Blockchain Disruption and Smart Contracts*

Lin William Cong† Zhiguo He§

September 26, 2017

Abstract

Distributed ledger technology embodied in blockchains feature decentralized consensus as well as low-cost, tamper-proof, and algorithmic executions, and consequently enlarge the contracting space through smart contracts. Meanwhile, the process of generating decentralized consensus which involves information distribution necessarily alters the informational environment. We analyze how decentralization improves consensus effectiveness, and how the quintessential features of blockchain reshape industrial organization and the landscape of competition. Smart contracts can mitigate information asymmetry and deliver higher social welfare and consumer surplus through enhanced entry and competition, yet blockchains may also encourage collusion due to the irreducible distribution of information. In general, blockchains can sustain market equilibria with a larger range of economic outcomes. We further discuss anti-trust policy implications targeted to blockchain applications such as separating consensus record-keepers from users.

Keywords: Competition, Distributed Ledger, FinTech, Anti-trust, Incomplete Contracts, Collusion, Information, Security Design, Enforcement.

*The authors thank Jingtao Zheng for excellent research assistance and discussions that helped shape an initial version of the paper. They are also grateful to Susan Athey, Tom Ding, Itay Goldstein, Brett Green, Campbell Harvey, Wei Jiang, Andrew Karolyi, Jiasun Li, Maureen O'Hara, Minyu Peng, David Yermack, and seminar and conference participants at the RFS FinTech workshop and CUHK Econ for very helpful discussions and comments.

†University of Chicago Booth School of Business. Email: will.cong@chicagobooth.edu.

§University of Chicago Booth School of Business; and NBER. Email: zhiguo.he@chicagobooth.edu.

1 Introduction

Blockchain, a distributed database and computing technology managed in a decentralized manner (often autonomously), first became well-known as the technology behind the cryptocurrency bitcoin in 2008. It has since emerged in various other forms, often allowing programs with logics stored and automated to trigger further recording and transactions on which the blockchain participants reach consensus. This has given rise to applications such as smart contracts – digital contracts with consensus terms and are self-enforcing and tamper-proof through automated execution. In the past few years, blockchain technology has been believed to potentially disrupt business and the financial services in a way similar to how the internet disrupted offline commerce.[1] Figure 1 displays Google searches showing the rising popularity of the blockchain technology in the past half decade, as well as the recent trend in new open-source projects that are related to blockchain and smart contract.

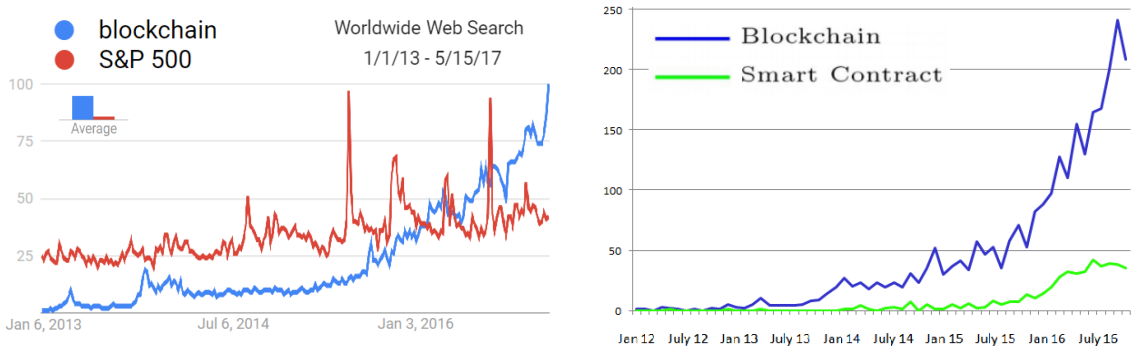


Figure 1: **Trends about Blockchain and Smart Contracts**. The left panel displays relative search interests and plots each relative to its peak (normalized to 100) for the given region and time. The right panel shows the number of blockchain and smart contract projects hosted on Github, a major open-source development platform for coding programs around the world.

In this paper, we argue that despite a plethora of definitions, descriptions, and applications of blockchain and decentralized ledger, the technology and its various incarnations share a core functionality in providing "decentralized consensus." From an economist's perspective, the backbone of the blockchain technology can be described as a robust consensus

[1]Bloomberg Markets featured in its Oct 2015 cover, "It's all about the Blockchain." The Economist ran a cover story around the same time "the Trust Machine" which argued that "the technology behind bitcoin could change how the economy works." The financial services industry rebranded the blockchain more generally as a form of distributed ledger technology. Marc Andreessen, the cocreator of Netscape, even exclaimed "This is the thing! This is the distributed trust network that the Internet always needed and never had."

generation scheme that leads agents with divergent perspectives and incentives to accept and act upon as if it is the "truth." For example, on the bitcoin blockchain, agents can check and verify transactions with one another, eliminating "double-spending" the digital currency and freeing everyone from the need of a centralized trustworthy arbitrator or third party.

Blockchains reach decentralized consensus via contacting decentralized recordkeepers. Two economic forces naturally arise in this setting. First, greater decentralized consensus makes contracting on contingencies easier; second, achieving such consensus requires distributing information to some degrees for verification. Hence blockchain applications typically feature a fundamental tension between *decentralized consensus* and *distributed information*. The former enhances contractibility and is welfare-improving, while the latter could be detrimental to the society.

We provide a simple framework to think about the process of reaching decentralized consensus on a blockchain. Most blockchains have a community of recordkeepers which typically overlap with users of the blockchain. Similar to third-party arbitrators in the real world, they receive (noisy) signal on the true state of the world and may have incentives to misreport (tamper or manipulate) or risks of being hacked or attacked. With the help of decentralized record-keepers, fast developing real-time communication technologies, and careful protocol designs such as proof-of-work or weighted aggregation, blockchains reduce individual's incentive to manipulate and misreport, allowing individual information to be aggregated more efficiently. Compared to the traditional world, blockchains produce consensus more cheaply and more accurately, which, together with automated execution, enables effective contractibility on contingencies that were difficult to contract on previously. However, effective consensus is predicated on decentralized record-keepers' observing and receiving greater amount of information. Many existing blockchains achieve this by making more information (potentially encrypted) available to the recordkeepers, or reaching out to large numbers of record-keepers. This information distribution process changes the informational environment, and hence the economic behaviors of the participants on the blockchain.

Armed with this insight, we then analyze the impact of blockchain technology on competition and industrial organization. Specifically, there are two incumbent sellers known to be authentic, and an entrant who only has some probability of being authentic (otherwise she is fraudulent). Authentic sellers will deliver the goods, while the fraudulent ones cannot. In each period, buyers as a group show up with a constant probability, and shop the sellers based on price quotes, then exit the economy. Each seller only observes whether customers

come to her; she observes neither the price quotes other sellers offer, nor whether other sellers get customers. This captures in the spirit of Green and Porter (1984) imperfect monitoring in repeated games, but in a setting in which sellers compete on price rather than quantity.

In the traditional world without blockchain and smart contract, due to contract incompleteness sellers cannot offer prices contingent on the success of delivering the goods. Entry does not occur due to the lemons problem, and two incumbents might engage in collusion in equilibium. Because incumbent sellers cannot differentiate the event of no buyers showing up from the event of the other seller stealing her market share, aggressive price wars which are the perfect competitive strategy occur too often, making it relatively hard to sustain collusion among incumbent sellers.

In contrast, blockchains, via decentralized consensus, enable agents to contract on service outcomes and automate contingent transfers. Hence the authentic entrant is now able to signal her authenticity fully and enter the market. This decimates information asymmetry as an entry barrier, leading to enhanced competition that improves welfare and consumer surplus. We further show in an extension that the mitigation of information asymmetry and improvement in consumer surplus and welfare hold even when sellers all have private type of service quality.

However, as mentioned before, generating decentralized consensus also inevitably leads to greater observability of aggregate service activity recorded on the blockchain, and we show this increased observability can foster tacit collusion among sellers. Different from traditional world and the environment in Green and Porter (1984), now sellers—by serving as recordkeepers as well—effectively observe the aggregate service activities on the blockchain and hence are able to detect deviations perfectly in any collusive equilibrium. Consistent with this intuition, we show that with only permissioned blockchain among the incumbents, there is always weakly more collusion equilibria than those that are sustainable in the traditional world.

Our model thus features the trade-off between enhanced competition and aggravated collusion, both arising from the blockchain technology. Under fairly weak conditions, with blockchain and smart contracts the set of possible dynamic equilibria is strictly bigger than that in a traditional world, leading to social welfare and consumer surplus that could be higher or lower than before.

In practice, there is indeed a wide-spread concern that blockchains may jeopardize market competitiveness in a serious way; this becomes especially relevant for permissioned

blockchains with powerful financial institutions as exclusive members.[2] Our paper highlights one particular economic mechanism through which blockchain facilitates collusion, and we explore regulatory implications of the model that policies that aim to improve consumer surplus. For instance, an oft-neglected regulatory solution is to separate usage and consensus generation on blockchains, so that sellers cannot use the consensus-generating information for the purpose of sustaining collusion.

By providing a conceptual description of blockchain and smart contracts from an economic and financial perspective, our analysis aims to demonstrate that blockchains are not merely database technology that reduces the cost of storing or sharing data, but have profound economic implications on consensus generation, industrial organization, smart contract design, and anti-trust policy.

Related Literature Our paper adds to the emerging literature on blockchains, which thus far has mainly come from computer scientists. Among studies on the applications and economic impact of the technology, Harvey (2016) briefly surveys the mechanics and applications of crypto-finance, especially Bitcoin.[3] Yermack (2017) evaluates the potential impacts of the blockchain technology on corporate governance: for managers, institutional investors, small shareholders, auditors, etc. Raskin and Yermack (2016) push further to envision that the central banks might use the technology to launch their own digital currencies. Complementary to our discussion on smart contracts, Bartoletti and Pompianu (2017) empirically document how smart contracts are interpreted and programmed on various blockchain platforms. We add by examining arguably the most defining features of blockchain, and how they interact with information asymmetry and affect market competition — both are important, general issues in economics.

Related are studies on the underlying mechanism for generating decentralized consensus. Kroll, Davey, and Felten (2013) note that Bitcoin miners are playing a strategic game and the "Longest Chain Rule" should be a Nash equilibrium. Biais, Bisiere, Bouvard, and Casamatta (2017) formalize this notion but show deviations from the rule can arise on the equilibrium path. Eyal and Sirer (2014) study "selfish mining" in bitcoin blockchain

[2]See, for example, "Exposing the 'If we call it a blockchain, perhaps it wont be deemed a cartel?' tactic," by Izabella Kaminska, Financial Times, May 11th, 2015.

[3]Other papers on more specialized applications include Catalini and Gans (2016) point out the blockchain technology can reduce the cost of verification and the cost of networking. Malinova and Park (2016) study the design of the mapping between identifiers and end-investors and the degree of transparency of holdings in a blockchain-based marketplace. Khapko and Zoican (2017) argue that blockchain allows for flexible settlement of trades, and the optimal time-to-settle trades off search costs and counter-party risk, creating vertical differentiation.

in which miners launch block-withholding attacks. Nayak, Kumar, Miller, and Shi (2016) discuss "stubborn mining" that generalize and outperform "selfish mining" attacks. For challenges facing the incentive and governance issues of maintaining decentralized ledgers, see Evans (2014).[4] Instead of focusing on strategic behaviors of the recordkeepers under specific blockchain protocols such as that of Bitcoin miners, we theorize the generation of decentralized consensus—a feature universal to all blockchains, and then examine its implications for contracting and competition.

Our analysis on collusion adds to the large literature on industrial organization and repeated games with monitoring (see, e.g., Tirole (1988)). Our model ingredients are similar to Porter (1983) and Green and Porter (1984), who study collusion in Cournot setting with imperfect public monitoring. Abreu, Pearce, and Stacchetti (1986) and Abreu (1987) generalize the results further to consider additional types of equilibria. Our analysis of sustainable equilibria is related to Fudenberg and Maskin (1986). Our paper examines Bertrand competition, and links the additional observable or contractible information to the type of monitoring in repeated games under the technological innovation.

Our discussion on the application of blockchain and smart contract in financial services and transactions is broadly linked to the literature on financial technology and innovation. Philippon (2015) points out intermediation cost has been 1.5-2 percent of intermediated assets for a long time, despite dramatic technology improvement. We provide a cautionary tale that blockchain technology, while holding great potential in mitigating information asymmetry and encouraging entry, can also lead to greater collusive behavior.

The rest of the paper is organized as follows: Section 2 provides institutional details on blockchains, smart contracts, and their applications; Section 3 develops a simple framework for understanding the key economic trade-offs surrounding decentralized consensus and information distribution; Section 4 takes the core functionality of blockchain as given to analyze dynamic industrial equilibria in both traditional and blockchain worlds, in order to demonstrate that blockchain technology facilitates entry and cartelism; Section 5 discusses policy implications and extensions on private types, imperfect consensus, and smart contract forms; Section 6 concludes.

[4]Along that line, several studies examine the organization and compensation of miners. Kiayias, Koutsoupias, Kyropoulou, and Tselekounis (2016) show that when the computational power of a miner is large, Nash equilibria unanticipated by the bitcoin designer arise. This fundamental tension between concentration of computation power and system stability is also studied in Baldimtsi, Cong, He, and Li (2017), who theoretically and empirically analyze mining pool formation and compensation contracts in an attempt to shed new lights on the theory of firm.

2 Blockchain, Smart Contracts, and Applications

In this section, we first provide an overview of the blockchain technology, highlighting decentralized consensus as its backbone. We then formally define smart contracts, before discussing various real world business applications of blockchains and smart contracts in the financial industry.

