Blockchain Economics[*]

Joseph Abadi and Markus Brunnermeier[†]

May 1, 2018

Abstract

When is record-keeping better arranged through distributed ledger technology (DLT) than through a traditional centralized intermediary? The answer depends on users' incentives to abandon an established ledger in favor of a competitor. Network externalities amplify centralized intermediaries' ability to extract rents from "anchored" users who have stakes in their ledgers. DLT allows for the replication of information on a competing ledger and removes impediments to switching, which is especially welfare-improving when network externalities are strong. Blockchain, a type of DLT, accomplishes this information replication through "fork competition." Entry competition among blockchain miners promotes fork competition that benefits users. In a repeated game setting, free entry implies blockchain miners are incentivized statically whereas centralized intermediaries face dynamic incentives for good behavior. 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

Fintech has the potential to revolutionize finance. Some observers even argue that blockchain technology is an invention as groundbreaking as the invention of double-entry bookkeeping in fourteenth-century Italy. Blockchains could drastically change record-keeping of financial transactions and ownership data. Traditionally, centralized entities have been responsible for maintaining records. A centralized intermediary implements rules for the operation of its ledgers to maximize profits, which distorts the incentives of the ledger's users and imposes potentially significant economic costs. These intermediaries typically have a monopoly on the information written on the ledger, preventing their rents from simply being competed away. The importance of coordination among the ledger's users cements the intermediary's stranglehold over them.

The emerging fintech industry has provided us with a radical alternative to record information: *distributed ledger technology.* Blockchains are a particular type of distributed ledger written by decentralized, usually anonymous groups of agents rather than known centralized parties. Free entry of ledger writers implies that competition will drive writers' profits to zero, meaning there are no incentives to set distortionary rules that benefit writers of the ledger. Achieving consensus on such a ledger is important given that, in principle, anyone may write essentially anything on it. Consensus is reached by making the ledger publicly viewable and verifiable. Because of the public nature of the information on the blockchain, any group of parties that wants to operate under a different set of rules for recording information may seamlessly replicate the information and start a new blockchain at no cost. Competition among blockchain ledgers is therefore far more intense than competition among traditional ledgers.

However, as anonymously written ledgers, blockchains require identity management in the form of high computational costs. That is, those who write on the ledger must perform a computationally expensive task in order to do so. Otherwise it would be possible for individual entities to masquerade as a large number of entities, subverting the democratic nature of the distributed ledger. A critical issue is whether the benefits of eliminating monopolistic rent extraction are worth the computational costs required to run a viable blockchain.

The key question we address is for which ledgers it is more economically efficient to use a monopolistic intermediary and and for which it is better to use a blockchain. In financial markets, centralized entities such as banks are currently in charge of recording payments, exchanges intermediate securities issuance and trade, and governments oversee land registries. It is often argued that "tech giants" that operate centralized platforms are the rent-extracting monopolists of the 21st century. In the tech space, platforms such as Alibaba record users' credit histories and retailers like Amazon maintain ratings for vendors. Writers of a given type of ledger can be incentivized to use rules that benefit the ledger's readers (users) by other competing writers or by the readers who take actions in response to the proposed rules. Writers may compete by implementing policies under which they earn lower rents, but the ultimate power to discipline the writing community lies with the readers. If readers are unhappy with the operation of the ledger, they may simply desert it. That is, they can stop taking actions in response to the reported history, at which point any attempt to extract profits from them becomes irrelevant.

Blockchains expand the scope of these incentives by allowing for "fork competition," in which writers replicate all the information contained on the existing blockchain but change its rules. This type of competition is far more intense than traditional competition among currencies à la Hayek. In particular, for cryptocurrencies, each owner of a unit of the established currency is entitled to a unit of a new forked currency, so there are no impediments to switching. In a competition between traditional fiat currencies, owners of the established currency must find someone to take the other side of the trade if they wish to move into the new currency, so it is impossible to simply transfer one's "stake" to the new currency.

We model readers' and writers' ledger choice problem as a global game in the spirit of Carlsson and van Damme (1993) and Morris and Shin (1998). In both a setting with a blockchain and one with a traditional ledger, readers attempt to coordinate by choosing between two ledgers. The distinction between a blockchain and a traditional ledger amounts to a restriction on who may write on a given ledger and what information is contained on each ledger. For example, for a monopolistic intermediary, the intermediary is the sole writer whereas the readers are its clients and possibly a regulatory agency. For a cryptocurrency, the writers who perform the required expensive computations (miners) may enter freely and the readers are the users of the currency. When agents choose between traditional ledgers, readers are anchored to an "established" ledger because they are averse to losing the information it contains by switching to the competing ledger. On the other hand, when agents choose between two branches of a blockchain "fork," information is perfectly replicated across branches, meaning readers have no fundamental preference for one branch over the other except to the extent that the policies imposed on each branch are different.

Our main result is that with traditional record-keeping, readers are anchored to an established ledger by the cost of losing the information it contains. Network externalities that arise due to the coordination motives among readers amplify this effect. A *no-entry condition* implies that competitors to the established ledger never enter when readers' stakes in an established ledger are large or network externalities are strong. In this case, a traditional ledger is maximally costly in the sense that the monopolist is able to extract rents from all readers of the ledger. Furthermore, even when entry is possible, readers' stakes in the established ledger give that ledger a competitive advantage over the entrant's. There is therefore a gap between the fee charged by the incumbent and the fee charged by the entrant. However, in the presence of even partial competition, the role of coordination is vastly different. With an entrant, network externalities serve as a disciplining device– *both* intermediaries charge lower fees because they fear the domino effect that occurs when even a single marginal reader switches to the competing ledger.

