Blockchain Economics*

Joseph Abadi and Markus Brunnermeier†

December 19, 2018

Abstract

When is record-keeping better arranged through a blockchain than through a traditional centralized intermediary? The ideal qualities of any record-keeping system are (i) correctness, (ii) decentralization, and (iii) cost efficiency. We point out a *blockchain trilemma*: no ledger can satisfy all three properties simultaneously. A centralized record-keeper extracts rents due to its monopoly on the ledger. Its franchise value dynamically incentivizes correct reporting. Blockchains drive down rents by allowing for free entry of record-keepers and portability of information to competing "forks." Blockchains must therefore provide static incentives for correctness through computationally expensive proof-of-work algorithms and permit record-keepers to roll back history in order to undo fraudulent reports. While blockchains can keep track of ownership transfers, enforcement of possession rights is often better complemented by centralized record-keeping.

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

*We are grateful for an insightful discussion by Gur Huberman and helpful comments from Zhiguo He, Stephen Morris, Ulrich Müller, Wolfgang Pesendorfer, and seminar participants at the NYU Five Star Conference, the SEC, the University of Chicago, the Wharton School at the University of Pennsylvania, the Princeton Department of Computer Science, the St. Louis Fed, the Princeton Department of Economics, the NYU Intermediation Conference, and the BIS.

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

1 Introduction

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

It is of paramount importance that a ledger record all information correctly, i.e., it should be devoid of fraud. In this paper we point out a "blockchain trilemma": a ledger's correctness requires either the remittance of rents to a centralized entity or a pure waste of physical resources. Hence, it is impossible for any ledger to simultaneously satisfy the following three properties (depicted in Figure 1): (i) correctness, (ii) decentralization, and (iii) cost efficiency.

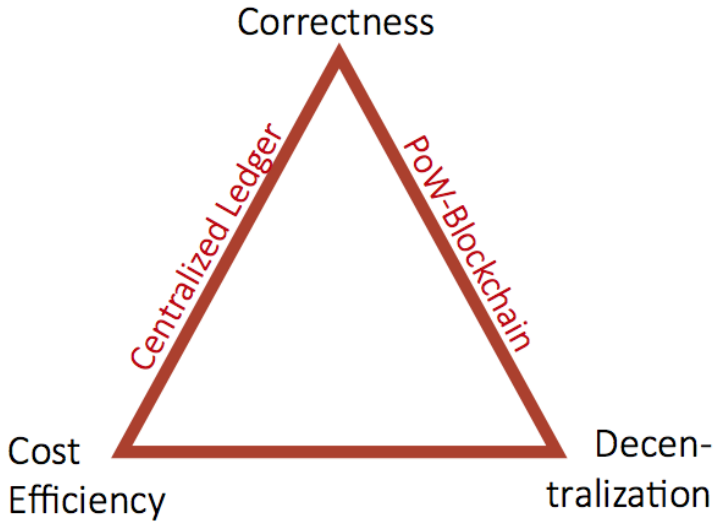

Figure 1: The blockchain trilemma.

Centralized record-keeping systems provide incentives to report honestly through the managing entity's franchise value. That is, this entity is *dynamically* incentivized to report honestly because it fears jeopardizing its future profits. As a result, centralized ledgers are often managed by monopolists that extract distortionary rents from the ledger's users.

Centralized record-keeping limits competition in two ways. First, only one entity is permitted to write on the ledger, so record-keepers cannot enter freely. Thus, the record-keeper cannot be disciplined internally: its incentives must stem from users' ability to switch to an outside ledger. Second, information is not portable: an entrant is not necessarily able to introduce a competing ledger containing all the information in the incumbent's. The lack

of information portability induces increased switching costs for users. These two features protect centralized record-keepers and promote the extraction of rent.

Blockchains are a type of *distributed ledger* written by decentralized and usually anonymous groups of agents rather than known centralized parties. This novel method of record-keeping has introduced two economic innovations that overcome the two limitations of competition among centralized ledgers. There is free entry of record-keepers: any agent may write on the ledger so long as they follow a certain set of established rules. Furthermore, information on an existing blockchain is portable to a competing one. A software developer can propose to "fork off" an existing blockchain to establish one with different policies while retaining all the information contained in the original blockchain. Fork competition eliminates the inefficiencies arising from switching costs in centralized record-keeping systems.

The features of blockchains that allow for enhanced competition are a double-edged sword. While centralized record-keepers are incentivized dynamically through their rents, free entry erodes these rents and necessitates a *static* mechanism to discipline bad actors. Since anybody can become a record-keeper (or miner) for a public blockchain, a consensus mechanism is needed to determine the true history written on the ledger (and reject fraudulent reports). Applying a majority rule is complicated by the fact that individual entities can masquerade as a large number of entities for free, subverting the democratic nature of the blockchain. Incentives to report honestly are instead provided through the imposition of a physical resource cost to write on the blockchain, making it costly to distort the ledger. Record-keepers must typically perform computationally expensive tasks in order to record information and validate others' reports. This consensus algorithm, known as proof-of-work (PoW), is responsible for the recent emergence of blockchains. It accomplishes the reduction of distortionary rents only by replacing those rents with socially wasteful electricity costs.

The PoW consensus algorithm is facilitated by the third economically important innovation of blockchains: the ability to roll back history. Those who attempt to commit fraud, censor information on the blockchain, or unilaterally change the rules are disciplined by other record-keepers, who may refuse to validate their reports. The community of record-keepers simply needs to identify the time of the attack and ignore all information reported after that time, thereby reverting the blockchain to a previous state. Such discipline may even be applied after a "51% attack" in which a record-keeper with a large amount of computational power attempts to report a valid history of events that conflicts with previous reports. This ability to roll back can be viewed as a conceptually distinct type of forking, where instead of creating new rules, the blockchain community creates a fork to undo the effects of a distortion of the ledger. The ease of rebuffing an attack provides strong static incentives to follow protocol. This internal discipline of dishonest reporting is significantly more robust than the

security guarantees of centralized ledgers, which rely on the existence of some external option to which users may switch if the ledger writer acts dishonestly.

In addition to analyzing the fundamental differences between blockchains and centralized ledgers, we apply our model of ledger competition to evaluate the benefits of blockchain in several applications of interest. While in most of our analysis we assume that users use only one ledger for analytical tractability, our qualitative results extend to an environment in which users can exchange several different currencies. We show that fork competition among cryptocurrencies is much fiercer than traditional currency competition (à la Hayek). Fork competition endows users of an established cryptocurrency with equal amounts of the newly created one, giving them a greater incentive to acquire information about and adopt the competing currency. By contrast, if users of an established currency are required to purchase the new currency from an outside issuer as in Hayekian competition, there is less of an incentive to adopt, which renders coordination on the new currency more difficult.

In between the polar cases of completely centralized traditional ledgers and completely decentralized blockchains, there is a third type of ledger called a "permissioned" blockchain that shows promise in many practical settings. The record-keepers of a permissioned blockchain are known agents rather than anonymous miners, so proof-of-work is unnecessary. Permissioned blockchains then seemingly break the trilemma: they allow for fork competition, like anonymous blockchains, but completely eliminate the waste of resources. We show that the impediments to entry of record-keepers on a permissioned blockchain substantially weaken fork competition. Permissioned record-keepers have franchise values and therefore can collude to prevent competing forks from surviving, whereas dynamic punishment schemes that sustain collusion are impossible when there is free entry of record-keepers.

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

Blockchains have applications that reach far beyond the realm of cryptocurrencies and

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

Related Literature. Our paper is related to the emergent literature on the economic properties and implications of blockchains. The paper most closely related to ours is that by Biais et al. (2017), which studies coordination among miners in a blockchain-based system. They show 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. Record-keepers' payoffs when forking depend exogenously on the number of record-keepers who choose a given branch of the fork. In our model, we instead place most of the focus on coordination among users. Record-keepers' payoffs are determined by users' actions, and a global games refinement of the game played among users puts more discipline on exactly how and when a fork may occur. Budish (2018) studies the costs of incentivizing honesty for cryptocurrency blockchains in isolation, whereas our work compares the cost of securing a blockchain to that of securing a centralized ledger. Cong and He (2017) focus mostly on the issue of how ledger transparency leads to a greater scope for collusion between users of the system. In contrast, we consider collusion between the blockchain's record-keepers rather than between users and show that collusion can occur only when entry of record-keepers is constrained.

Our framework uses a global game of the type pioneered by Carlsson and van Damme (1993) in order to select a unique equilibrium of the coordination game played by users. Rather than review the massive literature on global games here, we refer the reader to Morris and Shin (2001) for an extensive and general analysis of the global games framework. We use techniques from the more recent literature on global games with non-Gaussian private values pioneered by Sakovics and Steiner (2012) and advanced by Drozd and Serrano-Padial (2017). This framework is extended and generalized in Serrano-Padial (2018). 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 cryp-

tocurrency 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.

Some of the recent literature on blockchains in economics focuses on the security and the costs of the system. Huberman, Leshno, and Moallemi (2017) study transaction fees in Bitcoin and compare that environment to one with a monopolistic intermediary, as we do. They also emphasize the role of free entry and conclude that the blockchain market structure completely eliminates the rents that a monopolist would extract in an identical market. 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. 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.

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

Recent computer science literature has studied blockchain security extensively. Most papers in computer science, such as that by 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 record-keepers must always play honestly in order for the network to be secure. In contrast, we do not assume any record-keepers are compelled to play honestly and study a slightly more general class of attacks in which one record-keeper overwhelms the rest of the network. Our model shows that the cost of operating a blockchain is intrinsically linked to the cost of preventing attacks, no matter what they may be. Furthermore, the ability to

roll back history acts as an additional safeguard against attacks: the community of users and record-keepers may coordinate to ignore an attacker's reports even if those reports follow the protocols outlined in the blockchain's source code. That is, even if a given reported history is cryptographically sound, it may be the case that market participants recognize that everyone would be better off reverting to an earlier state.

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 may be 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 and illustrate how rents emerge under centralized record-keeping without information portability. We analyze a specific example where agents choose between two branches of a blockchain fork and another example in which agents choose between traditional ledgers in order to spell out the tradeoff between decentralization and cost efficiency. Section 4 discusses how free entry of record-keepers benefits competition. Section 5 presents a model of ledger security and investigates the incentive mechanisms required to ensure correctness for both a centralized ledger and a blockchain. Section 6 examines the distinction between ownership and possession as well as the feasibility of enforcing contracts recorded on a blockchain. Section 7 concludes.

2 Blockchain Technology

In this section we outline how blockchains work and discuss the distinguishing features of blockchains with anonymous record-keepers.

2.1 What is a blockchain?

A blockchain is a ledger in which agents known as nodes (or miners, or record-keepers) take turns recording information sequentially in data structures known as blocks. This information could consist of payment histories, contracts outlining wagers between anonymous parties, or data on ownership of domain names, among other applications. In principle, it is possible to use a blockchain in any application where it is necessary to record information. As will be discussed later, there are many possible algorithms to select the current record-keeper. The ledger consists of a tree of blocks that contains all the information recorded by miners 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").

On any blockchain, there are some rules that users and record-keepers tacitly agree to follow. These rules are written into the code distributed by the software developers for that blockchain. They include (but are not limited to) the cryptographic standards that must be followed when recording information, the types of information that may be recorded in a block, and the miners' compensation scheme. In general, blockchain security algorithms make it inexpensive for record-keepers to confirm that the rules are being followed. Nakamoto (2008) specifies that all other record-keepers should ignore any chain containing a block that does not conform to the proposed rules.

A chain of blocks leading back to the genesis summarizes a state. Ledger users and record-keepers must reach a consensus about which state is considered to be the current state. Hence, the rules followed by nodes must constitute a consensus algorithm. The consensus algorithm must coordinate nodes even when the network has some latency and not all nodes receive messages in the same order. Typically, the community coordinates on the longest valid chain of blocks as the current state, as suggested in Nakamoto (2008). Each record-keeper is periodically allowed to add a block to the tree. Record-keepers usually extend only the consensus chain, and users will act only in response to events on that chain. A record-keeper's decision to extend a given chain can be seen as a signal that the record-keeper accepts that chain as valid; that is, it is a tacit "vote" in favor of that chain. Record-keepers are rewarded for achieving consensus through users' acceptance of the chain extended. In general, record-keepers accrue rewards (stemming from the system) and record-keeping fees (paid by users) for each block added to the tree, so these rewards are realized only if those fees are on the consensus chain.

In principle, however, it is possible for users and record-keepers 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 a competing blockchain that builds off of the old chain, but they ignore any record-keepers who do not follow the new rules. Similarly, record-keepers who use the old rules will ignore all record-keepers who use the new ones, so the blockchain effectively forks and becomes two blockchains: one blockchain that uses the established rules and a competing blockchain that uses the new ones. The data contained in the established blockchain up until the point of the competing blockchain's birth is included in the competing blockchain, 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 ported to a

competing ledger.

Several important hard forks have taken place in cryptocurrency blockchains. 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. Other examples include the hard forks that created Bitcoin Gold and Bitcoin Cash, two cryptocurrencies that modified the rules governing the original Bitcoin blockchain.

In our model we also consider attacks on the blockchain. An attack on a blockchain involves the addition of blocks that are somehow invalid or reverse previous accepted transactions. 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 users and record-keepers go along with the attack, but these attacks are usually detected automatically by all users of the system. It is perhaps less obvious why an attacker would want to add valid blocks somewhere other than at 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 network's computing power, eventually the length of the attackers' chain will exceed that of the other chain. At this point it becomes the longest valid chain. All record-keepers, 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." 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 one of the largest security concerns 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 embeds double spending, but it encompasses a broader class of attacks.[1]

[1]For instance, one could also consider an attack on light clients (such as the ones used on mobile devices) that do not store and thus are unable to verify the entire history of transactions. An attacker could, in theory, bombard these clients with fraudulent messages that they cannot detect as invalid.

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 record-keeper. The users 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 record-keeper is disciplined entirely by the users, who may decide to punish the record-keeper in some way if the record-keeper 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 ledger's 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 record-keepers take turns adding blocks to the chain according to a predefined algorithm, so again costly identity management is unnecessary. The record-keepers on a permissioned blockchain are disciplined by users, just as in a private blockchain, but they are also disciplined by other record-keepers. If one record-keeper deviates and begins validating fraudulent ledger entries by including them in his block, other record-keepers may ignore him and refuse to extend his chain. If a record-keeper proposes a change to the blockchain's policies, other record-keepers may prevent such a change by keeping records 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. Record-keepers are disciplined exactly as in permissioned blockchains. All users of the network are anonymous. However, when record-keepers 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 record-keepers 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 and is used by most

major cryptocurrency blockchains such as Bitcoin, Ethereum, and Litecoin. In order to incentivize record-keepers 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 record-keepers' rewards tend to be economically large. For example, the Bitcoin blockchain currently uses more electricity than Hungary.

2.3 Differences between blockchains and centralized ledgers

The computer science literature stresses several differences between public blockchains and traditional systems of record-keeping. For one, public blockchains tend to be more transparent than centralized ledgers given that anyone can read or write on the blockchain.[2] Another advantage of blockchains is that they have the potential to be more efficient at achieving consensus than existing systems (e.g., settlement in financial markets). Public blockchains are also immutable in the sense that it is costly to rewrite history on a blockchain, and they are typically shielded from operational risk because there is no single point of failure.

Many of these features of public blockchains, however, are incidental. For example, there is no reason why PoW (or another identity management system) should be necessary in order to achieve fast settlement or transparency. These features could be implemented simply by moving existing systems over to publicly viewable electronic databases maintained by a known group of entities.

In contrast to the previous literature, our paper emphasizes three *economically* relevant differences between public blockchains and traditional centralized ledgers. First, record-keepers enter freely: any agent who wishes to write on the ledger may do so by following a given set of rules. Free entry of anonymous record-keepers is "trustless" and thus requires identity management. Public blockchains typically solve the trust problem by forcing record-keepers to pay a cost to record information and requiring that future record-keepers validate those reports. By contrast, record-keepers in centralized systems are known, and their privileged positions allow them to extract rents. Inasmuch as they fear losing the benefits stemming from record-keeping, centralized record-keepers are "trusted."

Second, the information in a public blockchain is portable: an agent who wishes to launch a competing ledger may propose a hard fork, thereby replicating all the information in the existing blockchain.[3] Importantly, to the extent that the information in the ledger is used to coordinate users' actions (as in the case of cryptocurrency and other applications), it is

[2]This assertion comes with a caveat. Recent developments in cryptography have led to proposals of blockchains that operate using "zero-knowledge proofs" that permit information to be verified without being made public.

critical that users can be confident that *all* of the information in the ledger is ported over to the competing blockchain, not just their own personal information.[4] The guarantee that all information is portable to a new blockchain will imply that users can respond to others' actions in the same way as they did when using the old blockchain.

Third, any group of record-keepers that detects fraudulent activity may roll back the history of the blockchain to a state preceding the attack and create a fork in which the attack never occurred. While the first two technological differences we highlighted are applicable to almost any type of distributed ledger, the rollback feature is specific to blockchains. The ability to revert to a previous state is a result of the lack of *finality* characteristic of blockchains. We show that a lack of finality can, in fact, enhance a blockchain's security. It is a desirable feature of blockchains rather than a bug.

In our model, we focus on these three economically relevant differences and explore their implications for competition between ledgers.

3 Model of Ledger Choice

In this section, we present a general static model of ledger choice as a coordination game. Our objective is to capture a variety of settings in which users choose among competing ledgers with different rules or policies. The main results explain how the features of a record-keeping system influence users' ability to coordinate on a ledger and the intensity of competition among ledgers. Hence, the static model we present in this section describes the emergence of rents. Our leading example applies our model to the study of 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. The specific examples of competition between different types of ledgers illustrate how decentralization attenuates rent extraction.

We focus on the importance of coordination *among users* because many types of ledgers are useful only if they are widely used.[5] For example, agents will be willing to accept a fiat currency only if it is accepted by others. Another situation in which coordination is important is when the ledger contains information about users' creditworthiness (such as Alibaba's Sesame credit score system). Users will not have an incentive to build up their

[4]While it would in principle be possible for a private blockchain to permit forking, we show that it would be strictly suboptimal for the agent who owns the blockchain to do so, meaning that lack of ownership of the blockchain is necessary for portability of information to be relevant.

[4]This distinguishes our notion of information portability from the notion of data ownership referenced in proposals such as the General Data Protection Regulation (GDPR). By the same token, it is not necessarily crucial that the information in the ledger be publicly viewable. The data in the ledger could be encrypted, as with cryptocurrencies that use zero-knowledge proofs. All that is needed, in principle, is some method of using the ledger to prove a minimal set of statements that sustain an equilibrium of a game played among users.

credit score if there are no lenders. Throughout, we will abstract from the specific details of the coordination motive.

There are two ledgers, A and B. There is a continuum of agents $i \in [0,1]$ known as users, who are users of the service offered by the ledger. There is a set of agents $j \in \mathcal{M}$ known as record-keepers. 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 record-keeper. 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 centralized ledger, the proposer is also the record-keeper. That is, the monopolist who runs the ledger also decides on the rules. In what follows, we will allow for the possibility that some record-keeper $j \in M$ is also one of the proposers.

Users must choose the extent to which they participate on ledgers A and B. User i chooses an action $\varphi_i = (\varphi_i^A, \varphi_i^B) \in [0,1]^2$, where φ_i^l represents the intensity of i's usage of ledger $l \in \{A, B\}$. Users will desire to coordinate with other users, so their utility functions will depend on $\phi^l \equiv \int_0^1 \varphi_i^l di$ for $l \in \{A, B\}$.

