

MIT Open Access Articles

Some simple economics of the blockchain

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Catalini, Christian and Joshua S. Gans. "Some simple economics of the blockchain." Communications of the ACM 63, 7 (June 2020): [dx.doi.org/10.1145/3359552](https://doi.org/10.1145/3359552). © 2020 Owner/Author

Published Version: <http://dx.doi.org/10.1145/3359552>

Publisher: Association for Computing Machinery (ACM)

Permanent Link: <https://hdl.handle.net/1721.1/130500>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: <http://creativecommons.org/licenses/by-nc-sa/4.0/>



Some Simple Economics of the Blockchain

Christian Catalini (MIT) and Joshua S. Gans (University of Toronto)*

April 20, 2019

Abstract

We build on economic theory to discuss how blockchain technology can shape innovation and competition in digital platforms. We identify two key costs affected by the technology: the cost of verification and the cost of networking. The cost of verification relates to the ability to cheaply verify state, including information about past transactions and their attributes, and current ownership in a native digital asset. The cost of networking, instead, relates to the ability to bootstrap and operate a marketplace without assigning control to a centralized intermediary. This is achieved by combining the ability to cheaply verify state with economic incentives targeted at rewarding state transitions that are particularly valuable from a network perspective, such as the contribution of the resources needed to operate, scale, and secure a decentralized network. The resulting digital marketplaces allow participants to make joint investments in shared infrastructure and digital public utilities without assigning market power to a platform operator, and are characterized by increased competition, lower barriers to entry, and a lower privacy risk. Because of their decentralized nature, they also introduce new types of inefficiencies and governance challenges.

Keywords: blockchain, digital platforms, advertising, early adopters, market design, tokens, initial coin offerings, smart contracts, distributed ledgers.

*Christian Catalini is the Theodore T. Miller Career Development Professor at MIT, Associate Professor of Technological Innovation, Entrepreneurship, and Strategic Management at the MIT Sloan School of Management. Founder, MIT Cryptoeconomics Lab: catalini@mit.edu. Joshua S. Gans is a Professor of Strategic Management and holder of the Jeffrey Skoll Chair in Technical Innovation and Entrepreneurship at the Rotman School of Management, University of Toronto: joshua.gans@rotman.utoronto.ca. We are thankful to Al Roth, Muneeb Ali, Naval Ravikant, Nicola Greco, Tim Simcoe, Scott Stern, Catherine Tucker, and Jane Wu for helpful discussions.

1 Introduction

In October 2008, a few weeks after the Emergency Economic Stabilization Act rescued the U.S. financial system from collapse, Satoshi Nakamoto (Nakamoto 2008) introduced a cryptography mailing list to Bitcoin, a peer-to-peer electronic cash system *“based on cryptographic proof instead of trust, allowing any two willing parties to transact directly with each other without the need for a trusted third party.”* With Bitcoin, for the first time, value could be reliably transferred between two distant, untrusting parties without the need of an intermediary. Through a clever combination of cryptography and game theory, the Bitcoin ‘blockchain’ – a distributed, public transaction ledger – could be used by any participant in the network to cheaply verify and settle transactions in the cryptocurrency. Thanks to rules designed to incentivize the propagation of new legitimate transactions, to reconcile conflicting information, and to ultimately agree at regular intervals about the true state of a shared ledger (a ‘blockchain’)¹ in an environment where not all participating agents can be trusted, Bitcoin was also the first platform, at scale, to rely on decentralized, internet-level ‘consensus’ for its operations: Without involving a central clearinghouse or market maker, the platform was able to settle the transfer of property rights in the underlying digital token (bitcoin) by simply combining a shared ledger with an incentive system designed to securely maintain it.

From an economics perspective, this new market design solution provides some of the advantages of a centralized digital platform (e.g. the ability of participants to rely on a shared network and benefit from network effects) without some of the consequences the presence of an intermediary may introduce such as increased market power, ability to renege on commitments to ecosystem participants, control over participants’ data, presence of a single point of failure, etc. As a result, relative to existing financial networks, a cryptocurrency

¹See Appendix for more details.

such as Bitcoin may be able to offer lower barriers to entry for new service providers and application developers, and an alternative monetary policy for individuals that do not live in countries with trustworthy institutions: Key commitments encoded in the Bitcoin protocol are its fixed supply, predetermined release schedule, and the fact that rules can only be changed with support from a majority of participants. While the resulting ecosystem may not offer an improvement for individuals living in countries with reliable and independent central banks, it may represent an option in countries that are unable to maintain their monetary policy commitments. Of course, the open and “permissionless” nature of the Bitcoin network, and the inability to adjust its supply also introduce new challenges, as the network can be used for illegal activity, and the value of the cryptocurrency can fluctuate wildly with changes in expectations about its future success, limiting its use as an effective medium of exchange.

In the paper, we rely on economic theory to explain how two key costs affected by blockchain technology – the *cost of verification* of state, and the *cost of networking* – change the types of transactions that can be supported in the economy. These costs have implications for the design and efficiency of digital platforms, and open opportunities for new approaches to data ownership, privacy, and licensing; monetization of digital content; auctions and reputation systems.

While the reduction in the cost of verification has economic consequences mostly on the intensive margin of production (improving existing applications), on the extensive margin (new applications), the reduction in the cost of networking is more consequential: Bitcoin was the first digital platform to be bootstrapped in a decentralized fashion without resorting to investments by an intermediary or planner. As early adopters and investors experimented with the cryptocurrency in the hope that the network would increase in users, security²

²In a proof-of-work blockchain such as the one used by Bitcoin, the security of the public ledger depends on the amount of computing power that is dedicated to verifying and extending the log of transactions over time (i.e. that is dedicated to “mining”).

and value, the underlying token appreciated, generating the positive feedback loop needed to attract subsequent batches of users. This organic diffusion process uses high-powered incentives similar to the venture capital model to reward early adopters for taking risks and dedicating their time, effort, and capital to a new platform. The same incentive system is now used by startups to raise capital and lower switching costs for the user base and developer community of entrenched digital incumbents. This allows them to compete in a context where network effects are strongly in favor of established players.

Whereas the reduction in the cost of verification is what allows Bitcoin to settle transactions without an intermediary, the reduction in the cost of networking is what allowed its ecosystem to scale in the first place: Within eight years, the digital, scarce token native to Bitcoin went from having no value to a total market capitalization of \$180B,³ and is considered by investors to be part of a new asset class and a novel type of store of value.

Beyond the idiosyncratic market design choices behind Bitcoin, the ability to track transaction attributes, settle trades and enforce contracts across a wide variety of digital assets is what makes blockchain technology a general purpose technology. Entries on a distributed ledger can represent ownership in currency, digital content, intellectual property, equity, information, contracts, financial and physical assets. As a result, the scaling model pioneered by Bitcoin has been adopted by open source projects and startups interested in creating platforms for the exchange of other types of scarce, digital goods. For example, Ethereum used its own token, Ether, to bootstrap a decentralized marketplace for computing power and applications, Filecoin for data storage, BAT for digital advertising, Blockstack for digital identity, etc.

The new types of networks that can be created using the technology challenge the business

³The market capitalization is calculated as the number of tokens in circulation (approximately 16.8M bitcoin) times the value of each token (the Bitcoin to USD exchange rate was \$10,633 in January 2018). The second largest cryptocurrency, Ethereum, had a \$94B market cap (source: <https://coinmarketcap.com/> - accessed 01-22-2018).

models of incumbent digital platforms and financial institutions, and open opportunities for novel approaches to the exchange of digital assets, data ownership and monetization, information licensing, and privacy. Whereas the utopian view has argued that blockchain has the potential to transform every digital service by removing the need for intermediaries, we argue that it is more likely to change the nature of intermediation by reducing the market power of intermediaries, and by progressively redefining how they add value to transactions.⁴ This transformation will unfold slowly because even in sectors that are well-suited for a more decentralized exchange of digital assets such as finance, there are currently substantial legal and regulatory frictions to adoption. While blockchain allows for the costless verification of state when all relevant information is born digital, most markets also rely on external information – including information about identity – to ensure safe and compliant exchanges. As a result, ‘last mile’ frictions limit the conditions under which blockchain-based networks can replace existing infrastructure, as complementary innovations are needed to ensure that the shared data managed through a consensus protocol is kept in sync with critical offline information and events.

The paper proceeds as follows: We first review the literature in Section 1.1. In Section 2, we discuss the effects of the reduction in the cost of verification. Section 3 focuses on the reduction in the cost of bootstrapping and operating a network. Section 4 concludes.

1.1 Literature

The paper contributes to the nascent literature on blockchain by providing an economic framework for understanding how the technology changes the types of transactions and networks that can be sustained in the economy. By focusing on the two key economic

⁴For example, while financial intermediaries are currently able to charge relatively high fees for cross-border payments and remittances, this revenue stream will disappear if blockchain-based payment networks diffuse and commodify the transfer of value. At the same time, this does not mean that intermediaries will not be able to provide and charge fees for added value services on top of a basic payments layer (e.g. fraud protection for merchants, dispute resolution, etc.).

costs the technology influences, the paper abstracts away from some of the idiosyncratic choices different protocols make (e.g. in terms of privacy, consensus algorithms, presence of mining versus not, etc.), and surfaces high-level dimensions that have implications for market structure and competition with existing digital platforms. This level of analysis allows us to highlight commonalities between protocols that may be different at a more fine-grained technical level, but ultimately share a similar trust and competition model, and will thus have a similar impact on how rents are allocated between users, developers and nodes providing resources to a network. The Appendix provides additional technical details on how some of the most popular cryptocurrencies work, and a taxonomy of transactions that the technology can support (e.g. auctions, smart contracts, digital identity and property rights, audit trails etc).