2.1 Blockchain as Decentralized Consensus

The work on a cryptographically secured chain of blocks dates back to 1991 by Stuart Haber and W. Scott Stornetta, but it was only until 2008 that the first blockchain was conceptualized by Satoshi Nakamoto, and was implemented and popularized through the cryptocurrency **Bitcoin** (Nakamoto (2008)).[5] Its simplest form entails a distributed database that autonomously maintains a continuously growing list of public records in unit of "blocks", secured from tampering and revision. Each block contains a timestamp and a link to a previous block. Other forms of blockchains emerged subsequently with different designs on exclusivity, transparency, and maintenance of the records.

All blockchains–with varying degrees–aim at creating a database system in which parties can jointly maintain and edit in a decentralized manner, with no individual exercising central control. One defining feature of blockchain architectures is thus their ability to maintain, in a relatively cheap way, a uniform view on the state of things and the order of events – a consensus, a type of information that leads individuals or organizations with divergent perspectives and incentives to accept and act upon as if it is truth.

The generation of consensus is essential to many economic and social functions. Its benefits and empowerment for everyone sharing and trusting the same ledger–the so-called decentralized ledger–are clear. Settlements no longer take days, lemon's markets and frauds disappear, and the list goes on. Traditionally, court, government, notary agencies, etc., provide such consensus, but in a way that is labor-intensive, costly, tamper-prone, local, and centralized.

[5] Böhme, Christin, Edelman, and Moore (2015) surveys Bitcoin's design principles and properties, risks, and regulation. Narayanan, Bonneau, Felten, Miller, and Goldfeder (2016) is an in-depth introduction for the technical details of Bitcoin blockchain.

Decentralized Consensus

Blockchain aims to produce ***decentralized consensus***, a specific state or set of information to be agreed upon by *all agents* via rules and protocols, without the need to trust or rely upon a centralized authority. This supposedly makes the consensus more secure and tamper-proof. Moreover, it rewards a community for properly maintaining the consensus, allowing greater recording and processing power in an incentive-compatible and typically competitive manner. Although maintaining decentralized consensus is also costly, as seen in the electricity consumption in Bitcoin mining, it is reducible through alternative designs, and is more cost-effective relative to the traditional world.

In neither aspect is the technology in its current forms perfect, but it has improved quickly and substantially enough that the challenges are not insurmountable, and the basic functionality that blockchain provides is clear. While there have been several hacking incidents on blockchains (the most notable one being the hack on **Decentralized Autonomous Organization** (DAO) on **Ethereum** blockchain) and Bitcoin mining seems to waste electricity, these can be addressed by improving the protocol design.[6] Storing wealth in centralized database is even riskier, and going to court or arbitration to reach consensus is likely to cost a magnitude higher for parties involved in a dispute on payment. Rather than analyzing the technical details and refinements of various protocols, this paper focuses on the economic implications of decentralized consensus, and the natural process that accompanies it — information distribution.

Information Distribution

The economic tradeoffs involving information distribution in generating decentralized consensus is interesting both from a practical perspective and for fundamental economic understanding. On Bitcoin, the consensus is reached and maintained through distributing all transaction information (with public-key-encrypted owner addresses) to the entire population on the blockchain, so all transaction details (except for identities) recorded on the consensus are public information. One obvious issue that arises when pushing for real-world blockchain applications is business privacy. For instance, financial institutions are typically sensitive to reveal the details the transaction to other unrelated parties; and traders may want to hide their identities to prevent front-running (Malinova and Park (2016)). At the aggregate level,

[6] **Lightning**, which builds on the Bitcoin blockchain, reduces the amount of information that has to be recorded on the blockchain to increase processing power. **Phi** from **String Lab** builds on Ethereum to ensure higher security and execution speed.

there could also be unintended consequences that greater information distribution brings about; we analyze one example in the industrial organization context of Green and Porter (1984) in Section 4.

Facing this fundamental trade-off, there are many proposals on better encryption which effectively masks sensitive information in the process of consensus generation. Other straightforward compromise is to reach decentralized consensus only on a subset of important states of world, or requesting verification from fewer nodes (recordkeepers) in the blockchain network.[7] In what ways does information distribution matter beyond privacy concerns? Will these reduce the effectiveness of blockchain consensus? Extant theory tells us very little.

Before answering these questions more formally in Section 3, we discuss how decentralized consensus enables the use of smart contracts, and various real-world blockchain applications.

2.2 Smart Contracts

In recent years, the development of blockchain technology has allowed customizable programming logic to be stored in a decentralized way. This course of development has revived the notion and facilitated the creation of smart contracts, originally envisioned by Szabo in 1994 (e.g., Tapscott and Tapscott (2016)):

"A smart contract is a computerized transaction protocol that executes the terms of a contract. The general objectives are to satisfy common contractual conditions (such as payment terms, liens, confidentiality, and even enforcement), minimize exceptions both malicious and accidental, and minimize the need for trusted intermediaries. Related economic goals include lowering fraud loss, arbitrations and enforcement costs, and other transaction costs."

While a consensus definition (no pun intended) for smart contracts has yet to be reached, their core functionality is clear — contracting on contingencies on a decentralized consensus, and on low-cost, algorithmic execution. To achieve decentralized consensus, a distributed ledger is needed, which also has to be self-executing. Contingencies (including allocation of property and control rights) in a smart contract should be codified, so that automated execution is feasible, which reduces enforcement cost.

Given the aforementioned facts, we provide a functional definition of smart contracts:

[7]For example, Aune, O'Hara, and Slama (2017) discuss the use of first-stage hashing to secure time-priority without revealing detailed information and revealing information later, in order to prevent front-running a transaction before it is recorded on a block on distributed ledgers. Directly related is the so-called "Zero Knowledge Proof" in computer science; in layman's language, participants can agree on certain facts without revealing useful information.

Smart contracts *are digital contracts allowing terms contingent on decentralized consensus and are self-enforcing and tamper-proof through automated execution.*

Our definition is consistent with and nests the definitions commonly seen in the legal circle (Lauslahti, Mattila, Seppälä, et al. (2016)), and in Szabo (1998) and Szabo (1997). It is important to note that smart contracts are not merely digital contracts (many of which rely on trusted authority for reaching consensus and execution), nor are they entailing artificial intelligence (they are rather robotic, on the contrary).[8]

Without decentralized consensus, the party holding the centralized consensus often enjoys huge bargaining power and monopoly rent (for example, a third party with data monopoly). And, traditional resolutions by third parties such as the court or an arbitrator do not fit well as they involve high degrees of human intervention that are less algorithmic, and could be costly because the resolutions are typically less deterministic and thus is costly to risk-averse agents. By enabling the use of smart contracts, the blockchain technology increases contractibility and enforceability in contingent contracts that facilitate exchanging money, property, shares, service, or anything of value in an algorithmically automated and conflict-free way.

That said, the added contractibility comes at the expense of greater information distribution as discussed earlier, and the overall impact on the economy is far from obvious. It is helpful to understand the impact better by first seeing some real-world applications.

2.3 Blockchain Applications in the Financial Industry

The applications of blockchain technology and smart contracts are broad, sometimes even beyond the Fintech industry.[9] Yet because of the tamper-proof nature, and the ease to automate rule-based monetary transfers, smart contracts are especially appealing for financial services and trading. Therefore we describe below existing uses of of blockchain and smart contracts in two financial industries, which are also the predominant applications

[8]Smart contracts do not even have to be implemented through blockchain; any decentralized consensus system would suffice. An interesting example is **Metromile**, a pay-per-mile car insurance company that aims to achieve consensus on miles driven using a device that synchronize data on driving. Another example is **Waze**, a driving app that aggregates real-time traffic information. What is missing in these apps is the tamper-proofness and security that blockchain technology brings about.

[9]Bartoletti and Pompianu (2017) analyze 834 smart contracts from Bitcoin and Ethereum with 1,673271 transactions. They find five main categories of uses (financial, notary, games, wallet, and library), three of which are related to monetary transfers and transactions, with the remaining two related to recording consensus information. More than two-thirds of the uses are on managing, gathering, or distributing money.

of the technology; importantly, they are not merely proofs of concept.

Trusted Payments

Suppose Alice in Chicago wants to make a payment to Bob in Africa but does not want Bob to cash it out immediately. She may have to go to a large bank and worry about issues such as whether Bob is the real Bob and the bank information he provides is accurate. The Society for Worldwide Interbank Financial Telecommunications (SWIFT) can mitigate the problem, but would require Bob's bank to be in the society, and may take a long time to verify the bank has not turned fraudulent or bankrupt, and the bank has to do the same verifications on Bob's account, before authorizing the transfer. If Alice is charged an unreasonably high exchange rate, or the transfer did not go through, resolving the issue may incur further cost and delay. This concern becomes even more severe with digital payments, where "double-spending" (spending the digital currency or using the credit more than once) issues abound.

Bitcoins as a form of cryptocurrency were first invented to offer a potential solution to the "double-spending" problem (Nakamoto (2008)). It enables peer-to-peer transactions recorded on the Bitcoin blockchain that is secure and time-stamped to make it tamper-proof, and by broadcasting all candidate transactions publicly and having "miners" constantly competing for the recording right and validating new blocks, its distributed ledger provides an almost real-time decentralized consensus on whether a transaction has taken place. If both Alice and Bob use bitcoins, they can make the transfer directly.

By design, maintaining the decentralized consensus record on Bitcoin blockchain requires the miners to solve difficult NP-complete computational problems (i.e., mining, a form of "proof-of-work") which is costly and limited in capacity, making it unsuitable for large volumes of financial transactions. Subsequent platforms such as **Lightning** (built on the Bitcoin blockchain) and **Stellar** (separate blockchain) help improve the processing capacity through local channels and multisignature accounts so that unnecessary information does not have to be part of the decentralized consensus. **Counterparty** also builds on the Bitcoin blockchain, but allows for more flexible smart contracts and maintains consensus through "proof-of-burn," i.e., fees paid by clients are destroyed, and nodes are rewarded for validation from the inflation of the currency.

That said, these blockchains' scripting language is limited (basic arithmetic, logical, and crypto operations such as hashing, verification of digital signatures), and only a small fraction of mining nodes can process more complex script by signature verification, making them

less useful for general smart contracts. **Ethereum** – then second largest blockchain platform by market capitalization after the Bitcoin blockchain – allows the use of Turing-complete language and permits more complex contingent operations (Turing (1937)), providing the archetypal implementation of smart contracts (Buterin (2014)). All and only the valid updates to the contract states are recorded and ensured with automated execution. A group of voluntary participants (Ether miners) maintain a decentralized consensus recording of the states, and other interacting parties utilize the consensus information to automate executions of contract terms. Additional applications such as **Monax** and **Phi (String Lab)** build on Ethereum to enrich and optimize its smart contract functionalities and processing power, just like websites build on the Internet protocol.

Traditional players in the financial industry are also actively adopting the blockchain technology to address the payment problem. Originally known as Ripple Labs, **Ripple** was founded in 2012 to provide global financial transactions and real-time cross-border payments. It has since been increasingly adopted by major banks and payment networks as the settlement infrastructure technology. A (typically large) set of validating nodes achieve decentralized consensus using the **Ripple** transaction protocol RTXP—an iterative consensus process as an alternative to Proof-of-Work, in which transactions are broadcasted repeatedly across the network of validating nodes until an agreement is reached. Digital transfers are then automated through connecting electronically to bank accounts or using its native cryptocurrency Ripples (XRP).

Now a system like **Ripple**, equipped with decentralized consensus (and almost real-time because it takes a few seconds per transaction), alleviates the concerns about Bob or Bob's bank's authenticity and functionality, and the use of contingent transfers ensures that if Bob violates the agreement, Alice's fund is reverted back.

Trade Finance

Another related and important application is trade finance, which accounts for more than USD 10 Trillion annually according to a WTO report in 2015. Despite technological advent in many areas of financial services, trade finance remains a largely paper-based, manual process, involving multiple participants in various jurisdictions around the world, and prone to human error and delays along the supply chain. An importer may fail to strike a deal because she cannot obtain a letter of credit from his/her bank, or the importer's bank offering the letter of credit is not as well-known in the exporter's country. An exporter may fail to get advanced financing because the exporter's bank worries about whether the goods can be

successfully and timely delivered, whether the quality of goods meets the requirement, and whether payments from the importer can be secured.

It is clear that the blochchain technology can help alleviate (and to a large extent, resolve) the aforementioned frictions in trade finance. A decentralized ledger can better track goods during the process in which goods are shipped and delivered; by giving all parties equal access to the transaction record, it also facilitates faster verification and authentication, thereby reducing shipment and financial transaction times and uncertainties. Moreover, with the help of the so-called "oracles," i.e., feeders of information from the offline world, smart contracts drastically reduce costs in executing transfers contingent on real outcomes such as a shipment's having arrived at an intermediate port.

In 2016, **Barclays** and Fintech start-up **Wave** claim to have become the first organization to complete a global trade transaction using distributed ledger/blockchain technology. The letter of credit (LC) transaction between Ornua (formerly the Irish Dairy Board) and Seychelles Trading Company is the first to have trade documentation handled on the new Wave platform. Software giant **IBM** has also been spearheading the application of blockchain and smart contracts to trade finance, launching solutions for Indian Mahindra Group in December 2016, and in partnership with Danish shipping behemoth Maersk. In early 2017, IBM has ventured further by rolling out the **Yijian Blockchain Technology Application System** for the Chinese pharmaceutical sector. It has also collaborated with a group companies to develop a blockchain-based crude oil trade finance platform.[10] Other blockchain-based platforms to support lending, issuing letters of credit, export credit and insurance include **HK Blockchain** for trade finance, **TradeSafe**, and **Digital Trade Chain (DTC)**.