Users are heterogeneous in their *fundamental* preferences for ledgers. Each user is assigned a type $s_i = (s_i^A, s_i^B)$. Here, s_i^l is meant to represent the *stake* that user i has in ledger l. The stake that a user has in a given ledger should be interpreted as the amount of information pertaining to that user that is encoded in the ledger. For any ledger that keeps track of asset holdings, a user's stake is simply the set of assets held by that user, with larger asset holdings being interpreted as a higher stake. However, a user's stake does not necessarily have to represent the market value of some asset. A user with a high stake may also be a consumer who has built up a high credit score in a credit registry or a financial institution with a complex set of contracts with other institutions written down in a particular ledger. Stakes could even be interpreted as ratings on e-commerce websites or as contacts on social media platforms. We denote the population CDF of stakes s by $Q(s)$.

There is also a common value component ζ in users' preferences. This component can be understood as parametrizing differences between fundamental features of the two ledgers,

[5]Other papers in the literature examine coordination among record-keepers instead. There is a vast computer science literature that analyzes the equilibria of the game played by blockchain miners, and Biais et al. (2017) provide a proof in a formal game-theoretic model that there is an equilibrium in which all miners choose the longest chain.

such as technological efficiency. When $\zeta > 0$, the common value induces a preference for A among all users, and when $\zeta < 0$, users prefer B. While parameter ζ itself does not play a significant role in our analysis, the introduction of uncertainty about this parameter will lead to equilibrium selection.

Each proposer P^l chooses a fundamental parameter $L^l \in \mathcal{L}^l$ determining the revenues earned by record-keepers and charged to users. A simple way of thinking about L^l is as an explicit fee charged to users by the record-keeper(s) of the ledger, but more broadly L^l could be interpreted as an implicit fee. Such implicit fees could arise, for instance, if a monopolist who runs a ledger chooses to sell users' data to an outside party. The fundamental parameter L^l 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). Parameter L^l could also be interpreted as a choice of mechanism that leads to a certain utility for users and record-keepers. Henceforth, for ease of exposition we will refer to L^l as a fee.

Record-keeper j chooses a ledger $l \in W_j \subset \{A, B\}$ and takes an action $a_j \in \mathcal{A}_j$ to write on the ledger. In our applications, the action a_j will refer either to an expenditure of computational resources to ensure cryptographic security of the ledger (proof-of-work), or an action taken by a dishonest record-keeper to distort the ledger (such as a double-spend or outright fraud by a monopolist). That is, the actions taken by record-keepers will capture two points of the blockchain trilemma: cost efficiency and correctness.

A user's utility depends on (i) her own action φ, (ii) her stakes s, (iii) others' participation choices $\phi = (\phi^A, \phi^B)$, (iv) the common value ζ, (v) the fundamental parameter $L = (L^A, L^B)$ chosen by proposers, and (vi) actions a taken by record-keepers. In general, a user's utility function can be written as

$$u(\varphi, s, \phi, \zeta, L, a).$$

A record-keeper's utility when writing on ledger l depends on the action a_j^l taken by that record-keeper, the actions a_{-j}^l taken by other record-keepers on the same ledger, usage of the ledger ϕ^l, and the fee L^l. The corresponding utility function will be denoted by

$$w^l(a_j^l, a_{-j}^l, \phi^l, L^l).$$

Finally, the utility of proposer P^l is given by a function $v(\phi^l, L^l)$.

The game is played in periods $t = 0, 1, 2, \ldots$. Each period has three subperiods $\tau = 0, 1, 2$. The timing of the game is as follows:

- $\tau = 0$: Proposers P^A and P^B choose L^A and L^B, respectively.

- $\tau = 1$: Users first observe previous actions and their own types s_i. They then choose actions φ.
- $\tau = 2$: Record-keepers observe previous actions, choose a ledger $l \in W_j$, and take actions $a_j \in \mathcal{A}$.[6] Payoffs are realized.

For the remainder of this section, we focus on the static Nash equilibria of the stage game played in each period t. Later, we will analyze the equilibria of the full dynamic game.

3.1 Characterization of equilibrium with arbitrary competing ledgers

We now prove properties of equilibrium that will hold in all of the settings we consider. First, we show that as noise about the common value vanishes, users' play is uniquely pinned down in equilibrium. We then characterize the way in which users coordinate and connect our results to the rents that can be extracted from them. 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 user to Fudenberg and Tirole (1991).

In the following analysis, we impose the restriction that, when choosing a ledger, users must choose either A or B, not both. In the notation of the previous section, they choose $\varphi \in \{(1,0),(0,1)\}$. Under this assumption, it is natural to consider

$$\tilde{u}(s,\phi,\zeta,L,a) \equiv u(A,s,\phi,\zeta,L,a) - u(B,s,\phi,\zeta,L,a)$$

(with some abuse of notation). Users choose A when they expect $\tilde{u}$ to be positive and B when they expect $\tilde{u}$ to be negative. For ease of notation, we will henceforth denote a choice of A as $\varphi = 1$ or $\varphi = A$ and set $\phi = \int_i \varphi_i di$ to be the proportion of users who choose A.

We will also take record-keepers' actions (as a function of ϕ and L) to be known by users at $\tau = 1$. That is, users take as given a function $a(\phi, L)$. In our main applications, users will know the actions record-keepers must take in equilibrium at $\tau = 2$ by backward induction. In this section, then, we focus on the equilibrium of the game played by users. We defer the discussion of the equilibrium among proposers until later, as it will depend on the particular application. Finally, we render users' individual types one-dimensional rather

[6] Although it would be natural in some cases to assume that record-keepers move before users or that the two groups move simultaneously, most of our results are robust to these alternative specifications. We impose this sequence of moves for ease of exposition and because we wish to focus on how coordination among users of a blockchain system differs from coordination among users of a centralized ledger. Furthermore, it is often reasonable to assume that record-keepers may condition on users' actions (such as with blockchain systems where miners condition effort on cryptocurrency prices).

than two-dimensional by assuming that u depends on the stakes s_i^A and s_i^B only through their difference $s_i^A - s_i^B$. Henceforth, we denote this difference by s_i.[7]

The assumptions that users may choose only one ledger and that users' stakes are fixed and nontransferable are stark and merit further discussion. Although users may choose only one ledger in our benchmark model, in the Appendix we extend the model to allow users to use both ledgers simultaneously, and most of our main results carry over in that framework. For ease of exposition, we assume that users choose only one ledger.

The assumption that users' stakes are nontransferable is suitable in some cases, such as when the ledger contains information about the types or reputations of users. For example, vendors on online platforms cannot trade their reputation scores to other vendors, and users of credit registries cannot transfer their credit scores among themselves. However, in the case of currency, this assumption is completely counterfactual: if two people with stakes in different currencies trade those currencies with each other, they relinquish their stakes in one ledger to obtain a stake in another. We remedy this issue by later presenting a more explicit model in which the ledger contains currency holdings, showing that, although the results are somewhat more complicated, the main intuition regarding the importance of information portability and free entry of record-keepers continues to hold.

In the following results, we will focus on the case where network externalities are strong in the sense that the only stable equilibria are those in which all users choose A or all choose B.[8] These cases are particularly important because there are multiple equilibria, so we will need to develop a method of selecting among them. In our applications of the model, we will also consider cases in which network externalities are weak, but in those cases the equilibrium will be unique regardless, meaning equilibrium selection is irrelevant.

Concretely, we focus on regions of the parameter space where if all users play the same action, it is individually rational for any type s to follow suit. Furthermore, mixed equilibria where some users play A and others play B are not robust to small perturbations of strategies.

To select among equilibria, we use a global games framework and introduce incomplete information about the common value ζ. Formally, we assume that each user i receives a signal $x_i = \zeta + \sigma\eta_i$, where η is uniformly distributed on $[-\frac{1}{2}, \frac{1}{2}]$. We typically work in the limit $\sigma \to 0$, so there is an arbitrarily small amount of noise in agents' signals.[9] Incomplete information about this value could be motivated by, for example, uncertainty about the properties of the ledger's technology. With incomplete information about ζ, users' types become two-dimensional. An individual user's type can be summarized by $\theta_i = (x_i, s_i)$.

In order to prove statements about equilibria in this environment, it will be useful to make the following assumption:

[7] This assumption will apply to all models in our applications.

Assumption 1. *There exist $\overline{\zeta}$ and $\underline{\zeta}$ such that it is strictly dominant for all types s_i to play A whenever $\zeta > \overline{\zeta}$ and B when $\zeta < \underline{\zeta}$.*

This type of assumption is typical in the global games literature. It states that there exist "dominance regions" in which fundamentals favor one action so heavily that agents find it optimal to play that action even if no other agent does so.

The main property of equilibria that we can prove at this point is that equilibria will take a "cutoff" form: there will be threshold values $k(s)$ such that all agents with $x_i < k(s_i)$ choose ledger B and all users with $x_i > k(s_i)$ choose ledger A. These cutoffs will be decreasing in s, meaning agents with larger stakes in ledger A will be more likely to choose A. This is true as long as the actions taken by record-keepers are the same on ledgers A and B. That is, users 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 users will switch to B.

Proposition 1. *Under Assumption 3.1, for sufficiently small σ, there is an essentially unique equilibrium of the game played by users at $\tau = 1$, holding fixed the actions of record-keepers at $t = 2$. There exist weakly monotonically decreasing cutoffs $k(s)$ such that all users with $x_i > k(s_i)$ choose $\varphi_i = A$, and all users with $x_i < k(s_i)$ choose $\varphi_i = B$.*

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

We can provide a sharper characterization of equilibrium when network externalities are strong. Before doing so, we define an important symmetry property of users' utility functions.

Definition 1. *Let $\chi = (\zeta, L, a(\phi, L))$. Define $\tilde{u}$ to be **symmetric** at χ if $\tilde{u}(s = 0, \phi, \chi) = -\tilde{u}(s = 0, 1 - \phi, \chi)$ for all (φ, ϕ). Define $\tilde{u}$ to be **asymmetric towards** A (resp. B) at χ if there exists a symmetric utility function $\tilde{u}' \neq \tilde{u}$ such that $\tilde{u} \geq \tilde{u}'$ (resp. $\tilde{u} \leq \tilde{u}$).*

[9] Users' priors over ζ become irrelevant in this limit. See Frankel, Morris, and Pauzner (2003).

When $\tilde{u}$ is symmetric, it simply means that users consider the fundamentals of the two ledgers to be equivalent, so differences in the utilities they obtain from using the two ledgers are due only to differences in stakes and participation on each ledger. By contrast, when $\tilde{u}$ is asymmetric toward A, users prefer the fundamentals of A to those of B.

Equipped with this definition of symmetry, we may prove our main result regarding the properties of the equilibrium played by users.

Proposition 2. *When network externalities are strong, all types $s \in \operatorname{supp} Q$ share the same cutoff k. This cutoff satisfies two properties.*

1. *Let $\chi = (\zeta, L, a)$ and consider the type $s = 0$ (types with the same stake in A and B). If u is symmetric at χ, then when $\operatorname{supp} Q = \{0\}$, $k = \zeta$.*

2. *If $\tilde{u}$ is symmetric, $\operatorname{supp} Q \subset \mathbb{R}_+$, and $Q(0) < 1$, then $k < \zeta$.*

The characterization of equilibria in games with strong network externalities demonstrated by Proposition 2 is actually quite natural. First, the proposition states that when network externalities are strong, all agents must use the same cutoff k (so that all users choose the same ledger in equilibrium). If there exist fundamentals (ζ, L, a) of the two ledgers such that users would be indifferent between them in the absence of network externalities, the cutoff k is equal to ζ as long as all users' stakes are equal to zero. Put simply, users coordinate on the ledger with more favorable fundamentals when their stakes do not anchor them to one ledger. Hence, there is essentially perfect competition between ledgers in this setting. On the other hand, when users' stakes in A are larger than their stakes in B, $k < \zeta$, meaning users may coordinate on A even when the fundamentals (ζ, L, a) favor B.

When users have larger stakes on one of the two ledgers, that ledger has a competitive advantage in the coordination game played by users. For example, if A is an established ledger and B is a newly proposed one, A will have an advantage over B unless the information from A is ported over to B. When A has a competitive advantage, it will be able to win out over B even when it extracts larger rents from users. Later, in our analysis of the differences between blockchain and centralized systems of record-keeping, we will emphasize the enhanced portability of information allowed by blockchain. In that sense, blockchain will also alter the equilibrium of the game played by users. However, the differences between the two systems are not limited to the equilibrium of the game played by users: there will be differences in the play of record-keepers and proposers as well.

3.2 Traditional competition between centralized ledgers

In this section, we analyze a competition between an outside ledger and a centralized ledger maintained by an incumbent monopolist. The differences between this setting and one with two branches of a blockchain fork will clarify exactly how fork competition differs from standard competition, i.e., what exactly is accomplished by decentralization. We will show that a lack of information portability will anchor users to the incumbent's ledger and give rise to rent extraction.

We first begin by assuming that the monopolist is the incumbent, in the sense that users have a stake in the monopolist's ledger but not the outside ledger. There are just two record-keepers: the incumbent monopolist $\mathcal{M}$ on ledger A and an outside record-keeper (entrant) $\mathcal{O}$ on ledger B. In this case, the record-keepers are also the proposers $P^A = \mathcal{M}$ and $P^B = \mathcal{O}$. Each record-keeper may write only on her own ledger. At $\tau = 0$, the incumbent may choose a fundamental parameter $L^A \geq 0$ and the entrant chooses $L^B \geq 0$. User i has stakes $s_i^A \geq 0$ on ledger A and $s_i^B = 0$ on ledger B. Users have no stake in the outside record-keeper's ledger, and that record-keeper is unable to replicate the stakes in the monopolist's ledger due to information frictions. Record-keepers do not take actions at $\tau = 2$.[10]

Users have preferences summarized by

$$\tilde{u}(s, \phi, \zeta, L) = s^A - s^B + \kappa(2\phi - 1) + \zeta - \alpha(L^A - L^B). \tag{1}$$

Users' types s_i have a cross-sectional distribution $Q(s)$ that is uniform on the interval $[S - \frac{d}{2}, S + \frac{d}{2}]$. Here, S is the average stake and d is the dispersion of stakes. It is important to distinguish between situations in which network externalities are strong enough to generate multiplicity and situations in which network externalities are weak. There is multiplicity if and only if $d > \kappa$. We will henceforth assume in this section that the true realization of the common value ζ is zero, but that this is unknown to users.

The monopolist receives a fee L^A from each user who participates. The monopolist's objective function is

$$\max_{L^A \geq 0} \phi L^A,$$

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

$$\max_{L^B \geq 0} (1 - \phi) L^B.$$

In order to proceed, we must determine how the record-keepers' choices of L^A and L^B map

[10] We postpone the discussion of record-keepers' honesty to Section 4.

to participation ϕ. Proposition 3 provides an answer.

Proposition 3. *When $d > \kappa$, all users for whom*

$$\frac{1}{2} + \kappa^{-1}\left(s_i - \alpha(L^A - L^B)\right) > Q(s_i)$$

choose to remain on ledger A, and all other users choose ledger B. When $d < \kappa$, all users choose A if $\alpha(L^A - L^B) < S$ and B if $\alpha(L^A - L^B) > S$.

Proposition 3 illustrates that coordination among users is inertial. Users' initial stakes in the incumbent's ledger anchor them to it, and network externalities amplify the initial anchoring effect. As a result, the entrant is at a disadvantage. Even if the entrant charges a lower fee, it will not necessarily claim a larger market share than the incumbent. Indeed, if network externalities are strong enough, it is possible for the entrant to charge a lower fee than the incumbent and still completely fail to capture any users.

Now that we have determined how users coordinate given the fees proposed, we may solve for the equilibrium of the game played by the incumbent and the entrant. In Appendix C.1, we will show that there are two types of equilibria: when network externalities are weak, the entrant manages to capture positive market share, whereas when network externalities are strong, the entrant does not capture any share of the market, but the market remains contestable in the sense that the incumbent sets its fee low enough to keep the entrant out of the market. Proposition 4 characterizes the fee charged by the incumbent in these different types of equilibria.

Proposition 4. *The incumbent monopolist charges fees*

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

These fees are increasing in the mean stake S.

The main insight of Proposition 4 is that the equilibrium fee charged by the incumbent is monotonically increasing in the average stake S users have in its ledger. This relationship is most pronounced when network externalities are strong enough that the entrant is unable to attract any user in equilibrium.

The comparison between centralized ledger competition and fork competition highlights one of the central tradeoffs of the blockchain trilemma: the tension between cost efficiency and decentralization. When two forks of a blockchain compete, two forms of decentralization aid

coordination on an efficient outcome. Portability of information nullifies the anchoring effect on the established ledger, so network externalities play no role in amplifying inertia. Hence, portability of information eases coordination on the new ledger *among users*. Competition among record-keepers eases coordination *among record-keepers* in the sense that it ensures they will write on the ledger preferred by users. The combination of these two features allows new, competing ledgers to emerge in equilibrium. Welfare losses in blockchain-based record-keeping systems come mostly from the waste of computational resources.

Under traditional centralized-ledger competition, by contrast, even when there is the possibility of entry, both the incumbent and the entrant may charge high fees. If there is no possibility of entry, strong network externalities protect the incumbent and increase distortionary rents. The incumbent further enjoys high rents because of its monopoly on information, which is detrimental to users' welfare. Taken together, these results suggest that fork competition is particularly beneficial when coordination motives among users are strong or when switching costs in a traditional setting are high.

Although we stress the emergence of new ledgers in our analysis of fork competition, there is another side of fork competition that is potentially just as important: the ability to enshrine a certain set of rules for the operation of the ledger. That is, if a new, inferior ledger is proposed, our model suggests that users will tend to coordinate on the existing ledger and ignore the new one, keeping the old rules intact. While our formal model of static centralized ledger competition does not allow us to analyze the dynamics of incumbent's fees, intuitively it stands to reason that they may rise over time as users become more anchored to its ledger. Thus, under centralized record-keeping it is not always possible to ensure that existing policies will remain in place.

3.3 Fork competition

In this section, we present our baseline model of competition between blockchain ledgers and analyze the equilibria of the stage game. In reality, this competition corresponds to a "hard fork," in which some of the blockchain's record-keepers decide to build their own blockchain with new protocols off of a previously existing (parent) blockchain. There is a critical distinction between a hard fork in which new rules are proposed and a fork that is created in an attack on the blockchain, such as a double-spend. The former is important in understanding the relationship between decentralization and competition among ledgers, which is our focus in this section. The latter relates to the security and correctness of the ledger, which we will discuss in Section 4.

Critically, a hard fork preserves all of the data in the parent blockchain. This observation will be crucial for our conclusions: the ability of record-keepers to change the rules of the

blockchain but keep users' stakes in the network intact will allow for perfect competition between ledgers. There will be no inertia in switching ledgers because users will not stand to lose their stakes by doing so. We will show that in an environment with information portability and free entry of record-keepers, competition among ledgers is essentially perfect. Record-keepers will no longer be able to earn rents. Blockchains will enhance competition between ledgers, but they will come at the cost of proof-of-work, i.e., the cost of decentralization. This example will thus illustrate the tradeoff between decentralization and cost efficiency postulated in the trilemma.