Previous research in this emerging area has focused on providing an overview of Bitcoin and its operations (Böhme, Christin, Edelman, and Moore 2015, Narayanan, Bonneau, Felten, Miller, and Goldfeder 2016); has combined theory and data to explain the velocity of Bitcoin and its use across countries as an investment vehicle, for gambling and illegal online markets (Athey, Parashkevov, Sarukkai, and Xia 2016); and has studied the role early adopters play in the diffusion and use of Bitcoin within a large-scale, field experiment (Catalini and Tucker 2017).

Researchers have also examined competition between alternative cryptocurrencies and their differences (Gandal and Halaburda 2014, Gans and Halaburda 2015, Dwyer 2015, Halaburda and Sarvary 2016); the changes they entail for trading behavior (Malinova and Park 2016); their integration with fiat-based currencies and direct use for providing citizens with central bank money (Raskin and Yermack 2016, Seretakakis 2017, Bordo and Levin 2017) and alternative payment systems (Beck, Czepluch, Lollike, and Malone 2016, Rysman and Schuh 2017); implications for regulation and governance (Wright and De Filippi 2015, Davidson, De Filippi, and Potts 2016, Kiviat 2015, Walport 2016); and the privacy trade-offs cryp-

tocurrencies and digital wallets introduce for consumers (Athey, Catalini, and Tucker 2017).

From a business perspective, scholars have compared the transformation brought about by blockchain to the introduction of communication protocols such as TCP/IP (Iansiti and Lakhani 2017, Ito, Narula, and Ali 2017), and have explored applications to digital platforms beyond finance and implications for the boundaries of the firm (Catalini 2017a, Catalini 2017b).

2 Cost of Verification

Markets facilitate the voluntary exchange of goods and services between buyers and sellers. For an exchange to be executed, key attributes of a transaction need to be verified by the parties involved. When an exchange takes place in person the buyer can usually directly assess the quality of the goods, and the seller can verify the authenticity of the cash. The only intermediary involved in this scenario is the central bank issuing and backing the fiat-currency used in the exchange. When a transaction is performed online instead, one or more financial intermediaries broker it by verifying, for example, that the buyer has sufficient funds. Intermediaries add value to marketplaces by reducing information asymmetry and the risk of moral hazard through third-party verification. This often involves imposing additional disclosures, monitoring participants, maintaining trustworthy reputation systems, and enforcing contractual clauses. As markets scale in size and geographic reach, verification services become more valuable, as most parties do not have preexisting relationships, but rely on intermediaries to ensure the safety of transactions and enforce contracts. In the extreme case where verification costs are prohibitively high, markets unravel, and beneficial trades do not take place.

In exchange for their services, intermediaries typically charge a fee. This is one of the costs buyers and sellers incur when they cannot efficiently verify all the relevant transaction

attributes by themselves. Additional costs may stem from the intermediary having access to transaction data (a privacy risk), and being able to select which transactions to execute (a censorship risk).

These costs are exacerbated when intermediaries gain market power, often as a result of the informational advantage they develop over transacting parties through their intermediation services (Stiglitz 2002). Transacting through an intermediary always involves some degree of disclosure to a third-party, and increases the chance that the information will be later reused outside of the original contractual arrangement. Moreover, as an increasingly large share of economic and social activity is digitized, keeping data secure has become more problematic and information leakage more prevalent. Classic examples are the theft of social security numbers (e.g. Equifax hack) and credit card data (e.g. Target data breach), or the licensing of customer data to advertisers. Blockchain technology can prevent information leakage by allowing market participants to verify transaction attributes and enforce contracts without exposing the underlying information to a third-party.⁵ This allows an agent to verify that some piece of information is true (e.g. good credit standing), without full access to all background information (e.g. past transaction records): i.e., the technology allows for the verification of transaction attributes in a privacy-preserving way.

Digitization has pushed verification costs for many types of transactions close to zero. When the relevant information is digital, blockchain technology contributes to this process by allowing for *costless verification*.⁶ Of course, at the interface between an offline record and its digital representation blockchain applications still face substantial frictions and “last mile” costs (Tucker and Catalini 2018). This explains why, despite claims by technology enthusiasts about the value of using the technology across a variety of applications including

⁵This is achieved by combining a distributed ledger with zero-knowledge cryptography. Examples include cryptocurrencies such as Zcash and Zcoin.

⁶In practice, verification costs will never be exactly zero. What we mean by ‘costless’ is low enough to be irrelevant from an economic perspective relative to the value of the transaction.

supply chain monitoring and digital identity, use cases outside of cryptocurrency and fintech (settings where key information and assets are digital) have been extremely limited. The link between online “on-chain” activities recorded on a blockchain and offline “off-chain” events introduces major challenges which cannot be overcome without complementary innovations. For example, a blockchain such as the Bitcoin one can be used to cheaply verify ownership and exchanges of its native digital asset. While this technically allows anyone to send and receive bitcoin globally without using an intermediary or being censored, actually being able to spend bitcoin to buy goods and services offline still runs into last mile issues. Hence, while Bitcoin has been used in countries with hyperinflation to escape devaluation, its use as a medium of exchange has been limited, and governments can still shape how these digital assets are used at the interface between the digital and the physical world. Similarly, information about identity is often used to increase the safety of market interactions, reduce fraud and build robust digital reputation systems, but being able to link an online action and digital record on a blockchain to an offline individual or entity is as expensive with blockchain technology as it would be with more traditional solutions. This drastically limits the benefits blockchain and smart contracts can bring in the absence of complementary technology (e.g. a tamper-proof GPS sensor), firms and institutions that can help ensure that the digital records are accurate to begin with.

The high-level process of verification is described in Figure 1: When a digital transaction is born, it immediately inherits some basic attributes, such as the fact that it exists and when it was created, information about the seller and buyer involved and their credentials, etc. We typically rely on these attributes to perform subsequent actions (e.g., once funds are transferred, the seller may ship the goods). Some of these actions take place every time (e.g. settlement), whereas others are only triggered by specific events. A particularly interesting subset of future events are those that require additional verification. For example, a problem may emerge, and transaction attributes may need to be checked through an audit. The

audit could range from actual auditors accessing the relevant logs or requesting additional information from market participants, to the execution of an internal process designed to handle the exception. Such processes tend to be costly, may involve labor and capital, and may require a third-party to mediate between buyer and seller. The ideal outcome of an audit is the resolution of the problem that emerged.




Figure 1: Costly Verification Through an Intermediary (Audit) versus Costless Digital Verification on a Blockchain

Blockchain technology affects this flow by allowing, when a problem emerges, for the costless verification of digital information. Any transaction attribute or information on the agents and goods involved that is stored on a distributed ledger can be cheaply verified, in real time, by any market participant. Trust in the intermediary is replaced with trust in the underlying code and consensus rules.⁷ These rules define how a distributed network reaches agreement, at regular intervals, about the true state of the shared data it needs to maintain to operate a well-functioning marketplace. At a minimum, such shared data can represent

⁷If we think of the audit capability of the third-party that intervenes when a problem emerges in a traditional market as surveillance or monitoring, blockchain technology can deliver “sousveillance” (Mann, Nolan, and Wellman 2002), i.e. an audit that is embedded within the rules of the marketplace.

past transactions and outstanding balances in an underlying, cryptographic token (i.e. it could be a snapshot of the ownership rights in the token). In more complex applications, the shared data can also cover the rules and data required to perform a specific operation (i.e. to run an application, verify that a contract clause is enforced). These operations, often referred to as ‘smart contracts’,⁸ can be automated in response to new events, adding flexibility to the verification process. For example, on a shared ledger used to exchange financial assets, transacting institutions can agree, ex-ante, on the rules for the settlement and reconciliation of trades, as well as on the process they will follow and third-parties they will involve if an audit is necessary or a dispute emerges. Trusted, independent oracles can also be incorporated to ensure that such financial contracts can respond to market conditions and new information (e.g. to implement a weather derivative, a smart contract can aggregate information across multiple weather sources to assess if a payout has to be made).

