

# References


[1] Rossella Argenziano. Differentiated networks: Equilibrium and efficiency. *RAND Journal of Economics*, 39(3):747–769, 2008.

[2] Bruno Biais, Christophe Bivière, Matthieu Bouvard, and Catherine Casamatta. The blockchain folk theorem. Working Paper, 2017.

[3] Hans Carlsson and Eric van Damme. Global games and equilibrium selection. *Econometrica*, 61(5):989–1018, 1993.

[4] Christian Catalini and Joshua Gans. Some simple economics of the blockchain. 2017.

[5] Jonathan Chiu and Thorsten Koeppl. The economics of cryptocurrencies– bitcoin and beyond. Working Paper, 2017.

[6] William Lin Cong and Zhiguo He. Blockchain disruption and smart contracts. Working Paper, 2017.

[7] William Lin Cong, Ye Li, and Neng Wang. Tokenomics: Dynamic adoption and valuation. Working Paper, 2018.

[8] Douglas Diamond. Financial intermediation and delegated monitoring. *The Review of Economic Studies*, 51(3):393–414, 1984.

[9] Lukasz Drozd and Ricardo Serrano-Padial. Credit enforcement cycles. Working Paper, 2017.

[10] David Easley, Maurenn O'Hara, and Soumya Basu. From mining to markets: The evolution of bitcoin transaction fees. Working Paper, 2017.

[11] David Frankel, Stephen Morris, and Ady Pauzner. Equilibrium selection in global games with strategic complementarities. *Journal of Economic Theory*, 108(1):1–44, 2003.

[12] Drew Fudenberg and Jean Tirole. Perfect bayesian equilibrium. *Journal of Economic Theory*, 53(2):236–260, 1991.

[13] Arthur Gervais, Ghassan Kharame, Karl Wüst, Vasileios Glykantzis, Hubert Ritzdorf, and Srdjan Capkun. On the security and performance of proof of work blockchains. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016.

[14] BitFury Group. Public versus private blockchains. Whitepaper, 2015.

[15] Gur Huberman, Jacob Leshno, and Ciamac Moallemi. Monopoly without a monopolist: An economic analysis of the bitcoin payment system. Working Paper, 2017.

[16] June Ma, Joshua Gans, and Rabee Tourky. Market structure in bitcoin mining. NBER Working Paper, 2018.

[17] Paul Milgrom and David Roberts. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, 58(6):1255–1277, 1990.

[18] Stephen Morris and Hyun Song Shin. Unique equilibrium in a model of speculative currency attacks. *American Economic Review*, 88(3):587–597, 1998.

[19] Stephen Morris and Hyun Song Shin. Global games: Theory and applications. 2001.

[20] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. 2008.

[21] Emiliano Pagnotta and Andrea Buraschi. An equilibrium valuation of bitcoin and decentralized network assets. Working Paper, 2018.

[22] Joszef Sakovics and Jakub Steiner. Who matters in coordination problems. *American Economic Review*, 102(7):3439–3461, 2012.

[23] Linda Schilling and Harald Uhlig. Some simple bitcoin economics. Working Paper, 2018.

[24] Michael Sockin and Wei Xiong. A model of cryptocurrencies. Working Paper, 2018.