The model of blockchain competition falls within the general class of models of ledger competition described earlier. In the game, users 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 fees L^A and users have stakes s_i^A on that branch. That is, we constrain the proposer P^A to choose L^A. This proposer can be thought of as one of the original developers of the blockchain. The proposer on branch B may choose a new fee $L^B \geq 0$ after observing ζ at $\tau = 0$ (so proposers have perfect information about ζ).[11] Furthermore, in a hard fork, all of the information on the original blockchain is *carried over to the new blockchain*, meaning that users' stakes on ledger B are $s_i^B = s_i^A$ for all i, so $s_i = s_i^A - s_i^B = 0$. In this sense, the established ledger A has no informational advantage over the entrant ledger B. 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 ϕ^B, P^B receives a payoff $v(\phi, L^B)$, where v is a decreasing function of L^B and an increasing function of ϕ. The proposer's payoff is assumed to come from an appreciation of the developer's stake when the proposed ledger is adopted, so proposers' incentives are aligned with users'.[12]

In this setting, the set M of record-keepers 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. Record-keepers are responsible for cryptographically securing the ledger, and they are given some surplus for contributing computing power to the blockchain. At $\tau = 2$, record-keeper j chooses a ledger $l_j \in \{A, B\}$ and an amount of computational power $c_j \leq 1$ to contribute to that ledger. We assume that record-keepers can observe users' actions before making a decision because, in practice, this is often exactly what happens. Cryptocurrency "mining pools" are set up to

[12]Here we assume that P^B may choose any $L^B \geq 0$. That is, because it is not a focus of our paper, we assume away the issue of implementability of L^B via some mechanism. For a discussion of how mechanisms map to fees in centralized and decentralized payments systems, see Huberman, Leshno, and Moallemi (2017).

[12]The assumption that the proposer's incentives are aligned with those of users is not overly restrictive in this context. As we will show, free entry of blockchain record-keepers implies that no matter what the fee structure, record-keepers will not earn positive profits, so there is no way the proposer could benefit record-keepers by choosing a higher fee.

automatically mine on whatever blockchain yields the highest profits at that moment. To the extent that the token price on a blockchain proxies for participation on that blockchain, mining pools essentially condition their decisions on users' actions.

Record-keepers pay a linear cost $f(c) = c$ of generating computational power. Let $C^l \equiv \int_{l_{j'}=l} c_{j'} dj'$ be the total computational power contributed to branch l of the fork, and denote the participation on that fork by ϕ^l. Then a record-keeper's net profits when contributing computing power c_j to branch l are

$$w(c_j, C^l, \phi^l, L^l) = \frac{c_j}{C^l} \phi^l L^l - c_j$$

when $C^l > 0$ and $-c_j$ otherwise. The record-keeper's revenues are proportional to participation and the fundamental parameter L^l but are inversely proportional to the computational power contributed by other record-keepers. This revenue function captures two features shared by most blockchains. Namely, the total rewards given to record-keepers are fixed, and those rewards tend to be more valuable when the blockchain has been adopted by a larger group of users. Record-keepers' rewards thus have a natural interpretation as a *coin* associated with a particular blockchain.[13]

The first important feature of the equilibrium with blockchain competition is that record-keepers' actions are pinned down by the free-entry condition. Optimizing w with respect to c_j, we see that the optimal computing effort c_j^* is given by

$$c_j^* = \begin{cases} 1 & C^l < \phi^l L^l \\ \in [0,1] & C^l = \phi^l L^l \\ 0 & C^l > \phi^l L^L \end{cases}.$$

Hence, in equilibrium it must be that record-keepers break even:

$$C^l = \phi^l L^l. \tag{2}$$

We later show that this result extends to all equilibria of the full dynamic game. Intuitively, with free entry, record-keepers will compete to write on the ledgers used by users and do not leave opportunities for profit on the table. Free entry thus distinguishes a public blockchain from a permissioned blockchain, which, as we will show later, can feature equilibria in which record-keepers collude in order to suppress competing ledgers.

[13] Almost all blockchains reward record-keepers internally with coins that exist as data on the blockchain. This reward scheme is common even for blockchains that are used mainly for purposes other than the exchange of cryptocurrency. The provision of rewards within the scope of the blockchain further incentivizes record-keepers to maintain the blockchain's integrity.

Users prefer ledgers that are cryptographically secure. In order to restrict attention to the competitive benefits of decentralization, we effectively abstract from record-keepers' endogenous choices to attempt to distort the ledger in this section. Instead, we assume an exogenously given relationship between the probability of attacks on the blockchain and the proof-of-work cost. Their preferences for cryptographic security are parametrized by a bounded utility function $\tilde{u}(s, \phi, \zeta, L, C)$ (where $C = (C^A, C^B)$) that depends on record-keepers' actions only through the ratios $\frac{C^l}{\phi^l}$,[14] is increasing in the ratio $\frac{C^A}{\phi^A}$, and is decreasing in the ratio $\frac{C^B}{\phi^B}$. That is, users value security in terms of the amount of computational power committed to the blockchain per user, so they prefer ledger l when more computing power is committed to that ledger. Hence, there is a sense in which users prefer high fees to a certain extent: fees provide incentives for record-keepers to secure the ledger. For now, we keep the dependence of preferences on computational effort exogenous and discuss the benefits of fork competition. In our discussion of attacks on the blockchain we outline how it can be endogenized and discuss in greater detail the tradeoff between free entry of record-keepers and costly proof-of-work.

In what follows, it will be convenient to adopt a particular form for users' utility function $\tilde{u}$ in order to derive analytical results. We assume

$$\tilde{u}(s, f, \phi, \zeta, L, C) = \underbrace{s}_{\text{stake}} + \underbrace{\kappa(2\phi - 1)}_{\text{network externalities}} + \underbrace{\zeta}_{\text{fundamentals}} - \Bigg(\underbrace{\alpha(L^A - L^B)}_{\text{fees}} - \underbrace{(g(\frac{C^A}{\phi}) - g(\frac{C^B}{1-\phi}))}_{\text{crypto security}} \Bigg). \tag{3}$$

Here, the utility function is simply represented as a linear sum of the five components of users' preferences. Parameter κ governs the strength of network externalities, and parameter α determines users' aversion to fees. Function g (which is, at this point, exogenous) relates users' utility to the cryptographic security of the ledger. Note that this utility function is symmetric when $\zeta = 0$, $L^A = L^B$, and $C^l = \phi^l L^l$ as defined in Definition 1. Throughout, we will continue to use a similar utility function in order to illustrate properties of equilibrium in our applications.

Now that we have set up the blockchain game, we may prove our main result about equilibria of the static blockchain competition game.

Proposition 5. *Suppose that network externalities are strong and that there exists $\bar{\zeta}$ such that $\tilde{u}$ is symmetric, taking as given equal fees on the two ledgers, $L^B = L^A$, and optimal equilibrium play by record-keepers, $C^l = \phi^l L^l$. Then if there exists $L^B \geq 0$ such that $\tilde{u}$ is*

[14] Of course, some value must be assigned to the function at $\phi^l = 0$. This value can be essentially arbitrary so long as users' utility when $C^l = 0$ does not change with ϕ^l. That is, when no computing power is contributed, users' utility is independent of participation.

asymmetric toward B at $\bar{\zeta}$ with fees $L = (L^A, L^B)$ and optimal equilibrium play by record-keepers, there exists a unique equilibrium of the stage game when $\zeta \leq \bar{\zeta}$. In this equilibrium, proposer P^B announces

$$\tilde{L}^B = \arg\min_{L^B} \tilde{u}(s = 0, \phi = 0, \zeta, L, C(\phi, L)),$$

all users and record-keepers choose ledger B, and record-keepers break even.

Proposition 5 is a remarkable result. It states that in a setting in which there is an opportunity to fork a blockchain, users will always choose the branch of the fork that has more favorable fees, and proposers (developers) will propose rules that are beneficial to users rather than record-keepers.[15] Of course, the result that proposers suggest protocols beneficial to users depends partly on the assumption that proposers' incentives are aligned with those of users. But in a setting with free entry of record-keepers, this assumption is not overly restrictive. Record-keepers always make zero profits, so proposing a ledger that increases record-keepers' revenues is pointless. Furthermore, users 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, users' stakes impede switching to a ledger where record-keepers earn lower revenue.

Proposition 5 highlights the benefits of a blockchain. When all users' fundamental preferences for an alternative ledger are identical, the absence of switching costs induces full coordination on the competing ledger. There is perfect competition among ledgers in that as long as it is feasible to make ledger B even slightly more desirable than ledger A, the competing ledger will win out over the existing one. Remarkably, there is perfect competition between ledgers. Coordination inefficiencies are precluded under these assumptions, but in the Appendix we will discuss how coordination can break down and multiple ledgers are used in equilibrium when users have heterogeneous fundamental preferences.[16]

Popular discussion has largely focused on the ways in which blockchains can decrease essentially exogenous costs, such as by inducing faster consensus about a ledger's contents. This result shows there is an *endogenous* channel through which blockchain reduces the

[15] Note that the hypothesis $\zeta \leq \bar{\zeta}$ 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 $\zeta = \bar{\zeta}$.

[16] The results in the Appendix are reminiscent of most hard forks, in which part of the community chooses to adopt the new fork of the blockchain and the other chooses the established blockchain. However, in many of these cases (such as with Ethereum or any of the forks of Bitcoin) the vast majority of blockchain users choose the same branch of the fork. We interpret our results as making predictions regarding which branch of the fork will be adopted by the larger group.

cost of maintaining a ledger: the synergy between *portability of information* and *competition among record-keepers*. When information can be ported to an outside ledger, users will want to use that ledger if record-keepers are paid lower fees (as long as the ledger is cryptographically secure). Individually, record-keepers are better off writing on a ledger with high fees, but competitive forces drive record-keepers to undercut each other by writing on the ledger with lower fees. Record-keepers know that all users will use the outside ledger when there are enough record-keepers to secure it, so the end result is that all record-keepers must switch to the outside ledger. The downside of a blockchain is that, while in a traditional setting, record-keepers' fees simply represent a (possibly distortionary) transfer, in the case of blockchain record-keepers' fees are a pure waste of resources. Later, we will examine under what conditions a traditional ledger maintained by a monopolist induces a large extraction of rents.

3.4 Currency competition

Although cryptocurrencies are perhaps currently the most important application of blockchain technology, the assumptions of our benchmark model are not directly applicable to cryptocurrencies. After all, there is no need for agents to use only one currency. Furthermore, in a model of money, agents should be able to exchange their currency holdings with one another. In this section, we briefly describe a formal monetary model of currency competition and show that, relaxing the restrictive assumptions made in our benchmark model, the main intuition that information portability enhances competition between ledgers goes through.

Our model of currency competition is based on the Lagos-Wright (2005) model, which is a workhorse in monetary economics. In the model, time is discrete and runs forever. All goods j are produced by labor, $Y_j = L_j$. Agents' utility of consumption is $u(c)$, the disutility of labor is l, and the per-period discount factor is δ. Within each period, there are two subperiods: day and night. During the day, *specialized* goods are traded in a decentralized market. Anonymous agents meet in pairs, and a double-coincidence-of-wants problem forces them to use currency in order to transact. Buyers make take-it-or-leave-it offers to sellers in which they propose an exchange of a given quantity of goods for money.[17] At night, agents trade *general* goods in a centralized market. Agents have homogeneous preferences for general goods.

We assume there are two currencies, A and B, and that users (agents) must pay a participation cost χ^l to adopt currency l. Adopting currency l enables a user to accept and use it in decentralized market transactions, as well as to trade it in the centralized market without

[17] We make this assumption in order to keep the characterization of agents' value functions simple, since it allows us to assume that buyers capture the entire trade surplus.

incurring a deadweight loss τ per unit of currency exchanged. For simplicity, we assume $\tau = 1$, so users who do not adopt a currency are completely unable to exchange it. Users are permitted to adopt one currency, both, or neither. The participation cost can be thought of as the effort expenditure required in order to recognize counterfeits, collect information about temporary price fluctuations, or understand the monetary policy of a given currency, for example.[18] For simplicity, we will assume $\chi^A = 0$. That is, users do not have to pay a cost to adopt the established currency and hence always adopt it. The only decision they make is whether to adopt B.

In this class of models, an agent's utility consists of two terms: the value of the agent's wealth (in monetary holdings) and the continuation value obtained through future trades in the decentralized market. First consider a user who does not choose to adopt currency B. If this user has initial holdings m^A of currencies A, the value function can be written as

$$W(m^A|A) = \psi^A m^A + \frac{\delta}{1-\delta} E[S|A], \tag{4}$$

where ψ^A denotes the price of a unit of currency A in terms of goods and $E[S|A]$ denotes the agent's expected surplus in future decentralized market meetings given that only A was adopted. For a user who adopts both currencies, the value function is

$$W(m^A, m^B|A, B) = \psi^A m^A + \psi^B m^B + \frac{\delta}{1-\delta} E[S|A, B] \tag{5}$$

if that user holds m^A units of currency A and m^B units of currency B. Here, ψ^B and $E[S|A, B]$ are defined analogously to ψ^A and $E[S|A]$.

We examine two distinct settings. With traditional Hayekian competition, the new currency B is issued by some outside agent who sells the currency to users in exchange for some of their holdings of A. With fork competition, on the other hand, all agents with stakes in the original currency are initially endowed with identical stakes of the new currency B. There is no need to purchase currency from any external issuer.

A user's decision to adopt currency B affects her utility through two channels: it may yield a discrete jump in the value function if she is initially endowed with some of currency B (only with fork competition), and it increases the surplus that she can expect to earn in future meetings if B has some transactional benefits relative to A (e.g., a lower cost of carry due to inflation). What matters for our application is that the increase in utility from adopting

[18] More realistically, users who decide not to accept or use the new currency could conceivably sell their entire initial stake in the new currency immediately. It is not obvious that doing so would entail any sort of substantial effort cost, but when users are unaware of the intricacies of a given currency it is possible that they will trade that currency at a loss due to an asymmetric information problem.

B depends on the fraction of others who choose to adopt B. In particular, we focus on the case where there are positive network externalities: that is, $W(m^A, m^B|A, B) - W(m^A|A)$ is increasing in the adoption of B.[19]

There are positive network externalities for two reasons. First, when other users adopt currency B, the probability of meeting a seller who accepts B in the decentralized market increases. This effect reduces the inflation cost of carrying the new currency because the inflation cost is incurred only when agents are unable to meet a seller who accepts the currency. Therefore, the marginal surplus of holding an additional unit of B increases, making it relatively more desirable to adopt B. Second, the increased liquidity of B induces agents to give it a greater portfolio weight, which increases the real value of the currency. This increase in the value of B, in turn, may increase the value of the initial stakes that users might have in currency B and thus the desirability of adoption.

As before, we select an equilibrium by using our global games framework. When there is incomplete information about χ^B, as long as network externalities are sufficiently strong this refinement implies the existence of a cutoff χ^{B*} such that when $\chi^B < \chi^{B*}$, all agents will choose to adopt B.

Proposition 6. *When the utility functions in Equations 4 and 5 describe a supermodular game with strong network externalities, there exists a cutoff χ^{B*} such that all users adopt B when $\chi^B < \chi^{B*}$. This cutoff is lowest when users have no initial stake in currency B.*

Proposition 6 states that the cutoff participation cost below which users adopt currency B is higher when users have an initial stake in the ledger. That is, if users are initially endowed with some of the new currency (as they would be after a blockchain fork, but not after an initial coin offering) they are more likely to expend effort to understand how it is traded. Even those who do not foresee themselves accepting the currency in future transactions have an incentive to do so in order to avoid losses when selling their initial stake or when purchasing goods. The knowledge that others have an incentive to pay this effort cost eases coordination on the new currency. This endowment effect, therefore, could help a currency that is superior to an established one along some dimension (e.g., lower inflation) to gain widespread recognition.

[19] We show in the Appendix that, under a reasonable calibration of the steady state of a Lagos-Wright-style model, there are indeed positive network externalities in the adoption of B. These network externalities are strong enough that the only equilibria are ones in which either all users adopt B or none do so.

4 Free Entry and Dynamic Ledger Choice

To fully realize the benefits of fork competition, a ledger must allow for free entry of record-keepers as well as portability of information. All of our substantive results on the benefits of fork competition so far, however, have relied only on the assumption that information is portable from one branch to the other. In our static model, free entry of record-keepers serves only to pin down the strength of the ledger's cryptographic security. Here, we show that in a dynamic setting, free entry is essential in fostering competition among ledgers.

We first show that in a setting with restricted entry, even when information portability is permitted, it may be possible for an incumbent who manages a centralized ledger to prevent an entrant from competing. The incumbent can threaten to dynamically punish the entrant by cutting its fees and driving rents down. Thus, perfect competition between ledgers will not obtain. We also briefly address the case of *permissioned* blockchain. We show that fork competition in a permissioned setting is limited by the restrictions on entry for record-keepers. It will be possible for collusion between permissioned record-keepers to prevent low fees from emerging. This example will further validate our claim that free entry is an economically relevant feature of public blockchains.

We then show that with a proof-of-work blockchain, this is not the case. We derive the equilibrium of the *repeated* blockchain fork game of Section **??**. We show that, remarkably, users and record-keepers must play the static equilibrium of Proposition 5 in every period of the blockchain ledger-choice game. In short, this is because the free-entry condition guarantees that record-keepers cannot be rewarded or punished by any dynamic scheme. Therefore, record-keepers will not be able to collude with each other on an outcome that is beneficial to them. The benefits of fork competition for public blockchains will thus be present in our dynamic setting as well.

4.1 Dynamic ledger choice with restricted entry

In order to highlight the essentiality of free entry, we consider a setting in which information is portable but entry is restricted. In particular, we examine competition between an incumbent centralized ledger writer and a single entrant who is able to replicate all of the information in the incumbent's ledger. The entrant's ability to port information over to its ledger can be thought of, perhaps, as the result of some policy that weakens the stranglehold that monopolistic record-keepers have on their data. We do not specify the mechanism that allows for portability of information because this example is simply an exercise to illustrate the consequences of restricted entry.

To analyze this problem, we will lay out a simple dynamic model of ledger competition.

Intuitively, restricted entry will hinder competition because the entrant will need to be guaranteed rents in equilibrium in order for the decision to enter to be ex-post optimal. Hence, the incumbent will be able to *dynamically punish* the entrant by threatening to act more competitively and drive rents to zero. This threat can prevent the entrant from attempting to compete in the first place. By contrast, in a setting with free entry, there will be no rents in equilibrium, so such threats have no bite.

Consider the repeated version of the ledger choice model presented in Section 3. There are two centralized entities that act as proposers and record-keepers: an incumbent P^A and an entrant P^B who discount payoffs at rate $\delta < 1$. The static ledger choice game is played in periods $t = 1, 2, \ldots$. There is only one difference: in subperiod $\tau = 0$ of period $t = 1$, the entrant may pay a cost $c > 0$ to enter. This allows the entrant to propose fees $L_t^B \in \mathbb{R}_+$ in all future periods. Furthermore, P^B is allowed to replicate the information in ledger A, so if P^B enters, users' stakes are identical on the two ledgers. As argued previously, when information is portable, users will simply choose the ledger with the lower fee.

If P^B chooses not to enter, users choose between ledger A and an outside option (called "autarky") instead of between A and B. We assume that the optimal fee charged by A when users face autarky as their outside option is L^{A*}.[20]

In this model, entry is restricted in the sense that only one entity is allowed to compete with the incumbent. In reality, it must be that the potential set of entrants is limited. It is not possible to allow any anonymous agent to write on the ledger without some form of identity management. Our assumption that only one entity is eligible to enter the market can be viewed as an extreme version of this observation. This entity must have some "reputation capital" that enables it to write the ledger, and, presumably, its reputation allows for the possibility to extract rents. While the costs of entry can be viewed as physical resource costs required to set up a new ledger, they can also be viewed as an opportunity cost that this known entity incurs by using its reputation capital in this market rather than another one.

We now show that in a subgame-perfect equilibrium, it is often possible for the incumbent to prevent the entrant from competing. The incumbent does so by threatening to drive rents to zero conditional on entry. The equilibrium is characterized by the following strategy profile:

- At $t = 1$, the entrant chooses not to enter. The incumbent sets $L^A = L^{A*}$.
- For $t \geq 2$, if the entrant chose to enter at $t = 1$, both the incumbent and entrant set their fees to zero: $L^A = L^B = 0$.
- For all t, if the entrant chose to enter at $t = 1$, users choose the ledger with the lower fee. If the fees are equal, they choose each ledger with equal probability.

[20] The existence of such an optimal fee is guaranteed by the analysis in Section 3.