As with past improvements in information and communication technology, reductions in the cost of verification enable the unbundling of services that were previously offered together, as part of the steps traditionally performed by an intermediary can now be delivered through a shared ledger. This allows these steps to be collectively owned and managed by a broader group of ecosystem stakeholders, in a way that resembles collaboration among competitors and complementors in standard setting organizations (Bekkers, Catalini, Martinelli, Righi, and Simcoe 2019) or open source foundations. The effects of this change have been mostly

⁸In 1996, Nick Szabo defined smart contracts as: *“The basic idea of smart contracts is that many kinds of contractual clauses (such as liens, bonding, delineation of property rights, etc.) can be embedded in the hardware and software we deal with, in such a way as to make breach of contract expensive (if desired, sometimes prohibitively so) for the breacher. A canonical real-life example, which we might consider to be the primitive ancestor of smart contracts, is the humble vending machine. Within a limited amount of potential loss (the amount in the till should be less than the cost of breaching the mechanism), the machine takes in coins, and via a simple mechanism, which makes a beginner’s level problem in design with finite automata, dispense change and product fairly. Smart contracts go beyond the vending machine in proposing to embed contracts in all sorts of property that is valuable and controlled by digital means. Smart contracts reference that property in a dynamic, proactively enforced form, and provide much better observation and verification where proactive measures must fall short.”*Source: http://www.fon.hum.uva.nl/rob/Courses/InformationInSpeech/CDROM/Literature/LOTwinterschool2006/szabo.best.vwh.net/smart_contracts_2.html

felt on the intensive margin of production (i.e. on improving the efficiency of pre-existing use cases), as firms are experimenting with moving different types of transactions to blockchain-based systems to reduce settlement and reconciliation costs.⁹

As a consequence, applications resulting from the reduction in the cost of verification have been complementary to incumbents, as they improve existing value-chains by lowering the cost of tracking ownership and trading digital assets without reducing the market power of existing players. Furthermore, even when verification can be automated, intermediaries can still add value and retain influence over a market by supporting regulatory compliance, market safety, handling edge cases (e.g. a chargeback), and certifying information that requires labor-intensive, offline forms of verification. This explains why implementations of the technology targeted at identity and provenance have been slower to diffuse: While the verification of digital attributes can be cheaply implemented on a blockchain, the initial mapping between offline events and their digital representations is still costly to bootstrap and maintain. Therefore, as digital verification costs fall, key complements to it that can improve the process of offline verification become more valuable.

On one extreme, blockchain technology can be used to settle trades of digital assets that are completely self-contained within a shared ledger (e.g. bitcoin, ether). The consensus rules established in the code define how tokens are created and earned, and how the network reaches agreement about the true state of ownership over time.¹⁰ The cost of verifying transaction attributes and enforcing simple contracts for self-contained tokens can be extremely low. This is what allows for value to be transferred through Bitcoin across the globe at

⁹Ripple, for example, uses the technology to allow for cross-border payments between existing financial institutions; Digital Asset Holdings is using the same approach to enable more efficient trades of financial assets; Western Union has invested in the space to lower costs in the remittances market; NASDAQ deployed a solution to track equity in privately held companies; the state of Delaware is moving incorporation data on a distributed ledger, etc.

¹⁰Changes in the rules are implemented through a voting process similar to standard setting negotiations, and disagreements can lead to part of the network forking the codebase and current state of the shared ledger to launch a competing platform with different market design rules.

a relatively low cost. Of course, compliance with Know-Your-Customer (KYC) and Anti-Money-Laundering (AML) rules may require individuals and firms to sustain additional costs to credibly link their offline identities with their Bitcoin ones, but as long as individuals agree that the underlying token has value, using it as a store of value and medium of exchange is possible. Similarly, a native crypto token can be used to facilitate low cost transactions of digital resources such as computation (Ethereum), data storage (Filecoin), bandwidth or electricity, as in all these cases verifying the exchange of a resource is not too expensive.

On the other extreme, when entries on a shared ledger are digital representations of offline identities, products, services and related transactions, costless verification is difficult to achieve. Under this scenario, the reduction in the cost of verification is contingent on maintaining a credible link between offline events and their online record. This link is cheaper to establish when offline attributes are easy to capture and expensive to alter or fake: e.g., in the case of diamonds, Everledger uses the physical properties of the gems as a digital fingerprint that can be recorded and tracked on a blockchain as the products move through the supply chain. In many cases, maintaining a robust link between offline events and distributed ledgers is very expensive, and may require not only one or more trusted intermediaries, but also multiple parties to agree on rules for secure data entry and sharing. In the absence of a strong link between offline and online events, asymmetric information and moral hazard will be an issue in these markets. In this context, Internet of Things devices are instrumental in expanding the set of contracts that can be automated on a blockchain because they can be used to record real-world information (e.g. through sensors, GPS devices, etc.) and substitute labor intensive verification with inexpensive hardware.

Overall, when last mile problems are limited – such as in the case of digital assets that are native to a blockchain – decentralized verification goes from being costly, scarce and prone to abuse, to being cheap and reliable. While this process is unlikely to be more efficient on a per transaction basis than verification through a centralized intermediary, the ability to

perform it without trusting a third-party can lead to savings from increased competition, the absence of centralized control, higher privacy and censorship resistance, and the removal of single points of failure. At the same time, when frictions between offline events and their digital representations are high, these improvements are unlikely to materialize in the absence of complementary innovations, as intermediaries will still be able to control key existing complements to digital verification and use them to exert influence over market participants.

As decentralized verification becomes cheaper, the scale at which it can be efficiently implemented also drops: On a distributed ledger, data integrity can be built, from the ground up, from the most basic transaction attributes to more complex ones. For example, a robust reputation system can be constructed from the full set of interactions an economic agent has throughout the economy, increasing transparency and accountability. Expensive audits and due diligence can be progressively substituted with more frequent and fine-grained verification to ensure market safety and reduce the risk of moral hazard. A lower cost of verification also makes it easier to define property rights at a more granular scale than before, as any digital asset (or small fraction of it) can be traded, exchanged or tracked at a low cost on a shared ledger.¹¹

3 Cost of Networking

The ability to verify state (e.g. the current ownership status of a digital asset) at a lower cost because of the reduction in the cost of verification allows a blockchain protocol to not only reach consensus about the history and proposed evolution of a digital asset, but also to define rules for state transitions that are particularly valuable from a network perspective.

¹¹In the same way that Twitter, because of the 140 character limitation, enabled new forms of communication, the ability to implement costless verification at the level of a single piece of information has the potential to change how information markets, digital property rights and micropayments are designed.

These transitions can be used to reward participants for performing actions that accelerate adoption and increase network value and welfare. For example, the protocol can be used to incentivize behavior that builds network effects (both in terms of users and applications), ensures the network has sufficient resources available to meet demand, guarantees its security, encourages savings or spending behavior, etc. Taken together, these incentives lower the cost of networking, i.e. the cost of bootstrapping, operating and scaling an economic network.

Whereas a reduction in the cost of verification is a necessary condition for a reduction in the cost of networking – as it is the ability to verify state that allows economic agents to establish property rights on network resources and define incentives without relying on an intermediary – it is not a sufficient condition, as implementations can take advantage of the former without the latter. In particular, when a blockchain protocol is permissioned and the entities developing it retain control over which participants can update and verify state, transitions are not fully defined by code and self-contained within the system, but rather can be influenced by external parties through fiat. As a result, from an economics perspective, the network will operate under constraints similar to those of traditional digital platforms, and participants will have to trust the platform architect and core constituents through formal and relational contracts, past reputation, etc. This tension is an important one from an organizational perspective, as it determines if a blockchain network can be considered a novel organizational form versus not (Catalini and Boslego 2019).

A permissionless blockchain protocol, instead, allows a network of economic agents to agree, at regular intervals, on the true state of a set of shared data without assigning residual rights to trusted entities. The flexibility in terms of what such shared data represents across settings (e.g. currency, intellectual property, financial assets, contracts, etc.) makes it a general purpose technology (GPT). GPTs typically take a long time to diffuse through the economy, but also lead to productivity gains across multiple industries (Bresnahan and Trajtenberg 1995, Helpman 1998, Rosenberg and Trajtenberg 2001, Moser and Nicholas

2004). Classic examples of GPTs include the steam engine, electricity, and the internet. While permissionless networks have been compared to communication protocols such as TCP/IP – which focus on how *information* is packetized and routed through the internet – they fundamentally differ from them because they allow for the secure provision, transfer and enforcement of *property rights*. On these networks, trust in a platform operator is replaced by trust in the underlying incentives, code and consensus rules. As a result, market power of the intermediary, privacy risk and censorship risk can be potentially reduced. The switch in the trust model also introduces new challenges, as bugs in the code can leave participants with little recourse beyond trying to coordinate a hard fork of the network. Issues with the new trust model have resulted from benign programming mistakes (such as the Parity wallet library removal),¹² from deliberate attempts at defrauding investors by promising high returns in the absence of any real technical or business plan (as in the case of fraudulent initial coin offerings), as well as from malicious attacks (such as the DAO hack, which led to a split of the Ethereum network).¹³ Similarly, while blockchain protocols can be designed to offer participants a high degree of privacy (e.g. Zk-Stark, Zcash, Monero, etc.), and users can take additional measures to protect their privacy from the public (e.g. using a mixing service, not reusing addresses), many shared ledgers such as the Bitcoin one are pseudonymous,¹⁴ allowing third-parties to deanonymize transactions and trace movements of funds over time.

Whereas permissioned networks only take advantage of the reduction in the cost of verification, permissionless ones build on the first by adding a self-contained incentives system to also deliver a decrease in the cost of launching and operating a network without relying on trusted intermediaries. The effects of this reduction in the cost of networking are felt both in the phase of bootstrapping a new platform, and in the phase of operating it. In

¹²See <https://www.parity.io/a-postmortem-on-the-parity-multi-sig-library-self-destruct/>

¹³See <https://www.bloomberg.com/features/2017-the-ether-thief/>

¹⁴Like a writer writing a book under a pseudonym, if a Bitcoin user is ever tied to a specific address, the entire history of her transactions with that address can be read on the public Bitcoin blockchain.

the first phase, a native token can be used to create incentives for adoption and to fund the development and scaling of the network, for example by having mining rewards or by raising capital through an initial coin offering (ICO). In the second phase, market design is used to define the conditions under which participants can earn tokens for contributing resources to the network (e.g. computing power in the case of Bitcoin, computing and applications for Ether, disk storage for Filecoin, digital content and advertising in the case of the Basic Attention Token, etc.).