By contrast, we find that a blockchain is beneficial to readers because of a synergy between *replicability of information* and *competition among writers*. Given that information can be replicated across ledgers, readers' stakes in the system no longer anchor them to the established ledger. The lack of a fundamental anchor eliminates the amplification role played by network externalities. When a fork (i.e. new set of rules) of the blockchain is proposed, writers replicate the information on the blockchain and compete to write on the branch of the fork that readers prefer. Readers then see that there are enough writers to cryptographically secure the new ledger and attempt to coordinate on the ledger that they prefer. If there is complete information about readers' preferences, there are multiple equilibria, and in some cases readers coordinate the established ledger even though they all prefer the new one. When there is even an arbitrarily small amount of incomplete information about readers' preferences, though, there is a unique equilibrium in which readers select the ledger they prefer. On a blockchain, the equilibrium of the ledger choice game always yields the policies most beneficial to readers.

Although blockchain technology makes it easier for competing ledgers to succeed over existing ones, blockchain may make switching to a competitor so easy that inefficient miscoordination will occur in equilibrium. When readers disagree about how the blockchain should operate, it may be the case that some readers switch to a competing ledger whereas others stick with the existing blockchain. In fact, this situation is quite common empirically. Given that readers do not internalize network externalities, such a split of the blockchain community may be detrimental to readers' welfare. We highlight that this inherent instability of blockchain is most pronounced when readers' preferences are sufficiently heterogeneous or when network externalities are weak, which again suggests that blockchains are most useful when coordination motives among readers are strong.

In addition to the two polar cases of centralized and decentralized record-keeping that we consider, there is a third type of ledger called a permissioned blockchain. Writers on a permissioned blockchain are granted special writing permissions, and the reading and writing protocols provide incentives that lie somewhere in between those faced on anonymous blockchains and those faced by centralized entities. The writers on a permissioned blockchain are known, so there is no need for costly identity management. Nevertheless, a permissioned blockchain can be forked with all of the information in the established ledger carried over to the new one, just like an anonymous (permissionless) blockchain. We find that in a dynamic setting where the ledger choice game is played repeatedly, a permissioned blockchain does not always give rise to rules that are most beneficial to readers. Writers earn profits in equilibrium, so they can collude via a dynamic reward and punishment scheme to prevent deviations to writing on a competing ledger with lower fees. On a permissionless blockchain, on the other hand, this sort of dynamic collusion scheme is impossible. The free entry condition implies zero profits, which in turn means writers must play the unique static equilibrium in every period.

We study an extension of our model in which we allow writers to distort the ledger by "cheating" and breaking the proposed rules. We highlight two distinct channels through which the writer(s) can be disciplined. Writers have *static* incentives to act honestly because readers may discover they are making fraudulent reports. Readers then could immediately cease to act in response to those reports, thereby nullifying any gain from dishonest reporting. If writers have franchise values, they also have *dynamic* incentives to comply with the rules: readers may threaten to permanently exit the ledger, which destroys any franchise value the writers might have.

We show that the tradeoff between provision of static and dynamic incentives is precisely the tradeoff between a blockchain and a traditional ledger in the context of cheating. When writers distort a blockchain ledger, readers will prefer to create a fork of the blockchain on which the fraud never occurred. Knowing this, writers will compete to create such a fork, rendering the previous distortion of the ledger useless. When a traditional intermediary distorts its ledger, there is no scope for competition by other writers. The readers must discipline the writer by abandoning the ledger, which destroys the writer's franchise value. We derive a minimum required fee that an intermediary must earn in order to ensure it does not distort its own ledger. This fee is higher when it is unlikely that deviations by the intermediary will be discovered, meaning the intermediary has a greater incentive to cheat. With a blockchain, the security criterion is sufficient decentralization of writing. The network will be secure only when it would take a large group of writers to overpower the rest. Intuitively, when there are many writers involved in an attack, the probability of detection is high because each one has incentives to steal too much (just as each producer in an oligopoly has incentives to produce too much). Sufficient decentralization implies a bound on fees that is increasing in the required number of writers.

Finally, we informally make the important point that in many settings transfers of possession as well as ownership must be guaranteed. 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 he holds the deed, his ownership of the house will be enforced. This concern is especially important in developing countries, where tracking and enforcing property rights is often an issue. Broadly speaking, blockchain can help to solidify property records when existing institutions are weak, but it does nothing to prevent an overtly corrupt government from refusing to enforce contracts.

Related Literature. The paper most closely related to ours is Biais et al. (2017), which studies 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 determined by readers' preferences and the forks that do occur are unambiguously beneficial. Cong and He (2017) focuses 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. 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. 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. 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.

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 even if there are no writers who are compelled to play honestly, the network still becomes secure when there is a sufficient number of writers. This result obtains because the readers can threaten to abandon the ledger, rendering any attempt to steal useless. 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. 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. 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.

An important question is whether we actually need PoW in order for the blockchain to function correctly. After all, it would be desirable to have a secure method of record-keeping that allows for competition between ledgers without incurring the substantial costs associated with PoW. While some claim that any attack on a blockchain is unprofitable because readers would detect the attack and ignore updates to the ledger, in this case it would be possible to secure the ledger using a single writer, meaning PoW provides a superfluous layer of security. That is, when readers are able to provide sufficient discipline, having a mechanism by which anonymous writers discipline each other is pointless. By the same token, we will find that if readers themselves are able to discipline a monopolist enough to prevent the monopolist from extracting rents, there will be no need for a PoW blockchain. The mainstream models of blockchain security, such as Gervais et al. (2016), assume that as long as no single entity has a majority of the computing power, the blockchain will be secure because no writer will ever assist another in an attack. This logic embeds the assumption that writers do not collude with each other. However, if this were the case, again PoW would be rendered pointless. A permissioned blockchain with a sufficient number of writers would suffice to secure the blockchain under these conditions. Some in the blockchain community have expressed doubts that a permissioned blockchain could be trusted because of the possibility of collusion among writers, which is where a PoW-based system finally becomes important. We show that the free entry condition implies there will be no dynamic collusion between writers in any equilibrium on a PoW-based blockchain, which is not necessarily the case for a permissioned blockchain because there is no entry. In sum, PoW is useful when (1) readers cannot adequately discipline writers, and (2) writers are able to collude with one another.