In order to verify that this is an equilibrium, we need only check that it is not strictly optimal for the entrant to enter at $t = 1$. In all future periods, the incumbent and the entrant play a static Nash equilibrium in which they charge fees equal to zero.[21] When the entrant chooses to enter, all users choose ledger B as long as $L^B < L^{A*}$, so the entrant's profits are at most L^{A*} in the first period. Both the incumbent and the entrant earn no profits after $t = 1$, so it is optimal to enter only if $L^{A*} \geq c$. Otherwise, the strategy profile outlined above constitutes an equilibrium.

Note the discrepancy between the threshold cost for entry and the value of the rents earned by the incumbent. The incumbent would be willing to pay up to $\frac{L^{A*}}{1-\delta}$ for the right to remain in operation, but the entrant is willing to pay no more than L^{A*} for the right to enter. That is, for entry to occur, the costs must be only a fraction of the profits that can be earned, meaning that efficiency is severely compromised. The model also suggests another possibility: even if there were no fixed cost of entry, the incumbent would likely be willing to pay off any other entity that gained the ability to enter its market. Contracting with competitors in this way is possible only when competitors are non-anonymous. In doing so, the incumbent would preserve its future profits.

Our results are summarized in the following proposition:

Proposition 7. *In the dynamic centralized ledger-choice game, if information is portable between ledgers, there is a threshold cost of entry c^* above which entry can never occur:*

$$c^* = L^{A*}.$$

The threshold cost is less than the present value of the incumbent's profits in the case without entry.

Our results on ledger competition without free entry of record-keepers can also be used to understand the limitations of permissioned blockchains. A permissioned blockchain can be thought of as a blockchain run by a consortium of M record-keepers. Given that there is no need for identity management, there are no proof-of-work costs in this setting. However, in principle, the blockchain can be forked in order to create a competing ledger.

Without free entry, the possibility of collusion via dynamic reward and punishment schemes will again impede competition among ledgers. It will not always be optimal for a group of permissioned record-keepers to create a competing ledger that undercuts the fees charged by the existing blockchain, Specifically, there are equilibria in the game played among permissioned record-keepers in which a subset $N < M$ of record-keepers can be prevented from creating a competing blockchain. The strategy profiles are analogous to the one

[21] If there were variable costs of writing the ledger in each period, they would instead set fees equal to cost.

described in the game between centralized record-keepers: when a group of record-keepers defects and undercuts the blockchain's fees, the existing blockchain will cut its fees as much as possible, thereby lowering the rents that can be earned by the record-keepers who defected. Formally, in the Appendix we present a model of a permissioned blockchain and show that defection can be deterred if

$$\frac{1}{1-\delta}\frac{L^{A*}}{M} > \frac{L^{A*}}{N},$$

where L^{A*} is the fee charged by the existing blockchain. This formula is the analogue of that in Proposition 7: here the cost of creating a new blockchain is simply the present value of the flow of rents $\frac{L^{A*}}{M}$ accruing to permissioned record-keepers. Again, the key point is that non-competitive behavior is sustained by collusion, and the threat of future punishments that destroy rents precludes perfect competition among ledgers. This type of collusion will be impossible with free entry.

Our results on permissioned blockchain tie into the blockchain trilemma. To restate the main point, there is nothing inherent in the blockchain data structure itself that impedes rent-seeking behavior. Adding a costly identity management system to allow for free entry of record-keepers, in fact, increases the costs of using the ledger for a *given* set of policies. However, perfect competition among record-keepers, 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 users.

4.2 Permissionless blockchain competition in a dynamic setting

We now analyze dynamic fork competition and show that free entry guarantees perfect competition among ledgers. The repeated game with a permissionless blockchain is played in periods $t = 1, 2, \ldots$. In each period, proposers, users, and record-keepers play the stage game. Users are short-lived and die after one period, but record-keepers and proposers P^A, P^B live forever and discount payoffs at rate δ.

The observable quantities are the fork $L^B \geq 0$ that was proposed, how many users chose branch A, and how much computing power was committed to each branch in each period. The ledger l chosen by the majority of users at $\tau = 1$ becomes the reference parameter on branch A in the next period. That is, when users choose a particular branch of the blockchain fork, that chain is extended and becomes the default for developers to build off of if they want to fork in the future. Users observe their own private signals, and record-keepers observe the entire public history.

We define subgame-perfect equilibrium in the usual way. We now show that in any SPE of the repeated game, record-keepers 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. In that equilibrium, the outcome most favorable to users always obtains, so competition among ledgers is indeed perfect in this dynamic setting.

Proposition 8. *In any SPE of the repeated game, record-keepers make zero profits. The unique SPE is the equilibrium of Proposition 1 played in every period t.*

Free entry of record-keepers forces their rents to zero. It is therefore not possible for record-keepers to collude on an outcome in which they maintain only an inferior ledger that users dislike because such a collusion scheme would require dynamic rewards or punishments, neither of which is possible when record-keepers break even in every period.[22] Therefore, record-keepers always write on the ledger preferred by users, and users are able to coordinate on that ledger without fearing it will lack cryptographic security. Proposers then face the same actions by users and record-keepers as in the static game, so they play the same action as in the static game.

Although the lack of viable dynamic punishment schemes intensifies competition among ledgers, we will show in the next section that it poses a problem for the incentivization of accurate record-keeping. The mechanisms used to incentivize centralized record-keepers to report honestly typically involve dynamic punishments in which their future rents are reduced. These mechanisms will be infeasible in a setting with free entry because record-keepers always break even.

5 Incentivizing Honesty

Traditional ledgers have been criticized for being vulnerable to tampering by a malicious record-keeper. One of the principal advantages of blockchain protocols is that the ledger is resilient to fraud by a single bad actor. Record-keepers are able to discipline each other. Additionally, the identity management mechanism (proof-of-work) makes it expensive to attack the blockchain by rewriting history but comes at the cost of a waste of resources.

We will now analyze the security of blockchains as well as traditional ledgers maintained by centralized entities. We outline a simple model of blockchain security and compare the security of a blockchain to that of a ledger written by a monopolist. At one extreme, we show that while centralized intermediaries have fully *dynamic* incentives not to distort their ledgers.

[22] More broadly, our point extends to a situation in which record-keepers do earn some rents (e.g., because there is a fixed cost of computing power that must be paid upfront). The key intuition is that a record-keeper's future rents are low relative to the static profits that could be made if the majority of record-keepers refused to support a given branch of the fork.

They can always make a short-term profit by reporting dishonestly because records written on the ledger are *final*, but in doing so they stake their entire future stream of profits. At the other extreme, blockchain record-keepers' incentives must be completely *static* because they earn no rents. The *reversibility* of the blockchain disciplines attackers because any fraudulent actions they take may be undone once discovered by the community.

5.1 Centralized ledger security

We first analyze the case where a monopolist is able to distort its own ledger while facing competition from a fixed outside ledger. There is a monopolist who discounts payoffs at rate δ, a manager of the outside ledger, and a continuum $i \in [0,1]$ of users who live for one period. We outline the timeline of the model and then specify agents' preferences.

Timeline: In each period t at $\tau = 0$, the monopolist proposes a fixed fee L^A and the outside proposer announces a fixed L^B.[23] Ledger B can be thought of as an exogenously given outside option. Users choose a ledger at $\tau = 1$. At $\tau = 2$, each record-keeper chooses its own ledger.

The monopolist is able to take an action to distort the ledger at $\tau = 2$ of each period. The monopolist chooses an action $h_t \in [0, \overline{h}]$ at $\tau = 2$. The action h_t represents the record-keeper's *honesty*: a smaller h_t represents a larger distortion of the ledger attempted by the monopolist.

There are public signals that permit users to detect distortions of the ledger. These signals could correspond to news media revealing that fraudulent activity has occurred, or large numbers of people realizing that their accounts on the ledger have been compromised and spreading word of the attack. In each period t, a public signal $y_t \in \{0,1\}$ is observed at $\tau = 2$ with probability $\Pr(y_t = 1|h_{t-1}) = p(h_{t-1})$.[24] We assume $p(\overline{h}) = 0$, so that no deviation is detected when the monopolist is completely honest, and that $p'(h) < 0$, so that when the monopolist's distortion is severe, it is more likely to be revealed to the public.

Preferences: Users' fundamental preferences for ledger A are given by $\tilde{u}_{i,t} = s_i + \zeta - \gamma E[(\overline{h} - h_t)|\{y_s\}_{s=1}^t]$, where s_i is user i's stake on ledger A and users receive signals $x_i = \zeta + \sigma\eta_i$ as usual. A user's utility is decreased when the monopolist distorts the ledger and plays $h < \overline{h}$. Given that the monopolist's actions to distort the ledger are final, users' decisions are forward-looking: their decision to continue to use ledger A depends on their beliefs about the monopolist's future behavior.

[23] We abstract from proposers' choices of L^A and L^B in order to focus on the problem of incentivizing honesty rather than the problem of competition among ledgers.

[24] Here the assumption that y_t takes only the values 0 or 1 restricts the set of possible equilibria. We impose this assumption so that we may analyze an equilibrium in which users inflict the harshest possible punishment on the monopolist upon detecting dishonesty.

All records written on the monopolist's ledger are final, so the profits earned through a distortion of the ledger cannot be nullified. When the monopolist distorts the ledger at time t, he immediately receives $(L^A + \overline{h} - h_t)\phi_t$, where ϕ_t is users' participation on ledger A. That is, profits from dishonesty are realized immediately, so a dynamic punishment scheme is needed to incentivize honesty.

Equilibrium: There will be multiple equilibria because there is no mechanism to pin down users' expectations of future play. However, we can establish a lower bound on the fee required by the monopolist to ensure that $h = \overline{h}$ 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 users 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 users, it suffices to assume that there is an action $\hat{h}$ the monopolist can take so that $\max_i s_i - \gamma(\overline{h} - \hat{h}) - \alpha(L^A - L^B) < 0$, meaning even the type who is most anchored to ledger A by a personal stake in the system prefers to leave the ledger when users expect $\hat{h}$ to be played going forward. The expectations that justify this equilibrium, then, are

$$E[\overline{h} - h_t | \{y_s\}_{s=1}^t] = \begin{cases} 0 & y_s = 0 \ \forall \ s \leq t \\ \overline{h} - \hat{h} & \exists \ s \leq t, \ y_s = 1 \end{cases}.$$

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

$$\max_h \ \overline{h} - h + \delta(1 - p(h))V.$$

The first-order condition is

$$1 = -p'(h^*)\delta V.$$

When the monopolist plays h^*, we have $V = \frac{L^A + \overline{h} - h^*}{1 - \delta(1 - p(h^*))}$. A sufficient condition for a unique optimum $h^* \in [0, \overline{h}]$ to exist is then just $\frac{d^2}{dh^2} - \frac{(L^A + \overline{h} - h)}{1 - \delta(1 - p(h))} < 0$. This condition is similar to the increasing hazard rate assumption that will be made in the next section. To ensure the monopolist plays $h^* = \overline{h}$, we need

$$1 < -\frac{\delta}{1 - \delta} p'(\overline{h}) L^A. \tag{6}$$

The monopolist trades off the marginal gain of dishonesty, which is a unit payoff in the current period, against the marginal cost. The marginal cost is the increase in the probability of detection, $-p'(\overline{h})$, times the discounted value of *all* future profits, $\frac{\delta}{1-\delta}L^A$. This condition is usually somewhat weak: in order for the monopolist to be dishonest, the marginal benefit of stealing must be on the order of his franchise value as long as the marginal increase in the probability of detection is not too low. Therefore, it tends to be cheap to provide a monopolist with incentives for honesty as long as the harshest punishment scheme can be implemented. This result is restated in Proposition 9.

Proposition 9. *Assume that the monopolist's objective function (under the harshest punishment scheme) is concave, i.e.* $\frac{d^2}{dh^2} - \frac{(L^A+\overline{h}-h)}{1-\delta(1-p(h))} < 0$. *As long as*

$$L^A > -(\frac{1}{\delta} - 1)\frac{1}{p'(\overline{h})}$$

the monopolist never finds it optimal to distort the ledger.

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 that result from users' stakes on the ledger. The less likely the monopolist is to be detected in its deviations, the higher this bound must be. It is worth noting that the fee charged by a monopolist may be far from this bound. If the rents earned by a monopolist through his normal activities are large, there is no competitive force to push the fee it charges down to the level derived in Proposition 9.

Although the fees required to incentivize the monopolist to report truthfully are not necessarily large, we highlight that the harshest punishment scheme is only one of many possibilities in equilibrium. In general, it will be more difficult to incentivize the monopolist. First, in a setting with a traditional ledger, equilibrium is not unique, so while it may be the case that under the harshest possible punishment scheme it is not necessary to pay an intermediary large fees to obtain ledger security, the equilibrium fee required to ensure good behavior may be much higher. Second, in this case the assumption that signals y_t are either zero or one is not without loss of generality. A richer signal structure would lead to even greater multiplicity that would allow the monopolist to "nickel and dime" users by proving to them that, although he is distorting the ledger, he is not doing so to the extent that users would prefer to switch to a competitor and lose their stakes in the established ledger.

5.2 Blockchain security and rollback

In this section, we present a model of blockchain security. There will be occasional opportunities for record-keepers with large computing capacities to attack the blockchain. The dynamic mechanism used by centralized record-keeping systems, in which a record-keeper's future rents are reduced when a distortion of the ledger is detected, is not feasible with free entry because rents are driven to zero. The incentives for these attackers to refrain will be static: they will trade off the benefits of distorting the ledger against the one-time cost of conducting an attack. The key technological feature of blockchains that will allow for a static incentive mechanism is *rollback*. Attacks that are detected can be immediately reversed by the community. In this sense, blockchain's lack of finality is a feature rather than a bug. Rollback provides an additional defense against attacks that goes above and beyond the onerous cost to acquire the computational power necessary to overwhelm the rest of the network.

Whereas we focused on hard forks that changed the blockchain's policies in Section 3, here we consider forks that roll the blockchain back to a previous state in the wake of an attack but keep policies intact. When a malicious record-keeper attacks the blockchain, such as by attempting a double-spend, the blockchain typically forks. The attacker creates one branch of the blockchain in secret while other record-keepers, being initially unaware of the attack, extend another branch. When the attacker manages to create a longer chain and reveals his branch, most users mechanically go along with it because their software looks for the longest chain. It may become clear only after some time that a malicious attack has occurred, at which point part of the community may propose a rollback of the blockchain to a state preceding the attack.[25] In our model, we will interpret this type of attack as lowering users' stakes on the attacker's chain. A rollback will permit users to recover the stakes that they lost at the cost of undoing some transactions. Figure 2 illustrates this process.

We will analyze several features of blockchain security. We differ from previous literature on the subject (e.g., Biais et al. (2017)) by focusing on how coordination among users of the system and the possibility of rollback influence the success of an attack. We show that, as argued by Budish (2018), proof-of-work may be expensive because the incentives of record-keepers are *static*, so the cost of providing incentives is a *flow* cost. Our novel result is that the possibility of rollback facilitates coordination among users to discipline bad actors. In

[25] While the attack on the Ethereum blockchain in 2016 was not a double-spend, it remains by far the largest attack on a cryptocurrency blockchain. It indeed resulted in a hard fork that rolled back the blockchain and undid certain transactions, although a minority of the community continued to use the blockchain containing the entire history of transactions after the attack. Other, smaller, double-spend attacks have occurred on cryptocurrency blockchains that typically elicit some sort of response to improve security, but often these attacks are not severe enough to merit a rollback.

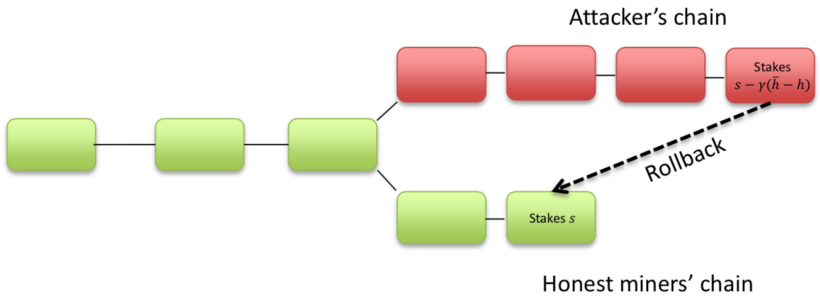

Figure 2: An illustration of a rollback response to a double-spend attack. The blocks mined by the attacker on the longer chain are marked in red, and, as described in the model, users' stakes on this chain are reduced. The community may roll back to the chain mined by honest record-keepers (in green) to recover users' lost stakes.

that sense, blockchain security is more robust than centralized ledger security. Our analysis thus highlights a tradeoff suggested by the Blockchain trilemma: cost inefficiency gives rise to a stronger form of ledger correctness.

The model of blockchain security is based on the repeated ledger-choice game. We first describe the timeline and then agents' preferences.

Timeline: As before, there are two proposers, P^A and P^B. In subperiod $\tau = 0$ of period t, Proposer P^A suggests extending the current longest chain, and proposer P^B suggests rolling back to the state at time $t-1$. Neither suggests a change of rules, so fees on both ledgers are equal to some fixed L.

There is a continuum of users $i \in [0,1]$ who live for one period. At $\tau = 1$ of period t, users choose between ledger A, the longest chain, and ledger B, the rolled back chain. Record-keepers who recorded information at time $t-1$ receive payoffs once users make their decisions.

The main difference between this model and our benchmark dynamic model of permissionless blockchain is in the modeling of record-keepers. In each period, there is a continuum of infinitesimally small record-keepers $j \in [0, M]$ with computing capacity dj. With probability $\mu > 0$, a single large record-keeper J is permitted to attempt an attack on the blockchain. For simplicity, we will assume $\mu \approx 0$, although all of our results will extend to the general case. This record-keeper has computing capacity M, so his computing power is enough to match the rest of the network's capacity. This assumption is meant to capture "51% attacks" in which an entity or mining pool able to control a majority of a blockchain's computing power mounts a malicious attack on the network in order to reap financial gains.[26] Record-keepers

who write on the ledger in period t receive payoffs at $t+1$ and then shut down operations.[27]

In subperiod $\tau = 2$ of period t, the large record-keeper may race the small record-keepers to create the longest chain. All record-keepers choose to write on the ledger chosen by users at $\tau = 1$, and then the large record-keeper J can attempt to distort the ledger by contributing at least as much computing power as the small record-keepers. Formally, if the computation power contributed by the large record-keeper is c_J and small record-keepers exert computational effort $c_j dj$, record-keeper J succeeds in creating the longest chain if $c_J \geq \int_j c_j dj$. The attacker may also attempt to distort the ledger by taking an action $h \in [0, \overline{h}]$, as in the model of centralized ledger security. Again, action $\overline{h} - h_J$ represents the record-keeper's dishonesty, which here can be thought of as the severity of a 51% attack. The type of attack modeled here is one in which the large record-keeper maliciously creates a fork of the blockchain on which he reverses previously accepted reports while small record-keepers write on the other branch. Users are initially fooled by large record-keepers' reports[28] and transact according to the his chain because it has greater proof-of-work.[29]

The structure of public signals is similar to that in the case with a centralized ledger. In each period $t > 0$, a public signal $y_t \in \{0, 1\}$ is revealed. The signal takes value 1 with probability $p(h_{t-1})$, where h_{t-1} is the action h played by the large record-keeper at $t-1$ (if an attack occurred and the attacker created the longest chain). We assume that

$$\frac{d}{dh} - \frac{p(h)}{1 - p(h)} < 0,$$

i.e., the hazard rate is decreasing, and again that $p(\overline{h}) = 0$, so no deviation is detected if the record-keeper is completely honest. In this setting, the assumption $y \in \{0, 1\}$ will be without loss of generality. The equilibrium will be the same regardless of whether users can perfectly observe $h_{J,t-1}$ with some probability.