Since during the bootstrapping phase the actual utility the network can deliver to users is limited by its small scale, and network effects work against users switching from existing alternatives, this phase relies on contributions from early adopters and investors with positive expectations about the future value of the network. As in open source projects (Von Hippel 2002, Von Hippel and Von Krogh 2003, Von Hippel 2005), early adopters may be willing to dedicate time and effort to support a new network because they want to create a viable alternative to established products or they derive utility from advancing the underlying technology (e.g. consumption utility from early access, from working on novel, complex problems, job-market signalling). Investors, instead, as in traditional equity finance, may come in early because they expect the token to appreciate in value and reward their investment (Catalini and Gans 2018). Of course, individuals can be simultaneously early adopters and investors, and contribute both effort and capital to these projects. For this set of individuals, the presence of a native token serves a similar purpose to founder and early-employee equity in startups, and allows these projects to attract talent without raising investment from traditional angels and venture capitalists. Since it only takes a few lines of code to write a smart contract for an initial coin offering, open source codebases can be forked or imitated at a low cost, and regulation is still uncertain in many jurisdictions, the ability to profit from launching a new cryptocurrency or manipulating its trading have attracted a large number of bad actors and speculators. While lower entry barriers and the

presence of technical investors could in theory open up capital for new types of entrepreneurs and ideas that traditional investors may be more reluctant to fund, the absence of regulation and oversight also allows fraudulent projects to blend in with legitimate ones and raise capital from unsophisticated investors. Combined with the fact that the value of a new token is, in most cases, purely based on expectations about its future success, and that such expectations, because of technical, regulatory and market uncertainty can rapidly turn when new information emerges or sentiment evolves, the valuations of cryptocurrencies have been extremely volatile. The resulting turmoil and speculative bubbles have made it more difficult for investors to identify high quality projects and teams, have attracted speculators and low quality entrants, and have shifted attention from technology R&D to short-term speculative returns.

If in the first phase of growth of a blockchain-based network incentives are predominantly targeted at accelerating adoption, in the second phase the key challenges from a market design perspective are ensuring that the incentives continue to support contributions of key resources to the ecosystem, and avoiding a tragedy of the commons. By design, the protocol layer is a shared resource among all network participants, and everyone benefits from investments in it – from better security to removing technical constraints on throughput, latency or liveness. At the same time, because of the public good nature of these improvements, in the absence of proper governance, a blockchain-based network may fail to invest enough resources on them. From a valuation perspective, whereas the bootstrapping phase of a new token is associated with extremely high volatility, as uncertainty around a network's potential is resolved, it should enter a more stable growth trajectory.¹⁵

Overall, relative to blockchain implementations that only take advantage of the reduction

¹⁵This is similar to the process of early-stage startup funding and growth. Within the cryptocurrency space, token sales and ICOs have shortened the time it takes for a developer team to raise capital. For instance, while Bitcoin – which was bootstrapped through a slow process of word-of-mouth within the cypherpunk community – took four years to reach a \$1B market cap, Ethereum took only two years to reach the same milestone.

in the cost of verification (e.g. permissioned networks), those that also benefit from the reduction in the cost of networking (e.g. permissionless ones) are different on at least four dimensions. First, they are less likely to leave market power in the hands of their founders or early participants. This limits the ability of any party to unilaterally censor transactions or exclude participants from the network, and removes single points of failure, as the network does not depend on the availability of one or a few key players to operate.¹⁶

Second, they are less reliant on off-chain governance, relational contracts and laws to support their operations, as by design, to take advantage of the lower cost of networking they need to embed as much as possible of the incentives and governance rules required for their operations into the protocol. Of course, permissionless networks still need off-chain governance and coordination between their key stakeholders to execute a hard fork, implement controversial changes, or respond to an attack, but relative to more closed networks that rely on trusted intermediaries they leave less discretion to any single party, and end up codifying more of their rules into their codebases.

Third, they involve a lower privacy risk, as no single entity (or group of entities) has preferential access to or visibility over the information generated by the network.¹⁷ In traditional platforms, the privacy risk is particularly salient in markets where consumers pay for services by allowing intermediaries to access and monetize their data, an issue that is increasingly relevant because of the role such data can play in the training of AI algorithms (Agrawal, Gans, and Goldfarb 2016). Whereas the trend of consumers relinquishing private

¹⁶The censorship risk is visible when an intermediary *revokes* a participant's access to the marketplace and digital assets through fiat; when it *degrades* access (e.g. in terms of speed, features, etc.) to some participants to reduce their ability to effectively compete; and when it *loses* control over the marketplace because of an attack or technical failure. All three cases have been observed in online platforms, which are heavily concentrated markets because of network effects and economies of scale in data collection, storage, and processing. This not only gives market power to a small number of entrenched incumbents, but also makes the underlying services less resilient to targeted attacks and errors.

¹⁷Although privacy may still be a concern if the public nature of the shared data exposes critical information about market participants and their transactions. For example, third-parties can use data from the public Bitcoin ledger to deanonymize transactions and identify entities (Athey, Catalini, and Tucker 2017).

information in exchange for free or subsidized digital services is unlikely to change because of blockchain technology – as small incentives and frictions can be used by digital platforms to persuade even privacy sensitive individuals to relinquish sensitive information (Athey, Catalini, and Tucker 2017) – startups in this space are experimenting with approaches that give users greater control over how, when and why their private data is accessed and monetized.

Fourth, blockchain implementations that take advantage of the lower cost of networking inevitably induce architectural changes in how firms create and capture value within markets. Architectural innovations, by destroying the usefulness of the assets and accumulated knowledge of incumbents (Henderson and Clark 1990), open opportunities for entrants to reshape the dimensions firms compete on, and experiment with new business models. In particular, by allowing for the separation of some of the benefits of network effects from the costs of market power – since even in the absence of a platform architect participants in a blockchain network are able to rely on shared infrastructure – the technology offers new ways to reward contributors, allocate rents in a marketplace, and build applications on top of shared data while preserving the privacy of the underlying information. In traditional digital marketplaces, platform operators have wide visibility over all interactions that take place on their networks, and users are unable to directly custody or control the digital assets they use or create while transacting on them. This is a direct result of the inability of these systems to generate and trade scarce, digital assets and establish digital property rights without also assigning control over them to a third-party (usually the platform operator). Before Bitcoin, for example, a central clearing house of some type was necessary to prevent the copying and double spending of digital cash. Bitcoin solves this problem by allowing users to self-custody digital tokens and exchange them without relinquishing control over them to a third-party. This reduces switching costs between digital wallets, and offers users a higher degree of privacy from service providers. Interestingly, while blockchain technology provides individuals and organizations with the opportunity to self-custody and exchange digital assets with-

out the need for traditional intermediaries such as banks, significant work is needed before users can reap the full benefits of this change – such as greater privacy, higher portability between service providers, and increased competition – as many implementations lack the convenience and usability of the centralized solutions consumers are used to. For example, while Bitcoin users can store and protect their own private keys, a large number of them rely on third-party wallets to do so, essentially trusting these entities with their funds as in traditional systems.

4 Conclusions

The paper focuses on two key costs that are affected by blockchain technology: the cost of verification, and the cost of networking. For markets to thrive, participants need to be able to efficiently verify and audit transaction attributes, including for example, the credentials and reputation of the parties involved, characteristics of assets exchanged, and external events and information that have implications for contractual arrangements.

Outside the boundaries of an organization, this is typically achieved by relying on trusted intermediaries. In exchange for their services, intermediaries charge fees and capitalize on their ability to observe all transactions taking place within their marketplaces. This informational advantage, combined with network effects and economies of scale, gives them substantial market power and control over market participants. Consequences of market power include higher prices, user lock-in and high switching costs, the presence of single points of failure, censorship risk, barriers to innovation, and reduced privacy.

Blockchain technology, by reducing the costs of running decentralized networks of exchange, allows for the creation of ecosystems where the benefits from network effects and shared digital infrastructure do not come at the cost of increased market power and data access by platform operators. This reduction in the cost of networking has profound con-

sequences for market structure, as it allows open-source projects and startups to directly compete with entrenched incumbents through the design of platforms where the rents from direct and indirect network effects are shared more widely among participants (e.g. users, application developers, investors, etc.), and no single entity has full control over the underlying digital assets.

Because of the absence of a central clearing house or market maker, these novel networks, when permissionless, exhibit low barriers to entry and innovation. As long as applications are compatible with the rules of the protocol, they can be deployed without permission from other participants, and compete for market share. This reduces the expropriation risk application developers face when building on top of traditional digital platforms. Furthermore, since contributors can participate in governance in a way that is often proportional to their stake in the system, these networks can democratically evolve over time to accommodate changes that are beneficial to the majority of their constituents.¹⁸

From a talent acquisition perspective, unlike open source projects, the digital platforms built on top of crypto tokens do not have to rely solely on pro-social contributions of time and labor and job market signalling (Lerner and Tirole 2002) to support their development. Using a native token, they can directly incentivize early contributions by developers, investors and early adopters. This novel source of funding combines crowdfunding with the simultaneous crowdsourcing of key resources needed to scale a platform and attract both developer and user activity on to it. Because of the reduction in the cost of verification, this model also allows for equity in the system to be defined at a much narrower scale, and to be allocated to a wider population of participants in response to verifiable contributions of resources.