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. An example of a setting in which agents need to coordinate on a ledger is when a fork in a blockchain arises. We then apply the model to two distinct situations: one in which a ledger is maintained on a blockchain and one in which a ledger run by a monopolist competes against an outside option.

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 four periods, $t = 0, 1, 2, 3$. 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).

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 π_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

$$\theta_i = \theta_{i,A} - \theta_{i,B} = (s_{i,A} + \epsilon_{i,A}) - (s_{i,B} + \epsilon_{i,B})$$

where in our applications θ_i is drawn from some distribution that induces beliefs $Q(\theta_i) \equiv \Pr(\theta_{i'} \leq \theta_i | \theta_i)$.[1] Here $s_{i,k}$ is meant to represent the *stake* that agent i has in ledger k and ϵ_i is an idiosyncratic preference for ledger k (typically assumed to be small but used for equilibrium selection). 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. The population CDF of types is denoted $P(\theta)$. Writers may also be assigned types $\theta_j^w \in \Theta^w$ that provide them with information about readers' preferences.[2]

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 privately informed about their stakes when they receive their types θ_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

[2] We assume $\tilde{Q}(\theta'|\theta) \equiv \Pr(\theta'' \leq \theta'|\theta)$ is uniformly continuous for technical reasons.

[2] The structure of writers' information is specific to each application and will be described in each example below.

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: After observing proposals at $t = 0$, writers are assigned types θ_j^w and choose a ledger $w(j) \in F_j \subset \{A, B\}$ on which to write.

t=2: Readers first observe writers' choices and their own types θ_i. They then choose a ledger $r(i) \in \{A, B\}$ in which to participate.

t=3: Writers on ledger $k \in \{A, B\}$ take actions $a_j \in \mathcal{A}(\pi_k)$ and 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 taken by writers at $t = 3$. The actions taken by writers at $t = 3$ may be important to readers for several reasons. If the action at $t = 3$ 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 $t = 3$ corresponds to a distortion of the ledger, readers will prefer ledgers that have not been distorted. Let π_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 $t = 3$. A reader who chooses ledger k obtains utility $u(\theta_{i,k}, \pi_k, L_k, a_k)$. We assume that u takes the form

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

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

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

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

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

We will define $\hat{\pi}(\theta, a, L_a, L_B)$ to be the π_A such that $\Delta = 0$ when a reader's type is θ and the fundamental parameters of the ledgers are L_A, L_B. We will usually suppress the dependence on a, L_A, L_B. According to this definition,

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

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

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

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 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 when Condition SC holds, 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 when Condition SC holds, equilibria will take a "cutoff" form: there will be a threshold value $\hat{\theta}$ such that all agents with $\theta_i < \hat{\theta}$ choose ledger B and all readers with $\theta_i > \hat{\theta}$ choose ledger 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 exists $\hat{\theta}$ such that all readers with $\theta_i > \hat{\theta}$ choose $r(i) = A$ and all readers with $\theta_i < \hat{\theta}$ choose $r(i) = B$ whenever Condition SC holds.*

The proof of Proposition 1 relies on standard techniques from the global games literature. The logic behind the proof is as follows. In this setup, there are certain types θ 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.

Condition SC is critical for the proof of Proposition 1. Intuitively, it guarantees that low types θ (who prefer B) believe that there are many readers who like B even more than they do, and high types θ (who prefer A) believe that there are many readers who like A even more. This condition is necessary for the contagion argument outlined above to work. When a type θ who likes B believes there are many others who like B more, agents of type θ will want to choose B as long as all other lower types choose B as well, so the contagion that begins in the dominance region reaches type θ. There are two ways to get Condition SC to hold. The first is to assume readers' preferences are heterogeneous. In our analysis of traditional ledgers run by centralized intermediaries, we simply assume readers' stakes in the ledger are sufficiently dispersed. When we study readers' choices between two branches of a blockchain fork, stakes will play no role because each reader will have the same stake on both ledgers. In that case, we get Condition SC to hold by introducing incomplete information as in Carlsson and van Damme (1993). Each reader will receive a noisy signal of others' preferences. When this signal is very precise, each agent will believe that half of all other readers received a lower signal and half received a higher one, meaning Condition SC will be satisfied.

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 type θ, generically there are three equilibria: one in which all agents 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.

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 π_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 $t = 1$, each writer j chooses ledger $w(j) \in \{A, B\}$. At $t = 3$, writer j chooses an amount of computational power $c_j \leq 1$ to contribute to that ledger. 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 π_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. Writers also pay an arbitrarily small cost ν if they choose a ledger k at $t = 1$ such that $\pi_k = 0$. This cost can be thought of as a small cost of producing code to write on a blockchain that is never used in the future.

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, but in our discussion of attacks on the blockchain we outline how it might be endogenized.

There is incomplete information about readers' preferences. Each reader has an idiosyncratic fundamental preference ϵ_i for ledger A. Preferences ϵ_i are independently and identically distributed uniformly on the interval $[\bar{\theta} - \sigma, \bar{\theta} + \sigma]$, where $\bar{\theta}$ is random and we take the limit $\sigma \to 0$. The value of θ 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. This small amount of noise in preferences gives rise to a type distribution $\theta_i \sim U[\bar{\theta} - \sigma, \bar{\theta} + \sigma]$, 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. To see this, note that when σ is sufficiently small, Condition SC is satisfied: $Q(\theta_i) \to \frac{1}{2}$ as $\sigma \to 0$.[4] Writers also have incomplete information about readers' preferences. Each writer j receives a signal $\theta_j^w \sim U[\bar{\theta} - \sigma, \bar{\theta} + \sigma]$.

Readers' preferences are summarized by

$$\Delta(\pi) = \big(\frac{1}{2}\kappa\pi_A + \theta_i - g(L_A) - h(\frac{C_A}{\pi_A})\big) - \big(\frac{1}{2}\kappa\pi_B - g(L_B) - h(\frac{C_B}{\pi_B})\big)$$

since each reader's stake is the same on both ledgers. Here κ 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

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

where π represents participation on ledger A, as before.