3 Decentralized Consensus and Information Distribution

In all the aforementioned applications, smart contracts have augmented contractibility and enforceability on certain contingencies, be it the lock-in requirement for Bob's fund withdrawal, or the automated payment upon an importer's successfully receiving the goods. Moreover, transaction and contingency information on the ledger, at whatever form, is distributed and observable to many agents on the blockchains. For example, though the payment identity is confidential on **Ripple** and hence it is difficult for anyone to associate transaction information with any specific user or corporation, transaction information is

[10] http://www.coindesk.com/ibm-blockchain-platform-oil-trade-finance/

public.[11] Trade finance blockchains typically require information from shipment ports to be distributed to get the whereabouts of shipments.

In this section, we model the mechanism of consensus generation, highlighting the role of recordkeepers and the inevitable nature of information distribution. In doing this, we bridge the gap between the research centered on the keepers' strategic behaviors (e.g., Eyal and Sirer (2014)) and that on the economic functionality of blockchains (e.g., Yermack (2017)).

3.1 Keepers and Verification

In all the applications, we observe that to reach decentralized consensus, blockchains contact a set of recordkeepers, or simply ***keepers***—who are typically dispersed blockchain participants and hence the name "decentralized"—for verification. Cryptocurrency mining as a way to maintain consensus record is a prominent feature for the likes of **Bitcoin** and **Ethereum**. **Ripple** and **R3 CEV** use their own consensus process but also relies on a community of keepers. The process typically entail both competition and assignment for recordkeeping as well as post-block validations, and is an interesting industry on its own. Table B1 in the Appendix contains a list of further examples of keepers, including finer descriptions of their roles.

These keepers observe some public information regarding the contingency to be recorded on the blockchain. The public knowledge includes both the information broadcasted on the blockchain and some information fed from the off-chain world through "oracles" — decentralized eternal actors or applications that feed data onto the blockchain. For **Bitcoin**, miners observe all transactions (before they are put into a block) on the blockchain without off-chain oracles. Many trade finance blockchains mentioned above use information from local ships, ports, banks, and border customs to track delivery status, though transaction details may not be fully public (e.g., **Corda** or some **Hyperlydger** blockchains). They may also receive some extra information about the transaction upon contact. For example, **IBM** currently works on trade-finance blockchains that provide keepers additional information about shipment status because to generate consensus record on whether the goods has been delivered requires off-chain collaborations and cross-validations with shipping companies and import-export controls.

These observations point to the fact that information distribution is often necessary in order to obtain valid consensus record. For instance, **R3's Corda**'s validating model (as

[11] For details, see https://ripple.com/insights/what-an-open-network-means-for-banks-market-makers-and-regulators/.

opposed to its non-validating model which does not provide validation consensus) requires distributing private information to the notaries, in order to prevent Denial-of-Service (DoS) attacks.[12]

The concern about information distribution and privacy is voiced by practitioners, among which **R3 CEV** – an active blockchain consortium – has been outspoken. R3's **Corda** system sets out to tackle the challenge that the only parties who should have access to the details of a financial transaction are those parties themselves and others with a legitimate need to know.[13] Even with that, the request (thus a form of information) for proving transaction uniqueness is distributed to some independent observers, changing the information environment of this economy at least partially.

While these measures potentially ensure confidentiality, two important economic insights are missing from current discussions. First, contacting less recordkeepers may reduce the effectiveness of the consensus; second, no news is news – even encrypted data are still data, as the mere act of verification request still informs recordkeepers something about the state of the world. As we demonstrate shortly, contacting less record-keepers does not solve the problem either without compromising the quality of consensus.[14] In other words, there is an irreducible lower bound on information distribution in order to achieve consensus.

Our model below focuses on the above mentioned tension: creating decentralized consensus leads to greater contractibility but necessitates greater information distribution.

3.2 A Simple Model of Decentralized Consensus

Suppose a smart contract references a contingent outcome $\widetilde{\omega}$ such as a bitcoin transfer, an arrival of a shipment, or a completion of international bank transfer.In general, $\tilde{\omega}$ could involve a sequence of events over a time period; but for simplicity we just focus on one contingency, which we refer to as "successful delivery" of service or goods, in the context of our main model in Section 3.

Denote the decentralized consensus on $\widetilde{\omega}$ on a blockchain by $\widetilde{z}$. To achieve this, suppose that the blockchain protocol contacts a set of μ potential keepers. Keepers in the set

[12]A node knowingly builds an invalid transaction consuming some set of existing states and sends it to the notary, causing the states to be marked as consumed.

[13]Similar to many other blockchains, a transaction on **Corda** requires both validity and uniqueness. What is different is that consensus over the transaction validity is performed only by parties to the transaction in question. So we do not have consensus at the ledger level. Consensus over uniqueness is customizable and involve independent observers, typically random and pluggable.

[14]For more details, please see https://docs.corda.net/key-concepts-notaries.html and https://docs.corda.net/key-concepts-consensus.html.

$\mathbb{K} \equiv \{1, 2, \cdots, K\}$ are homogeneous, and for simplicity we model the **effectiveness** of the consensus for contracting (and potentially other purposes) by $-Var\left(\widetilde{\omega} - \widetilde{z}\right)$.[15] An effective consensus is the cornerstone for the trust that many Fintech firms so extensively purport.

Upon contact, each keeper $k \in \mathbb{K}$ submits $\widetilde{y}_k$, yielding a collection of reports $\mathbf{y} \equiv \{\widetilde{y}_k\}_{k\in\mathbb{K}}$. Depending on the specific blockchain protocol, the consensus $\widetilde{z}(\mathbf{y})$ is then a transformation of inputs collected from these contacted keepers. For ease of illustration, we assume

$$\widetilde{z}(\mathbf{y}) = \frac{1}{K}\sum_{k} \widetilde{y}_k, \tag{1}$$

i.e.,the decentralized consensus is a simple average of all selected reports.[16]. We focus on how the metric of decentralization μ affects the quality of decentralized consensus and the system-wide information distribution.

3.3 Information Set of Recordkeepers

Consistent with the discussion at the beginning of this section, we assume that each keeper on the blockchain has a private signal $\widetilde{x}_i = \widetilde{\omega} + \widetilde{\eta}_i$, where for simplicity $\widetilde{\eta}_i$ are i.i.d. with zero mean and variance σ_η^2. $\widetilde{\eta}_i$ captures noisy observations of the true state based on public information and off-chain information available on blockchain, as well as additional information keepers have when generating consensus. As discussed, For **Bitcoin**, $\widetilde{\eta}_i = 0$ because the transaction information is publicly broadcast, rendering it noiseless in the verification on the existence and validity of blocks and proof-of-work, and the consistency of timestamps. For trade-finance applications, keepers such as shipping companies and import-export controls observe the shipment status for the trade.

Denote by $\mathbf{1}_k$ the event of keeper k being contacted, upon which his/her signal turns to $\widetilde{x}_k = \widetilde{\omega} + \widetilde{\eta}_k$, where $\widetilde{\eta}_k$'s are with zero mean and variance σ_K^2. We have $\sigma_K \leq \sigma_\eta$, thanks to the additional (potentially encrypted) information. To sum it up, the entire information distrubution on the blockchain can be written as a tuple of $\left\{R, \{\widetilde{x}_i\}_{i\notin\mathbb{K}}, \{\widetilde{x}_k, \mathbf{1}_k\}_{k\in\mathbb{K}}, \widetilde{z}\right\}$.

[15] In reality, the effectiveness depends on the purpose and use of consensus on each specific blockchain. Our specification qualitatively captures the universal feature that a consensus is not effective even when it is unbiased if it is uncertain to always reflect truth accurately. Our results are robust to introducing penalty terms for bias as well as high moments of $-Var\left(\widetilde{\omega} - \widetilde{z}\right)$.

[16] It is easy to show that our results are robust to heterogeneous and stochastic weights on signals, say $z(\vec{y}) = \frac{1}{R}\sum_r \widetilde{w}_r\widetilde{y}_r$ with $\sum_r \widetilde{w}_r = 1$. This specification includes certain well-known blockchains such as Bitcoin, in which the miner who solves a hard NP complete problem first (which is completely random if miners have homogeneous computation power) gets to make the record block. In the language of our model, the blockchain protocol randomly chooses one report from all contacted keepers (all miners).

We highlight that, although in many applications the extra information about the transaction is partially or fully encrypted, the mere fact being contacted reveals certain information; at least, the contacted keeper learns that a transaction has occurred. Our later analysis of blockchain on industrial organization hinges on this minimum information distribution necessitated uniquely by the formation of decentralized consensus.

3.4 Misreporting and Manipulation

Record-keepers may have incentives to misreport. For example, in trade-finance applications if record-keepers are also parties involved in the transaction; bitcoin miners may hide report through "selfish-mining", or double-spend certain coins, or get contaminated by hackers (in which case the distortion incentive comes from the hackers). Such incentives also exist in traditional economies. Business arbitraters may favor a client; double-spending was the issue in traditional online payments that originally inspired the creation of Bitcoin. In fact, media reports and practitioners' discussions have largely centered on how blockchain helps reduce tampering, manipulation, and hacking.

We assume that each risk-neutral keeper who submits a report of y_k derives a utility of

$$U(y_k; \mathbf{y}) = \widetilde{b}_k \cdot (\widetilde{z}(\mathbf{y}) - \widetilde{x}_k) - \frac{1}{2h} (y_k - \widetilde{x}_k)^2 . \tag{2}$$

The first coefficient $\widetilde{b}_k \equiv \widetilde{b} + \widetilde{\varepsilon}_k$ is keeper k's bias in misreporting, which is known to the keeper k before submitting his/her report. Here, the common bias $\widetilde{b}$ (among contacted keepers) is with zero mean and variance σ_b^2, capturing the common bias on the blockchain, which can be interpreted as one institutional transaction party choosing validators within its proprietary network (peer selection on **Ripple** and notary choice on **Corda**), an attempt by holders of the crypto-currency to slow down the creation of inflation of the native currency, and/or a system-wide hacking motive.[17] The idiosyncratic part $\widetilde{\varepsilon}_k$ is i.i.d., zero mean, and with variance σ_ε^2.

The second term captures the private cost of manipulation, where h parametrizes how quickly the cost rises with the magnitude of manipulation.[18] For example, for a bitcoin block

[17] Such common bias is not alien in the traditional economy – arbitrators in business arbitration are only rewarded if they are chosen by their clients and may systematically cater to major clients.

[18] In Equation (2), we could introduce a fixed cost of manipulation, as hacking or falsifying record could incur fixed cost on each record-keeper. The only effect that has is that exactly no manipulation would occur if K is large enough, but does not have to be infinity. We could include another system-wide value from the quality of consensus, $\kappa(K)$, as a function of K. Without much specifics of the details of the blockchain protocols, not much can be said about the net benefit of consensus provision κ. For Bitcoin and Ethereum,

that differs drastically from other miners' record, it takes longer to be confirmed and has higher probability to be reversed, which are costly to the miner recording that block.

3.5 Information Distribution and Quality of Consensus

Each individual contacted keeper chooses y_r to optimize U in (2), which gives

$$\widetilde{y}_k^* = \widetilde{\chi} + \widetilde{\eta}_k + \frac{h}{K}\widetilde{b}_k. \tag{3}$$

The equilibrium consensus then is (recall $\widetilde{b}_k = \widetilde{b} + \widetilde{\varepsilon}_k$)

$$\widetilde{z} = \frac{1}{K}\sum_k \widetilde{y}_k = \widetilde{\omega} + \frac{1}{K}\sum_k \widetilde{\eta}_k + \frac{h}{K}\left(\widetilde{b} + \frac{1}{K}\sum_k \widetilde{\varepsilon}_k\right), \tag{4}$$

with the resulting quality of the decentralized consensus:

$$-Var(\widetilde{\omega} - \widetilde{z}) = -\left[\underbrace{\frac{\sigma_K^2}{K}}_{\text{signal quality}} + \underbrace{\frac{h^2}{K^2}\left[\sigma_b^2 + \frac{\sigma_\varepsilon^2}{K}\right]}_{\text{manipulation}}\right]. \tag{5}$$

Public information disclosure policy on any blockchain will likely affect the keeper's signal quality σ_K^2, thereby affecting the quality of decentralized consensus in (5). But the unique benefit blockchain and decentralization bring derives from how the size of contact pool K improves the quality of consensus, which we focus on highlighting.

The first term is related to signal quality per se. For instance, contacting for verification, via sharing some details of the transaction information, may reduce σ_K and hence is quality-improving. Another evident channel in the first term in (5) is that the average over a greater sample size K smooths out the observation noises $\widetilde{\eta}_k$'s, and hence a better consensus.

The second channel is more novel and is rooted in the process of decentralized consensus generation. When the blockchain contacts more and more keepers, i.e., a greater K, each understands that each individual has less influence on the final consensus outcome. The

κ is the mining reward (coin plus fee) lest the cost (electricity). More generally, motivated by the fact that for typical blockchains, competition among record keepers drives the net profit from consensus generation to a lower value. But this term $\kappa(K)$ is less of our focus since our paper is not about optimal design of the blockchain protocol. Even though Bitcoin mining consumes a large amount of electricity, it does not have to be a social waste and can be channeled for scientific computing, as startup firms such as BOINC has demonstrated.

resulting reduced manipulation in report $\widetilde{y}_r^*$ in (3) translates to a higher consensus effectiveness. This effect is reflected in the scaling of $1/K^2$ in the terms in "manipulation" in (5). This is the key economic reason why blockchain is deemed more secure, in addition to its technical improvements on cyber-security. Of course, aggregation certainly helps reach a better consensus by reducing the idiosyncratic components of misreporting, as reflected in the denominator of the second term in "manipulation" in (5).

Equation (5) shows that soliciting more reports improves the quality of decentralized consensus; in particular, consensus becomes perfect, i.e., i.e. $\widetilde{z} \approx \widetilde{\omega}$, as $K \to \infty$. We focus on this case later in the context of an industrial organization framework analyzed in Section 3, and show how decentralized consensus improves the contractibility and enhances entry (hence competition).