Preferences: Users' preferences are as before. Their fundamental preferences for each

[27] Even in a setting where the primary purpose of the blockchain is not to facilitate the exchange of crypto-assets, 51% attacks would likely be financially profitable because almost all blockchains compensate record-keepers internally with coins. Thus, blockchains will typically contain some sort of financial asset that can, in principle, be double-spent.

[27] We make this assumption in order to abstract away from dynamic punishments for large record-keepers who can attack the network. This assumption is reasonable because large record-keepers 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 even if the blockchain completely shut down, these record-keepers could simply choose to attack another blockchain.

[29] If users 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. The success of any attack relies on users being fooled to a certain extent.

[29] A 51% attack works because users' software looks for the longest chain of blocks. Even though small record-keepers are sending reports as well, these reports are initially ignored by users.

branch of the fork are given by $u_{i,t} = \zeta - \gamma E[\overline{h} - h_{t-1}|y_t]$. The term ζ is a (vanishingly small) preference for the longer chain, reflecting the fact that users prefer a ledger that does not omit the most recent information. The term $-\gamma E[\overline{h} - h_{t-1}|y_t]$ corresponds to the fact that users' stakes are impacted by the attack and they can essentially reverse their losses from the *previous* distortion of the ledger by forking and rolling back to a point in the blockchain before the distortion occurred.

These preferences differ from those in the example with a traditional ledger in an important way. Users' stakes in the ledger at $t-1$ (which are related to distortions in previous periods) are now relevant. This is because users have the option to fork off to a ledger on which the distortion that occurred at $t-1$ never happened. Whereas in the model with a monopolist users' play was affected by public signals only because it was informative about the monopolist's continuation play, in this model public signals matter because they materially affect users' payoffs when they choose to roll back. In this sense, whereas users' actions in the example with a monopolist were *forward-looking*, in this example their actions are *backward-looking*. Users receive a noisy signal $x_i = \zeta + \sigma\eta_i$ of the common value as before, and their only signals of h_{t-1} are the public signals y_t. Users' types are hence given by $\theta_{i,t} = (x_i, y_t)$.

Given that all users have identical fundamental preferences for rolling back the blockchain, they will all coordinate on the same ledger (as in Proposition 5). There are two cases. If no attack is detected in period t ($y_t = 0$), users will prefer the longer chain because $\zeta > 0$ and the possibility that an attack occurred at $t-1$ is unlikely. If an attack is detected ($y_t = 1$), users will prefer to roll back the blockchain because their preference for the longer chain is small and their expectations of the severity of the attack, $\overline{h} - h_{t-1}$, are strictly positive.

Record-keepers are compensated for their computational effort at time t only if they contributed to the longest chain and users select the longest chain (ledger A) at $t+1$. If an attacker does not manage to overwhelm the rest of the network at time t, small record-keepers receive $\frac{c_j}{\int_{j'} c_{j'} dj'} \phi_{t+1} L \equiv \frac{c_j}{C} \phi_{t+1} L$ at $t+1$, where ϕ_{t+1} is the fraction of users who choose A at $t+1$. On the other hand, if an attacker does manage to overwhelm the network at t, the attacker receives $(L + \overline{h} - h)\phi_{t+1}$ in period $t+1$. Small record-keepers get nothing because they did not contribute to the longest chain.

Under the assumption that an attack is unlikely ($\mu \approx 0$), the aggregate computational effort of small record-keepers will always be equal to L, the fees that can be earned through honest record-keeping. Hence, in order to conduct a successful 51% attack, the large record-keeper will need to pay at least L.

Equilibrium: Having pinned down the behavior of users and small record-keepers, we may now analyze the large record-keeper's choices at $\tau = 2$ when an attack is possible. Of course, the only interesting case is the one in which the large record-keeper chooses $c_J = L$

and $h < \overline{h}$. We look for conditions under which he never does so in equilibrium. We have argued that whenever $y_t = 1$, the attack is not successful. When the attack is successful, the large record-keeper gets fees L plus the revenue $\overline{h} - h$ from the distortion. Therefore, the large record-keeper must solve

$$\max_h (1 - p(h))(L + \overline{h} - h).$$

The first-order condition implies

$$1 = -\frac{p'(h^*)}{1 - p(h^*)}(L + \overline{h} - h^*). \tag{7}$$

This equation illustrates the static nature of the attacker's incentives. The attacker cheats just enough so that the marginal gain from an increase in dishonesty, $1 - p(h^*)$, is equal to the marginal cost, which is just the profits $L + \overline{h} - h^*$ times the change in the probability of failure, $-p'(h^*)$.

Equation 7 implies that an equilibrium in which the attacker attempts to steal may exist only when the hazard rate $-\frac{p'(h)}{1-p(h)}$ is large enough. By our earlier assumption that the hazard rate is decreasing, a sufficient condition to ensure $h^* = \overline{h}$ is

$$1 < -p'(\overline{h})L. \tag{8}$$

When the hazard rate is large, the probability of detection is high enough to completely dissuade the large record-keeper from even attempting an attack. Note that this condition is satisfied for sufficiently large L, meaning that when the fee earned by blockchain record-keepers is high, even agents with the ability to subvert the network prefer not to attack it because they stand to lose the fee they would earn through honest record-keeping. Proposition 10 summarizes these findings:

Proposition 10. *The large record-keeper chooses not to attack the blockchain if and only if*

$$\max_h (1 - p(h))(\overline{h} - h) - p(h)L \leq 0.$$

A sufficient condition that guarantees this inequality will hold is

$$L > -\frac{1}{p'(\overline{h})}.$$

This bound on L characterizes the tradeoff between decentralization and cost efficiency

required to maintain correctness. It furthermore clarifies the nature of static incentives: the one-time cost of an attack L must be sufficiently high to deter malicious record-keepers. Relative to the case of a centralized ledger, security is expensive because it is based on a *flow* cost rather than a stock. The intuition for this result is simple and similar to that in Budish (2018): while a blockchain record-keeper is punished for misbehavior only through nullification of the profits obtained by attacking the blockchain, a monopolist is punished via the destruction of its franchise value, which consists of all future fees earned through honest play.

Proposition 10 has a striking implication. When the probability of detection is sufficiently large, it is unnecessary to set up an expensive fee structure for record-keepers that leads to a large waste of computational resources. Record-keepers will abstain from distorting the ledger regardless because each marginal unit of computational power spent on an attack earns less on average than one spent on keeping records honestly. The cost of conducting an attack, which is exactly equal to the fee earned in equilibrium, acts as further protection against attacks. The fact that incentives must be static coupled with the need for a pure waste of resources makes incentivizing honesty expensive for a blockchain.

There are several advantages that make the security of a blockchain more robust than that of a centralized ledger, however. Crucially, in this framework the equilibrium is unique: users always abandon the ledger after detecting an attack. The uniqueness of equilibrium is a direct consequence of fork competition. If an attack makes all users worse off they will coordinate on an alternative ledger on which the attack never happened, but the rest of the information on the ledger remains intact. Competition among record-keepers will cause them to coordinate on that ledger as well, and the attacker will get nothing. Hence, rollback is an effective security feature because it reduces the cost of negating an attack. As we will show, this static mechanism is quite different from that securing a centralized ledger. Also, the signal structure $p(h)$ may well be more revealing for a blockchain than for a traditional ledger, since blockchains are designed specifically to provide transparency about attacks on the ledger. Fork competition is thus important in securing a ledger as well as in forcing competition among record-keeper compensation schemes. Hence there is a sense in which, relative to blockchains, centralized ledgers trade away correctness in order to achieve greater cost efficiency (due to the fact that franchise values can be used to incentivize honesty).

5.3 Why proof-of-work?

One may wonder why we focus on PoW given that it is only one of many possible consensus algorithms. Popular alternatives that bear mentioning are proof-of-stake (PoS). These algorithms provide a way to achieve consensus without forcing record-keepers to solve difficult

computational problems and hence avoid the waste of resources that proof-of-work entails. However, the physical resource expenditures involved in PoW actually provide a distinct advantage in achieving consensus among a decentralized group of record-keepers: PoW makes it costly to extend invalid chains of blocks. With PoS, for example, a record-keeper may record information whenever a coin that she owns is randomly selected. That is, a record-keeper's voting power is proportional to the number of coins she owns rather than the computational effort she expends. If a record-keeper is randomly selected to record information on an invalid chain, however, there is very little to stop her from doing so since it costs her nothing. This impediment to consensus is known as the *nothing-at-stake problem* and is an active area of research in computer science.[30]

The PBFT algorithm essentially achieves consensus by selecting a "validating committee" and relaying transaction information among validators until a supermajority agree on the set of transactions to include in the next block. PBFT-style algorithms are often used by permissioned blockchains such as Ripple. It avoids the problem of forking by forcing validators to update the current state rather than allowing them to choose to roll back to a previous state. The lack of forking is detrimental in the sense that PBFT makes it more difficult to create a competing ledger that contains the information in the existing one. The main difficulty with PBFT, however, lies in how to select the validating committee without a system of identity management. There is currently a lack of consensus in the computer science community regarding how this might be done in a permissionless setting.

6 Enforcement

In this section, we discuss some practical matters and present results related to the application of blockchain when enforcement is necessary. The issue with blockchains is that, though 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 record-keepers 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

[30] While there are proposals that aim to curb the severity of the nothing-at-stake problem (such as Ethereum's Casper), to our knowledge most of these algorithms are based on the idea of punishing record-keepers who use the same address to validate transactions on two different chains. One can circumvent this security measure by splitting coins up over multiple addresses, thus minimizing the possibility that any one address is selected to record information on both chains.

when those contracts involve the transaction of physical assets. This type of enforcement would likely be the role of some centralized entity, 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 blockchain raises two issues. First, while several commentators claim that blockchain 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 blockchain is the solution. If the government is overly bureaucratic and incapable of setting up good institutions to track property rights, then a blockchain 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 blockchain will be useless. Again, the users of the ledger are the ultimate source of discipline, so a blockchain is useful only insofar as it helps them to discipline a corrupt government (through greater disclosure of information, most likely). If the enforcer is itself a private firm, such as a bank that enforces debt obligations, it may be optimal for the enforcer to maintain the ledger as well. The enforcer will have an incentive to fulfill its obligation for fear of losing the privilege of maintaining the ledger.

The second issue is that the enforcer must choose which forks of a blockchain to support. An enforcing entity cannot simply commit to enforce contracts on all forks because the same physical asset may be promised to two different individuals on different forks of the blockchain. The enforcer could say it will enforce all contracts so long as certain policies are followed, which prevents hard forks that change the 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 were to occur, such as the one on the Ethereum blockchain in 2016, the enforcer would have enormous power to resolve the issue in its own favor.

We have two formal results corresponding to these two issues. We fully lay out the model with enforcement in the Appendix, but present the main elements here. The only change to the benchmark dynamic model is that a fourth subperiod $\tau = 3$ is added to each day on which an agent known as the "enforcer" takes an action $e \in [0, 1]$. In each period, the enforcer earns a fee proportional to the participation in its ledger that does not depend on the action e. Users prefer the enforcer to play larger values of e, but by playing e the enforcer incurs a utility cost of e. Deviations from $e = 1$ are detected on the next day with probability $q(e)$. Users have the option to abandon the ledger and get zero utility at any time.

First, we show that there is a synergy between writing the ledger and enforcing its contents. An agent who both writes and enforces the ledger may distort the ledger by choosing an action h as in Section 4.4 and choosing an enforcement level e. The probabilities of detection

of these two actions are independent. We can then compare this situation to one in which there is a continuum of blockchain record-keepers in charge of reporting the ledger's contents and a separate enforcer. Proposition 11 summarizes the incentive compatibility constraints in the two situations.

Proposition 11. *When the enforcer also writes the ledger, the fee L earned by this entity must satisfy*

$$L \geq -\left(\frac{1}{\delta} - 1\right) \min\left\{\frac{1}{p'(0)},\ \frac{1}{q'(1)}\right\}, \tag{9}$$

where p is the detection function for distortion of a centralized ledger. On the other hand, when the ledger is written by a blockchain and enforced by a centralized entity, the fees L_W and L_E earned by blockchain record-keepers and the enforcer, respectively, must satisfy

$$L_W \geq -\frac{1}{p_B'(0)},\ L_E \geq -\left(\frac{1}{\delta} - 1\right)\frac{1}{q'(1)}, \tag{10}$$

where p_B is the detection function for distortions of the blockchain.

Clearly, the bound on the fee L derived in (9) is less than $L_E + L_W$ in (10) when δ is reasonably large and $p_b \approx p$. When distortion of the ledger and lack of enforcement are strategic substitutes, bundling record-keeping and enforcement together reduces the fee required to incentivize honest behavior. The fee only needs to be large enough that the intermediary will not choose the more attractive of the two deviations, after which the incentive compatibility condition for the other type of deviation is automatically satisfied. Furthermore, bundling is beneficial because it eliminates the waste of resources required by a blockchain. The only case in which it may be beneficial to keep record-keeping and enforcement unbundled is if the p function is much more sensitive for blockchains than standard ledgers, and the cost of enforcement $\frac{1}{q'(1)}$ is small relative to the cost of honest reporting.

Our second result is that when there is a centralized enforcer, there is no longer a unique equilibrium with a blockchain in which the ledger that is best for users is always selected. There are "dictatorial" equilibria in which the enforcer effectively decides which branch of a fork is chosen.

Proposition 12. *With a centralized enforcer, there is always an equilibrium in which the ledger preferred by the enforcer is chosen by all users.*

We defer the proof of this proposition to the Appendix, but the intuition is simple. The enforcer just threatens not to enforce contracts on any branch of a fork that uses policies of which it does not approve. Users will not want to coordinate on any ledger that the enforcer will ignore, so they use the one selected by the enforcer. The existence of a centralized entity

that is indispensable for the proper functioning of the ledger destroys the benefits that come with decentralizing the record-keeping function.

7 Conclusion

We present a general model of ledger competition and apply it to understand when a blockchain is more economically beneficial than a centralized ledger managed by a centralized intermediary. Our analysis of the tradeoffs between centralized and decentralized record-keeping is guided by the blockchain trilemma. First, we examine the competitive benefits of decentralization. We show that two key technological features of blockchains that promote competition between ledgers are information portability and free entry of record-keepers. The portability of information allowed by fork competition reduces switching costs for users when a new ledger emerges to compete with an established one. Thus, fork competition eases coordination among users on an entrant ledger with better protocols and policies. Free entry is critical for fork competition as well because it forces record-keepers' rents to zero. In particular, the absence of rents prohibits record-keepers from coordinating on an outcome detrimental to users. Such coordination requires dynamic punishment schemes, but in an environment without rents, such punishments are infeasible. Record-keepers simply compete to write on the ledger preferred by users. By contrast, in an environment with portability of information but no free entry of record-keepers (e.g., a permissioned blockchain), it need not be the case that the policies most beneficial to users emerge in equilibrium.

We then explicitly examine the costs of incentivizing record-keepers to report honestly (correctness). We show that on the one hand, centralized intermediaries are incentivized dynamically. Any distortion of a centralized ledger is final, so when fraud is discovered, the record-keeper must be punished by a reduction in future profits. Centralized record-keepers are incentivized by ensuring that the future profits they will earn are large enough to guarantee they do not want to risk losing them. On the other hand, blockchain record-keepers do not earn rents, so they must be incentivized statically. Rollback, the third important technological feature of blockchain, enables static incentivization by making attacks on the blockchain reversible. Attackers hence weigh the probability that a 51% attack will succeed against the one-off cost of conducting such an attack. Correctness is incentivized by raising the *flow* cost of proof-of-work to the point that attacks become unprofitable. It is expensive to incentivize blockchain record-keepers with flow rewards, but the security provided by blockchains is more resilient to misbehavior because of the rollback feature. Blockchains hence sacrifice cost efficiency for decentralization and more robust correctness.

We highlight the important distinction between ownership and possession. Blockchains

can effect transfers of ownership, but when enforcement of possession rights is required it is often more efficient to bundle record-keeping and enforcement duties in a single centralized intermediation structure.

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

References

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

[2] Eric Budish. The economic limits of bitcoin and the blockchain. Working paper, 2018.

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

[4] Jonathan Chiu and Thorsten Koeppl. The economics of cryptocurrencies– bitcoin and beyond. 2017.

[5] William Lin Cong and Zhiguo He. Blockchain disruption and smart contracts. 2017.

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

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

[8] Lukasz Drozd and Ricardo Serrano-Padial. Financial contracting with enforcement externalities. *Journal of Economic Theory*, 178:153–189, 2018.

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

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

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

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

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

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

[15] Ricardo Lagos and Randall Wright. A unified framework for monetary theory and policy analysis. *Journal of Political Economy*, 113(3):463–484, 2005.

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

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

[18] Stephen Morris and Hyun Song Shin. Global games: Theory and applications. In Mathias Dewatripont, Lars P. Hansen, and Stephen J. Turnovsky, editors, *Advances in Economics and Econometrics: Theory and Applications, Eighth World Conference*, Cambridge, 2003. Cambridge University Press.

[19] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. Whitepaper, 2008.

[20] Emiliano Pagnotta and Andrea Buraschi. An equilibrium valuation of bitcoin and decentralized network assets. 2018.

[21] Jean-Charles Rochet and Jean Tirole. Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4):990–1029, 2003.

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

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

[24] Ricardo Serrano-Padial. Coordination in global games with heterogeneous agents. Working paper, 2018.

[25] Michael Sockin and Wei Xiong. A model of cryptocurrencies. 2018.

[26] Karl Wüst and Arthur Gervais. Do you need a blockchain? *IACR Cryptology ePrint Archive*, 2017.

A Proofs

Proof of Proposition 1:

Proof. This proposition is an immediate consequence of Lemma B.1 in Appendix B. □

Proof of Proposition 2:

Proof. This proposition follows from Theorem B.1 in Appendix B. □

Proof of Proposition 3:

Proof. The equilibrium follows from Propositions B.2 and B.3. □

Proof of Proposition 4

Proof. This proposition follows from equations 17, 23, and 21. □

Proof of Proposition 5:

Proof. We prove the proposition by backwards induction.

τ**=2**: At $t = 2$, record-keepers know the value of ϕ^l, $l \in \{A, B\}$. We show that $C^l = \phi^l L^l$ in equilibrium. Suppose first that $C^l < \phi^l L^l$. Then since M is large, there exists a record-keeper j such that $c_j < 1$, but record-keeper j could make profits by setting $c_j = 1$ because

$$\frac{1}{C^l}\phi^l L^l - 1 > 0$$

Now suppose $C^l > \phi^l L^l$. This means that any record-keeper j for whom $c_j > 0$ would benefit by setting $c_j = 0$, since

$$\frac{c_j}{C^l}\phi^l L^l - c_j = (\frac{1}{C^l}\phi^l L^l - 1)c_j < 0$$

Hence $C^l = \phi^l L^l$.

τ**=1**: We will guess and verify that in any equilibrium, $\tilde{u}(s = 0, \phi, \zeta, L, C)$ is asymmetric towards B. Record-keepers' optimal play at $\tau = 2$ implies that $\frac{C^l}{\phi^l} = \underline{C}$ for each branch of the fork. Then it must be that P^B selects L^B such that

$$\tilde{u}(s = 0, \phi = \frac{1}{2}, \zeta, L, C = (\phi L^A, (1 - \phi)L^B)) \leq 0$$

This utility function is isomorphic to a symmetric utility function with positive stakes on B by setting

$$\tilde{s}_i = g(L^B) - g(L^A) - \alpha(L^B - L^A) \equiv \tilde{g}(L^B) - \tilde{g}(L^A)$$

We may now apply Theorem B.1 given that we have a symmetric utility function with positive stakes on one of the two ledgers. According to Theorem B.1, then, type $\tilde{s}_i$'s cutoff signal $k(s_i)$ is nonnegative, so as long as $\zeta \leq 0$, all users will have such signals when σ is sufficiently close to zero. Therefore all users play B.