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 writers W_k who chose ledger k at $t = 1$ for $k \in \{A, B\}$. At $t = 3$, players observe π.

[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 $\bar{\theta}$ is smooth and has full support to guarantee uniform convergence of $Q(\theta)$ to $\frac{1}{2}$.

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 $\bar{\theta} \leq 0$. 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.*

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.

This result shows that there is an *endogenous* channel through which blockchain reduces the cost of maintaining a ledger: the synergy between *replicability of information* and *competition among writers.* When information can be replicated on 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.

3.3 Competition between traditional ledgers

In this section, we analyze a competition between a ledger maintained by a monopolist and an outside ledger. 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 $\mathcal{O}$ on ledger B. In this case, the writers are also the proposers $P_A = \mathcal{M}$ and $P_B = \mathcal{O}$. Each writer

[5] Note that the hypothesis $\bar{\theta} \leq 0$ 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 $\bar{\theta} = 0$.

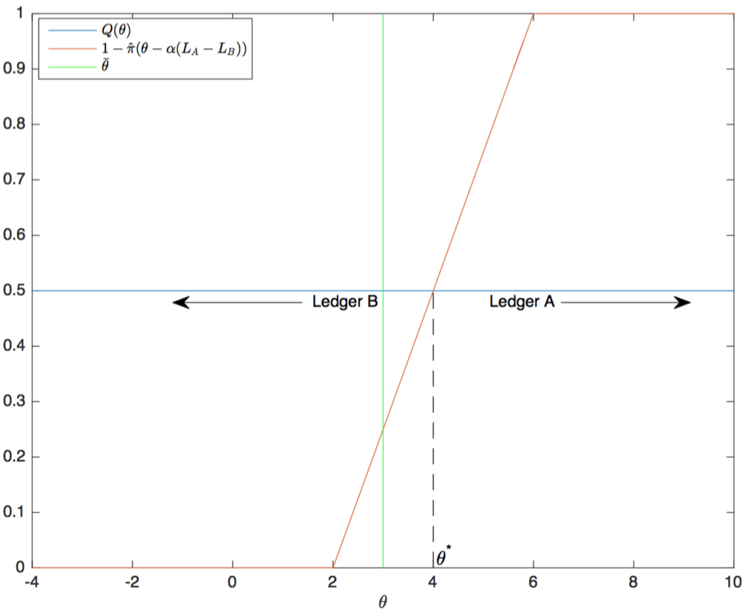

Figure 1: An example of an equilibrium of the blockchain game. Here $\bar{\theta} = 3$, $g(L) = \alpha L$, $L_A = 50$, $L_B = 10$, and $\kappa = 4$. The green line represents the actual CDF of types, which is concentrated in a small interval around $\bar{\theta}$.

may only writer on her own ledger. At $t = 1$, the monopolist may choose a fundamental parameter $L_A \in \mathcal{L}$, where $\mathcal{L} \subset \mathbb{R}_+$ is a compact set, but the outside writer is constrained to choose a fixed L_B. The monopolist and outside writer choose stakes $\hat{S}$ and 0, respectively. The restriction that the outside writer 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 $t = 3$.

Readers have preferences summarized by

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

Here we use a linear function αL 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 θ_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 $[S - \frac{d}{2}, S + \frac{d}{2}]$.[6] Here S is the average stake and d is the dispersion in stakes. Readers could have arbitrarily small fundamental preferences for ledger A as in the case of blockchain, but when there is dispersion in stakes this type of incomplete information is irrelevant. It is therefore sufficient to assume that S is large enough that Condition SC is satisfied. This condition is simply $d > \kappa$.

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

$$\max_{L_A \in \mathcal{L}} \pi L_A$$

where π denotes participation on ledger A. By Proposition 1, when the monopolist selects L_A, all readers for whom $1 - \hat{\pi}(\theta) > Q(\theta) = \frac{\theta - S}{d} + \frac{1}{2}$ choose to remain on ledger A. Figure 2 illustrates this situation. To find the cutoff type θ^* who is indifferent between remaining on the monopolist's ledger and leaving, we solve

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

which implies

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

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

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

[6] We could also derive similar results under the assumption that the cross-sectional distribution of s differs from the conditional quantile $Q(s)$. What is critical is that the conditional quantile satisfies Condition SC. Numerical results are similar even when Q is non-uniform.

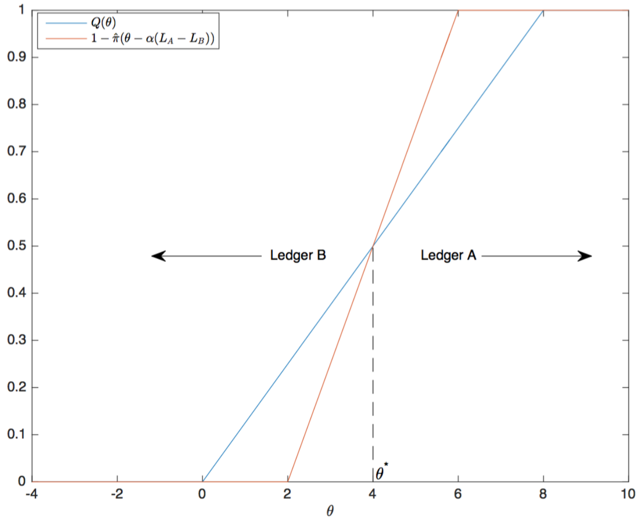

Figure 2: Participation on the monopolist's ledger when stakes are distributed uniformly on $[S-\frac{d}{2},S+\frac{d}{2}]$ with $S=4$, $d=8$, $\kappa=4$, $L_A=40$, $L_B=0$, and $\alpha=0.1$.

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

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

From equation (4), several results are immediate.

