

 Open access • Journal Article • DOI:10.2139/SSRN.3160070

Blockchain: The birth of decentralized governance — [Source link](#)

[Benito Arruñada](#), [Benito Arruñada](#), [Luis Garicano](#), [Luis Garicano](#)

Institutions: [Barcelona Graduate School of Economics](#), [Pompeu Fabra University](#), [IE University](#), [Economic Policy Institute](#)

Published on: 10 Apr 2018 - [Social Science Research Network](#)

Related papers:

- [Governance in Decentralized Networks](#)
- [Decentralized Blockchain Technology and the Rise of Lex Cryptographia](#)
- [A novel framework for policy based on-chain governance of blockchain networks](#)
- [Blockchain Technology and Decentralized Governance: The Pitfalls of a Trustless Dream](#)
- [The Best of Both Worlds: A New Composite Framework Leveraging PoS and PoW for Blockchain Security and Governance](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/blockchain-the-birth-of-decentralized-governance-2syi9w7h0k>



**Universitat
Pompeu Fabra**
Barcelona

Department
of Economics and Business

Economics Working Paper Series

Working Paper No. 1608

**Blockchain: the birth of decentralized
governance**

Benito Arruñada and Luis Garicano

April 2018

Blockchain: The Birth of Decentralized Governance

Benito Arruñada* and Luis Garicano**

April 10, 2018***

Abstract

By allowing networks to split, decentralized blockchain platforms protect members against hold up, but hinder coordination, given that adaptation decisions are ultimately decentralized. The current solutions to improve coordination, based on “premining” cryptocurrencies, taxing members and incentivizing developers, are insufficient. For blockchain to fulfill its promise and outcompete centralized firms, it needs to develop new forms of “soft” decentralized governance (anarchic, aristocratic, democratic, and autocratic) that allow networks to avoid bad equilibria.

Keywords: blockchain, platforms, networks, hold-up, coordination, relational capital, incomplete contracts, decentralized governance.

JEL codes: D23, L12, L22, L86.

* Pompeu Fabra University and Barcelona GSE. E-mail: benito.arrunada@upf.edu.

** IE Business School, London School of Economics and CEPR. E-mail: luis.garicano@gmail.com.

*** We thank Eric Brousseau, Miguel Espinosa, and Spyros Terovitis for valuable comments; Jin Lee, Luis Rayo and participants at the Workshop on Economic Governance of Data-driven Markets held at TILEC (Tilburg University), the Blockchain, Law & Policy Workshop held at the University of Amsterdam, and seminars at CUNEF (Madrid), the JRC of the European Commission (Seville) and the Department of Economics (University of Vigo) for useful discussions; and the research assistantship of Juan Barenys. This work received support from the Spanish Government through grant ECO2017-85763-R.

1. Introduction

The contractual technology on which the market economy is based has evolved over a period of two thousand years. Blockchain promises to overhaul it, by allowing for the creation of open, distributed, secure, encrypted and programmable digital ledgers and enabling secure and fully decentralized “P2P” trade.

Fulfilling this promise requires decentralized blockchain platforms to provide operating standards, a central reputation for the whole ecosystem, and a communication protocol that allows applications (“Apps”) to communicate with each other. But the success of such decentralized platforms hinges crucially on finding good governance solutions to align the incentives of all participants.

Specifically, all platforms, centralized or not, must solve two classes of problems: protecting the investment of platform participants from hold-up and coordinating adaptation and change. In this paper we discuss how decentralized, blockchain-based, platforms solve these two problems compared to traditional centralized ones, and evaluate the governance solutions currently being proposed by different platforms.

Hold up problems are particularly prevalent in centralized platforms. Once participants have committed to the platform, switching is costly. This allows the “network architect” (e.g., Apple in iOS Apps) to hold up application developers, for instance, by implementing changes that make their previous investments obsolete. The risk of being held up leads participants to underinvest. When the network architect has a sufficiently long horizon (again, like Apple in iPhone Apps or Google in Android), it has an incentive to develop a reputation for protecting the investments of its application developers (and not holding them up) that enforces the relational contract. Decentralized platforms such as blockchain drastically reduce the potential for hold up, since network members can always refuse to introduce any change. This reduction, however, comes at the cost of increasing coordination problems.

In both types of platforms, coordination problems result from the fact that platforms must ensure that core software and application developers have the right incentives to cooperate and adapt together to evolving needs. In centralized platforms, the decision to change is made by the architect, who ignores the private costs of the developers and, as a result, tends to introduce too much change. In a decentralized platform the opposite happens: individual application developers (and other platform members) have some veto power through their freedom to update. They must therefore agree to any change. This poses a problem for changes that are not win-win, that is, any change which is good for the whole network but causes losses to some players. Decentralized platforms will find this type of change hard to implement.