However, contacting more keepers affects the information environment in which the agents reside on the blockchain. First, depending on detailed blockchain protocols, soliciting reports involves transferring certain transaction information to contacted record-keepers (and σ_K changes); recall the example of **Corda**'s validating model in Section 3.1. Second, even with encrypted content information, the act of contacting conveys information (denoted by $\mathbf{1}_k$). In the context of an industrial organization framework analyzed in Section 4, this renders the aggregate economic activities public information if all agents are contacted, which opens greater room for collusion and potentially jeopardizes competition.

In general, the quality of consensus and the amount of information distribution on blockchains depend on their specific protocols. There is a great diversity of algorithms for building consensus based on requirements such as performance, scalability, consistency, data capacity, governance, security, and failure tolerance. Moreover, the underlying cryptographic mechanism for consensus generation is complex and still under development. A number of papers make incremental progresses in improving upon the Bitcoin protocol (e.g., Kroll, Davey, and Felten (2013) and Nayak, Kumar, Miller, and Shi (2016)), but detailing the various consensus mechanisms or deriving the "optimal" blockchain design is beyond the scope of this paper.

3.6 Relating to the Literature on Information Economics

We conclude this section by highlighting the difference between our analysis and the extant literature on information economics. Earlier studies on information disclosure typically concern transparency, which does not qualitatively change the primary market function in facilitating trading and exchange. Namely, even without pre- and post- transparency re-

quirements on TRACE, trading and aggregation can still take place. In contrast, without information distribution, blockchains cannot perform their core function of generating decentralized consensus and tamper-proofness. In traditional settings, though greater public information may be detrimental, regulators or agents can opt to distribute no information.[19] But for blockchains to generate decentralized consensus, which is the key feature of the technology, an irreducible level of information distribution is required. What is more, protocols are typically designed to facilitate adoption, and also managed in a decentralized manner, rendering it extremely difficult for a centralized regulator or agent to easily alter the informational environment.

We next take the defining features of the blockchain technology as given, and study its impact on the real economy, especially the contracting and competition in various markets.

4 Blockchain Disruption and Industrial Organization

We use a standard dynamic industrial organization model similar to Green and Porter (1984) to analyze the impact of blockchain technology on contracting and business competition. Decentralized and smart contracts helps entry which promotes competition, but greater information distribution may foster collusion which hurts competition.

4.1 The Setting

Consider a risk-neutral world in which time is infinite and discrete, and is indexed by t, $t = 1, 2, \cdots$. Every agent has a discount factor δ. In every period $t \geq 0$, with probability λ a unit measure of buyers show up, each demanding a unit of goods, which could be a service such as fund transfer. Buyers (if present) only live for one period and exit the economy. We use $\mathbb{I}_t$ to denote the aggregate event whether buyers show up in period t. Throughout, we use "buyers", "consumers", and "customers" interchangeably.

There are three long-lived sellers who produce and sell the goods, and are either authentic or fraudulent. Sellers should be broadly interpreted as large financial institutions providing goods or services. A fraudulent seller is unable to deliver the goods, while the authentic one always delivers. At the start of the game $t = 0$, two of them, A and B, are incumbents known to be authentic (who have already established a good reputation). There is also

[19] Transparency in trading bonds is a good example. See, e.g., Goldstein, Hotchkiss, and Sirri (2006) and Bessembinder and Maxwell (2008). In particular, Bloomfield and O'Hara (1999) also find that market makers can use trade information to maintain collusive behavior.

a new entrant C who privately knows her authenticity, but others only have the common prior belief that C is authentic with probability π — later referred to as C's reputation. In every period $t \geq 0$, each seller gets an i.i.d. draw of the quality q_i, $i \in \{A, B, C\}$ of the goods they offer, which is the expected utility a buyer can get conditional on delivery of the goods. A buyer gets zero utility otherwise. Denote the cumulative distribution function and probability density function of quality distribution by $\phi(q)$ and $\Phi(q)$, and its support by $[\underline{q}, \bar{q}]$. It costs a seller μ to produce the goods, where $\mu < \underline{q}$ to reflect that transaction with an authentic seller is welfare-improving.

The quality profile $\mathbf{q} = (q_A, q_B, q_C)$ is realized at the beginning of each period and is publicly observable, capturing temporal differences among sellers. We discuss the case when quality is the seller's privately information in Section 5.3. For exposition, we denote the elements in $\mathbf{q}$ in decreasing order by $q^{(1)}$, $q^{(2)}$, and $q^{(3)}$ respectively. Without loss of generality, we treat q in the remainder of the paper as the probability of successful rendition of service by an authentic seller, which delivers one unit utility to a buyer.

C can potentially enter by paying an arbitrarily small cost of $\epsilon > 0$; hence C enters only if she can ever make strictly positive profit in this market after entry.[20] This allows us to focus on information asymmetry of seller authenticity as the relevant entry barrier. We further assume that before getting customers the entrant has no loss-absorbing capacity, for example, due to endogenous borrowing capacity, so that potential entrants cannot take aggressive penetration pricing schemes.[21]

In the context of financial industry, one can think of the buyers as bank customers, and the goods demanded as a certain type of financial service. The incumbents then represent well-established financial institutions with high reputation, and the entrant represents new service providers such as PayPal in its early days. In this example, μ represents the cost entailed in performing the service, and q the quality of service in terms of customer experience or speed or probability of completion conditional on seller's authenticity.

[20] Whether the entry decision is made before the quality q_C realization or not is immaterial, given the arbitrary small entry cost.

[21] A sufficient condition to rule out aggressive penetration pricing (in which entrants suffer huge losses in order to enter). This is realistic because without accumulation of service profit over time, the entrant typically does not have large initial capital (deep pocket) to undercut price aggressively. In fact, all we need is that C's tolerance for loss, L, is no more than $[\underline{q} - \pi\bar{q}]^+$.

4.2 Traditional World

Contracting Space and Information

We now give the key assumption on contracting space and information environment that are available in the traditional world.

Assumption 1. *In traditional world, no payment can be contingent on whether service delivery occurs or not. Each seller can only observe her own buyers and associated transaction information.*

The first part of Assumption 1 reflects certain contract incompleteness in real life that either limits the effectiveness of consensus or makes contracting on it too costly; for a good reference on the costs of writing and enforcing complete contracts, see Tirole (1999). In our context, this implies that the sellers first quote price $p_i(\mathbf{q})$ privately to buyers;[22] then, payoff-maximizing buyers choose one of the sellers, pay the offered price, and wait for the service to be delivered.

The second part of Assumption 1 implies that in the traditional world sellers do not observe others' price quotes. This assumption plays a role when we solve for the sellers' collusion equilibrium, and is similar to the assumption in Green and Porter (1984) and Porter (1983).

Bertrand Competition and Entry

Let us first consider a competitive equilibrium, in which sellers will keep lowering their offered prices until their competitors quit. Suppose that C enters. If $\pi q_C < \max\{q_A, q_B\}$, at least one of the incumbents always competes to lower the price to μ to get the customer this period and prevent the enhanced future competition they face had C entered in this period. Without a reputation of being authentic, C only stands a chance of getting a customer if buyers show up and $\pi q_C \geq \max\{q_A, q_B\}$.[23] The next proposition follows,

Proposition 4.1. *In a competitive equilibrium, the first time C can serve customers is in period* $\tau \equiv \min\{t \geq 0 | \pi q_{C,t} \mathbb{I}_t \geq \max\{q_{A,t}, q_{B,t}\}\}$ *or later. Consequently, C never enters if* $\pi\overline{q} < \underline{q}$.

[22] That sellers make offers is realistic in many applications where the customers or buyers are short-lived and dispersed. For example, banks typically quote the fee for making an international transfer, and customers can decide which bank to go to. Our main results are robust to this particular trading protocol.

[23] Even so, C may not get a customer if the incumbents use predatory pricing. Note that when $\pi\overline{q} < \underline{q}$, no matter what $\mathbf{q}$ is, C cannot enter even with penetration pricing because the maximum loss C can afford is less than $\underline{q} - \pi\overline{q}$.

Basically, an entrant can get customers only when her perceived quality is higher than the incumbents. If she does not expect to get customers, she never enters.

In the remainder of the paper, we focus on the case $\underline{q} > \pi\overline{q}$; in other words, the entrant C's reputation is sufficiently low that no entry ever occurs in the traditional world.[24] Obviously, the presence of fraudulent sellers causes this inefficiency, and traditionally the market relies on sellers' reputation, i.e., the probability being authentic, to mitigate it. We shall show later that this problem can be fully or better resolved by smart contracts with blockchain technology offering decentralized consensus.

Welfare and Consumer Surplus

With $\underline{q} > \pi\overline{q}$, C never enters. The expected future consumer (buyer) surplus and social welfare at any time s are, respectively,

$$\Pi_{buyer} = \mathbb{E}_s \left[\sum_{t=s+1}^{\infty} \delta^{t-s} \mathbb{I}_t \left(\min\{q_{A_t}, q_{B_t}\} - \mu \right) \right] = \frac{\delta\lambda}{1-\delta} \mathbb{E} \left[\min\{q_A, q_B\} - \mu \right] \tag{6}$$

and

$$\Pi_{total} = \mathbb{E}_s \left[\sum_{t=s+1}^{\infty} \delta^{t-s} \mathbb{I}_t \max\{q_{A_t}, q_{B_t}\} \right] = \frac{\delta\lambda}{1-\delta} \mathbb{E} \left[\max\{q_A, q_B\} \right]. \tag{7}$$

Collusive Equilibria

Besides the competitive equilibrium derived, there may exist collusive equilibria in this economy. Recall that sellers cannot make contingent contracts, and entrant cannot enter ($\pi\overline{q} < \mu$); we only need to examine potential tacit collusion among the incumbents.

We restrict each seller's strategies to the standard supergame strategies discussed in, for example, Green and Porter (1984). Specifically, consider the following strategy for A and B to collude. There are two phases:

1) *Collusion phase*: Every period, after the realization of types, A charges price q_A and B charges price q_B. A and B gets $\mathbb{I}_t f(q_A, q_B)$ and $f(q_B, q_A) = \mathbb{I}_t(1 - f(q_A, q_B))$ fractions of buyers, respectively. Here $f(x, y) \in (0, 1)$ is the proposed anonymous allocation function, potentially as a function of realized types. For example, the sellers can split the total

[24] Together with $\underline{q} > \mu$, this case highlights that authenticity matters more than the dispersion in service quality: we would rather make an international transfer at reputable banks despite the differential customer service they have, than entrust the money to a random person on the street who is polite and offers to make the transfer for me.

customers by setting quota on service to be delivered. This allocation function f includes the case where sellers always equally split buyers, and the case where buyers all go to the better seller.

2) *Punishment phase*: The punishment phase is triggered once one of the sellers does not have any buyers.[25] More specifically, the punishment phase can be triggered either by i) buyer not showing up this period or ii) one of the seller deviates by quoting a cheaper price to get all the buyers. Once triggered, A and B are engaged in Bertrand competition for a fixed T period.

Recall that the sellers do not observe other sellers' price quotes, but observes their own customers. However, A and B's private signals are always correlated (i.e., observing either no customers or all customers), making this repeated game with private monitoring essentially a game with imperfect public monitoring. It is imperfect in the sense that punishment could be triggered even when no one deviates. The equilibrium notion corresponding to the above strategies is thus akin to public perfect equilibrium.

A standard result in the literature of dynamic repeated games is that sustainable equilibria crucially depend on the discount factor δ, with the Folk Theorem as the best-known example. We therefore proceed to find the lower bound of discount factor, denoted by $\delta_{(T,f)}$, above which an equilibrium with a specified T and $f(x,y)$ exists.

Lemma 4.2. *A collusion strategy with (T,f) as described above is an equilibrium, if*

$$\frac{\lambda\delta\left(1-\delta^{T}\right)}{1-\lambda\delta-(1-\lambda)\delta^{T+1}} \geq \frac{M_3}{M_1-M_2} \tag{8}$$

where $M_1 = E[f(q)(q-\mu)], M_2 = E[(q_i - q_{-i})^+], M_3 = \max_q\{(1-f(q))(q-\mu)\}, f(q_i) = E_{q_{-i}}[f(q_i,q_{-i})]$.

We note M_1 is a seller's expected payoff in each stage game in the collusion phase, M_2 is that in the punishment phase, and M_3 is the maximum gain from deviating. To sustain a collusion, we basically need the incentive for one time deviation from collusion to be relatively small compared with the punishment going forward. Note that the LHS of (8) is increasing in δ and in T. Therefore, there exists a $\delta_{(T,f)}^{Traditional}$ above which the collusive equilibrium is sustained. Moreover,

Proposition 4.3. *The discount threshold $\delta_o^{Traditional} \equiv \inf_f \frac{1}{\lambda}\frac{M_3}{M_1+M_3-M_2}$ is well-defined and positive. When $\delta < \delta_o^{Traditional}$, no collusion equilibrium exists for any (T,f).*

[25] We could alternatively allow punishment to be triggered with some probability, which is similar to shortening the punishment phase. This does not affect our main results and is left out for exposition clarity.

With sufficiently small discount factor, no collusion can be sustained because sellers value future cost of punishment too lightly and prefer the one-time deviation gain in the current stage.

The welfare under (T, f) collusion is determined by f, and consumer surplus by both (T, f) and colluding price. One feature that stands out is that the consumer surplus depends on the length of punishment period T because buyers earn positive surplus only when the punishment phase is triggered due to absence of buyers in the economy (which occurs with probability $1 - \lambda$ in each period).

Now consider a collusion where sellers charge a lower colluding price (less than q_i), it easily follows that (8) is relaxed. Therefore if a collusive equilibrium extracting all rent in collusion phase is sustainable, an equilibrium with the safe (T, f) but lower colluding prices is also sustainable. This implies that equilibria with consumer surplus ranging from the competitive level and the (T, f) collusion level are all sustainable.