Proposition 4. *The optimal fee L_A chosen by the monopolist is*

1. *Increasing in the mean stake S on its ledger;*
2. *Increasing in the dispersion of stakes d;*
3. *Increasing in the fees charged on the outside ledger L_B;*
4. *Decreasing in the strength of the coordination motive κ;*
5. *Decreasing in the sensitivity of readers' utility to fees α.*

All readers participate on the monopolist's ledger when $S+\alpha L_B\geq\frac{3}{2}(d-\kappa)$.

The logic behind this proposition is straightforward. 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 is detrimental to the monopolist's business. When 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. Proposition 4 then suggests that a blockchain will be most valuable in breaking up monopolies in which (1) readers have large and heterogenous stakes, or (2) participation in the ledger is relatively unimportant for readers. Another way to state this second result is that network externalities work as a disciplining device against the monopolist. Competition between writers is able to lower readers' costs when the coordination motive cannot be used as a disciplining device against a monopolist.

3.4 Traditional competition with an entrant

Here we analyze a situation in which an entrant monopolist who maintains a traditional ledger competes against a fixed outside ledger. The monopolist is an entrant in the sense that readers have stakes on the outside ledger. The model is exactly as in the previous section except that the monopolist $\mathcal{M}$ chooses L_B and L_A is held fixed. The stakes on ledger A are $\hat{S}$, as before, and the stakes on the monopolist's ledger are equal to zero. Now the monopolist solves

$$\max_{L_B \in \mathcal{L}} (1-\pi)L_B$$

Rearranging an expression from the previous section, we have that when stakes are distributed uniformly in the interval $[S-\frac{d}{2}, S+\frac{d}{2}]$

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

so

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

The entrant's problem is then

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

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 $-\frac{3}{2}(d-\kappa) \leq S - \alpha L_A \leq \frac{1}{2}(d-\kappa)$.[7]

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 κ, the mean stake on the incumbent's ledger S, and readers' sensitivity to fundamentals α. 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.

3.5 Competition between two active monopolists

Here we analyze competition between two monopolists, an incumbent and an entrant, who manage traditional ledgers. The model is identical to the one studied in the previous section, but neither L_A nor L_B is fixed. We denote the monopolists by $\mathcal{M}_A$ and $\mathcal{M}_B$. At $t = 0$, each monopolist k proposes L_k. As before, readers' stakes on ledger A are uniformly distributed on the interval $[S - \frac{d}{2}, S + \frac{d}{2}]$ and the stakes on ledger B are equal to zero.

The logic of the optimization problems in the previous sections still holds because each monopolist takes the other's choice as given. Therefore, in order to find an equilibrium we need to simultaneously solve equations (4) and (5). This yields

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

[7] If $S - \alpha L_A$ is greater than the upper bound of that interval, the incumbent extracts such small rents from readers that the entrant could not attract any readers even by setting $L_B = 0$. If $S - \alpha L_A$ is greater than the upper bound, the entrant finds it optimal to capture the entire market.

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 $0 \leq \pi_A, \pi_B \leq 1$. 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,

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

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 5. *The no-entry bound on the average stake S is decreasing in the strength of the coordination motive κ 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 κ.*

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 κ 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 $\mathcal{M}_A$ can extract higher rents than monopolist $\mathcal{M}_B$ because readers' stake in ledger A acts as an inertial force preventing them from leaving.

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 replicability of information and competition between writers drives fees down as far as they can go while still providing sufficient incentives for writers to contribute enough computing power to make the network secure. The equilibrium outcome is independent of the distribution of readers' stakes and their desire for coordination. Welfare losses come mostly from the waste of computational resources. Under traditional monopolistic competition, both the monopolists may charge high fees, especially if network externalities are weak. The incumbent further enjoys high rents because of its monopoly on information, which is detrimental to readers' welfare.

3.6 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.[8] 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

[8] 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.

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. 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.

3.7 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.

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

$$\theta_i = \eta_i + \epsilon_i$$

where $\eta_i \in \{0, \eta\}$. The type η_i reflects a preference for forking: readers with $\eta_i = \eta$ dislike all forks equally, and readers with $\eta_i = 0$ are not averse to forking the existing blockchain.[9] Types η_i are independently and identically distributed across readers with $\Pr(\eta_i = \eta) = \mu$. Types ϵ_i are distributed uniformly in the interval $[-\sigma, \sigma]$ as before. Readers observe both η_i and ϵ_i. Under these conditions, the function $Q(\theta)$ is not well-defined because the distribution of types is two-dimensional. Nevertheless, we can define the function $Q(\theta|\eta_i)$ as

$$Q(\theta|\eta_i) = \begin{cases} \frac{1}{2}\mu & \eta_i = \eta \\ \frac{1}{2}(1-\mu) & \eta_i = 0 \end{cases}$$

which denotes the mass of readers of type η_i whose types are expected to be below θ in the eyes of a reader of type θ.

Note that if there exists $L_B \in \mathcal{L}$ such that $L_B \geq \underline{C}$ and $g(L_B) - g(L_A) > \eta$, 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 $g(L_B) - g(L_A) < \eta$. We find that when μ lies in a certain range, we again obtain a unique equilibrium among readers at $t = 2$. This result is summarized in Proposition 6.

Proposition 6. *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 μ of readers are of type $\eta_i = \eta$.*

[9] We adopt this specification for simplicity. Allowing for η_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.

Then if $\eta \geq \frac{\kappa}{2}$ and $\mu \in [1 - 2\kappa^{-1}(g(L_A) - g(L_B)), 2\kappa^{-1}\big(\eta - (g(L_A) - g(L_B))\big)]$, there is a unique equilibrium starting from $t = 2$ in which all readers of type $\eta_i = \eta$ play on branch A and all others play on branch B.

This proposition essentially shows that when (1) readers' fundamental aversion to forking is strong relative to the coordination motive and (2) the proportion of readers who dislike forking is in some intermediate range, the blockchain is vulnerable to a hard fork that splits the community. As long as those conditions are satisfied, such a split is the unique rationalizable outcome. Intuitively, when network externalities are weak relative to some readers' dislike of forks, readers who dislike forks will still prefer not to leave the existing ledger even if all other readers join the new fork. Put another way, network externalities 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.