τ**=0**: Now we confirm our guess that the proposer selects such a value L^B if possible. The equilibrium derived above shows that whenever $\tilde{g}L^B < \tilde{g}(L^A)$, the proposer obtains a payoff of $\tilde{u}(s = 0, \phi = 0, \zeta, L, C(\phi = 0, L))$. It is never possible for the proposer to obtain a higher payoff by choosing L^B such that $\tilde{g}(L^B) < \tilde{g}(L^A)$. Then it must be that the proposer chooses the lowest possible L^B in order to maximize payoffs, so

$$L^B = \arg\min_{L^B} \tilde{u}(s = 0, \phi = 0, \zeta, L, C(\phi, L))$$

□

Proof of Proposition 6

Proof. When $\chi^A = 0$, it is always a dominant strategy to adopt currency A given that there is uncertainty about the adoption of B. Then this game reduces to one with two possible choices: to adopt B or not to adopt B. When χ^B is negative, it is dominant to adopt B. When χ^B is large and positive, it is dominant to not adopt B, so we have shown the existence of dominance regions. Then by Lemma B.1, there is a unique cutoff equilibrium. The proof of Theorem B.1 shows that the cutoff χ^* is smallest when all users' initial stakes are equal to zero. □

Proof of Proposition 7

Proof. The proof is outlined in Section 4.1. □

Proof of Proposition 8

Proof. First we show that in subperiod $\tau = 2$ after any history, on either branch l of the fork, the total computing power contributed by record-keepers must be $\phi^l L^l$. Suppose that $C^l < \phi^l L^l$. Then there must be some record-keeper j who contributes $c_j < 1$. By deviating to $c_j = 1$ on branch l, this record-keeper can achieve positive profits in the current period. Furthermore, this record-keeper's deviation does not affect any publicly observable signal in

the future history, since the record-keeper is of measure zero. An analogous argument shows that C^l cannot be greater than $\phi^l L^l$, so $C^l = \phi^l L^l$ at any history.

Second, we must check that proposers play static best responses. Given that both users and record-keepers play the same strategies that they do in the stage game, a proposer can maximize her flow of payoffs by playing the same L^B as in the stage game. □

Proofs of Proposition 9 and Proposition 10:

Proof. Proposition 9 follows from the analysis in Section 5.1. Proposition 10 follows directly from optimality condition 7. □

Proof of Propositions 11 and 12:

Proof. See Appendix C.4. □

B Global Games with Heterogeneous Preferences

The majority of the proofs in this section are adapted from Drozd and Serrano-Padial (2017) and Serrano-Padial (2018). There is a continuum of players $i \in [0, 1]$ who play a one-shot coordination game in which they choose between two options, A and B. Players' fundamental preferences consist of heterogeneous private values $\theta_i \in \mathbb{R}$ for choice A. There is also a common value $\tau \in [\tau_l, \tau_u]$ that affects players' preferences for A. We assume that θ is iid across players with distribution $F(\theta)$. For now, we assume F is a discrete distribution with finite support but later take the limit of a continuous distribution F. Players' preferences can be described by the function

$$v(\theta, \tau, \phi)$$

where ϕ is the proportion of players who choose A. That is, there is a coordination motive. When $v(\theta, \tau, \phi) > 0$, it is a best response for a player of type θ to choose A. Conversely, a player of type θ should choose B if $v(\theta, \tau, \phi) < 0$. We additionally impose a symmetry assumption. Assumption B.1 summarizes the restrictions on v:

Assumption B.1. *We make the following assumptions about v:*

1. *Function $v(\theta, \tau, \phi)$ is increasing in θ;*
2. *Function $v(\theta, \tau, \phi)$ is increasing in τ;*
3. *Function v is **symmetric** in the sense that $v(0, 0, \phi) = -v(0, 0, 1 - \phi)$.*

4. *There exist dominance regions for both actions:* $\exists\, \overline{\tau} < \tau_u, \underline{\tau} > \tau_l$ *such that* A *is dominant for all* θ *whenever* $\tau > \overline{\tau}$ *and* B *is dominant for all* θ *whenever* $\tau < \underline{\tau}$.

Henceforth we will assume that players have incomplete information about the common value τ. We assume players have a uniform prior over τ[31] and receive signals $s_i = \tau + \sigma\eta_i$ ($\sigma > 0$), where η_i is iid across players and independent of τ. The noise term η_i is distributed with CDF $H(\eta)$ with support on the interval $[-\frac{1}{2}, \frac{1}{2}]$. In what follows, we will frequently consider the limit $\sigma \to 0$.

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

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

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

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

Proof. The assumption that there exist dominance regions for actions A and B mean we may focus on signals $s \in [\underline{\tau} - \frac{\sigma}{2}, \overline{\tau} + \frac{\sigma}{2}]$ (since agents with signals outside that interval know τ lies

[31] In the limit $\sigma \to 0$, the results would be unchanged if we were to assume an arbitrary prior on τ with full support. See Frankel, Morris, and Pauzner (2003).

in a dominance region). For small σ, the density of τ conditional on s is $h(\frac{s-\tau}{\sigma})$. We have

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

□

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

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

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

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

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

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

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

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

Proof. Begin by defining "virtual types" $\delta(s, \theta) = s - k(\theta)$. This reduces the two-dimensional

type space to a one-dimensional one. Agents play A whenever $\delta(s,\theta) > 0$ and B when $\delta(s,\theta) < 0$. With this definition,

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

For brevity, we will denote $\delta(s_i,\theta_i)$ by δ_i. First we find the marginal density of (s,θ) conditional on $\theta \in \Theta$ after integrating out τ. Then we show that $\Pr(\psi(\tau,\Theta) \leq z|\delta(s_i,\theta_i) = 0) = z$. This property is due to Morris and Shin (2003).

We have

$$\Pr(s,\theta,\tau|\Theta) = \Pr(s|\theta,\tau,\Theta)\Pr(\theta|\tau,\Theta)\Pr(\tau|\Theta)$$

Recalling that the prior on τ is uniform, this yields

$$\Pr(s,\theta,\tau|\Theta) = \frac{1}{\sigma}h\left(\frac{s-\tau}{\sigma}\right)\frac{f(\theta)}{\sum\limits_{\theta'\in\Theta} f(\theta')}$$

whenever $s \in [\underline{\tau} + \frac{\sigma}{2}, \overline{\tau} - \frac{\sigma}{2}]$. Cutoffs $k(\theta)$ must always lie in this range as argued earlier. When $s \in [\underline{\tau} - \frac{\sigma}{2}, \overline{\tau} + \frac{\sigma}{2}]$, the marginal density is then

$$\Pr(s,\theta|\Theta) = \int\limits_{s-\frac{\sigma}{2}}^{s+\frac{\sigma}{2}} \frac{1}{\sigma}h\left(\frac{s-\tau}{\sigma}\right)\frac{f(\theta)}{\sum\limits_{\theta'\in\Theta} f(\theta')}dF(\tau) = \frac{f(\theta)}{\sum\limits_{\theta'\in\Theta} f(\theta')}$$

Define $\tilde{\eta}_i = \frac{\delta_i - \tau}{\sigma}$, and denote the distribution of $\tilde{\eta}$ conditional on $\theta \in \Theta$ by $\tilde{H}_\Theta$. Observe that

$$\tilde{H}_\Theta(\tilde{\eta}) = \sum_{\theta\in\Theta} \Pr\left(\frac{s-k(\theta)-\tau}{\sigma} < \tilde{\eta}\right) = \sum_{\theta\in\Theta} \Pr\left(\epsilon < \tilde{\eta} + \frac{k(\theta)}{\sigma}\right)$$

Hence $\tilde{\eta}_j$ is iid across players and independent of τ.We have

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

Now to complete the proof, observe that

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

Given the uniform prior over τ, the information environment is translation-invariant, so

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

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

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

as desired. □

Up until this point, none of the results have depended on taking the limit $\sigma\to 0$; we just needed σ to be small enough that $\underline{\tau}-\frac{\sigma}{2}>\tau_l$ and $\overline{\tau}+\frac{\sigma}{2}<\tau_h$. Now we specialize to the case considered in the text where σ becomes arbitrarily small and define $\mathbf{k}^\sigma$ to be the threshold equilibrium played for variance parameter σ. Correspondingly, we denote a specific type θ's cutoff by $k^\sigma(\theta)$. We then define

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

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

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

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

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

Proof. Fix $\tilde{\sigma}>0$ and define a partition of types $\Theta_1,\dots,\Theta_S$ by placing two types θ,θ' in the same equivalence class whenever $|k^{\tilde{\sigma}}(\theta)-k^{\tilde{\sigma}}(\theta')|<\tilde{\sigma}$. Define $Q_\theta^{\tilde{\sigma}}(\chi|\mathbf{k}^{\tilde{\sigma}},z)=\Pr(\tau\leq\chi|s=$

$k^{\tilde{\sigma}}(\theta), \psi(\tau, \Theta_i) = z)$ (for $\theta \in \Theta_i$) to be type $k^{\tilde{\sigma}}(\theta)$'s belief about τ conditional on the event that a proportion z of players in the same equivalence class of the partition play B. We have

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

The term in the integrand is bounded by $v\Bigg(\theta, k^{\tilde{\sigma}} \pm \frac{\tilde{\sigma}}{2}, 1 - 2\sum_{\Theta_j, j<i} f(\theta) - 2z\sum_{\Theta_i} f(\theta)\Bigg)$, so

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

and

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

Note that as $\tilde{\sigma} \to 0$, the right-hand side of B converges to the right-hand side of B as long as dA_θ is bounded, which is shown in Lemma 10 of Drozd and Serrano-Padial (2017).

Now, for each i, take some arbitrary $\theta_i \in \Theta_i$ and set $k_i = k^{\tilde{\sigma}}(\theta_i)$. As σ is taken to zero from $\tilde{\sigma}$, set cutoffs $\hat{\mathbf{k}}_\sigma^{\tilde{\sigma}}$ so that $\Delta_{\theta_i, \theta_i'} = \frac{k_i - \hat{k}_\sigma^{\tilde{\sigma}}(\theta_i')}{\sigma} = \frac{k_i - k^{\tilde{\sigma}}(\theta_i')}{\tilde{\sigma}}$ for all $\theta_i' \in \Theta_i$. Note that $A_\theta(z|\hat{\mathbf{k}}_\sigma^{\tilde{\sigma}}, \Theta_i)$ is constant as $\sigma \to 0$ under these transformed cutoffs. Then as $\sigma \to 0$,

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

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

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

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

k_i was arbitrary, this proof has the additional implication that $k^\sigma(\theta) \to k_i$ uniformly for all $\theta \in \Theta_i$. The uniqueness of the partition is immediate from the uniqueness of equilibrium in this setting. □

In order to proceed, we will need to define stable and unstable equilibria of the complete information game.

Definition B.1. *An equilibrium strategy profile $a(\theta)$ constitutes an **unstable** equilibrium of the complete information game if for all $\mu > 0$, if a fraction μ of players' choices are changed from the prescription of $a(\theta)$ (from A to B), all players wish to play B, and if a fraction μ of actions are changed from B to A, all players wish to play A. An equilibrium profile $a(\theta)$ constitutes a stable equilibrium if it is not an unstable equilibrium.*

We now prove the main theorem.

Theorem B.1. *Suppose that in the complete information game, for any value of τ there are at most two pure strategy (stable) equilibria: one in which all players choose A and one in which all players choose B. In the limit $\sigma \to 0$, the equilibrium strategies are given by a cutoff k common to all players such that*

1. *When $\theta_i = 0$ for all i with certainty, $k = 0$;*
2. *When $F(0) = 0$ and $F(\theta) < 1$ for some $\theta > 0$, $k < 0$.*

Proof. First note that as $\sigma \to 0$, the cutoff used by all agents must be the same. Suppose that there are two groups of types Θ_1 and Θ_2 such that types in Θ_1 use k_1 and types in Θ_2 use $k_2 > k_1$ as $\sigma \to 0$. Then there would be an equilibrium of the complete information game in which types in Θ_1 choose A and types in Θ_2 choose B when $\tau \in [k_1, k_2]$. By assumption, this equilibrium must be unstable. That is, for all $\tau \in [k_1, k_2]$ and any $\mu > 0$, all agents would prefer to play A if the fraction of agents playing A were $Q(\Theta_1) + \mu$, and all agents would prefer to play B if the fraction of agents playing A were $Q(\Theta_1) - \mu$.

For $s \in [k_1 - \frac{\sigma}{2}, k_1 + \frac{\sigma}{2}]$, there is uncertainty about whether other types in Θ_1 will play A. In fact, the belief constraint implies that some cutoff type must believe at least half of all others in Θ_1 will play B (or at least half of all others in Θ_2 will play A). By the definition of an unstable equilibrium, there must then always be some cutoff type who wants to deviate from the equilibrium cutoff strategy (since this type believes that the fraction of agents who play A is strictly less than $Q(\Theta_1)$). Hence an unstable equilibrium cannot be the equilibrium selected in the limit, meaning there must be a cutoff common to all agents.

The rest of the proof follows from Lemma B.3. The cutoff k must satisfy the limit condition

$$\int_0^1 v\Big(\theta, k, 1-2z\sum_{\Theta} f(\theta')\Big) dA_\theta(z|k) = 0$$

for all $\theta \in \Theta$. When $\Theta = \{0\}$ the belief constraint (Lemma B.2) implies that $dA_\theta(z|k) = dA(z|k) = zdz$. Then we have

$$\int_0^1 v(0, k, 1-2z) z dz = 0$$

From the symmetry assumption in Assumption B.1 it is immediate that this condition is satisfied when $k = 0$.

On the other hand, when there is a positive mass of agents with $\theta > 0$, we have

$$\sum_{\Theta} f(\theta) \int v\Big(\theta, k, 1-2z\sum_{\Theta} f(\theta')\Big) dA_\theta(z|k) = 0$$

by summing over θ. Then by the assumption that v is increasing in θ,

$$\begin{aligned}
\sum_{\Theta} f(\theta) \int v\Big(\theta, k, 1-2z\sum_{\Theta} f(\theta')\Big) dA_\theta(z|k) &> \sum_{\theta} f(\theta) \int v\Big(0, k, 1-2z\sum_{\Theta} f(\theta')\Big) dA_\theta(z|k) \\
&= \int v\Big(0, k, 1-2z\sum_{\Theta} f(\theta')\Big)\Big(\sum_{\theta} f(\theta) dA_\theta(z|k)\Big) \\
&= \int v\Big(0, k, 1-2z\Big) z dz
\end{aligned}$$

where the first equality is obtained by noting that the integrand is independent of θ and the second follows from the belief constraint. Note that if $k \geq 0$, then the right-hand side is nonnegative, meaning that $\int v\Big(\theta, k, 1-2z\sum_{\Theta} f(\theta')\Big) dA_\theta(z|k)$ must certainly be positive. Thus the limit condition can only be satisfied by some $k < 0$. □

We may specialize the results to a case in which the function v is linear in order to derive further analytical expressions characterizing the equilibrium. From now on, we assume that

$$v(\theta, \tau, \phi) = \theta + \tau + \kappa\phi$$

The following proposition describes the main results in this linear case.

Proposition B.1. *In the limit $\sigma \to 0$, the equilibrium strategies are given by a monotone*

partition $\Theta_1, \ldots, \Theta_S$ *of* Θ *and cutoffs* $k_1 > \cdots > k_S$ *such that*

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

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

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

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

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

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

a contradiction. From this it immediately follows that

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

which is point (ii).

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

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

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

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

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

Equipped with Proposition B.1, we may now prove some properties of equilibria when the distribution F satisfies certain conditions. We consider three scenarios:

1. F is continuous and $\theta + \kappa(1 - 2F(\theta))$ is monotonically increasing;

2. F has a symmetric, single-peaked density f and $\theta + \kappa(1 - 2F(\theta))$ is non-monotonic;

3. F is a two-point discrete distribution.

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

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

Proof. We show that the partition described in Proposition B.1 must consist of singletons in this case. Suppose that $\theta_1 < \theta_2$ are the boundaries of equivalence class i of the partition. By property (ii) of Theorem 4, we have

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

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

Proposition B.3. *Suppose F has a symmetric, single-peaked density f and $\theta+\kappa(1-2F(\theta))$ is non-monotonic. Let $\hat{\theta} = \arg\max_{\theta} f(\theta)$. The equilibrium is characterized by a parameter Δ such that*

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

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

Proof. Observe that under the assumptions on F, there must be only one interval $[\underline{\theta}, \overline{\theta}]$ where $\theta + \kappa(1 - 2F(\theta))$ is decreasing. All θ in this interval must belong to the same equivalence class of the partition described in Proposition B.1. We show this by contradiction. Suppose $\theta \in [\underline{\theta}, \overline{\theta})$ is at the upper boundary of an equivalence class Θ_i and $\theta' \in (\underline{\theta}, \overline{\theta}]$ is the lower bound of Θ_{i+1}. Then by point (ii) of Proposition B.1 we have

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

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

Hence the entire increasing region $[\underline{\theta}, \overline{\theta}]$ belongs to a single equivalence class of the partition. At the boundaries of the equivalence class containing that interval, $k(\theta)$ must be continuous (which follows by again applying the argument showing that there cannot be two equivalence classes containing points in the increasing region). By the argument in Proposition B.2, there cannot be an equivalence class of the partition containing only points in the decreasing region, so it must be that the partition consists of a single equivalence class $[\underline{\theta}, \overline{\theta}]$ containing all values of θ in the increasing region and singletons for all θ outside that interval.

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

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

Continuity of the cutoff at the boundaries of the interval implies

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

Rearranging these expressions, we find

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

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

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

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

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

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

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

Proposition B.4. *When F is a two-point distribution with support $\{\theta_L, \theta_H\}$ (and $\theta_L < \theta_H$)*

such that $\Pr(\theta = \theta_L) = \mu$, $\Pr(\theta = \theta_H) = 1 - \mu$, *the equilibrium cutoffs are*

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

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

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

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

Consider the equilibrium with finite, nonzero σ. We have

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

where $\Delta_{H,L} \equiv \frac{k(\theta_L) - k(\theta_H)}{\sigma}$ and the third line uses the fact that H is the uniform distribution on $[-\frac{1}{2}, \frac{1}{2}]$. Similarly, we find

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

Suppose that as $\sigma \to 0$, $\Delta_{H,L} \to \xi$. Then these equations imply

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

so

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

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

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

when the cutoffs are equal.

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

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

by part (iii) of Proposition B.1. Rearranging, we get

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

Given that $k(\theta_L) > k(\theta_H)$, we must have

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

which completes the proof. □

C Applications

C.1 Competition between centralized ledgers

Here we analyze the optimal fees charged by the two centralized record-keepers in Section 3.2. First, we consider the simpler case in which $d < \kappa$. In this case, network externalities are so strong that all users will end up choosing the same ledger regardless of how invested they are in ledger A. Proposition 3 shows that when network externalities are strong, the incumbent and entrant effectively compete à la Bertrand. Each will try to undercut the other as long as it is possible to do so. However, the incumbent has a competitive advantage corresponding to the average stake S users have in its ledger. Therefore, in equilibrium the entrant must choose $L^B = 0$, and the incumbent monopolist chooses L^A just small enough so that users do not switch to B. By Proposition 3, this yields

$$L^A = \frac{S}{\alpha} \tag{17}$$

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

Now we consider the case in which $d > \kappa$. By Proposition 3, when the monopolist selects L^A, all users for whom $\frac{1}{2} + \kappa^{-1}\Big(s_i - \alpha(L^A - L^B)\Big) > Q(s_i) = \frac{s-S}{d} + \frac{1}{2}$ choose to remain on ledger A. To find the cutoff type s^* who is indifferent between remaining on the monopolist's ledger and leaving, we solve

$$\frac{1}{2} + \kappa^{-1}(s - \alpha(L^A - L^B)) = \frac{s-S}{d} + \frac{1}{2}$$

which implies

$$s^* = \frac{d}{d-\kappa}\Big(\alpha(L^A - L^B) - \kappa(\frac{S}{d} - 1)\Big) \tag{18}$$

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

$$\phi(L^A, L^B) = 1 - Q(s^*(L^A, L^B)) = \frac{S + \frac{d}{2} - \frac{\kappa}{2} - \alpha(L^A - L^B)}{d - \kappa}$$

Then the monopolist's problem reduces to

$$\max_{L^A} \ \Big(S + \frac{d}{2} - \frac{\kappa}{2} - \alpha(L^A - L^B)\Big)L^A$$

which yields

$$L^A = \frac{S + \frac{d}{2} - \frac{\kappa}{2} + \alpha L^B}{2\alpha} \tag{19}$$

The rents extracted by the monopolist are increasing in the average stake on its ledger because when the average stake is higher, users 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 users, so an upwards adjustment of L^A does not result in a large exodus of users from ledger A. Finally, when the parameter L^B is large, users are reluctant to leave ledger A because they know that they will be charged high fees on the outside ledger regardless, so the monopolist enjoys higher profits.