Similarly, by allowing for the definition of scarce digital property rights, native tokens

¹⁸Minorities that disagree with a proposed change face reduced lock-in and hold-up risk because they can fork the existing codebase at any time and launch a separate, backwards-compatible platform under their preferred rules. At the same time, since forks introduce uncertainty and a network split may decrease the overall value of a platform, off-chain governance is needed to allow communities of developers to reach agreement about fundamental changes to market design without destroying confidence in the network.

allow decentralized networks of exchange to coordinate activity around shared objectives, and transact digital resources without assigning market power to a market maker. Through blockchain-based networks, individuals and organizations can source ideas, information, capital and labor, and enforce contracts for digital assets with substantially reduced frictions. These changes allow for the design of novel types of networks that blend features of competitive markets with the more nuanced forms of governance used within vertically integrated firms and online platforms.¹⁹

Whereas intermediaries will still be able to add substantial value to transactions by focusing on tasks that are complementary to digital verification (e.g. secure recording of offline events, curation, certification of identity and services, etc.), they are likely to face increased competition because of the ability to establish and exchange digital assets on decentralized open networks without them. This challenges some of their revenue sources and reduces their influence over markets, opening up opportunities for new business models and novel approaches to data privacy, ownership and portability, as well as to the regulation of networks that should be considered public utilities. By reducing barriers to entry within sectors that are currently heavily concentrated because of network effects and control over data, the technology may enable a new wave of innovation in digital services, and greater consumer choice.

For these changes to materialize, however, substantial hurdles will have to be overcome. First, the technology will need to reach a level of performance (e.g. throughput, latency, cost per transaction, etc.) comparable to traditional networks. While decentralization inevitably comes at a cost, the gains from greater competition, openness, privacy and censorship resistance will have to outweigh the lower efficiency of blockchain networks to make adoption

¹⁹For example, the hedge fund Numerai uses smart contracts to transparently reward contributions to its financial prediction model: Data scientists across the globe can collaborate through this new type of network knowing that their inputs will be rewarded – using the native Numeraire token – according to their impact on the performance of the fund.

worthwhile. Hybrid networks that embrace key features of permissionless systems – such as low barriers to entry and a competitive market for resources and applications – while initially borrowing trust from existing institutions to overcome scaling problems, may also provide a viable transition path when performance is an obstacle to adoption.

Second, regulatory frameworks will have to evolve to reduce uncertainty for founders and network participants, and to provide stronger protections for investors and early adopters. Because of their similarities but also their differences with equity finance (Catalini and Gans 2018), crypto tokens lend themselves to both legitimate fundraising activity by high quality entrepreneurs, as well as flagrant abuse by fraudsters (Catalini, Boslego, and Zhang 2018). As in other technological bubbles, this constitutes a challenge for the space, as investors have a difficult time separating projects worth supporting from the much larger number of low quality imitators, and entry by speculators has brought extreme price volatility and additional risks to the market.

Third, and possibly most important, blockchain technology, like other technological advancements, is not a panacea for every possible technical and market challenge a digital ecosystem may face. As discussed throughout the paper, the technology can add substantial value under fairly narrow conditions: 1) when last mile problems are not severe and digital verification can be implemented in a novel or more fine-grained way because of a reduction in the cost of verifying state without assigning control to an intermediary; 2) when the reduction in the cost of networking allows participants to allocate rents from a digital platform more efficiently between users, developers, and investors; 3) when the combination of a reduction in both costs (verification and networking) allows for the definition of new types of digital assets and property rights; 4) when there is a need for greater privacy and ability for users to control when and how their data is accessed and used. When none of these conditions are met instead, more centralized solutions that rely on traditional intermediaries and relational contracts are unlikely to be replaced, as the benefits of transitioning to a blockchain-based

system are unlikely to counterbalance the costs introduced by a decentralized infrastructure and governance, and the replication of state across the network.

References

- AGRAWAL, A., J. GANS, AND A. GOLDFARB (2016): “The simple economics of machine intelligence,” *Harvard Business Review*, 17.
- ATHEY, S., C. CATALINI, AND C. TUCKER (2017): “The Digital Privacy Paradox: Small Money, Small Costs, Small Talk,” *National Bureau of Economic Research Working Paper*.
- ATHEY, S., I. PARASHKEVOV, V. SARUKKAI, AND J. XIA (2016): “Bitcoin pricing, adoption, and usage: Theory and evidence,” .
- AUSUBEL, L. M., P. MILGROM, ET AL. (2006): “The lovely but lonely Vickrey auction,” *Combinatorial auctions*, 17, 22–26.
- BECK, R., J. S. CZEPLUCH, N. LOLLIKE, AND S. MALONE (2016): “Blockchain-the Gateway to Trust-Free Cryptographic Transactions.,” in *ECIS*, p. ResearchPaper153.
- BEKKERS, R., C. CATALINI, A. MARTINELLI, C. RIGHI, AND T. SIMCOE (2019): “Disclosure rules and declared essential patents,” Discussion paper, National Bureau of Economic Research.
- BÖHME, R., N. CHRISTIN, B. EDELMAN, AND T. MOORE (2015): “Bitcoin: Economics, technology, and governance,” *The Journal of Economic Perspectives*, 29(2), 213–238.
- BORDO, M. D., AND A. T. LEVIN (2017): “Central Bank Digital Currency and the Future of Monetary Policy,” *National Bureau of Economic Research Working Paper*.
- BRESNAHAN, T. F., AND M. TRAJTENBERG (1995): “General purpose technologies Engines of growth?,” *Journal of econometrics*, 65(1), 83–108.
- CATALINI, C. (2017a): “How Blockchain Applications Will Move Beyond Finance,” *Harvard Business Review*.
- (2017b): “How Blockchain Technology Will Impact the Digital Economy,” *Oxford Business Law Blog*.
- CATALINI, C., AND J. BOSLEGO (2019): “Blockchain Technology and Organization Science: Decentralization Theatre or Novel Organizational Form?,” *MIT Working Paper*.
- CATALINI, C., J. BOSLEGO, AND K. ZHANG (2018): “Technological Opportunity, Bubbles and Innovation: The Dynamics of Initial Coin Offerings,” *Working Paper*.
- CATALINI, C., AND J. S. GANS (2018): “Initial coin offerings and the value of crypto tokens,” Discussion paper, National Bureau of Economic Research.
- CATALINI, C., AND C. TUCKER (2017): “When early adopters don’t adopt,” *Science*, 357(6347), 135–136.

- DAVIDSON, S., P. DE FILIPPI, AND J. POTTS (2016): “Economics of blockchain,” *Working Paper*.
- DWYER, G. P. (2015): “The economics of Bitcoin and similar private digital currencies,” *Journal of Financial Stability*, 17, 81–91.
- EDELMAN, B., M. OSTROVSKY, AND M. SCHWARZ (2007): “Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords,” *American Economic Review*, 97(1), 242–259.
- GANDAL, N., AND H. HALABURDA (2014): “Competition in the Cryptocurrency Market,” *NET Institute Working Paper*.
- GANS, J. S., AND H. HALABURDA (2015): “Some economics of private digital currency,” in *Economic Analysis of the Digital Economy*, pp. 257–276. University of Chicago Press.
- HALABURDA, H., AND M. SARVARY (2016): *Beyond bitcoin: The economics of digital currencies*. Springer.
- HELPMAN, E. (1998): *General purpose technologies and economic growth*. MIT press.
- HENDERSON, R. M., AND K. B. CLARK (1990): “Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms,” *Administrative science quarterly*, pp. 9–30.
- IANSITI, M., AND K. R. LAKHANI (2017): “The Truth About Blockchain,” *Harvard Business Review*, 95(1), 118–127.
- ITO, J., N. NARULA, AND R. ALI (2017): “The Blockchain Will Do to the Financial System What the Internet Did to Media,” *Harvard Business Review*.
- KIVIAT, T. I. (2015): “Beyond Bitcoin: Issues in Regulating Blockchain Transactions,” *Duke LJ*, 65, 569.
- LERNER, J., AND J. TIROLE (2002): “Some simple economics of open source,” *The journal of industrial economics*, 50(2), 197–234.
- LUCA, M. (2017): “Designing Online Marketplaces: Trust and Reputation Mechanisms,” *Innovation Policy and the Economy*, 17(1), 77–93.
- MALINOVA, K., AND A. PARK (2016): “Market Design with Blockchain Technology,” *Working Paper*.
- MANN, S., J. NOLAN, AND B. WELLMAN (2002): “Sousveillance: Inventing and using wearable computing devices for data collection in surveillance environments,” *Surveillance & society*, 1(3), 331–355.
- MILGROM, P. R. (2004): *Putting auction theory to work*. Cambridge University Press.