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 δ. Histories of this game are defined recursively. Let $\mathcal{H}^1 = \{\emptyset\}$. Then define

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

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 writers chose branch A, 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 7. *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. 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. 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 $T = 1, 2, \dots,$ and each day consists of subperiods $t = 0, 1, 2, 3$ 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 δ 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 $t = 1$, writers announce which fork of the blockchain they will support. At $t = 2$, after learning writers' decisions, readers individually choose a fork $k \in \{A, B\}$ of the blockchain. Writing on each branch of the fork occurs at $t = 3$. Readers may choose a branch of the fork only so long as at least one writer supports that branch.

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 π_k, surplus parameter L_k, and W_k writers.

Now we show that when δ 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

$$\frac{L_A}{L_B} \geq \delta + (1-\delta)M \tag{8}$$

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 δ 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 reduces the costs of interacting with intermediaries. 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.

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 not susceptible to fraud by a single bad actor. However, in principle a large group of writers may conspire to (at least temporarily) fool others into accepting a ledger that is in some way invalid. In this section, we analyze the security of both traditional ledgers maintained by monopolists and blockchains. We outline a simple model of security in a blockchain and compare the security of a blockchain to that of a ledger written by a monopolist.

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 are large writers indexed by $J \in \{1, \ldots, M\}$, each of whom has computing capacity 1, and a continuum $j \in [0, M]$ of infinitesimally small writers with computing power dj. The large writers represent exactly half of the total computing power in the network. This assumption is meant to capture "51% attacks" in which an entity or group of entities controlling a majority of a blockchain's computing power mount a malicious attack on the network in order to reap financial gains. We will also assume the large writers live 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 $L = 2M$ so that in an equilibrium with no attacks, 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. Subperiods $t = 1$ and $t = 2$ are as in the benchmark model. Writers choose one branch of the fork at $t = 1$ after receiving signals and readers choose a ledger at $t = 2$ after learning their types.[10]

The main difference from the benchmark model is at $t = 3$. 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, large writers choose an action $a_J \in [0, \overline{a}]$ as well as computing power at $t = 3$. 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, all large writers must choose $a_J > 0$ (meaning each one joins the 51% attack) and $c_J = 1$, so that the computing power provided by large writers is sufficient to overwhelm the rest of the network. The type of attack modeled here is one in which large writers create an invalid fork of the blockchain on which they distort the ledger while small writers write on a valid fork. Readers are initially fooled by large writers' reports[11]and transact according to the invalid chain because it has greater proof-of-work.[12]

On each day $T > 0$, a public signal $y_T \in \{0, 1\}$ is revealed. The signal takes value 1 with probability $p\hat{a}_{T-1}$, where $\hat{a}_{T-1}$ is the average of the actions a_J played by writers at $T - 1$. 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 $\hat{a}_{T-1}$.

Readers' preferences are as before. Their fundamental preferences for each branch of the fork are given by $\theta_i = \tau - \gamma E[\hat{a}_{T-1} + \hat{a}_T | y_T] + \sigma\epsilon_i$, where $\gamma > 0$ is a constant. The term τ 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 $\hat{a}_{T-1}$ corresponds to the fact that readers can essentially reverse their losses from the *previous* distortion of the ledger by forking from a point in the blockchain before the distortion occurred. The term $\hat{a}_T$ is present because readers are also concerned that an attack may occur in the current period. However, in this setting where attacks occur infrequently, this term may be ignored. Again $\sigma\epsilon_i$ is an arbitrarily small noise term. Small writers receive revenues $\frac{c_j}{C_k}\pi_k L_k$

[10] 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.

[12] 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 in the first place.

[12] 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.

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 may only take one action, so we do not model their preferences. When large writers attack ledger A successfully at time T, they receive revenues $\frac{L+a}{M}\pi_{A,T+1}$ where $\pi_{A,T+1}$ denotes the participation on ledger A at time $T+1$. 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 $t = 3$, play at $t = 1$ and $t = 2$ 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 $t = 3$ of day $T - 1$, we must analyze the equilibrium after an attack at time T. Again, the equilibrium at $t = 1$ is the same because writers are unaware of the attack. At $t = 2$, however, the public signal 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 $a^* = E[\hat{a}_{T-1}|y_T = 1]$ and note that $a^* > 0$. Then when τ is sufficiently small, we have

$$\theta_i = \tau - a^* + \sigma\epsilon_i$$

so were are in the same case as Proposition 3 with $\bar{\theta} = \tau - a^* < 0$. 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 large writers' choices at $t = 3$ when an attack is possible. Of course, the only interesting case is the case in which they choose $c_J = 1$ and $a_J > 0$. We look for conditions under which they never do so in equilibrium. We will restrict attention to symmetric equilibria in which $a_J = \hat{a}$ for all J. We have argued that writers must solve

$$\max_a \frac{1}{M}\left(1 - \frac{p}{M}((M-1)\hat{a} + a)\right)(L + a)$$

The first-order condition and the symmetric equilibrium condition $a_J = \hat{a}$ imply

$$\hat{a} = \frac{1}{p}\frac{M}{M+1} - \frac{L}{M+1} \tag{9}$$

This equation implies that since $L = 2M$, a symmetric equilibrium in which writers attempt to steal may exist only when p is sufficiently low, i.e. $p < \frac{1}{2}$. When p is large, the probability of detection is high enough to completely dissuade writers from even attempting an attack.