On the other hand, a strong coordination motive can be detrimental to the monopolist's business. If the coordination motive is strong, when a single marginal user leaves the ledger it induces many other users 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 users' 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

depends only on users' fundamental preferences– the ledger that is best for users is chosen automatically. In the traditional environment, when even partial competition is possible, network externalities work as a disciplining device against the incumbent monopolist. That is, network externalities enhance the importance of ledgers' fundamental parameters when replication of information and perfect, blockchain-style competition between record-keepers is impossible.

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

$$\frac{\frac{d}{2} - \frac{\kappa}{2} - S + \alpha(L^A - L^B)}{d - \kappa}$$

The entrant's problem is then

$$\max_{L^B \geq 0} \left(\frac{d}{2} - \frac{\kappa}{2} - S + \alpha(L^A - L^B) \right) L^B$$

The first-order condition of this problem is

$$L^B = \frac{\frac{d}{2} - \frac{\kappa}{2} - S + \alpha L^A}{2\alpha} \tag{20}$$

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

Equation 20 shows that the entrant will extract high rents if the dispersion in users' stakes is large or if the incumbent also extracts large rents. When the dispersion in users' 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 users no matter which ledger they ultimately choose. When the fundamental parameter L^A on the incumbent's ledger is large, users 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 users' 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 users leaving ledger B. The fee charged by the entrant is also decreasing in the mean stake S on the incumbent's ledger because that stake gives the incumbent a competitive advantage, so the entrant must charge a lower fee in order to capture a significant segment of the market.

In order to find the equilibrium of the game between the incumbent monopolist and the entrant, we simply combine their first-order conditions. Hence we simultaneously solve

equations 19 and 20. This yields

$$L^A = \frac{1}{2\alpha}(d-\kappa) + \frac{1}{3\alpha}S,\ L^B = \frac{1}{2\alpha}(d-\kappa) - \frac{1}{3\alpha}S \tag{21}$$

Then participation on each ledger is

$$\phi^A = \phi = \frac{1}{2} + \frac{1}{3}\frac{S}{d-\kappa},\ \phi^B = 1-\phi = \frac{1}{2} - \frac{1}{3}\frac{S}{d-\kappa}$$

We need $0 \leq \phi^A, \phi^B \leq 1$ and $L^l \geq 0$ for $l \in \{A, B\}$. A necessary and sufficient condition is

$$S \leq \frac{3}{2}(d-\kappa) \tag{22}$$

This inequality is a *no-entry bound*. If this inequality does not hold, the incumbent A is in fact able to retain all users even when $L^B = 0$. That is, the stakes users have in ledger A endogenously prevent entry by even the most competitive entrant. While network externalities discipline the fees charged by the incumbent, inequality 22 shows that they actually impede entry by competitors as well. When the participation of others is important to users, it is difficult for a competitor to enter because it cannot attract enough users to get itself off the ground. On the other hand, when users' stakes on ledger A are dispersed, it is easier for the entrant to attract the users with the least to lose by switching, which in turn induces switching by other users. When the no-entry bound holds,

$$L^A = L^{NE} = \frac{1}{\alpha}\left(S - \frac{1}{2}(d-\kappa)\right) \tag{23}$$

The incumbent sets L^A to be the highest value such that all users participate in the ledger. We have the following results regarding the case with no entry.

Proposition C.1. *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. Users' welfare under the no-entry bound is decreasing in S, increasing in d, and decreasing in κ.*

Proof. These properties are a result of Equations 22 and 23. □

Now we turn to the case in which there is entry. Equation 21 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 user leaves a ledger, other nearly marginal users follow suit because of the importance of coordination. On the other hand, dispersion in stakes has the opposite effect. When users'

stakes are heterogeneous, only a small mass of users 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 users' stakes on ledger A is high, there is a competitive wedge between ledgers A and B. Monopolist $\mathcal{M}$ can extract higher rents than the entrant $\mathcal{O}$ because users' stake in ledger A acts as an inertial force preventing them from leaving.

C.2 Currency competition: Calibration

Here we briefly describe the parameters we use in our currency competition model and the system of equations we solve. We also show that under our calibration, the currency adoption decision exhibits positive network externalities.

First, when buyers make take-it-or-leave-it offers in the decentralized market, they are always able to purchase c units of goods when they hold real balances that would be worth c in the centralized market of that period. This is because buyers are always able to extract full surplus from sellers. Second, there are only two types of meetings: meetings in which the buyer encounters a seller who accepts only A (fraction ϕ^A of meetings) and meetings in which the buyer encounters a seller who accepts both A and B (fraction $\phi^{\cup} = 1 - \phi^A$ of meetings).

As is well known, all agents who make the same adoption decision will carry the same real balances into the decentralized market in this type of model. Let $c_{\cup}^A, c_{\cup}^B$ be the real balances of currencies A and B, respectively, held by agents who adopt B. Let c_A^A be the real balances of A held by agents who do not adopt B. In the centralized market, quasilinear preferences imply that agents always consume c^* such that $u'(c^*) = 1$. Hence we only need to pin down consumption in the decentralized market in order to solve the model. We focus on the region of the parameter space where agents' liquidity constraints bind in all meetings. For an agent who adopts B, the Euler equations for holdings of A and B are

$$u'(c^*) = \frac{\delta}{\pi^A}\left(\sigma(\phi^A u'(c_{\cup}^A) + \phi^{\cup} u'(c_{\cup}^A + c_{\cup}^B)) + (1-\sigma)u'(c^*)\right)$$

$$u'(c^*) = \frac{\delta}{\pi^B}\left(\sigma\phi^{\cup} u'(c_{\cup}^A + c_{\cup}^B) + (1-\sigma\phi^{\cup})u'(c^*)\right)$$

On the other hand, for an agent who adopts only A, the Euler equation is

$$u'(c^*) = \frac{\delta}{\pi^A}\left(\sigma u'(c_A^A) + (1-\sigma)u'(c^*)\right)$$

We assume $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ with $\gamma = 5$ and set $\delta = 0.96, \sigma = 0.05$.[32] We assume that

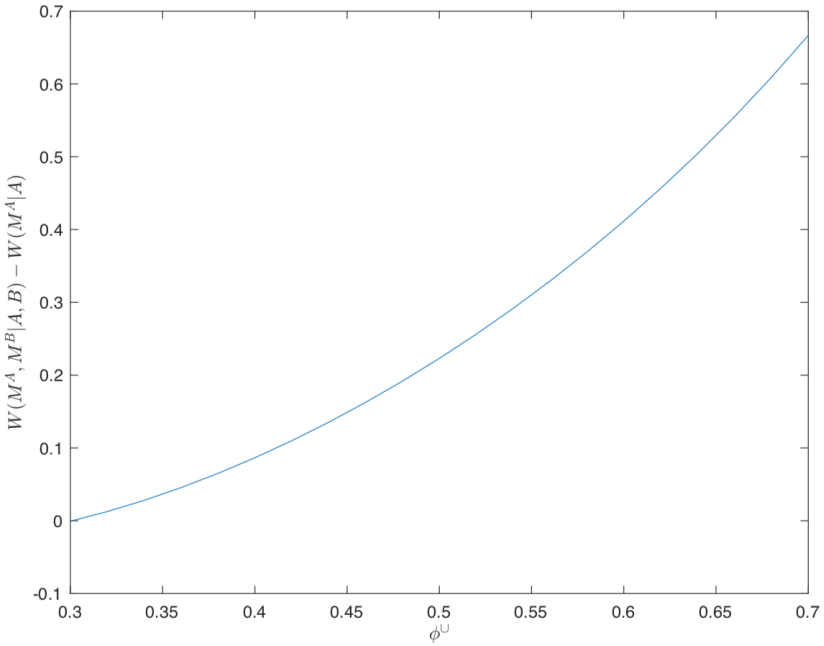

Figure 3: Coordination externalities in the model of currency competition. Parameters are given in Appendix C.2.

the established currency A is a high-inflation currency with $\pi^A = 25\%$ (e.g., the Argentine peso) and that currency B has no inflation, $\pi^B = 0$.[33] Using these parameters, the system of equations above can be solved for any values of $(\phi^A, \phi^\cup)$. We then show that there are strategic complementarities in the adoption of currency B by showing that the difference $W(M^A, M^B|A, B) - W(M^A|A)$ (defined in Equations 4 and 5) is increasing in the fraction $\phi^\cup$ of users who adopt B. Figure 3 illustrates this result.

C.3 Permissioned blockchain

We present a formal model of permissioned blockchain in order to study fork competition in this environment and prove the claims in Section 4.1. The model is similar to the baseline model with the exception that there is a finite number of record-keepers who do not incur computational costs. Play occurs in periods $t = 1, 2, \ldots$, and each day consists of subperiods $\tau = 0, 1, 2$, just as in the benchmark ledger-choice model. There are $M \in \mathbb{N}$ record-keepers

[33] As in Lagos and Wright (2005), the calibration of σ is somewhat arbitrary.

[33] In this class of models we may specify inflation directly by choosing the long-run rate of growth of the money supply.

who discount payoffs at rate δ. These record-keepers are split into two groups: group A (with N^A members) and group B (with N^B members). Each group l assigns a dedicated member P^l as a proposer. Proposer P^l may choose to announce a fee $L^l \in \mathbb{R}_+$ in subperiod $\tau = 0$ or to announce nothing. There is also a continuum of short-lived users $i \in [0, 1]$ who choose a ledger at $\tau = 1$ if a fork has been proposed. As permissioned blockchain does not require identity management, there are no computational problems to be solved. Record-keepers in group A do not take an action at $\tau = 2$. The only action record-keepers in group B take is to choose whether to continue to write on the established ledger A or the competing ledger B (if it has been proposed) at $\tau = 2$.[34]

Here branch A of the fork can be seen as the reference ledger, and branch B can be seen as a competing ledger that may emerge. Our main result will be that with a permissioned blockchain, it will be possible for record-keepers to prevent forking to branch B and thus perpetuate high user fees.

Given that information is portable from one branch of the fork to the other, users will simply coordinate on the ledger with the lower fees, as demonstrated in our analysis of the static ledger choice model. Our analysis hence focuses on record-keepers' actions.

The incidence of a fork in this setting should be thought of as a proposal to create a competing system (potentially made by one of the original record-keepers). A subset of the record-keepers may break off and begin writing on the new ledger, but the group of record-keepers who continue to write on the original ledger will be able to punish defectors by threatening to act more competitively and reduce the appeal of the new ledger relative to the old one.[35]

Now we show that when δ is sufficiently large or M is sufficiently small, there is an SPE of this game in which all record-keepers choose ledger A, and a new ledger with lower fees is never proposed. This is in contrast to the permissionless blockchain case, in which users and record-keepers would always coordinate on ledger B if the proposed fee L^B were lower than L^A. Consider the following equilibrium conjecture:

1. After any history in which all record-keepers chose A in all previous periods,

    - At $\tau = 0$, proposer P^A announces $L^A = L$ and proposer P^B announces nothing. All record-keepers and users choose A.

[34] We assume that if P^B chooses to announce a fork, he must write on ledger B at $\tau = 2$. This assumption is made to avoid a situation in which a fork has been announced, users choose that ledger, and no record-keeper writes on that branch of the fork.

[35] Another possible punishment is for record-keepers who break off to be sanctioned and forbidden from writing on the established ledger ever again. We do not consider this type of punishment, but the equilibria that could be sustained with this punishment are similar to those that can be sustained in our environment because both punishments essentially reduce to a threat to lower defectors' franchise values.

- At $\tau = 1$, and P^B announced a fork at $\tau = 0$ of the current period, all users choose A if $L^A \leq L^B$ and B otherwise.
- At $\tau = 2$, if a fork was announced at $\tau = 0$ and ϕ users chose A at $\tau = 1$, record-keepers in group B choose A if $\frac{\phi L^A}{N^A+1} > \frac{(1-\phi)L^B}{N^B}$ and B otherwise.

2. After any history in which P^B announced a fork in some previous period, proposer P^A sets fees equal to zero and P^B announces nothing. All record-keepers and users choose A.

This equilibrium conjecture is based on a grim trigger punishment scheme. Within a given period, if a fork has been proposed with $L^B < L$, record-keepers in group B have an incentive to choose ledger B because then all users switch to B and they obtain all the revenues on branch B. However, afterwards they receive payoffs equal to zero, and they cannot individually deviate to obtain higher payoffs because users will choose A (with $L^A = 0$) in every period.

The continuation play in the conjectured equilibrium depends only on the action taken by proposer P^B at $\tau = 0$. Otherwise, all agents play static best responses. Furthermore, the play after a deviation is a static Nash equilibrium. Hence we need only check the incentive compatibility condition at $\tau = 0$. If P^B deviates at $\tau = 0$, he (and the rest of group B) can obtain $\frac{L^B}{N^B}$ in the current period, but future payoffs are equal to zero. Formally, the incentive constraint that must be satisfied in order for proposer P^B not to deviate is

$$\frac{L^B}{N^B} \leq \frac{1}{M(1-\delta)} L^A$$

for all $L^B < L^A$. This inequality holds if and only if

$$\frac{M}{N^B} \leq \frac{1}{1-\delta}. \tag{24}$$

This inequality holds when δ is large or $\frac{M}{N^B}$ is small. When record-keepers are patient, they have an incentive to conform to equilibrium play because the threat of losing future payoffs is perceived to be harsh. When the total size of the record-keeping consortium M is small relative to the size of the deviating coalition N^B, there is also an incentive not to deviate because the static payoffs earned by continuing to write on the original ledger are large relative to those that can be earned by deviating to a competing ledger with fewer record-keepers.

C.4 Enforcement

Dictatorial equilibrium with an enforcer: In this section, we describe a model of ledger choice with enforcement issues. In addition to proposers, users, and record-keepers, there is a fourth type of agent known as an enforcer. This agent must take a costly action in a new subperiod $\tau = 3$ in order to enforce obligations written down in the ledger. Formally, the enforcer exerts effort e to enforce obligations and receives fees ϕL_E if ϕ users choose its ledger. For simplicity, in this section we assume that users choose between using a given ledger A and not using a ledger at all. Users get payoff $s_i + \kappa\phi - \alpha(L_W + L_E)$ when they use the ledger and the enforcer chooses to enforce at $t = 3$ and payoff $-\alpha(L_W + L_E)$ otherwise (where L_W is the fee earned by record-keepers). They may also choose not to use the ledger at all, in which case they get zero.

Of course, the enforcer has no incentive to take the costly action in the stage game, since the payment to the enforcer is not contingent on whether he takes the costly action. We assume, as in our dynamic setting, that the game is played on days $t = 0, 1, \ldots$. We now illustrate the presence of an enforcer can undermine the decentralization that is central to blockchain's benefits.

Intuitively, the need for enforcement (whether by a government or some other agency) is detrimental to decentralization because this entity can threaten to stop enforcing agreements unless the blockchain adopts certain protocols, thereby destroying the mechanism by which blockchain selects the outcome most beneficial to users. Blockchains that operate under conditions favorable to users will be dysfunctional because the enforcer will refuse to recognize them. In our model, this type of equilibrium can be captured in a simple way: suppose that there is a policy L_W^* preferred by the enforcer. The enforcer receives an additional benefit $V(L_W^*)$ whenever policy L_W^* is proposed. If the proposer in period t announces L_W^*, the enforcer takes the costly action to enforce at $\tau = 3$ of period t. Otherwise, the enforcer refuses to take the costly action in all subsequent periods, and users stop using the ledger.

Checking incentive compatibility of this strategy profile is simple. The enforcer may deviate either by choosing not to enforce in periods where he is supposed to or by enforcing when he is not supposed to. Clearly, the latter deviation is always suboptimal, since users do not use the ledger in any future period regardless of what the enforcer does. For it to be optimal to enforce along the equilibrium path, we must have

$$e \leq \frac{\delta}{1-\delta}(L_E - e + V(L_W^*))$$

That is, the one-shot benefit of shirking enforcement duties must exceed the enforcer's stream of future profits plus the benefits it receives when L_W^* is played on future days.

There is no incentive for any other agent to deviate from the prescribed strategy profile. Users play myopically, so they choose A whenever the enforcer is expected to take the costly action and exit otherwise. Record-keepers always make zero profits, so they have no incentive to deviate. Proposers get nothing if they propose anything other than L_W^*, so they also play according to the equilibrium prescription.

Of course, this equilibrium is not unique, but it illustrates that, unlike in the case without enforcement, there is not a unique equilibrium in which the best outcome for users is always realized. The equilibrium becomes dynamic rather than static because despite the fact that *record-keeping* is decentralized, *enforcement* is not decentralized. A single centralized enforcer that is crucial to the functioning of the blockchain may be able to subject the blockchain's policies to its will because, unlike other agents, it can threaten to prevent the blockchain from functioning properly.

Synergy between record-keeping and enforcement: We can also consider a situation in which the enforcer, record-keeper, and proposer are the same entity (analogously to the case without enforcement where monopolistic intermediaries were both record-keepers and proposers). The goal is to determine whether there is a synergy between record-keeping and enforcement in the sense that the rents required by an intermediary who performs both functions are less than those needed to compensate two separate entities who write on the ledger and enforce its contents.

The intermediary takes an action $h \in [0, \overline{h}]$ to distort the ledger and an action $e \in [0, 1]$ to enforce obligations. Deviations from the "honest" strategy $(h, e) = (0, 1)$ go undetected with probability $\tilde{p}(h, e) = (1 - p(h))(1 - q(e))$, where $p(h)$ is increasing in h and $q(e)$ is increasing in e. We further assume $p(\overline{h}) = 1$ and $q(1) = 1$, so $\tilde{p}(x, e) = 0$ when the intermediary plays honestly. The analysis in the previous section was equivalent to the assumption that $1 - q(e) = \mathbf{1}\{e < 1\}$, so the enforcer always chose $e \in \{0, 1\}$.

The intermediary obtains benefit $v(h, e)$ from taking action (h, e) in addition to a fee L that it receives every period (independent of its actions). We assume strategic substitutability between distortion of the ledger and lack of enforcement:

$$v(h, e) = \overline{h} - h - e$$

Under the harshest punishment, then, the intermediary solves

$$V = \max_{h,e} \tilde{V}(h, e) = \max_{h,e} \frac{L + \overline{h} - h - e}{1 - \delta(1 - p(h)))(1 - q(e))}$$

We may now derive bounds on the fee L such that the record-keeper-enforcer plays honestly.