- MOSER, P., AND T. NICHOLAS (2004): “Was electricity a general purpose technology? Evidence from historical patent citations,” *The American Economic Review*, 94(2), 388–394.
- NAKAMOTO, S. (2008): “Bitcoin: A peer-to-peer electronic cash system,” *White Paper*.
- NARAYANAN, A., J. BONNEAU, E. FELTEN, A. MILLER, AND S. GOLDFEDER (2016): *Bitcoin and Cryptocurrency Technologies: A Comprehensive Introduction*. Princeton University Press.
- RASKIN, M., AND D. YERMACK (2016): “Digital currencies, decentralized ledgers, and the future of central banking,” *National Bureau of Economic Research Working Paper*.
- ROSENBERG, N., AND M. TRAJTENBERG (2001): “A General purpose technology at work: the Corliss steam engine in the late 19th Century US,” *National Bureau of Economic Research Working Paper*.
- ROTH, A. E. (2002): “The economist as engineer: Game theory, experimentation, and computation as tools for design economics,” *Econometrica*, 70(4), 1341–1378.
- ROTH, A. E., AND A. OCKENFELS (2002): “Last-minute bidding and the rules for ending second-price auctions: Evidence from eBay and Amazon auctions on the Internet,” *The American Economic Review*, 92(4), 1093–1103.
- ROTHKOPF, M. H., T. J. TEISBERG, AND E. P. KAHN (1990): “Why are Vickrey auctions rare?,” *Journal of Political Economy*, 98(1), 94–109.
- RYSMAN, M., AND S. SCHUH (2017): “New innovations in payments,” *Innovation Policy and the Economy*, 17(1), 27–48.
- SERETAKIS, A. (2017): “Blockchain, Securities Markets and Central Banking,” *Working Paper*.
- STIGLITZ, J. E. (2002): “Information and the Change in the Paradigm in Economics,” *The American Economic Review*, 92(3), 460–501.
- TUCKER, C., AND C. CATALINI (2018): “What Blockchain Cant Do,” *Harvard Business Review*.
- VON HIPPEL, E. (2005): *Democratizing innovation*. MIT press.
- VON HIPPEL, E., AND G. VON KROGH (2003): “Open source software and the private-collective innovation model: Issues for organization science,” *Organization science*, 14(2), 209–223.
- VON HIPPEL, E. A. (2002): “Open source projects as horizontal innovation networks-by and for users,” .

WALPORT, M. (2016): “Distributed ledger technology: beyond block chain,” *UK Government Office for Science*.

WRIGHT, A., AND P. DE FILIPPI (2015): “Decentralized blockchain technology and the rise of lex cryptographia,” .

A Online Appendix

A.1 What is a Blockchain?

Using blockchain technology, a network of economic agents can agree, at regular intervals, about the true state of shared data. The flexibility in terms of what such shared data represents makes the technology extremely versatile, and allows distributed ledgers to track and settle exchanges across multiple types of digital assets. The rules through which the network reaches consensus about the state of the shared data over time are a key aspect of the market design of a crypto token, as they define the incentives for users and contributors of key resources to the platform. The extent towards which the data in a shared ledger is completely public and associated with pseudonyms (as in Bitcoin),²⁰ or can be strategically shielded for anonymity (as in Zcash) is also a market design choice of a specific token. Similarly, the frequency at which the network reaches consensus and the amount of data recorded in each period of time are features that vary across implementations.

The distributed ledger where the shared data resides is called a ‘blockchain’ because it typically constitutes a chain of blocks of transaction data (see Figure A-2). Each one of the blocks contains valid transaction records for a specific period of time and their attributes. A key attribute of each transaction (and each block) is its timestamp. Blocks are chained together by incorporating a digital fingerprint of the previous block (a hash) in the current block. Any change in the transaction information contained in a specific block would alter such a fingerprint, irreparably breaking the chain of consensus linking that block with all subsequent ones. As a result, one can think of a blockchain not only as a large-scale, distributed database, but also as an immutable audit trail where the ‘DNA’ of each block is incorporated in all following ones, making it impossible to alter history without being noticed.




Figure A-2: A Blockchain

²⁰Whereas it is often believed that Bitcoin transactions are anonymous, they are actually pseudonymous. Like a writer writing a book under a pseudonym, if a Bitcoin user is ever tied to a specific address, the entire history of her transactions with that address can be read on the public Bitcoin blockchain.

A.2 Proof of Work and ‘Mining’

In proof-of-work systems (PoW) such as Bitcoin or Litecoin, participants contribute to broadcasting and verifying transactions while “miners” take on the additional computational work required to assemble new, valid blocks and commit them to the shared ledger. The computationally costly tasks involved in mining are essentially part of a race between miners for the right to add the next block to the chain, and earn the associated reward. The more computing power a miner dedicates to mining, the higher the chances of winning the race by finding a valid solution for a new block first and broadcasting it to the rest of the network. Each time a miner commits a new block to the chain it can assign a predefined amount of the crypto token to itself as a reward (coinbase transaction). This reward, combined with the transactions fees participants may have included in their individual transactions to incentivize miners to prioritize them over others in the construction of the next block, serves as an incentive for miners for the work they perform. To incentivize a decentralized network of miners to contribute resources to secure and operate the network, blockchain protocols typically rely on a native, built-in “token” (in Bitcoin, this is represented by an unspent output on the ledger). This explains why Bitcoin and its blockchain are “joined at the hip”:²¹ for the network to operate in a decentralized way without trusted intermediaries, the process of maintaining the shared ledger must generate enough of an incentive in bitcoin for attracting miners.

Interestingly, in proof-of-work systems, mining does not serve the purpose of verifying transactions (as this activity is fairly light computationally), but of building a credible commitment against an attack: since blocks are chained together, the audit trail formed over time becomes more difficult to tamper with as more blocks are added, and computing power has been sunk to support it. Consensus about the true state of a distributed ledger therefore emerges and becomes stronger as time (and blocks) go by. If a bad actor wanted to reverse a past transaction (e.g. one that is stored n blocks in the past), it would have to spend a disproportionate amount of resources to do so. This is the result of the bad actor not only having to outpace the growth rate of the legitimate chain (which is still maintained by the rest of the network), but also of having to recompute all the blocks after the one that is being manipulated. Since the network always takes the longest valid chain as the true state of the ledger (i.e. as the “consensus”), the task of altering a past block of transactions and imposing it on the rest of the network becomes increasingly difficult as the chain is extended.²²

As a result, in proof-of-work systems, a blockchain is only as secure as the amount of computing power dedicated to mining it. This generates economies of scale and a positive feedback loop between network effects and security: as more participants use a crypto token,

²¹See <http://avc.com/2015/11/are-bitcoin-and-the-blockchain-joined-at-the-hip/> and <http://www.multichain.com/blog/2015/07/bitcoin-vs-blockchain-debate/> (accessed 11-01-2015).

²²Ironically, even if bad actors managed to control a disproportionate share of the computing power dedicated to securing a specific blockchain, it would be in their rational best interest to keep mining honestly (and earn the corresponding mining rewards and transaction fees), as tampering would be visible to others and would destroy the value of the underlying cryptocurrency.

its value increases, which in turn attracts more miners (due to higher rewards), ultimately increasing the security of the shared ledger. Similarly, if confidence in a crypto token drops (e.g. because of a security flaw or because of conflict between developers on its future direction), its value would drop, possibly triggering a negative feedback cycle where miners leave the ecosystem because of the lower rewards until the point where the ledger becomes insecure and is rendered useless. Whereas proponents of alternative consensus systems (such as proof-of-stake) criticize proof-of-work for being inherently wasteful (e.g. in terms of electricity, hardware), from a game theoretic perspective it is exactly the wasteful nature of the mining computations that defends the ledger from an attack: i.e. the sunk, irreversible commitment to the audit trail constitutes the cost a bad actor would have to sustain to manipulate it. If the output of mining was useful for some other purposes too (e.g. if the computations helped find large prime numbers), then the marginal cost of mining would be lower (as part of the cost would be absorbed by the benefits the miner can obtain from selling the solutions to these problems). This does not mean that the network would be less secure, as the marginal returns from mining would also increase and miners would dedicate more resources to it, leading to a corresponding increase in the difficulty of mining. As a result, the network would be as secure as before, and in a scenario in which there are frictions in the market for solutions, it could actually become harder for an attacker to accumulate enough resources to perform a 51% attack.²³

The process through which consensus on the true state of a distributed ledger is reached and secured over time has implications for market design. Depending on the degree of security needed for a specific transaction (e.g. buying a house versus buying coffee), participants will want to wait for a different number of blocks to be settled after the one containing their transaction. This means that the interval at which a new block is added to the chain and consensus is formed, together with the maximum number of transactions that can be included in a block (block size) endogenously determine the optimal transaction type on a specific blockchain. Whereas participants can include higher transaction fees to entice miners to grant them priority within the first available block (i.e. they can reward miners with a higher transaction fee to increase their priority in the queue of unsettled transactions), there is still a limited number of transactions that can be included in any single block.²⁴

²³Proof-of-stake, where the ability to extend a ledger depends on one's ownership stake in the currency, is among the proposed solutions to this trade-off between security and wasteful computations. See: <http://bitfury.com/content/5-white-papers-research/pos-vs-pow-1.0.2.pdf> and <https://bytecoin.org/blog/proof-of-activity-proof-of-burn-proof-of-capacity/> - accessed 09-07-2016.