When $p < \frac{1}{2}$, writers' expected revenues are

$$\frac{1}{M}(1 - p\hat{a})(L + \hat{a}) = \frac{1}{p}\frac{1}{(M+1)^2}(1 + pL)^2$$

where $L = 2M$ (i.e., fees scale with the size of the network). This quantity converges to a finite limit $4p$ as $M \to \infty$. Large writers' outside options are to either (a) sit out, or (b) write on the chain that small writers use and receive revenues $\frac{L}{M+1}$ (since all other large writers are attacking and do not claim rewards on that chain). Large writers' costs are always equal to 1, so sitting out dominates attacking if $4p < 1$, i.e. $p < \frac{1}{4}$. Note that writers' revenues are monotonically decreasing in M if $p < \frac{1}{2}$. Then for M sufficiently large, their expected revenues from attacking will be close to $4p$ and their expected revenues from writing on the valid chain that the small writers use will be close to $\lim_{M \to \infty} \frac{L}{M+1} = 2 > 4p$, so for large M an attack will never occur. These results are formalized in Proposition 8.

Proposition 8. *When the probability of detection p is sufficiently large ($p > \frac{1}{2}$) attacks on the blockchain never occur regardless of the number of large writers. For each $p \in [\frac{1}{4}, \frac{1}{2})$ there exists $\bar{M}$ such that writers always prefer to write on the valid chain rather than to attack when $M \geq \bar{M}$.*

The main mechanism at work in securing the blockchain is similar to the mechanism by which an oligopoly breaks down. In an oligopoly, each producer wants to extract rents from consumers, but given the inability to coordinate actions each producer will produce more than would be optimal for a monopolist in the same market. Prices are driven down because producers do not internalize the effect of their production on others' profits. Here, the same logic implies that writers do not internalize the effect of their distortion on other writers' revenues. Therefore, writers take actions a_J that are larger than what a single large writer with M units of computing power would choose, and the attack becomes unprofitable relative to other strategies.

The cutoff $p = \frac{1}{2}$ at which writers switch strategies in this example should not be taken literally. It comes from the assumption that $L = 2M$, i.e. that the computing power available to writers is exactly enough for all writers to break even. Nevertheless, it is not unreasonable to assume that $L = O(M)$ in general, in which case a similar cutoff would arise– there would still be a region in which writers do not steal and one in which they prefer to write on the valid chain when M is large.

Proposition 8 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 the invalid chain earns less on average than one spent on the valid chain. An important factor in this tradeoff is the fact that when a group of writers decides to attack, they reduce the computational power on the valid chain, so each individual writer stands to gain by defecting and moving to the valid chain where there are certain profits to be made. When the probability of detection is low, by contrast, the network is secure as long as it is sufficiently *decentralized*, meaning M must be large. Then the mechanism described above applies, which is necessary to reduce the profitability of ledger distortion. This is another sense in which competition between writers benefits readers.

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 δ, 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 $t = 3$ of each period. The monopolist chooses an action $a \in [0, \bar{a}]$ at $t = 3$ 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 $\Pr(y_T = 1|a) = pa_{T=1}$. 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 $\theta_i = s_i - \gamma E[a_T]$. These preferences differ from those in the blockchain security example in an important way. The action taken by the monopolist at $T-1$ is not relevant for readers. 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. 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 many 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 $s \leq T$. Then the monopolist's problem is

$$\max_a \ (L_A + a) + \delta(1 - pa)(L_A + a) + \delta^2(1 - pa)^2(L_A + a) + \dots$$

which is just

$$\max_a \ \frac{(L_A + a)}{1 - \delta(1 - pa)}$$

The first-order condition is

$$L_A \geq \frac{1 - \delta}{\delta p} \tag{10}$$

This yields a key result:

Proposition 9. *There is a threshold value of L_A such that the monopolist never distorts the ledger:*

$$L_A = \frac{1 - \delta}{\delta p}$$

Proposition 9 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.

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).

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. We focus the analysis of our static model on the issue of rent extraction. 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 replicability 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.

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.

In another extension, we show that with sufficient decentralization of the network, a blockchain will always generate a consensus about the true history. Hence, our argument relies on a novel mechanism: the disciplining of writers through static incentives. When writers do not discipline each other, competitive forces lead them to distort the ledger so much that readers almost certainly discover their misbehavior and abandon the ledger, rendering the fraud unprofitable. In markets where centralized intermediaries have weak dynamic incentives, this static competitive incentive is a more efficient way of securing the ledger because it imposes a weaker lower bound on the compensation of intermediaries.

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 mechanics securing 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

Proof of Proposition 1:

Proof. Let $D_0 = \sup\{\theta : 1 - \hat{\pi}(\theta) < 0\}$. Note that for all types $\theta < D_0$ it is dominant to play $r(i) = B$, since these types prefer to play on ledger B even if all other agents participate on ledger A. Fix $\delta > 0$ and choose η such that $\Pr(\theta' \leq \theta - \eta|\theta) \geq Q(\theta) - \delta$ for all (θ, θ'). Then recursively define

$$D_n = \{\theta : \theta \leq D_{n-1} + \eta\}$$

Suppose all $\theta \in D_{n-1}$ play $r(i) = B$ and that $Q(D_n) - \delta > 1 - \hat{\pi}(D_N)$. For all $\theta \in D_n$, it is dominant to play $r(i) = B$. This is because

$$E[1 - \pi|\theta] \geq Pr(\theta' < D_{n-1}|\theta) \geq Q(\theta) - \delta > 1 - \hat{\pi}(\theta)$$

For sufficiently small δ, $Q(\theta) - \delta$ crosses $1 - \hat{\pi}(\theta)$ only once by Condition SC. Hence for all types θ such that $Q(\theta) - \delta > 1 - \hat{\pi}(\theta)$, it is dominant to play B. An exactly analogous argument shows that for all θ such that $Q(\theta) + \delta < 1 - \hat{\pi}(\theta)$, it is dominant to play A. To obtain the desired result, simply take the limit $\delta \to 0$. □

Proof of Proposition 2:

Proof. If neither action is dominant for type θ, 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 θ, 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 θ agents are indifferent between A and B). □

Proof of Proposition 3:

Proof. We prove the proposition by backwards induction.