4.3 Blockchain World

The blockchain technology enables the consensus recording of success or failure of the service rendered by verifying and validating certain transactions, which, as detailed earlier in Sections 2.1 and 3, typically involves distributing information. Its algorithmic nature enables certain transfers to be automated based on consensus, reducing enforcement costs and mitigating contract incompleteness. For example, transaction amount is observable on **Ripple** network.

To highlight the economic force, we examine the case where the consensus provision regarding service delivery (but not the transaction term) is perfect ($K \to \infty$). This case captures many extant blockchains such as **Bitcoin**, **Ripple**, and **Symbiont**, where either the verification request or transaction information is distributed to sufficiently large numbers of people including major institutional participants such that consensus is near perfect. The basic tradeoff under imperfect consensus is qualitatively the same, which we discuss in Section 5.2.

Assumption 2. *The blockchain contacts all participants (including the sellers and the continuum of consumers) to generate effective decentralized consensus. More specifically, the blockchain consensus $\widetilde{z} = \widetilde{\omega}$ and a seller upon being contacted infers that customers are present.*

Recall $\widetilde{\omega}$ is the delivery outcome (whether successful or not). This assumption implies that

i) self-executed smart contracts can be perfectly contingent on service outcome consensus; and ii) the sellers observe whether there is aggregate activity on the blockchain. These are in sharp contrast to Assumption 1. For example, a bank can credibly offer a transfer between itself and a customer contingent on the service outcome. Two customers or two banks can arrange a credible transfer between them contingent on customers being serviced by all the banks. The list goes on. In addition, the sellers may have richer information, but our arguments require weak conditions and it suffices that they observe the aggregate activity.

In the rest of this section, we will first demonstrate how blockchain and smart contracts can enhance entry and competition, then show the same technology can lead to greater collusive behavior, before discussing regulatory implications.

Smart Contracts and Enhanced Entry

With blockchain, the entrant now can offer a price contingent on the success of service provision $\mathbb{P} = (p^s, p^f)$, where p^s and p^f are prices charged upon success and failure. An authentic entrant C can separate from her fraudulent peer by offering $(p^s, -\epsilon)$, where $p^s > 0$ and ϵ is infinitesimal. The fraudulent type gains nothing from mimicking: she knows that she can never deliver the service and hence never receive the payment.

Let us first analyze the equilibrium without potential collusion. We have

Proposition 4.4. *With smart contracts, the entrant C enters almost surely, and first gets customers in period* $\tau = \min\{t \geq 0 | q_{C,t}\mathbb{I}_t \geq \max\{q_{A,t}, q_{B,t}\}\}$ *or earlier.*

Smart contracts completely remove the reliance on sellers' reputation of being authentic, thus as long as the entrant's quality is higher than the incumbents, she can get customers. Thus she enters for sure.

In the world with blockchain and smart contracts, the expected future consumer surplus and total welfare at $t = s$ under a competitive equilibrium are, respectively,

$$\Pi_{buyer} = \mathbb{E}_s \left[\sum_{t=s+1}^{\infty} \delta^{t-s} \mathbb{I}_t \left(q^{(2)} - \mu \right) \right] = \frac{\delta\lambda}{1-\delta} \mathbb{E}\left[q^{(2)} - \mu \right] \tag{9}$$

and

$$\Pi_{total} = \mathbb{E}_s \left[\sum_{t=s+1}^{\infty} \delta^{t-s} \mathbb{I}_t q^{(1)} \right] = \frac{\delta\lambda}{1-\delta} \mathbb{E}\left[q^{(1)} \right]. \tag{10}$$

Compared to (6) and (7), we see that with smart contracts that facilitate entry and hence enhance competition, the economy becomes more efficient. Both consumer surplus (linear in the second order statistic), and welfare (linear in first order statistic) improve. Therefore, we clearly see that smart contracts can help improve consumer surplus and welfare. This welfare improvement due to enhanced entry is present in collusive equilibrium as well, as C always enters.

Enhanced Collusion under Permissioned Blockchain

While blockchain and smart contracts can improve both consumer surplus and welfare by encouraging entry and competition, they have a dark side and may result in dynamic equilibria with lower welfare or consumer surplus than in all the equilibria in the traditional world. To highlight the collusion-enhancing effect of blockchain, we first focus on permissioned blockchain for the incumbents which C cannot use (hence no entry), before discussing public blockchains that C can utilize.

Collusion using Smart Contract With blockchain and smart contracts, it is apparent that sellers can use the enlarged contingencies to facilitate collusion. For example, the sellers can form a coalition and sign on to a smart contract, which essentially uses side payments contingent on the service outcome to punish deviation.

Recall that Assumption 2 implies that the transaction information stored on the blockchain include whether buyers show up. Consider the following collusion with smart contract. All sellers collude to charge 1 dollar upon delivery, which effectively extract full rents from buyers. The sellers reach an agreement that they never serve all buyers (always leaving some strictly positive measure to other sellers); and if all the buyers go to seller i, which is a contractible contingency, then the smart contract automatically transfers all profit of seller i to other sellers. By imposing such automatic punishment upon deviation, the smart contract can potentially support any collusion, regardless of the discount factor.

Such explicit form of collusion using smart contracts is easy to detect and can be forbidden by anti-trust law (Section 5.1). The more relevant and interesting phenomenon is that even without explicit side payment through smart contracts, the blockchain still can facilitate greater collusion, which we discuss next.

Tacit Collusion with Permissioned Blockchain In the case of tacit collusion, the collusion and punishment phases as well as the allocation rule f are exactly the same as in

the traditional world. However, instead of triggering punishment upon deviating or receiving no buyers, punishment in the blockchain world can be further conditioned on whether buyers show up because participants upon being contacted for verification at least know that service requests are made, which allows the sellers to perfectly monitor deviation behavior by a colluding fellow.

In other words, the repeated game with traditionally imperfect public monitoring now achieves perfect public monitoring as deviations can be accurately detected using blockchain.[26] And punishing deviations more accurately, collusive equilibria become easier to sustain. Denoting the threshold discount factor above which collusion is sustained with permissioned blockchain by $\delta_{(T,f)}^{Blockchain2}$, we have

Proposition 4.5. *Compare the thresholds above with the specified collusion strategy is an equilibrium. We have*

$$\delta_{(T,f)}^{Blockchain2} < \delta_{(T,f)}^{Traditional} \tag{11}$$

When the discount factor $\delta \in [\delta_{(T,f)}^{Blockchain2}, \delta_{(T,f)}^{Traditional})$, the consumer welfare under collusion with blockchain is lower than that under competitive market without blockchain. In particular,

Corollary 4.6. *When* $\delta \in \left[\inf_f\{\delta_{(\infty,f)}^{Blockchain2}\}, \delta_o^{Traditional}\right)$*, there cannot be collusion without blockchain, but there could be with blockchain.*

Furthermore, if f does not always allocate all customers to the better incumbent, welfare is also lower.

Importantly, for general δ, any traditional collusion equilibrium with (T, f) has a corresponding sustainable blockchain collusion equilibrium with the same (T, f), but the latter extracts greater rents for the sellers. We note that the sellers can always return the same rents to the buyers by lowering the colluding price, or activating punishment phase even when there is no buyers in the system at all. Hence if a dynamic equilibrium with a certain welfare and consumer surplus is sustainable in the traditional world, it is also sustainable with blockchain. Yet, as demonstrated earlier, there could be dynamic equilibria with additional welfare and consumer surplus outcomes sustained with blockchain.

[26] With private monitoring with less correlated signals, to the extent that private signals are generated from a noisy signal of the true state of the world, having a consensus on the noisy signal increases private signals' correlation, which also makes the equilibrium more easily sustained. This is beyond our current discussion but constitutes an interesting future work. For more discussion on private monitoring, see for example Mailath and Morris (2002) and Hörner and Olszewski (2006).

4.4 Blockchain Disruption

While a permissioned blockchain does not facilitate entry, an authentic entrant can use smart contracts on public blockchain to separate from the fraudulent type (Section 4.3). Would the benefit of entry outweigh the cost of potential greater collusion? We now answer this question under the premise that there is a public blockchain that all three firms (incumbents A, B, and new entrant C) all have access to.

Consumer Surplus under Public Blockchain

Recall that Section 4.3 has solved the competitive equilibrium. To characterize other collusive equilibria in this economy, consider the following collusion strategy:[27]

1) *Collusion phase*: Every period, after the realization of types, each seller i charges 1 dollar contingent on success. Let $\hat{f}(q_i, q_j, q_k)$ be the fraction of the buyers that go to the seller with quality q_i where

2) *Punishment phase*: The punishment phase is triggered if one of the sellers does not have any buyers AND there are buyers showing up this period. In other words, the punishment phase is triggered only if there is some seller deviates. Once triggered, all sellers get involved in Bertrand competition for T periods.

Lemma 4.7. *With blockchain and smart contract, the above strategy is an equilibrium if the parameters satisfy*

$$\frac{\delta\lambda\left(1-\delta^T\right)}{1-\delta} \geq \frac{\hat{M}_3}{\hat{M}_1 - \hat{M}_2} \tag{12}$$

where $\hat{M}_1 = E[\hat{f}(q)(q-\mu)], \hat{M}_2 = E[(q_i - \max_{j\neq i} q_j)^+], \hat{M}_3 = \max_q\{(1-\hat{f}(q))(q-\mu)\}$.

The $\hat{M}$s have similar interpretations as in Lemma 4.2, but for three sellers instead of two. The LHS of equation 12 is also modified because with perfect public monitoring, the punishment is more accurately targeted.

We know smart contracts enhance competition by facilitating entry. Is it possible that even with entry, blockchain reduces welfare and consumer surplus? The answer is largely "yes," though the analysis needs extra care because with three sellers colluding, the customer-splitting rule $\hat{f}$ is necessarily different from f for the case of two incumbent sellers, making it

[27] Again, it suffices to examine the collusive behaviors that allow the sellers full rent, any equilibrium with the same allocation but higher consumer surplus can be achieved by lowering the collusion price. Also, because the incumbents would not do better by colluding among themselves when C is competitive, than by colluding altogether with C, we only examine collusion of all three sellers — a more severe case of collusion.

inappropriate to directly compare thresholds $\delta^{Traditional}_{(T,f)}$ and $\delta^{Blockchain3}_{(T,\hat{f})}$, where *Blockchain3* indicates the public blockchain with all three sellers. A series of formal results ensue.

Proposition 4.8. *The discount threshold* $\delta^{Blockchain3}_o \equiv \inf_{\hat{f}}\{\delta^{Blockchain3}_{(\infty,\hat{f})}\}$ *is well-defined and satisfies* $\delta^{Blockchain3}_o < 1$*. For all* $\delta > \delta^{Blockchain3}_o$*, there exists a collusion equilibrium with blockchain such that the consumer surplus is lower than that in any equilibrium in the traditional world.*

This proposition gives a sufficient condition on the discount factor δ so that even with entry, blockchain and smart contracts hurt consumers' surplus.

Corollary 4.9. *For* $m \geq n \geq 2$*, if* $\lambda < \frac{n-1}{n}$*, then* $\delta^{Traditional,n}_o > \delta^{Blockchain,m}_o$*, where* m *and* n *indicate the number of colluding sellers with and without blockchain respectively. Consequently for all* $\delta \in [\delta^{Blockchain,m}_o, 1)$*, there is no collusion in the traditional world with* n *incumbents, while there can be collusion with blockchain with* m *sellers that reduces consumer surplus.*

This proposition highlights that the way blockchain disruption could potentially hurt the consumer surplus is through bringing in new entrant only to collude with those incumbents, whereas in the traditional world the incumbents cannot sustain any collusion.

Dynamic Equilibria under Blockchain Disruption

More generally, in terms of welfare and consumer surplus, the set of equilibrium outcomes with blockchain disruption is a non-trivial superset of those in equilibria in the traditional world.

Theorem 4.10. *The discount threshold* $\delta^{Blockchain3}_a \equiv \sup_f\{\delta^{Blockchain3}_{(\infty,\hat{f})}\}$ *is well-defined and satisfies* $\delta^{Blockchain3}_a < 1$*. For all* $\delta > \delta^{Blockchain3}_a$*, any consumer surplus and welfare attainable in the traditional world can be attained with blockchain, and some additional equilibria with higher or lower consumer surplus or welfare can also be sustained.*

In this proposition, the subscript a in $\delta^{Blockchain3}_a$ stands for "all", indicating that if the discount factor is above $\delta^{Blockchain3}_a$, *all* collusion equilibria can be sustained.

Corollary 4.11. *The most collusive equilibrium with blockchain, which generates the highest payoff to the sellers, improves social welfare but results in strictly lower consumer surplus, compared to any equilibrium outcome under the traditional world.*

Finally, this corollary implies that if the sellers can initially coordinate to pick their favorable collusion, welfare improves but buyers are worse off. The intuition is that with new entry, welfare improves when the sellers allocate the business to the highest quality one in the most collusive equilibrium, but consumer surplus is lower than that in the traditional world because consumers always retain some rent in the traditional world due to the fact that punishment phase (competitive stage-game equilibrium) occurs in equilibrium.

5 Discussions and Extensions

5.1 Measures to Reduce Collusion on Blockchain

There is a wide-spread concern that blockchains can jeopardize market competitiveness; this concern becomes especially acute for permissioned blockchains whose members are powerful financial institutions. As described in a Financial Times article, "...the technology really facilitates is *cartel management* for groups that don't trust each other but which still need to work together if they are the value and stability of the markets they serve."[28] Our paper highlights one particular economic mechanism through which blockchain hinders competition, and provides the rigorous analysis on why and how collusion could occur. We now explore regulatory and market solutions to curb collusive behaviors in our framework.

Blockchain Competition

Although we focus on the case of a single blockchain on which multiple sellers compete, in practice there are likely to be multiple blockchains which both sellers and buyers can choose. Suppose that sellers only participate in a subset of blockchains in equilibrium; this can be justified by some fixed participation cost. Then the competition among blockchains seems to go against the collusive behaviors of sellers on one blockchain, because buyers can always pick the blockchain which offers the best price-adjusted service, which drives blockchains that are with a critical mass of sellers but are colluding out of market.