²⁴For example, Bitcoin currently adds a new block every 10 minutes, and blocks currently have a size of 1MB. The alternative cryptocurrency Litecoin was instead designed with shorter confirmation times (2.5 minutes): while this means that less computing work is done for each block (and therefore the sunk commitment and security per block is lower), the shorter time interval between blocks makes Litecoin more suited for smaller transactions. This basic trade-off between security and bandwidth also affects how different stakeholders within an ecosystem view scaling: in the case of Bitcoin, startups and users that see it predominantly as a cheap medium of exchange would rather have it process a large number of transactions per second and keep transaction fees low, whereas others that are interested in security because they see the token as a store of value ('digital gold') would rather have the incentives system drive out smaller transactions to other blockchains through fees and keep the same level of decentralization. Solutions like the Lightning Network

From a standards perspective, whereas there are advantages to being able to rely on a single blockchain because of economies of scale in security and direct and indirect network effects, it is clear that a single blockchain will not be able to perfectly accommodate every type of transaction (e.g. exchange of value versus the execution of software applications or of legal contracts). The size of a transaction, its attributes and functionality, and the related degree of security and privacy needed to execute it will push different marketplaces on different blockchains.²⁵

A.3 Permissionless versus Permissioned Blockchains and Privacy of Transactions

Bitcoin's market design choices were driven by a desire to make the cryptocurrency as decentralized and democratic in its governance as possible:²⁶ there are no trusted intermediaries, anybody can become a miner or add legitimate transactions, and nobody can block other participants' transactions. Whereas this makes Bitcoin extremely resilient to attacks and censorship, it also makes it less efficient, in its current form, than centralized payment networks.²⁷ Permissioned blockchains, which are distributed ledgers where participants typically need to be granted permission to add (or even view) transactions, can instead deliver higher bandwidth because they do not need to rely on proof-of-work for maintaining a shared ledger. When mining is completely absent from a private blockchain, the audit trail is not protected by sunk computational work, and if the trusted nodes are compromised (or if they collude to rewrite the ledger), the integrity of the chain is at risk.²⁸

Private blockchains are therefore very similar to the replicated, distributed databases already extensively used by corporations. The introduction of distributed ledger technology in this context is usually motivated by incentives to further standardize operations and increase compatibility across industry participants without, at the same time, changing the pre-existing market structure. It is important to note that while private blockchains

enable instantaneous transactions between users through bidirectional payment channels (as in correspondent banking). If successful, this approach would allow a large number of payments to be routed through this parallel network of two-parties ledger entries, drastically reducing the number of transactions that need to be recorded on the main ledger.

²⁵Solutions such as sidechains are being developed through which different blockchains could sync and exchange information seamlessly: e.g. daily microtransactions could take place on a sidechain with lower security but faster confirmation times, and end-of-the-day settlement takes place on the Bitcoin blockchain.

²⁶Whereas Bitcoin was designed to be fully decentralized (one cpu, one vote in the consensus process), economies of scale in mining have driven this activity towards centralization. In 2014, one single mining pool reached more than 50% of the network raising concerns about the integrity of the consensus process (as a miner with such a share could potentially censor transactions, revert them or perform double spending).

²⁷According to a 2014 stress test, the VISA network was able to handle at peak 56,582 transactions messages per second. As of this writing, Bitcoin can only handle approximately 7 transactions per second (source: <http://visatechmatters.tumblr.com/post/108952718025/56582-transaction-messages-per-second> and <https://en.bitcoin.it/wiki/Scalability> - accessed 09-08-2016).

²⁸This makes them less suited for problems where the integrity of the audit trail is paramount (e.g. for regulatory compliance, a network of banks should not be able to collude and revert the state of a distributed ledger ex-post).

benefit from costless verification, they do not take advantage of the reduction in the cost of networking, since control over transactions and assets is still in the hands of trusted nodes. Reliance on trusted intermediaries also comes with advantages, as these systems are more likely to be compatible from the start with pre-existing regulation.²⁹ Whereas this makes a distributed ledger more compatible with legacy systems, it also ties it back to traditional intermediaries as sources of trust. As a result, such a blockchain is unlikely to have a drastic effect on market structure and innovation in the long run.

Related to the issue of trusted intermediaries, is the question of how much privacy a particular blockchain needs to deliver to its users: patterns in a publicly available, distributed ledger can be used to de-anonymize transacting entities behind a pseudonym and gather useful information about the market (Athey, Parashkevov, Sarukkai, and Xia 2016, Athey, Catalini, and Tucker 2017, Catalini and Tucker 2017). To protect their privacy, users can adopt privacy enhancing techniques (e.g. use a new address for each transaction, obfuscate their transactions by mixing them with others), use a truly anonymous cryptocurrency (e.g. Zcash), rely on an intermediary (e.g. a digital wallet provider),³⁰ or use a system that separates basic information about a transaction (e.g. its existence and timestamp) from more sensitive attributes. Additional, sensitive information could be stored on a private blockchain (or database) and immutably linked to the public blockchain entry using a digital fingerprint.³¹ This would preserve the blockchain role as a time-stamping machine, since any tampering with the private record would irreparably break the cryptographic link between the two data sources.³²

While this is still an active area of research, new protocols are being developed to obfuscate transaction data, offer full anonymity to users through zero-knowledge cryptography, and implement different degrees of access to transaction information. Although perfect obfuscation might be not always possible to achieve,³³ it is clear that different cryptocurrencies will be able to compete also in terms of the privacy level they provide to their users (either at the protocol level, or through a trusted intermediary).

As discussed in the paper, costless verification can take place at the level of a single piece of information. When combined with privacy-enhancing measures, this can solve the trade-off between users' desire for customized product experiences (e.g. when using a virtual assistant like Siri), and the need to protect their private information (e.g. the queries sent to the service). If the sensitive data is stored on a blockchain, users can retain control of

²⁹For example, they can be designed to allow for ex-post editing of transactions through fiat, a feature that would undermine the very premise of a public, immutable blockchain, but that clearly has value for certain types of financial transactions.

³⁰Some digital wallet providers do not settle each transaction of their customers on a public blockchain, but only record aggregate inputs and outputs among all their users at regular intervals. These "off-chain" transactions offer a greater degree of privacy from the public, although all information is of course available to the digital wallet provider.

³¹For example, this could be achieved by applying a cryptographic hash function to the private part of the record and recording the output (typically a short string of characters) on the distributed ledger.

³²The blockchain entry would only act as "proof-of-existence" of the original transaction, and if the private record was lost or destroyed there would be no way from the public ledger to extract that information again.

³³See <https://www.iacr.org/archive/crypto2001/21390001.pdf> (accessed 08-01-2016).

their data and license it out as needed over time (e.g. Electronic Medical Records, etc).

A.4 Types of Transactions Enabled By Blockchain Technology

A.4.1 From Atomic Transactions and Immutable Audit Trails to the Exchange of Intellectual Property and Other Types of Digital Assets

An atomic transaction is a transaction that can be fully executed and enforced through a distributed ledger, and whose key attributes can also be verified through the same ledger at a very low cost without the need for an intermediary. Examples include transactions that imply an exchange of cryptocurrency between a buyer and a seller (e.g. a Bitcoin lending contract, a gambling contract) or an exchange between different crypto tokens.³⁴

If all we care about proving with certainty is if (existence) or when (timing) a certain transaction took place, then we can use a pre-existing blockchain to do so: e.g. we could rely on the Bitcoin blockchain to prove that we knew a certain piece of information at a specific point in time (proof of existence). Whereas we would not be able to directly embed the information on the Bitcoin blockchain, we could incorporate a digital fingerprint of it (e.g. a hash) inside a regular transaction. The digital fingerprint would then be secured by the proof-of-work done to maintain and extend the Bitcoin blockchain. At the verification stage, we could point the public to our transaction while at the same time revealing our private piece of information (e.g. the lab notes we wanted to timestamp) to prove the immutable link between the two. Without any additional infrastructure, a blockchain allows us to implement a “first to file” system based on a secure, historical record of timestamped, digital fingerprints.

Because of the ability to implement atomic transactions, build immutable audit trails, and simplify settlement and reconciliation across organizations, blockchain technology has seen fast adoption and experimentation within finance and accounting. Within these fields, the technology can be used to create more open, flexible and programmable exchange platforms, substantially extending the concept of double-entry bookkeeping (e.g. examples include Chain, Digital Asset Holdings, Blockstream etc.).³⁵ Beyond time, labor and cost savings, the development of more interoperable exchange platforms for digital assets substantially reduces entry cost for new players in these heavily regulated markets.

Applications also include novel forms of intellectual property registration and content licensing (e.g. Mediachain). Royalties for the use and remixing of IP or digital content can be tracked in a granular and transparent way on a blockchain by all market participants, which is likely to be particularly useful when different parties have conflicting incentives (e.g.

³⁴Online gambling is an interesting example because costless verification allows for the house to transparently demonstrate fair odds, as users can ex-post verify a dice roll or deck reshuffle was not manipulated to favor the house. Reputation of the gambling house would still be important, as a one-time defection would only be visible ex-post.

³⁵For example, the underlying structure and performance of mortgage-backed security can tracked on a blockchain and made accessible to relevant parties in real time (including regulators), and accounting records can be audited in an automatic fashion while preserving the privacy of the entities involved.

in a principal-agent relationship).³⁶ Pricing and digital privacy models are also becoming more flexible and granular thanks to crypto tokens: e.g. with micro-payments implemented within a browser (e.g. Brave), users can seamlessly pay for access to content behind a paywall, or for an ad-free or no-tracking experience. Similarly, content creators and advertisers can reward users for their attention or for revealing their preferences (e.g. BAT).