t=3: Let $W_k = \int \mathbf{1}\{w(j) = k\}dj$ denote the measure of writers on branch k for $k \in \{A, B\}$. At $t = 3$, writers know the value of π. We show that $C_k = \min\{\pi_k L_k, W_k\}$ in equilibrium. Suppose first that $C_k < \min\{\pi_k L_k, W_k\}$. Then there exists a writer j with $w(j) = k$ 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 $C_k = \min\{\pi_k L_k, W_k\}$.

t=2: We will guess and verify that in any equilibrium, $W_B \geq \underline{C}$ and $L_B < L_A$. Writers' optimal play at $t = 3$ implies that $\frac{C_B}{\pi_B} \geq \underline{C}$. We also have $h(\frac{C_A}{\pi_A}) \geq h(\underline{C})$. Then it must be that

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

in equilibrium. By Proposition 1, given that Condition SC holds, we have that all i with $Q(\theta_i) > 1 - \hat{\pi}(\theta_i)$ choose ledger B. When $\bar{\theta} = 0$ and $\sigma \to 0$, θ_i is sufficiently small for each i that $1 - \hat{\pi}(\theta_i) < \frac{1}{2}$. On the other hand, when $\sigma \to 0$, $Q(\theta_i)$ approaches $\frac{1}{2}$ for all i, so it must be that $Q(\theta_i) > 1 - \hat{\pi}(\theta_i)$ for all i, meaning all readers choose $r(i) = B$.

t=1: Now we confirm that $W_B \geq \underline{C}$ in any equilibrium. In fact, all writers will choose ledger B. First note that writers of type $\theta_j^w \in [\bar{\theta} - \sigma, \bar{\theta} + \sigma]$ do not believe they can make positive profits on ledger A. If they were to make positive profits in some state, the equilibrium at $t = 3$ implies that in that state $W_A < L_A$. If this were the case, we would have $W_B > \underline{C}$, so the $t = 2$ equilibrium would imply that in all states $\theta < g(L_A) - g(L_B)$ all readers choose B. Writers of type $\theta_j^w \in [\bar{\theta} - \sigma, \bar{\theta} + \sigma]$ then know that making positive profits on branch A is impossible when σ is small.

Let $\underline{\theta}(L_A, L_B) = g(L_A) - \frac{1}{2}\kappa - g(L_B)$. For all types $\theta < \underline{\theta}$, it is dominant to play B. Now consider a writer who receives a signal $\theta_j^w \leq \underline{\theta} - 2\sigma$. This writer knows that all readers are of type $\theta \leq \underline{\theta}$, so all readers must play B. Therefore, it is dominant for such writers to play B, since it is never possible to make profits on branch A but it is possible to pay a cost ν when $\pi = 0$. Now consider a writer of type $\underline{\theta} - 2\sigma + \eta$, where η is small relative to σ. This writer knows that a measure of writers larger than $\underline{C}$ are of type $\theta_{j'}^w \leq \bar{\theta} - 2\sigma$, so those writers choose B. In this case, all readers will choose B by the equilibrium at $t = 2$, so writers of type $\theta_j^w \leq \underline{\theta} - 2\sigma + \eta$ must choose B. By the same logic, writers of type $\theta_j^w \leq \underline{\theta} - 2\sigma + 2\eta$ will choose B, and so forth up through type $\bar{\theta} + \sigma$. Thus all writers $\theta_j^w \in [\bar{\theta} - \sigma, \bar{\theta} + \sigma]$ choose $w(j) = B$, validating our guess at $t = 2$.

When all writers choose $w(j) = B$, clearly, they must make zero profits (by the equilibrium at $t = 3$).

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. All five statements follow immediately from equation (4). □

Proof of Proposition 5:

Proof. The first point follows from inequality (7). The second follows from the formula for L^{NE}. □

Proof of Proposition 6:

Proof. Using the logic of the proof of Proposition 1, all types $\theta_i = \eta + \epsilon_i$ will choose to stay on ledger A as long as $Q(\eta|\eta) + 1 - \mu \leq 1 - \hat{\pi}(\eta)$. This inequality is equivalent to

$$\frac{1}{2}(1+\mu) \leq \frac{1}{2} + \kappa^{-1}\big(\eta - (g(L_A) - g(L_B))\big)$$

which can be rewritten as

$$\mu \leq 2\kappa^{-1}\big(\eta - (g(L_A) - g(L_B))\big)$$

On the other hand, all types $\theta_i = \epsilon_i$ will choose to switch to ledger B if $Q(0|0) \geq 1 - \hat{\pi}(0)$. This inequality is

$$\frac{1}{2}\mu \geq \frac{1}{2} - \kappa^{-1}\big(g(L_A) - g(L_B)\big)$$

which is just

$$\mu \geq 1 - 2\kappa^{-1}\big(g(L_A) - g(L_B)\big)$$

as desired.

Readers are not concerned with the possibility that $\frac{C_k}{\pi_k} < \underline{C}$ because there are sufficient writers on each branch to ensure that so long as writers break even, computational power contributed per reader on each branch will be at least $\underline{C}$. □

Proof of Proposition 7

Proof. First we show that at any history $h^{T,3}$, on either branch k of the fork, the total computing power contributed by writers must be $\min\{W_k, \pi_k L_k\}$. Suppose that $C_k < \min\{W_k, \pi_k L_k\}$. Then there must be some writer j on branch k who contributes $c_j < 1$. By deviating to $c_j = 1$, 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 $C_k = \min\{W_k, \pi_k L_k\}$ at any history.

Now we show that at $t = 1$, all writers must choose branch B if it has been proposed and $L_B < L_A$. Again, since each writer's action does not affect the observable history, all writers will play static best responses at $t = 1$. In the $t = 1$ part of the proof of Proposition 3, it was shown that all writers choose ledger B when $L_B < L_A$. Hence writers never make profits in equilibrium.

Finally, 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 $L_B = \min\{L : L \in \mathcal{L}, L \geq \underline{C}\}$. $\square$