Of course, this discussion begs another question: Why is it more difficult for blockchains to collude, at least relative to sellers on the same blockchain? We believe it is indeed the case, as blockchains are decentralized with many players, and information is not shared across blockchains.

[28] "Exposing the 'If we call it a blockchain, perhaps it wont be deemed a cartel?' tactic," by Izabella Kaminska, Financial Times, May 11th, 2015.

While blockchain competition may mitigate collusive behaviors on specific blockchains, in the long run if a single blockchain becomes dominant due to network effect, regulators have to step in to prevent collusions.

Regulatory Node and Design

In the traditional world, it helps for regulatory agency to observe and collect more information about the market, in order to better detect collusive behaviors. Similarly, adding a regulatory node in the blockchain, especially private blockchains, can help monitor the information distribution. To some extent, being part of the business ecosystem on the blockchain helps the regulators to reduce tacit collusion because they have the same information about other market participants on the blockchain as a colluding party.

The key difference from regulatory monitoring in the traditional world is that because blockchain records are real-time and tamper-proof, regulators do not have to worry about misreporting and time-delays, enabling the detection and containment of collusions and market malfunctions at short-horizons. Moreover, retrospective auditing is no longer prone to manipulation.

What is even more helpful is for the regulators to potentially participate in the protocol design. For example, the government can reserve access to certain encrypted information that is broadcasted to blockchain participants or recordkeepers. Not only does this direct access enables elimination of collusions using smart contracts (discussed in Section 4.3), but it also allows detections of tacit collusions based on statistical analysis of transaction and pricing behaviors. Alternatively, regulators can impose a separation of keepers and end users, which we discuss next.

Separation of Usage and Consensus Generation

In the model, sellers can use the information on the blockchain to punish deviations from collusion in a more accurate way. They observe the information because the information is distributed and recorded on the blockchain during the process of consensus generation. From this perspective, one obvious potential solution is to separate the players who help generate the decentralized consensus, from the users of that consensus. For example, if sellers can only use the blockchain for signing smart contracts with buyers, then they no longer have access to the aggregate activity information that facilitates collusion. On some blockchains such as **Symbiont**, keepers tend to be a rather separate group from the end users, though this resolution has not been sufficiently explored.

In our model, if the sellers are excluded from the recordkeeping activities and individual transactions are only accessible among the parties involved, then sellers would no longer accurately observe if other sellers have deviated, rending collusive equilibria less sustainable. In practice, this may be challenging because the parties that we should exclude from being contacted for recordkeeping are also likely the ones who are the most qualified to validate a record (for example by being in the same industry). Nevertheless, the separation of usage and consensus generation is novel and constitutes a promising direction for future regulatory policies targeted to blockchain applications.

5.2 Imperfect Consensus

In our discussion on industrial organization we have assumed that all participants on the blockchain are contacted as recordkeepers, which gives perfect consensus ($K \to \infty$) and observability for the sellers. This is consistent with the majority of extant blockchains in practice, and as we argue next, is welfare-maximizing (and hence Pareto-efficient) in our setting.

Suppose a finite subset $\mathcal{K}$ of blockchain participants serve as keepers, and $|\mathcal{K}| = K$. For simplicity, suppose any keeper correctly observes the service outcome with probability $\tilde{\psi} > \frac{1}{2}$, and let the resulting imperfect consensus be such that if a service is delivered, with probability ψ the blockchain correctly records it; similarly, failed delivery is also recorded as failure with probability ψ. The exact link between ψ, $\tilde{\psi}$, and K depends on the consensus protocol. For example, if the consensus is based on unanimity of contacted keepers, $\psi = \tilde{\psi}^K$; if it is based on majority rule, $\psi = \sum_{m=\lceil \frac{K}{2} \rceil}^{K} \binom{K}{m} \tilde{\psi}^m (1-\tilde{\psi})^{K-m}$. In either case, ψ captures in reduced-form the imperfectness of consensus generation.

In Proposition 4.1 we have shown that C never enters in the traditional world. Therefore, the only way for the authentic type to enter with the help of blockchain is through signaling with smart contract (p^s, p^f) and separating from (instead of pooling with) the fraudulent type. Recall that the entrant's capacity to bear initial loss is L and the service cost is μ, the authentic type in the separating equilibrium then solves the following:

$$\begin{aligned} &\max_{(p^s,p^f)} \quad \psi p^s + (1-\psi)p^f \\ s.t. \quad &\psi p^s + (1-\psi)p^f \geq \mu, \quad -p^f \leq L, \quad \textit{and} \quad (1-\psi)p^s + \psi p^f < 0, \end{aligned}$$

where the inequalities are the authentic type's participation constraint, limited loss capacity, and the fraudulent type's no-mimicking constraint, respectively. To understand the last

inequality, the fraudulent entrant whose probability of success is $1-\psi$ by mistake can now afford to rebate the consumers upon failure for an amount up to $-\hat{p}^f = \frac{1-\psi}{\psi}\hat{p}^s$ and still breaks even in expectation. As such, the authentic type has to separate with a smart contract $\left(p^s, -\frac{1-\psi}{\psi}p^s - \epsilon\right)$, for an arbitrarily small positive ϵ.

The program admits a solution when $\psi \geq \frac{\mu+L}{\mu+2L}$, which allows the authentic type to at least break even. We thus have

Proposition 5.1. *As long as the consensus quality is not too low ($\psi \geq \frac{\mu+L}{\mu+2L}$), the use of smart contract facilitates entry of the authentic type.*

The result also implies that for businesses that require high investment relative to the initial wealth of an entrant, smart contract under imperfect consensus may fail to mitigate informational asymmetry.

In our model, there is a continuum of consumers upon arrival, which implies that there is a continuum of transactions to be verified. If each verification process draws keepers in some independent way, then the law of large numbers across transactions reveals the aggregate state of customer arrival, even under imperfect consensus. Therefore, imperfect consensus does not affect the collusive equilibria supported. Overall, it weakly reduces entry and competition, and it is in this sense that it is weakly welfare improving to have perfect consensus.

What really matters for collusion is whether sellers learn the arrival of consumers in each period. One way to reduce this observability is by contacting a smaller number of keepers randomly, so that sellers are less likely being contacted to verify transactions. However, since imperfect consensus reduces efficient entry and allocation of business among sellers, this is not an ideal way for reducing collusion. A better design, as we described earlier, is to separate the sellers from the keepers and reduce directly contacting the former.

To model this exclusion of sellers, suppose for each service record, a seller is contacted with probability $\hat{\zeta}$, then the probability that a seller is completely unaware of the aggregate service activity conditional on consumers' arriving is $1-\zeta \equiv (1-\hat{\zeta})^n$ where n is the number of transactions. In the collusion-phase, a deviation is detected with probability of ζ instead of with probability one, triggering less punishment and making the collusion equilibrium harder to sustain. That said, if the number of transactions is large, the equilibrium approaches the one with perfect public monitoring unless the sellers are strictly prohibited from acting as keepers ($\hat{\zeta} = 0$).

5.3 Information Asymmetry and Private Qualities

In our analysis so far, $\mathbf{q}$ is publicly known, and many forms of smart contract can be used to solve the problem of inefficient entry (extensive margin of competition). This is obviously a strong assumption, without which the matching of consumers with incumbents could also be inefficient (intensive margin of competition). In this section we allow privately observed qualities. Collusion with private information in general is complex (Athey and Bagwell (2001) and Miller (2011)), therefore our focus is on the competitive equilibrium and competitive stage games in the punishment phase of a collusion.

We characterize how smart contracts can help mitigate allocative inefficiency beyond entry, and derive the equilibrium form of smart contracts under market equilibrium.

Allocative Inefficiency in the Traditional World

Suppose that in addition to uncertainty on authenticity, quality q_i is also only privately known to seller i. Without smart contracts, the entrant would always claim it is authentic and has high quality (cheap talk), and incumbents cannot separate themselves either because both A and B can claim that it is of higher quality. Following the same logic before, the entrant does not enter because its perceived quality $\pi\mathbb{E}[q_c]$ is below the incumbents' perceived quality $\mathbb{E}[q_c]$. In a competitive equilibrium among the incumbents, we have

Lemma 5.2. *In the traditional world, sellers post the same price $p_i = \mu$, and each buyer select (randomly) one of them. The expected buyer's surplus and social welfare per period is $\mathbb{E}[q] - \mu$.*

Lemma 5.2 shows that there is no separating equilibrium in this economy. The reason is simple. When payment cannot be made contingent on whether the transaction succeeds or not, there is no cost of misreporting: the seller can always misreport to get the highest (expected) upfront payment.

In the traditional world without smart contract, there are both welfare loss and market breakdown. In the case where $\mu \leq \mathbb{E}[q]$, there is social welfare loss since the transaction is implemented by a randomly selected seller instead of the highest type. With two incumbents, the consumer surplus is higher than the case with publicly-known $\mathbf{q}$. However, with more sellers, the mean is lower than the second order statistics and consumer surplus is lower than the case with publicly-known $\mathbf{q}$. Therefore in general, information asymmetry on seller qualities leads to lower welfare and consumer surplus.

World with Blockchain and Equilibrium Smart Contracts

Smart contracts enlarge the space of price quotes sellers can use. In general, sellers can offer $\mathbb{P} = (p^s, p^f)$, and type q upon getting customers earns $S_q(\mathbb{P}) = qp^s + (1-q)p^f - k$ from each buyer who in return gets a utility $B_q(\mathbb{P}) = q(1-p^s) + (1-q)(-p^f)$, where $1-p^s$ is the unit utility from successful service lest the payment.

Sellers' flexibility in offering contingent prices implies that buyers' choice generally depends on their beliefs regarding the smart contracts that each seller type submits in equilibrium, making smart contract offering a signaling game.[29] We further impose that $p^f \leq p^s$ so that $\mathbb{P}$ is higher upon success, a standard assumption in the security design literature (see, e.g., Innes (1990), Hart and Moore (1995), and DeMarzo and Duffie (1999)).

Sellers may offer a large variety of smart contracts. We show that only one particular class of contracts emerges in equilibrium, which is further characterized by the following proposition.

Proposition 5.3. *There is a unique competitive equilibrium for the stage game, and sellers offer smart contracts of the form $\mathbb{P}^* = (p, p-1)$. A seller of quality q offers $(p_q, p_q - 1)$, where*

$$p_q = 1 - q + \mu + \int_{\underline{q}}^{q} \left[\frac{\Phi(q')}{\Phi(q)}\right]^2 dq' \tag{13}$$

which is decreasing in q. Buyers all go to the highest-quality seller.

Recall that Φ is the cdf of q. We note that such a contract means buyers get utility $1-p$ regardless of the service outcome. The conclusion mirrors the well-known result in the literature of security design that the sellers would offer the least information-sensitive security ("flattest" security in the language of security-bid auctions, e.g. DeMarzo, Kremer, and Skrzypacz (2005)).[30]

In a competitive equilibrium, we essentially have a cash auction in which a bidder with quality q has a private valuation of his/her service opportunity $q - \mu$, and bids p. In equilibrium, buyers choose the highest quality seller, who gets the second highest valuation

[29] Under a market mechanism where buyers shop sellers and choose the most favorable one, our setup has a natural reinterpretation under informal first-price auctions with security bids. See, for example, DeMarzo, Kremer, and Skrzypacz (2005) and Cong (2016).

[30] We further remark that the result is robust to a payment rule that depends on all price quotes (not only the winning seller's). To see this, any combinations of the smart contract bids are still in the same contract class whose convex hull is itself. Therefore the proof still applies. This implies no matter whether the sellers quote final prices (first-price auction) or gradually out-compete other sellers (English auction) or choose a third-price auction, sellers always use quality-insensitive contracts.

$\mathbb{E}[q^{(2)} - \mu]$ in each period the buyers arrive, which can be seen by applying the revenue equivalence theorem. Notice that the economic outcomes are exactly the same as in the case where $\mathbf{q}$ is public ((9) and (10)). Therefore we have,

Corollary 5.4. *Smart contracts fully resolve informational asymmetry in a competitive equilibrium, and welfare and consumer surplus are independent of whether seller qualities are private or not.*

That said, one can show that restricting the form of smart contracts can potentially increase the consumer surplus in a way similar to security design affects issuer's payoffs. For regulators concerned with consumer surplus, collusion and smart contract forms should be jointly considered — a topic for future studies.

6 Conclusion

In this paper we argue that decentralized ledger technologies such as blockchains feature decentralized consensus as well as low-cost, tamper proof algorithmic executions, and consequently enlarge the contracting space and facilitate the creation of smart contracts. However, the process of reaching decentralized consensus changes the information environment on blockchain.

We analyze how this fundamental tension can reshape industry organization and the landscape of competition; it can deliver higher social welfare and consumer surplus through enhanced entry and competition, yet may also lead to greater collusion. In general, blockchain and smart contracts can sustain market equilibria with a larger range of economic outcomes. We discuss regulatory and market solutions to further improve consumer surplus, such as separating agents generating consensus from end-users.

To focus on the impact of blockchain and smart contracts on the financial sector and the real economy, we have modeled in reduced-form the universal feature of blockchains and the key tradeoffs of consensus generation and information distribution. With this in mind, designing a robust consensus protocol and providing the right incentives for maintaining consensus on specific blockchains constitute interesting future studies, which likely require the joint effort of computer scientists and economists.

References

Abreu, Dilip, 1987, *Repeated games with discounting: A general theory and an application to oligopoly* (University Microfilms).

——— , David Pearce, and Ennio Stacchetti, 1986, Optimal cartel equilibria with imperfect monitoring, *Journal of Economic Theory* 39, 251–269.