A useful extension of an atomic transaction is one that relies on an external source of information (e.g. weather data, exchange rate, price of a stock, outcome of public events) to execute a contract. Examples range from prediction markets (e.g. Gnosis) to betting denominated in a crypto token, to future contracts, mining pool contracts, escrow contracts etc. The external source of information (an ‘oracle’) could be a trusted intermediary, the aggregation of multiple sources (to avoid manipulation), a crowdsourced voting mechanism, or a trusted hardware device. A particularly interesting set of transactions is the one enabled by linking an IoT device to a cryptocurrency. If the hardware device is secure and cannot be tampered with, then the information it collects can act as the trusted arbiter in a digital transaction.³⁷ This allows new marketplaces to emerge where energy (e.g. from solar panels), bandwidth, access to resources and information, data processing through an API, or work performed by the crowd are priced in novel ways.³⁹

A.4.2 The Identity, Credentials and Provenance Verification Problem

The process of identity verification is central to all economic transactions. Each time we authenticate ourselves (or an entity we represent, or a device), we are essentially creating a transaction allowing a third-party to verify that we are authorized to perform an action. This transaction is usually what stands between a legitimate use and fraud, leakage of information, digital and physical theft. A well functioning market (and economy), relies on robust identity verification as well as on the ability to verify the goods and services being exchanged (e.g. in terms of their provenance, how they were changed through the supply chain etc.), and

³⁶For example, artists that license their music to Apple or Spotify could track how many times their songs are played by consumers, or seamlessly receive royalties from other artists for remixes that include parts of their songs according to a predetermined smart contract. Similarly, backers on a crowdfunding platform could obtain royalties each time a song they funded is played, artists could sell the rights to the first copy of a digital artwork, stock photography websites could certify legitimate uses of their content at a lower cost.

³⁷For example, a weather or pollution sensor³⁸ could capture local information and sell it back to the network for a price. IoT devices and robots, when combined with a cryptocurrency, can seamlessly earn, barter or exchange resources with other devices on the same network. In a futuristic scenario, a self-driving car could buy up lane space from surrounding vehicles on a highway for priority. If the IoT device also contributes to mining the underlying cryptocurrency (e.g. by dedicating computing cycles during idle time to securing a digital ledger), then this may also allow for new business models to emerge (e.g. a cellular phone’s plan could be partially subsidized through its mining chip).

³⁹Given current technology, users can already be paid instantaneously and with less frictions through a cryptocurrency to perform small tasks both offline and online across the globe (e.g. answering surveys, translating text or audio, writing a review, training machine learning algorithms, collecting offline prices etc.). Whereas payments from users to services online are pervasive, the reverse flow is substantially more rare (e.g. Amazon Mechanical Turk) and cumbersome (e.g. linking of a bank account). Crypto tokens, by enabling bidirectional, low friction exchange of value, can substantially expand these markets (e.g. 21.co).

the credentials of the parties involved (e.g. degrees on a curriculum vitae, professional licensing status, bad actor status, driving record etc.). Current solutions to the identity and credentials verification problem typically rely on insecure secrets and documents (e.g. social security number, passwords, passports, signatures, university transcripts etc.) or public-key encryption and hardware (e.g. multiple factor authentication, certificates). In many cases the intermediary is the government, although it can also be a consortium, or a private firm (e.g. Facebook Connect). This always involves some degree of information leakage and risk of reuse of private information outside of the designated transactions. Blockchain technology can reduce this risk by allowing for authentication without disclosure of sensitive information. The same way a distributed ledger can track the attributes of financial transactions, it can also track changes to an individual's status and credentials (or firm, good, service). An individual's ability to perform (or not) a certain action could be tracked on a blockchain and queried when needed without necessarily disclosing all underlying information (e.g. a bank could verify, after being authorized by a customer, a credit history). Similarly, access to medical records could be granted, revoked or ported between providers as needed.

From a privacy perspective, the ability to license out subsets of personal information for limited amounts of time and to seamlessly revoke access when necessary has the potential to not only increase security, but also to enable new business models where customers retain greater control over their data and firms can dynamically bid for access.

Attributes of digital and physical goods can also be tracked on a distributed ledger as they move through the economy, increasing our ability to verify their integrity, provenance, manipulation and status (e.g. warranties, food safety) over time. This is particularly powerful when immutable properties of a good (e.g. the properties of a diamond, art piece or geographic coordinates of a parcel of land) can be reliably recorded on a blockchain, i.e. when a unique, digital fingerprint can link ownership of a blockchain token to the underlying asset. When this is not possible, the problem of identity, credential and provenance verification will still require trusted intermediaries (or at least a secure IoT device or sensor) to reliably capture what is happening in the offline world and record it on a distributed ledger: Intermediaries and secure devices act as key complements to online forms of verification enabled by blockchain technology.

A.4.3 Online Reputation Systems

A key function of online intermediaries is to design and maintain a robust reputation system to facilitate transactions between buyers and sellers (Luca 2017). In this context, blockchain technology can be used to increase transparency, ensure that reviews and ratings are only produced after a verified purchase, and to build open reputation platforms. Advantages include the ability to port and use the resulting reputation scores across different services and contexts, increased transparency, and lower barriers to entry in markets currently dominated by a few intermediaries (e.g. Yelp, Airbnb, Uber). This has implications for how policymakers approach regulation, monitoring, and antitrust issues in these markets, as it gives a public entity the ability to enforce market design rules (e.g. safety standards, worker compensation, liquidity standards etc.) through a well-designed protocol.

A.4.4 Central Bank Money

A particularly interesting application is the development of a blockchain-based, fiat-endorsed digital currency. If a central bank were to switch from the current infrastructure to a cryptocurrency, it would be able to directly provide citizens with digital, central bank money. This would challenge some of the revenue models of commercial banks, as citizens may prefer the more secure central bank money to their traditional checking account. Startups could then compete in providing security and protection for consumer digital wallets, payments and billing services, etc. While the implications of such a switch are not the focus of this paper, the change would have broad implications for how governments implement taxation (because of costless verification), manage money supply and interest rates, deliver quantitative easing, on their ability to enforce financial sanctions on other governments, and more generally facilitate intertemporal transactions in the economy. Such a currency would also become an appealing alternative - because of its digital nature - for foreign citizens in countries facing currency devaluation or where trust in the government is low. Events such as India's demonetization of the 500 and 1000 rupee notes, and broader pushes towards greater traceability and government surveillance in transactions (e.g. by reducing the role of cash), are likely to increase consumers' interest in cryptocurrencies as a store of value and for privacy concerns (i.e. fiat-based currencies will have to increasingly compete with their decentralized counterparts). Recent moves such as the Chinese ban of initial coin offerings and Bitcoin exchanges foreshadow increasing tension between regulators and permissionless cryptocurrencies, possibly while the same governments consider adopting blockchain technology to lower cost and enable new types of services for their citizens through a fiat-based digital currency.

A.4.5 Auctions

Economists have made great strides in applying economic theory to the design of practical markets (Roth 2002). But issues remain and, apart from once-off auctions of public assets, the best designs are not often implemented. An example of this is the second-price auction developed by William Vickery (Ausubel, Milgrom, et al. 2006), where bidders submit their reservation price to an auctioneer, and the bidder with the highest bid wins the auction but only has to pay the second highest bid. This auction has the property that its outcomes are efficient (the auction winner is the agent with the highest valuation), and involves a straightforward bidding process since bidders can simply submit the highest amount they would be willing to pay. Nonetheless, it has found limited applicability in practice. A notable exception is Google's AdWords auction (Edelman, Ostrovsky, and Schwarz 2007). One of the reasons why market designs that require agents to submit their true valuation (or costs) do not actually emerge in practice is that there is a potential lack of trust in the intermediaries involved. One aspect of this is that a seller may use the fact that a bidder has a high willingness to pay for an object to somehow turn the tables on them in the auction.⁴⁰

⁴⁰For example, suppose there are two bidders for an object. One has a value of \$5 and another has a value of \$10. Suppose also that it turns out that the seller will keep the object if it does not attract more

Hence, they may choose not to do so and the value of the auction may be undermined. An open-cry auction may resolve this issue by forcing the seller to reveal when their reserve price is met but such auctions have their own costs (including having to assemble all bidders at the same time and location) and may not be practical online.

A distributed ledger solves these potential expropriation problems. For instance, eBay offers an automated bidder which allows people to submit their highest bid and then bids on their behalf. In effect, it is supposed to replicate a second-price auction. Often people do not actually use the automated bidder properly and wait until the last minute to bid (Roth and Ockenfels 2002). One reason could be some kind of mistrust or alternatively a concern that the bids will not be submitted properly. With a distributed ledger, the bids could be collected through a smart contract without ever exposing the information to the seller or a third party. When the auction closes, the contract would rank the offers, identify the winner, and destroy the information about the other bids. The smart contract could also ensure that the bidder has enough funds to make an offer and does not default (Milgrom 2004), reducing the worry that the auction will be re-run and the bidding information used against participants. Thus, we can see how the full verifiability that accompanies a blockchain can potentially render practical the full commitment assumptions required for efficient auction designs to be implemented.

than \$4 in the auction. In a second-price, auction where bidders bid their true values, the winning bidder would be the \$10 value bidder who would only have to pay a price of \$5. Suppose, however, that the seller does not reveal their reservation price. A concern might arise that they might see the bids and then claim the reservation price is \$7. In that situation, the bidders would face expropriation and a reduced surplus from bidding their true values. See (Rothkopf, Teisberg, and Kahn 1990) for an analysis. The authors also examine what might happen if truthful bids leak to third parties who can then exploit the bidders.