The derivatives of the intermediary's objective function satisfy

$$\frac{\partial}{\partial h}\tilde{V}(h,e) = -(1-\delta(1-p(x))(1-q(e))) - \delta p'(x)(L+\overline{h}-h-e)$$

$$\frac{\partial}{\partial e}\tilde{V}(h,e) = -(1-\delta(1-p(x))(1-q(e))) - \delta q'(e)(L+\overline{h}-h-e)$$

If the intermediary plays honestly, it must be that $\tilde{V}$ is nondecreasing in e and nondecreasing in h at $(h,e) = (\overline{h},1)$. Hence we must have

$$L \geq \left(\frac{1}{\delta}-1\right)\max\left\{-\frac{1}{p'(0)},\ -\frac{1}{q'(1)}\right\} \tag{25}$$

We can now compare this situation to one in which the ledger is written by a decentralized group of blockchain record-keepers and enforced by an outside entity. We assume that users abandon the ledger whenever any type of deviation is detected, whether it is by the record-keepers or the enforcer. The enforcer earns a fee L_E per period and solves

$$\max_e \frac{L_E - e}{1-\delta q(e)}$$

As in the benchmark model of blockchain security, large blockchain record-keepers solve

$$\max_x (1-p_B(h))(L_W+\overline{h}-h)$$

where L_W is the aggregate fee earned by blockchain record-keepers. The (necessary) incentive compatibility conditions that come out of the enforcer's and record-keepers' optimization problems are

$$L_E \geq -\left(\frac{1}{\delta}-1\right)\frac{1}{q'(1)},\ L_W \geq -\frac{1}{p'_B(0)}$$

respectively.

D Extensions and Additional Results

D.1 A realistic "hard fork"

In this section, we analyze a hard fork that is more realistic than the type highlighted in the preceding analysis where *all* users of the blockchain switch to one branch of the fork and the other is completely abandoned. In reality, hard forks usually lead to a split of the community. For example, the Ethereum community split after hackers stole cryptocurrency from a smart contract. Although the majority of the blockchain's users joined the segment

of the community that decided to fork, a significant percentage of users continued to use the original blockchain. The Bitcoin blockchain has also been forked by the (significantly less popular) cryptocurrencies Bitcoin Cash and Bitcoin Gold, both of which changed the rules of Bitcoin in order to benefit users. In these cases, many users of Bitcoin refused to actively use the new cryptocurrencies because they felt that the changes to the rules were actually detrimental or compromised the security of the blockchain. This section will focus on the tradeoff between fork competition and network externality inefficiencies, the second cost of decentralization. Although fork competition can benefit users, we will show that it can also lead to inefficient miscoordination, or "too many ledgers" in equilibrium.

The key mechanism that will underlie realistic hard forks in our model is preference heterogeneity. Although in the benchmark model agents have heterogeneous stakes, we have shown that when considering a blockchain fork, stake heterogeneity is irrelevant. We now consider a model identical to the benchmark with the exception of the specification of types. Users' types are given by

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

where $f_i \in \{0, f\}$. The type f_i reflects a preference for forking: users with $f_i = f$ dislike all forks equally, and users with $f_i = 0$ are not averse to forking the existing blockchain.[36] Types f_i are independently and identically distributed across users with $\Pr(f_i = \eta) = \mu$. Types x_i are distributed uniformly in the interval $[\zeta - \frac{\sigma}{2}, \zeta + \frac{\sigma}{2}]$ as before. Users observe both x_i and f_i.

We will use essentially the same utility function $\tilde{u}$ as in the previous section. We assume

$$\tilde{u}(s, f, \phi, \zeta, L, C) = f + \kappa(2\phi - 1) + \zeta - \left(\alpha(L^A - L^B) - (g(\frac{C^A}{\phi}) - g(\frac{C^B}{1-\phi}))\right) \tag{26}$$

Here the only difference from the previous section is that users' fundamental preferences depend on f_i rather than on a stake s_i.

Let $\tilde{g}(L) = g(L) - \alpha L$. Note that if there exists $L^B \geq 0$ such that $\tilde{g}(L^B) > \tilde{g}(L^A) + f$, we obtain the same result as in Section 3.2. Proposer P^B will propose such an L^B and all users will switch to branch B. In this case, there exists a feasible fee L^B that is better than L^A by such a wide margin that all users, including those who dislike forks, prefer ledger B with parameter L^B.

We therefore consider only the case in which all $L^B \geq 0$ satisfy $\tilde{g}(L^B) < \tilde{g}(L^A) + f$. In fact, the only situation in which multiple equilibria would arise under complete information

[36] 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.

is if

$$f + \tilde{g}(L^A) - \tilde{g}(L^B) \geq \kappa(1-2\mu) \geq \tilde{g}(L^A) - \tilde{g}(L^B)$$

We derive the unique equilibrium under these conditions. The results are summarized in Proposition D.1.

Proposition D.1. *Suppose users face ledgers with fees L^A, L^B, $W^l \geq L^l$ record-keepers commit to branch l at $\tau = 2$, and a fraction μ of users are of type $f_i = f$. Then if*

$$f + \tilde{g}(L^A) - \tilde{g}(L^B) \geq \kappa(1-2\mu) \geq \tilde{g}(L^A) - \tilde{g}(L^B)$$

the essentially unique equilibrium at $\tau = 1$ is of one of two types.

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

Proof of Proposition D.1:

Proof. These statements follow from Proposition B.4. □

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

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

$$\mu(f - \tilde{g}(L^A) + \kappa\mu) + (1-\mu)(-\tilde{g}(L^B) + \kappa(1-\mu))$$

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

$$(1-\mu)(\tilde{g}(L^A) - \tilde{g}(L^B)) - 2\kappa\mu(1-\mu)$$

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

$$2\kappa\mu > \tilde{g}(L^A) - \tilde{g}(L^B)$$

The results of the previous section and this one highlight the main tradeoff relevant for determining whether a blockchain is worthwhile. Although a blockchain greatly enhances competition between ledgers and lowers fees, it may also induce an undesirable breakdown of coordination. The possibility of miscoordination is especially strong when network externalities are weak but the welfare losses are large only when coordination matters, so miscoordination should be a concern when network externalities lie in an intermediate range.

D.2 Multi-Homing

Up until this point, for analytical tractability we have made the extreme assumption that users use only one of the two ledgers. We now extend the analysis to allow users to use both ledgers simultaneously. The main results of our analysis go through: users' stakes in a traditional centralized ledger anchor them to it and hinder competition, but the portability of information permitted by a blockchain removes this barrier to entry.

As in the benchmark model, there are two ledgers A and B and three time periods, $\tau = 0, 1, 2$. However, users are allowed to "multi-home" and choose to use both ledgers at $\tau = 1$. In particular, users now have four possible choices $\{\emptyset, A, B, A \cup B\}$ instead of just two. Here $\emptyset$ and $A \cup B$ represent the choice not to use either ledger and the choice to use both ledgers, respectively. A user who uses ledger l pays a cost αL^l. By choosing just one ledger (say A), a user obtains utility from interacting with all other users who chose A or $A \cup B$. If that user then chooses to participate in ledger B as well, the marginal gain in utility is just the utility obtained from interacting with users who chose B. That is, a user who chooses $A \cup B$ does not get utility from coordinating twice with other users who chose $A \cup B$; utility is gained only by interacting with a new set of users.

Formally, suppose that $\phi^\emptyset, \phi^A, \phi^B, \phi^{A\cup B}$ are the proportions of users who choose the strategies $\emptyset, A, B, A \cup B$, respectively. Then (neglecting cryptographic security) the utilities obtained by choosing $\emptyset, A, B$, and $A \cup B$ are given by

$$u^\emptyset = 0, \quad u^A = s_i^A + \kappa(\phi^A + \phi^{A\cup B}) - \alpha L^A$$

$$u^B = s_i^B + \kappa(\phi^B + \phi^{A\cup B}) - \alpha L^B, \quad u^{A\cup B} = s_i^A + s_i^B + \kappa(\phi^A + \phi^B + \phi^{A\cup B}) - \alpha(L^A + L^B)$$

If there is complete information about payoff parameters, there will be multiple equilibria

in this game. As such, we must introduce an arbitrarily small amount of noise in users' information structure in order to obtain unique predictions. We suppose users observe a signal ζ of the "fundamentals" of ledger B such that the expected cost of using B is $\alpha L^B + \zeta$. Again, users observe ζ with an arbitrarily small amount of noise. This information structure leads to the equilibrium described in Proposition D.2.

Proposition D.2. *For a given L^A, there is a cutoff value $\overline{L}^B(L^A)$ such that users all choose B whenever $L^B < \overline{L}^B(L^A)$, and no user chooses B otherwise. The cutoff value is*

$$\overline{L}^B(L^A) = Q(\alpha L^A)\left(\frac{1}{2}(1 - Q(\alpha L^A))\kappa + Q(\alpha L^A)(\alpha L^A - E[s|s < \alpha L^A])\right)$$

This proposition outlines two features of the equilibrium. First, in any equilibrium where some user uses B, *all* users use B. This is because users have no initial stakes in B, so their preferences for B are homogeneous. If one user finds that using B is worth the cost, then the same is true for all users. On the other hand, some users with large stakes in A will choose to continue to use A even after a large portion of the population has abandoned the ledger.

Second, users' stakes still anchor them to ledger A as in our previous results. As expected, the anchor on A is stronger when users' stakes in A are higher. The formula for the cutoff is also useful in a situation where users' stakes on both ledgers are the same (as in a blockchain). For example, if we simply take $s_i = 0$ for all i, in which case the formula simplifies to $\overline{L}^B(L^A) = L^A$, meaning that as in our benchmark model, the ledger more favorable to users is always selected. In order to obtain the result that competition among ledgers is restricted, however, we now must also assume that the proposer of ledger B cannot feasibly run a ledger with $L^B = 0$. That is, there must be some lower bound $\underline{L}^B$ on the fee charged by proposer B. As long as $\overline{L}^B(L^A) < \underline{L}^B$, ledger B will not emerge. There are several reasons why there may be a lower bound on L^B– there could be some economic cost to producing a viable ledger, or, as we discuss in the next section, incentive issues may require record-keepers to earn rents.

Here we prove the results. First we describe the equilibria of the complete information game.

Proposition D.3. *Suppose B is the entrant ledger in the sense that $s_i^B = 0$ for all i. There are two equilibria in the complete information game as long as $\min\{\alpha L^A, \alpha L^B\} < \min\{\kappa, \max_i s_i^A\}$ and $\kappa(1 - Q(\alpha L^A)) > \alpha L^A$: there is one in which all users choose only A, and there is one in which all users choose B and some users multi-home.*

Proof. First note that a strategy profile in which all users choose A is indeed an equilibrium:

we have

$$u^A = s_i^A + \kappa - \alpha L^A > \max\{u^\emptyset, u^B, u^{A\cup B}\}$$

for all i since $u^B = -\alpha L^B$ and $u^{A\cup B} = u^A - \alpha L^B$.

Now suppose that there is some equilibrium with $\phi^B > 0$. Then

$$u^B = \kappa(\phi^{A\cup B} + \phi^B) - \alpha L^B > 0$$

If some agents multi-home (i.e. $\phi^{A\cup B} > 0$), then it must be that

$$u^{A\cup B} = s_i^A + \kappa(\phi^A + \phi^{A\cup B} + \phi^B) - \alpha(L^A + L^B) > s_i^A + \kappa(\phi^A + \phi^{A\cup B}) - \alpha L^A = u^A$$

which implies

$$\kappa\phi^B - \alpha L^B > 0 \Rightarrow \kappa(\phi^{A\cup B} + \phi^B) - \alpha L^B > 0$$

Therefore, if some agents multi-home, then all agents must choose B. However, if no agents multi-home, then $\phi^{A\cup B} = 0$, so we still have that if some agents choose only B, then all agents must choose B.

Now we show that some agents must multi-home. Even if $\phi^A + \phi^{A\cup B} = 0$, we have

$$u^{A\cup B} - u^B = s_i^A - \alpha L^A > 0$$

for some i by assumption. Hence in any equilibrium where some agents choose only B, all agents must choose B or multi-home.

Another possibility is that some agents choose not to use a ledger. All agents with $s > \alpha L^A$ clearly must at least choose to use A. If it is the case that $\kappa(1 - Q(\alpha L^A)) > \alpha L^A$, then, no agent will choose not to use a ledger at all, since any agent with $s < \alpha L^A$ will still find it preferable to choose ledger A over choosing neither.

Finally, by the argument above it is evident that there can be no equilibrium in which all agents either choose A or multi-home. In this case, all agents would find it optimal to choose B, so all agents would multi-home, which cannot be an equilibrium. □

Proposition 1 shows that when ledger B is used by some, all users must use it. This may seem like a somewhat extreme result, but it can be modified by assuming that users have heterogeneous preferences for the properties of ledgers A and B. In such a case, there would be equilibria in which some users choose A, some users choose B, and some users multi-home, which is arguably closer to the empirically relevant case.

Note that there is a discontinuity at $L^A = L^B = 0$. In that case, the only trembling-hand perfect equilibrium is one in which all users multi-home. When the cost of using a ledger is

even $\epsilon > 0$, though, this ceases to be an equilibrium at all.

Since we want to obtain unique predictions about equilibrium play, we now turn to a setting with incomplete information. Suppose that users derive fundamental utility $-(\alpha L^B + \zeta)$ from using ledger B, where ζ is unknown. As in the benchmark model, users receive signals $x_i = \zeta + \sigma\epsilon_i$, where the distribution of ϵ_i has support on $[-\frac{1}{2}, \frac{1}{2}]$. We now consider the global games refinement.

Global games refinement: We consider a game in which users choose among their four possible strategies simultaneously. This game is more complicated because (1) we need to show it can be cast as a supermodular game, and (2) we need to determine what the cutoff strategies look like (since they are no longer a choice between just two options). We state without proof that the game is supermodular under the following parametrization: identify the set $\{0,1\}^2$ with users' choices by taking any element with first coordinate 1 to as a strategy where the user chooses A and any element with second coordinate 0 as a strategy where the user chooses B.

As the noise in the signals of ζ approaches zero, the equilibrium play of the complete information game is recovered. Therefore, there must be a cutoff $\overline{k}$ such that all users with $s_i > \alpha L^A$ choose only A when $x_i > \overline{k}$ and multi-home when $x_i < \overline{k}$, and users with $s_i < \alpha L^A$ choose only A when $x_i > \overline{k}$ and choose only B when $x_i < \overline{k}$.

We characterize the cutoff equilibrium and derive the fee L^B required for some users to choose ledger B.

Proof of Proposition D.2. At the cutoff, users with $s_i > \alpha L^A$ are indifferent between choosing only A and multi-homing. Hence

$$0 = \int_0^1 \Big((s + \kappa(1 - Q(\alpha L^A) + Q(\alpha L^A)(1-z)) - \alpha L^A) - (s + \kappa - \alpha(L^A + L^B) - \overline{k}) \Big) dA_s(z|\mathbf{k})$$

This yields

$$\alpha L^B + \overline{k} = \kappa Q(\alpha L^A) \int_0^1 z dA_s(z|\mathbf{k})$$

for $s > \alpha L^A$. This condition immediately yields $L^B \leq \frac{\kappa}{\alpha} Q(\alpha L^A)$. It also implies

$$(1 - Q(\alpha L^A))(\alpha L^B + \overline{k}) = \kappa Q(\alpha L^A) \int_{s > \alpha L^A} \int_0^1 z dA_s(z|\mathbf{k}) dQ(s) \tag{27}$$

By contrast, users with $s < \alpha L^A$ choose between using only A or only B at the cutoff.

Their indifference condition is

$$0 = \int_0^1 \Big((s + \kappa(1 - Q(\alpha L^A) + Q(\alpha L^A)(1 - z)) - \alpha L^A) - (\kappa z - \alpha L^B - \overline{k}) \Big) dA_s(z|\mathbf{k})$$

This condition implies

$$\alpha L^B + \overline{k} = \alpha L^A - s - \kappa \Big(1 - (1 + Q(\alpha L^A)) \int_0^1 z dA_s(z|\mathbf{k}) \Big)$$

for $s < \alpha L^A$. Integrating over $s < \alpha L^A$,

$$Q(\alpha L^A)(\alpha L^B + \overline{k}) = Q(\alpha L^A)(\alpha L^A - E[s|s < \alpha L^A] - \kappa) + \kappa(1 + Q(\alpha L^A)) \int_{s < \alpha L^A} \int_0^1 z dA_s(z|\mathbf{k}) dQ(s) \tag{28}$$

Combining Equations 27 and 28 and using the belief constraint $\int_s \int_0^1 z dA_s(z|\mathbf{k} dQ(s) = \frac{1}{2}$,

$$\alpha L^B + \overline{k} = Q(\alpha L^A) \left(\frac{1}{2}(1 - Q(\alpha L^A))\kappa + Q(\alpha L^A)(\alpha L^A - E[s|s < \alpha L^A]) \right)$$

so it any user uses B in equilibrium, L^B must be less than or equal to the right-hand side (assuming $\zeta = 0$). □

As expected, the anchor on A is usually stronger when users' stakes in A are higher (i.e., under a first-order shift upwards of $Q(s)$). The formula for the cutoff is also useful in a situation where users' stakes on both ledgers are the same (as in a blockchain). For example, if we simply take $s_i = 0$ for all i, the formula simplifies to

$$\overline{k} = \alpha(L^A - L^B) \tag{29}$$

Then when $\zeta \approx 0$ (as we assume throughout) we find that even in the presence of multi-homing, the ledger with the lower fee wins out.

D.3 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 of competition between two centralized entities is that the agent who proposes the fee structure for a blockchain does not care about the fees earned by

record-keepers because record-keepers always break even. Rather, the proposer's incentives are aligned with those of users. 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 record-keeper on ledger A, and $P^B = \mathcal{D}$ (for "developer") who is not a blockchain record-keeper. Proposers P^A and P^B choose parameters $L^A, L^B \geq 0$ at $\tau = 0$. The stakes on proposer P^A's ledger are uniformly distributed on $[S - \frac{d}{2}, S + \frac{d}{2}]$, and the stakes on P^B's ledger are all equal to zero. When a blockchain competes against a monopolist, there is still perfect competition between blockchain record-keepers, but the blockchain cannot replicate the information contained on the monopolist's ledger.

As in the baseline blockchain model, there is a continuum of record-keepers $j \in [0, M]$. However, there is no longer incomplete information. When users' 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 record-keepers will be the same as in the baseline model of a blockchain fork.[37] Furthermore, blockchain record-keepers cannot write on the monopolist's ledger, so they all must commit to ledger B at $\tau = 1$. The equilibrium at $\tau = 2$ 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 $\tau \geq 1$, so solving the model reduces to solving the proposers' optimization problems at $\tau = 0$. The monopolist behaves as if facing a fixed outside ledger with parameter L^B, so the optimal L^A is again given by 19. However, P^B has different preferences than an entrant monopolist. As in the baseline blockchain model, P^B's preferences are given by users' utility function $\tilde{u}$. If there is a value L^{B*} that is uniformly best for users (with utility function $\tilde{u}$) given optimal play by record-keepers, proposer P^B will choose it. 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

[37] Indeed, the $\tau = 2$ part of the proof of Proposition 5 is independent of the information structure so long as all record-keepers observe participation on the ledger.

computational power required to support a blockchain is large, the compensation necessary to attract record-keepers (and thus the minimum blockchain fee) will be higher, thereby dissuading users 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 21 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

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

A blockchain lowers costs for users when the computational expenditure required to placate users' need for cryptographic security is small, when the dispersion of users' 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 users. Hence when the incumbent already charges high fees, competition by a traditional intermediary should be enough to lower costs to users. Blockchain is useful primarily when entrants into the market have incentives to charge high fees. Free entry of blockchain record-keepers implies that there is no incentive for a proposer to choose a policy that gives record-keepers large fees because all record-keepers break even regardless. The feature of the blockchain that allows it to more effectively compete with traditional intermediaries is that it strips record-keepers of their market power.