Athey, Susan, and Kyle Bagwell, 2001, Optimal collusion with private information, *RAND Journal of Economics* 32, 428–465.

Aune, Rune Tevasvold, Maureen O'Hara, and Ouziel Slama, 2017, Footprints on the blockchain: Trading and information leakage in distributed ledgers, *The Journal of Trading*.

Baldimtsi, Foteini, Lin William Cong, Zhiguo He, and Jiasun Li, 2017, The creation and organization of firms: Theory and evidence from bitcoin mining pools, .

Bartoletti, Massimo, and Livio Pompianu, 2017, An empirical analysis of smart contracts: platforms, applications, and design patterns, *arXiv preprint arXiv:1703.06322*.

Bessembinder, Hendrik, and William Maxwell, 2008, Markets transparency and the corporate bond market, *The Journal of Economic Perspectives* 22, 217–234.

Biais, Bruno, Christophe Bisiere, Matthieu Bouvard, and Catherine Casamatta, 2017, The blockchain fold theorem, *Preliminary Work in Progress*.

Bloomfield, Robert, and Maureen O'Hara, 1999, Market transparency: who wins and who loses?, *Review of Financial Studies* 12, 5–35.

Böhme, Rainer, Nicolas Christin, Benjamin Edelman, and Tyler Moore, 2015, Bitcoin: Economics, technology, and governance, *The Journal of Economic Perspectives* 29, 213–238.

Buterin, Vitalik, 2014, Ethereum: A next-generation smart contract and decentralized application platform, *URL https://github. com/ethereum/wiki/wiki/% 5BEnglish% 5D-White-Paper*.

Catalini, Christian, and Joshua S Gans, 2016, Some simple economics of the blockchain, Discussion paper, National Bureau of Economic Research.

Cong, Lin William, 2016, Auctions of real options, *(NBER Working Paper)*.

DeMarzo, Peter, and Darrell Duffie, 1999, A liquidity-based model of security design, *Econometrica* 67, 65–99.

DeMarzo, Peter, Ilan Kremer, and Andrzej Skrzypacz, 2005, Bidding with securities: Auctions and security design, *American Economic Review* 95(4), 936–959.

Evans, David S, 2014, Economic aspects of bitcoin and other decentralized public-ledger currency platforms, .

Eyal, Ittay, and Emin Gün Sirer, 2014, Majority is not enough: Bitcoin mining is vulnerable, in *International Conference on Financial Cryptography and Data Security* pp. 436–454. Springer.

Fudenberg, Drew, and Eric Maskin, 1986, The folk theorem in repeated games with discounting or with incomplete information, *Econometrica: Journal of the Econometric Society* pp. 533–554.

Goldstein, Michael A, Edith S Hotchkiss, and Erik R Sirri, 2006, Transparency and liquidity: A controlled experiment on corporate bonds, *The review of financial studies* 20, 235–273.

Green, Edward J, and Robert H Porter, 1984, Noncooperative collusion under imperfect price information, *Econometrica: Journal of the Econometric Society* pp. 87–100.

Hart, Oliver, and John Moore, 1995, Debt and seniority: An analysis of the role of hard claims in constraining management, *American Economic Review* 85.

Harvey, Campbell R, 2016, Cryptofinance, *Available at SSRN 2438299*.

Hörner, Johannes, and Wojciech Olszewski, 2006, The folk theorem for games with private almost-perfect monitoring, *Econometrica* 74, 1499–1544.

Innes, Robert D, 1990, Limited liability and incentive contracting with ex-ante action choices, *Journal of economic theory* 52, 45–67.

Khapko, Mariana, and Marius Zoican, 2017, 'smart' settlement, *Working Paper*.

Kiayias, Aggelos, Elias Koutsoupias, Maria Kyropoulou, and Yiannis Tselekounis, 2016, Blockchain mining games, in *Proceedings of the 2016 ACM Conference on Economics and Computation* pp. 365–382. ACM.

Krishna, Vijay, 2009, *Auction theory* (Academic press).

Kroll, Joshua A, Ian C Davey, and Edward W Felten, 2013, The economics of bitcoin mining, or bitcoin in the presence of adversaries, in *Proceedings of WEIS* vol. 2013. Citeseer.

Lauslahti, Kristian, Juri Mattila, Timo Seppälä, et al., 2016, Smart contracts–how will blockchain technology affect contractual practices?, Discussion paper, The Research Institute of the Finnish Economy.

Mailath, George J, and Stephen Morris, 2002, Repeated games with almost-public monitoring, *Journal of Economic theory* 102, 189–228.

Malinova, Katya, and Andreas Park, 2016, Market design for trading with blockchain technology, *Available at SSRN*.

Miller, David A, 2011, Robust collusion with private information, *The Review of Economic Studies* 79, 778–811.

Nakamoto, Satoshi, 2008, Bitcoin: A peer-to-peer electronic cash system, .

Narayanan, Arvind, Joseph Bonneau, Edward Felten, Andrew Miller, and Steven Goldfeder, 2016, *Bitcoin and cryptocurrency technologies* (Princeton University Pres).

Nayak, Kartik, Srijan Kumar, Andrew Miller, and Elaine Shi, 2016, Stubborn mining: Generalizing selfish mining and combining with an eclipse attack, in *Security and Privacy (EuroS&P), 2016 IEEE European Symposium on* pp. 305–320. IEEE.

Philippon, Thomas, 2015, Has the us finance industry become less efficient? on the theory and measurement of financial intermediation, *The American Economic Review* 105, 1408–1438.

Porter, Robert H, 1983, Optimal cartel trigger price strategies, *Journal of Economic Theory* 29, 313–338.

Raskin, Max, and David Yermack, 2016, Digital currencies, decentralized ledgers, and the future of central banking, Discussion paper, National Bureau of Economic Research.

Szabo, Nick, 1997, Formalizing and securing relationships on public networks, *First Monday* 2.

——— , 1998, Secure property titles with owner authority, *Online at http://szabo. best. vwh. net/securetitle. html*.

Tapscott, Don, and Alex Tapscott, 2016, *Blockchain Revolution: How the Technology Behind Bitcoin Is Changing Money, Business, and the World* (Penguin).

Tirole, Jean, 1988, *The theory of industrial organization* (MIT press).

——— , 1999, Incomplete contracts: Where do we stand?, *Econometrica* 67, 741–781.

Turing, Alan Mathison, 1937, On computable numbers, with an application to the entscheidungsproblem, *Proceedings of the London mathematical society* 2, 230–265.

Yermack, David, 2017, Corporate governance and blockchains, Discussion paper, .

Appendix

A Derivations and Proofs

Proof of Proposition 4.1

Proof. In a competitive equilibrium, the sellers lower price until their competitors quit. If $\pi q_C < \max\{q_A, q_B\}$, at least one of the incumbents always competes to lower the price to μ to get the customer this period and prevent the enhanced future competition they face had C entered in this period. Without a reputation of being authentic, C only stands a chance of getting a customer if buyers show up and $\pi q_C \geq \max\{q_A, q_B\}$.

Because C does not have a capacity to bear loss at the point of entry, C cannot charge a penetration price below production cost μ and gets customers when $\pi q_{C,t}\mathbb{I}_t < \max\{q_{A,t}, q_{B,t}\}$. Even when $\pi q_{C,t}\mathbb{I}_t \geq \max\{q_{A,t}, q_{B,t}\}$, C may not be able to enter if the incumbents have deep pocket to do predatory pricing. □

Proof of Proposition 4.2

Proof. Let $V^+(q_i, q_{-i})$ be the present value of payoff to a seller with realized quality q_i in the collusion phase. In the collusion phase, buyers are indifferent between different sellers.

Let V^- be the present value of payoff to a seller before the realization of type in the first period of punishment phase. According to the collusion strategy, the continuation values satisfy:

$$V^+(q_i, q_{-i}) = \lambda(f(q_i, q_{-i})(q_i - \mu) + \delta V^+) + (1-\lambda)\delta V^- \tag{14}$$

$$V^- = \lambda E[(q_i - \max_{j \neq i} q_j)^+]\frac{1-\delta^T}{1-\delta} + \delta^T V^+ \tag{15}$$

For the strategy to be an equilibrium, we need to verify, by one-shot deviation principal, that a seller does not have incentive to unilaterally deviate. This is obvious in the punishment phase, since it is Bertrand equilibrium. In the collusion phase, to prevent deviation, we need

$$\forall q, V^+(q) \geq \lambda((q-\mu) + \delta V^-) + (1-\lambda)\delta V^- \tag{16}$$

Denote $V^+(q_i) = E_{q_{-i}}[V^+(q_i, q_{-i})], f(q_i) = E_{q_{-i}}[f(q_i, q_{-i})]$. Integrate (14), we have

$$V^+(q) = \lambda(f(q)(q-\mu) + \delta V^+) + (1-\lambda)\delta V^- \tag{17}$$

$$V^+ = \lambda(E[f(q)(q-\mu)] + \delta V^+) + (1-\lambda)\delta V^- \tag{18}$$

With (16) -(18), we have

$$\delta(V^+ - V^-) \geq (1 - f(q))(q-\mu), \forall q \in [\underline{q}, \bar{q}] \tag{19}$$

From (18) (15), we solve for $(V^+ - V^-)$, plug into the above equation, and get the range of discount factor

to support the collusion strategy as an equilibrium:

$$\delta\lambda\frac{\left(1-\delta^T\right)\left(M_1-M_2\right)}{1-\lambda\delta-(1-\lambda)\delta^{T+1}} \geq M_3 \tag{20}$$

where $M_1 = E[f(q)(q-\mu)], M_2 = E[(q_i - \max_{j\neq i} q_j)^+], M_3 = \max_q\{(1-f(q))(q-\mu)\}$. □

Proof of Proposition 4.3

Proof. Since M_1 is the expected stage game collusion rent to a seller, and M_2 is her payoff in a competitive stage-game equilibrium, we have $M_1 > M_2$. Moreover, $M_1 + M_3 > \mathbb{E}[q] - \mu$, thus $\frac{1}{\lambda}\frac{M_3}{M_1+M_3-M_2} > \frac{1}{\lambda}\frac{\mathbb{E}[q]-\mu-M_1}{\mathbb{E}[q]-\mu-M_2} > 0$. By the least-upper-bound property (and its implied greatest-lower-bound property) holds, the infimum exists.

□

Proof of Proposition 4.4

Proof. Since the payment can be contingent on completion of service, the authentic type can be separated out from fraudulent type by the following smart contract: the buyer pays the seller p^s conditional on the success of service, otherwise pays zero (or an infinitesimally small negative amount). The fraudulent type can ill-afford imitating the good type, since they cannot complete the service and get the payment anyway. So she does not enter and never gets any customer. For the authentic entrant to get buyers (if present), if $q_C \geq \max\{q_A, q_B\}$, she can charge payment $p^s = \frac{\mu+(q_C-\max\{q_A,q_B\})}{q_C}$ contingent on completion of service, and $-\epsilon$ upon failure, where ϵ is infinitesimal just to break the fraudulent type's indifference (alternatively we can assume a tiny cost for offering the contract and the fraudulent type would not bother to offer since she gets no customer anyway).

Given the smart contract allows authentic C to costlessly separate. A, B, and C are basically competing based on $\mathbf{q}$. Any predatory behaviors would only incur losses for the current period without improving future continuation value as future $\mathbf{q}$ is i.i.d.. Therefore there would not be any predatory (or penetration) pricing.

Finally for collusive equilibria, if A and B collude, they must be charging a weakly higher price, which enables C to first get customer earlier.

□

Proof of Proposition 4.5 and Corollary 4.6

Proof. It is easy to derive,

$$V^+(q_i, q_{-i}) = \lambda(f(q_i, q_{-i})(q_i-\mu) + \delta V^+) + (1-\lambda)\delta V^+ \tag{21}$$

$$V^+ = \lambda(E[f(q)(q-\mu)] + \delta V^+) + (1-\lambda)\delta V^+ \tag{22}$$

$$V^- = \lambda E[(q_i - q_{-i})^+]\frac{1-\delta^T}{1-\delta} + \delta^T V^+ \tag{23}$$

$$\forall q, V^+(q) \geq \lambda((q-\mu) + \delta V^-) + (1-\lambda)\delta V^+ \tag{24}$$

The collusion can be supported if

$$\frac{\delta\lambda\left(M_1-M_2\right)\left(1-\delta^T\right)}{1-\delta}\geq M_3 \tag{25}$$

where $M_1=E[f(q)(q-\mu)], M_2=E[(q_i-q_{-i})^+], M_3=\max_q\{(1-f(q))(q-\mu)\}$

Compared to tacit collusion without blockchain, the only difference in the above recursive equations is that the punishment phase is not trigged if the buyers do not show up, i.e., the corresponding part of the continuation value is $(1-\lambda)\delta V^+$ instead of $(1-\lambda)\delta V^-$.

We show that whenever (8) is satisfied, so is (25). This is equivalent to showing

$$1-\lambda\delta-(1-\lambda)\delta^{T+1}>1-\delta \tag{26}$$

which is equivalent to

$$\delta(1-\delta^T)(1-\lambda)>0 \tag{27}$$

Now for the corollary, note that there cannot be collusion when $\delta<\delta_o^{Traditional}$ is proven in Proposition 4.3.

To show there could be when $\delta\geq\inf_f\{\delta_{(\infty,f)}^{Blockchain2}\}$, we note again by the least upper bound property, $\inf_f\{\delta_{(\infty,f)}^{Blockchain2}\}$ is well-defined and positive. To show one collusion equilibrium exists, we only need to search within the class of f such that $f(q)$ is continuous function, i.e. $f\in\mathcal{C}([0,1])$. Because $\mathcal{C}([0,1])$ is a locally convex Hausdorff space that is complete, there exists a sequence of allocation functions that gets infinitely close to the infimum. This means for any $\delta\geq\delta_o^{Blockchain2}$, we can find a (T,f) that can be sustained. This holds true for our later discussions on infimum and supremum as well. □

Proof of Lemma 4.7

Proof.

$$V^+(q_i,q_{-i})=\lambda(\hat{f}(q_i,q_{-i})(q_i-\mu)+\delta V^+)+(1-\lambda)\delta V^+ \tag{28}$$

$$V^+=\lambda(E[\hat{f}(q)(q-\mu)]+\delta V^+)+(1-\lambda)\delta V^+ \tag{29}$$

$$V^-=\lambda E[(q_i-\max_{j\neq i}q_j)^+]\frac{1-\delta^T}{1-\delta}+\delta^T V^+ \tag{30}$$

$$\forall q, V^+(q)\geq\lambda((q-\mu)+\delta V^-)+(1-\lambda)\delta V^+ \tag{31}$$

Note $V^+-V^-=\frac{\lambda\left(\hat{M}_1-\hat{M}_2\right)\left(1-\delta^T\right)}{1-\delta}$. □

Proof of Proposition 4.8

Proof. We note $\hat{M}_2$ is simply the payoff to a seller in a competitive stage game, and is almost surely less than M_1 which is the expected stage game payoff under collusion. Therefore, $\inf_{\hat{f}}\{\delta_{(\infty,\hat{f})}^{Blockchain3}\}=\inf_{\hat{f}}\frac{\hat{M}_3}{\hat{M}_3+\lambda(\hat{M}_1-\hat{M}_2)}$. But $\frac{\hat{M}_3}{\hat{M}_3+\lambda(\hat{M}_1-\hat{M}_2)}\in(0,1)$ for all $\hat{f}$. Again, by the greatest-lower bound property of real-numbered set, the threshold is well-defined and smaller than 1.

When $\delta>\delta_o^{Blockchain3}$, the blockchain can support at least one collusion that fully extracts consumer surplus (with no punishment phase on equilibrium path). This is so because, again, there is a sequence of

allocation function $\hat{f}$ within in the complete function space $\mathcal{C}[0,1]$ that arbitrarily approaches the infimum. Without blockchain, consumer surplus is never zero as competitive stage game is always on equilibrium path (even with collusion, there has to be punishment on equilibrium path), thus consumer surplus is always positive. The conclusion follows. □

Proof of Corollary 4.9

Proof. For $m > 2$ in general, the previous proposition's proof still applies and $\delta_o^{Blockchain,m} < 1$. For $n \geq 2$, when $\lambda < \frac{n-1}{n}$, we have $\frac{1}{\lambda}\frac{M_3}{M_1+M_3-M_2} > \frac{1}{\lambda}\frac{\mathbb{E}[q]-\mu-M_1}{\mathbb{E}[q]-\mu-M_2} > \frac{1}{\lambda}\frac{\mathbb{E}[q]-\mu-\frac{1}{n}(\mathbb{E}[q]-\mu)}{\mathbb{E}[q]-\mu} > 1$. Therefore there cannot be collusion with $n \geq 2$ in the traditional world. The proposition follows. □

Proof of Theorem 4.10 and Corollary 4.11

Proof. Again, $\frac{\hat{M}_3}{\hat{M}_3+\lambda(\hat{M}_1-\hat{M}_2)} \in (0,1)$ for all $\hat{f}$. Therefore by the least upper bound property, $\sup_f\{\delta_{(\infty,\hat{f})}^{Blockchain3}\}$ exists and is less than 1. When $\delta > \sup_f\{\delta_{(\infty,\hat{f})}^{Blockchain3}\}$, any $(\infty,\hat{f})$ can be sustained, including the one allocating buyers to the highest quality seller and the one allocating to the lowest-quality seller. Note that for any realization of seller qualities, the best-quality (worst-quality) seller with blockchain is better (worse) than the best-quality (worst-quality) incumbent, we could attain higher or lower welfare. Moreover, since competitive stage game is always on equilibrium path without blockchain, consumer surplus is positive. With blockchain we can extract full rent, so lower consumer surplus is attainable. Moreover, by introducing some punishing on equilibrium path or lowering collusion price under blockchain, consumer surplus can be increased all the way to be higher than that in the traditional world (for example, under perfect competition). Thus consumer surplus thus can be higher too with blockchain.

Note for the corollary, the most collusive equilibria maximizes welfare but sellers fully extracts that. This equilibrium can be sustained and the results follows. □

Proof of Lemma 5.2

Proof. The information asymmetry here is that the buyer does not know a seller's type. Therefore the buyer makes his decision based on his perception of the type $\hat{q}_i$ and the price charged p_i. To be specific, the buyer maximizes his payoff by choosing the seller who can deliver the highest expected utility:

$$\max_i \hat{q}_i - p_i \tag{32}$$

If the payoff by choosing any seller is negative, the buyer will step out of the market.

Suppose there is a separating equilibrium where the pricing schedule is $p(q)$ and the probability for a seller with type q to be chosen is $f(q)$. For a seller with type q, she can pretend to be type $\tilde{q}$ by posting the price $p(\tilde{q})$. The seller's expected payoff by doing so is

$$f(\tilde{q})(p(\tilde{q}) - \mu) \tag{33}$$

Every seller will choose the same $\tilde{q}$ to maximize (33), which does not depend on q. Therefore, the separating equilibrium does not exist.

Since there is no separating equilirium, we consider the pooling equilbirium. Without a reputation system, the buyer's perception of each seller's type is the mean $\mathbb{E}[q]$.

This is similar to Bertrand competition. Suppose the lower price of the two firms is higher than μ, say, $p_1 > \mu$. Consider the deviation for the second firm to the price $p_2 = p_1 - \epsilon > \mu$, which increases the profit of the second firm. Therefore, in equilibrium, we must have $p_1 = p_2 = \mu$. Since we always assume the buyer's decision rule is non-discriminating, the tie is broken randomly. Therefore the ex-ante consumer surplus and social welfare is $E[q_i u - \mu]$, where the expectation is taken over the realization of q_i, which yields $\mathbb{E}[q] - \mu$. As a remark outside our parameter assumption, if the cost is so high that $\mu > \mathbb{E}[q]$, the ex-ante utility for buyer is negative, then the buyer will stay out of the market, i.e., the market breaks down. □

Proof of Proposition 5.3

Proof. We first show that using $\mathbb{P}^*$ is an equilibrium. We then prove that no other equilibrium exists. The proof resembles the argument in DeMarzo, Kremer, and Skrzypacz (2005) on how the flattest securities are always used in an equilibrium of informal auctions with security bids. However, because the sellers can always offer quality-insensitive smart contracts, we do not need to worry about equilibrium refinement. Readers who are familiar with DeMarzo, Kremer, and Skrzypacz (2005) should skip the detailed proof below.

With $\mathbb{P}^*$, buyers get utility $1 - p$ regardless of the service outcome; in other words, the smart contract is quality-insensitive. Conversely, any quality-insensitive smart contract has to be of the form $\mathbb{P}^*$. Given that the buyer taking an offer $(p, p-1)$ gets $1-p$ utility, the setup is equivalent to a first-price auction where the buyers are the auctioneers who allocates business opportunity, and sellers are bidders who bid cash $1-p$. The buyers go to the seller with the lowest p. We already know from the auction literature that a unique symmetric equilibrium with cash bids exists. Therefore, there is a unique equilibrium when restricting smart contracts to $\mathbb{P}^*$, implying that there is no profitable deviation using quality-insensitive contracts. The equilibrium offer of type q follows the solution of symmetric equilibrium of first price auctions (Krishna (2009)), and is given by p_q that solves

$$1 - p_q = \mathbb{E}[q^{(1),N-1} - \mu | q^{(1),N-1} < q] = q - \mu - \int_{\underline{q}}^{q} \left[\frac{\Phi(q')}{\Phi(q)}\right]^{N-1} dq' \tag{34}$$

where $q^{(1),N-1}$ is the highest realized quality among other $N-1$ sellers. We note the expression is increasing in q, thus buyers all choose the highest-quality seller. Substituting in $N = 3$ gives the expression in the proposition.

Now suppose this equilibrium breaks down when we allow for smart contracts beyond $\mathbb{P}^*$, then there must be a profitable deviation by a type q to a quality-sensitive smart contract $\mathbb{P}_q$ such that $Pr(B(\mathbb{P}_q))S_q(\mathbb{P}_q) > Pr(B_q(\mathbb{P}^*_q))S_q(\mathbb{P}^*_q)$, where $Pr(B)$ is the probability of getting customers when buyers believe that they can get utility B, and $B(\mathbb{P}_q)$ is the buyers' perceived value of the deviation contract. Denote the set of types that find it profitable to deviate to $\mathbb{P}_q$ by Q, then $B(\mathbb{P}_q) \in B(\mathbb{P}_q(Q))$. Therefore, $\exists q' \in Q$ (possibly q) such that $q' - S_{q'}(\mathbb{P}_q) = B(\mathbb{P}_q(q')) > B(\mathbb{P}_q)$. Consider the deviation by type q' to $(p', p'-1)$, where $p' = 1 - q' + S_{q'}(\mathbb{P}_q)$. Then the probability of winning is higher than q' deviating to use $\mathbb{P}_q$, and the payoff conditional on getting customers are both $S'_q(\mathbb{P}_q)$, implying that if it is profitable for q' to deviate to $\mathbb{P}_q$ (which is true since $q' \in Q$), it is also profitable for q' to deviate to a quality-insensitive contract $(p', p'-1)$. However this contradicts the fact that there is no profitable deviation using quality-insensitive contracts. Therefore we conclude that the equilibrium described in the previous paragraph is an equilibrium even when we allow general smart contract forms.

Next, we show that the above equilibrium is essentially unique for the game, i.e., all other symmetric

equilibria have the same payoffs.

We first argue that if a smart contract $\mathbb{P}$ is offered in an equilibrium and is quality-sensitive, then at most one type uses it. Suppose otherwise and more than one types use it. Let the lowest and highest types offering the smart contract be q_L and q_H, then $B(\mathbb{P}) = B_{q^*}(\mathbb{P}_{q^*})$ for some $q^* \in (q_L, q_H)$. However, $\mathbb{P}$ is increasing in quality because $p^s > p^f$, q_L would find it profitable to deviate to offering $(p, p-1)$ where $p = 1 - B(\mathbb{P})$, contradicting that in equilibrium both q_L offers $\mathbb{P}$. Therefore, at most one type uses $\mathbb{P}$.

Let the type be q, then $B(\mathbb{P}_q) = B_q(\mathbb{P}_q)$. This implies the allocation and payoffs are unaltered if type q replaces the offer by $(p_q, p_q - 1)$ where $p_q = 1 - B(\mathbb{P}_q)$. This is so because, $S_q(\mathbb{P}_q) = q - B_q(\mathbb{P}_q) = q - (1 - p_q)$.

Because each type q is solving the same optimization problem as in the case where we restrict to $\mathbb{P}^*$, we have shown that any unrestricted equilibrium are payoff equivalent to the unique and monotone equilibrium with restriction of smart contracts to $\mathbb{P}^*$.

Finally, the smart contract (p_q^s, p_q^f) used by type q in such an essentially unique equilibrium gives type q the same value as $(p_q, p_q - 1)$, i.e. $qp_q^s + (1-q)p_q^f = qp_q + (1-q)p_q$. Because in the equilibrium with $\mathbb{P}^*$, a seller's expected payoff is differentiable q for all q, by a standard envelope argument, taking derivatives in the unrestricted equilibrium yields $p_q^s - p_q^f = 0$. From this we conclude that all possible equilibria are payoff equivalent to the unique equilibrium when restricting smart contracts to $\mathbb{P}^*$, and the smart contracts used are also in $\mathbb{P}^*$. This basically means that no equilibrium exists other than the one described in the second paragraph of the proof.

□

B Examples of Keepers

Network	Token	Keeper	Function & Objective	Governance Role & Development State
Maker	MKR (gov.) Dai (tx)	Arbitrage seekers, Dai borrowers, CDP closers	Seek arbitrage opportunities and thereby maintain Dai stable.	Indirect – Keepers don't earn MKR directly but the biggest MKR holders are also Keepers.
Polkadot	DOT	Validators, Nominators, Fishermen	Validate transactions and calls across blockchains/smart contracts as well as police validation.	Direct but undefined currently as the appropriate governance mechanism is still being researched.
0x	ZRX	Relayers	Liquidity provision, market making, arbitrage seeking	Direct on-chain stake-voting of protocol updates. Final design and testing
Tezos	XTZ	Miners	Stake-mining and governance of the protocol	Direct on-chain stake-voting of protocol proposals including upgrades and new applications. Final stage of testing.
Dfinity	DFINITIES	Miners	Stake-mining and governance of the protocol through the Blockchain Nervous System (BNS)	Direct on-chain voting with AI inputs that evolves and learns over time.
Filecoin	FIL	Miners	File storage and retrieval considering persistence and latency	No direct governance of the protocol.
Raiden	RDN	Operators	Path-finding, channel monitoring, gateway, etc.	Governance mechanism is still being researched.
Truebit	TRU	Solvers & Challengers	Verifying computation and policing the system	Direct governance but undefined currently as the governance mechanism is still being researched.
1protocol	CRED	Operators	Stake-mining on behalf of Capitalists.	Direct but undefined currently as the governance mechanism is still being researched.
Cosmos	ATOM	Validators	Staking and validating transactions.	Direct on-chain stake-voting of protocol updates. Mid-stage design and testing.
Augur	REP	Reporters	Reporting on the outcome of events.	Indirect (for the time being).
Bitcoin	BTC	Miners	Confirming transactions and securing the network.	Indirect as Miners effectively have some control of the protocol updates (see BTC-BCH).

Table B1. A List of Keepers' Functions

Source: KeepersWorkers that Maintain Blockchain Networks, Ryan Zurrer, *Medium*, Aug 5, 2017. https://medium.com/@rzurrer/keepers-workers-that-maintain-blockchain-networks-a40182615b66