1.1. The roadmap

In Section 2 we study the key contracting difficulties in centralized platforms and their relational solutions. The network architect can unilaterally make the adaptation decisions of the platform. It may require excessive investments from partners and too frequent adaptation, since it does not fully internalize the costs of such decisions. However, a network architect with sufficient relational capital can use it to safeguard a governance structure that protects participants in the platform against its own hold up and thus provide them with adequate incentives to invest. To illuminate these problems, we briefly describe the governance of the Apple and Google platforms and their ability to protect developers from hold up and adequately respond to change.

In Section 3 we consider blockchain-based platforms. In a decentralized setting, each member of the network (individual developers, final users and other partners, such as blockchain miners) can refuse to change, which protects her from hold up. However, incentives to adapt in a coordinated manner are limited. There may exist multiple equilibria, as in the traditional network externality literature (e.g., Katz and Shapiro, 1985, 1986; Farrell and Saloner, 1985). Absent explicit monetary incentives or other governance solutions, the system is likely to yield too little adaptation or generate inefficient “splits” (where only part of the network adopts a particular protocol change). However, unlike in standards “wars”, where two standards battle in a winner-take-all confrontation, platforms can implement governance rules to handle the trade-off between hold-up and coordination as efficiently as possible.

For instance, consider the two main decentralized platforms: the ones developed on top of Bitcoin and Ethereum. While Bitcoin is a digital currency that can be used for structuring simple predefined transactions, such as escrow accounts or document signing, Ethereum was created with the aim of making even more complex transactions easier. In particular, its programming language (Solidity) supports complex logical functions like conditions, recursions, loops, go-tos, macros, etc. Codified Ethereum’s transactions can be called repeatedly, to send and receive funds, to act as new currencies, as well as to interact with other contracts. These features make it suitable for coding “smart contracts”—those which are enforced automatically by computer code.

The governance of both platforms is also very different. Bitcoin is an extremely decentralized platform with very little governance. Ethereum is a considerably more centralized platform where some additional governance aims to achieve some coordinated adaptation.

It is precisely the same ability to adopt changes individually—even to split from the main network—that protects App developers and others from expropriation what may lead to large scale coordination failures. In each of these decentralized platforms we study how they handled a situation where the implementation of a change led to a large coordination failure and a split in the network, in connection with what is called in the blockchain community a “hard fork”. We evaluate, in each case, the governance solutions that are being currently proposed by the different platforms, examine their weaknesses, and propose some further solutions.

While blockchain governance involves significant shortcomings, we study some of the new tools allowed by blockchain that may be used to solve incentive conflicts.

First, platform designers can rely on the complementarity between crypto currency and platforms. Specifically, “pre-mining” and other ways of allocating cryptotokens and cryptocurrencies to developers and, particularly, core developers, may alleviate coordination problems. We argue that this is only of limited value, and only early in the life of the platform. The impossibility of imposing “vesting periods” and other restrictions substantially weakens this tool.

Second, some networks are implementing solutions designed to facilitate coordination around the best equilibrium. We classify these emerging solutions as anarchic, aristocratic, autocratic and democratic, by analogy to political “hard” governance in the real world. Although these solutions “feel” like governance, ultimately network members always have the ability to implement or not these changes, and thus they are better understood as devices to facilitate the convergence of expectations than as real governance devices. Hence our use of the term “soft governance”.

In Section 4 we compare the relational capital needed by optimally-governed centralized and decentralized platforms. We show that centralized platforms require less relational capital when network members are highly heterogeneous, while decentralized ones are preferred when investments are hard to verify. This results point the way to the instances in which we expect one or the other solution to prevail.

The main precedent for our approach, using the economics of contracts and organization to study these platforms is Bresnahan and Greenstein (2014), who discuss informally mobile computing platforms, and suggest that the key tradeoff is that, while a hierarchical, or centralized platform is superior in terms of coordinating a response to change, a decentralized one is superior for exploration of new products and ideas. Our discussion is distinct in that we focus on a different tradeoff, between investment (hold up) and coordination; we analyze the problem in terms of relational contracting; we propose a formal model that we believe captures this tradeoff in the simplest possible way; and finally our case study discussion has to do with the blockchain platforms, not discussed in their paper.

There is also a burgeoning literature on bitcoin and blockchain. Böhme et al. (2015) focus on the economics of bitcoin. They note two key costs: one is the technological waste, and the second is the cost associated with the market concentration of intermediaries, notably of miners. Catalini and Gans (2016) study blockchain economics and argue that its main benefits are the reduction in these cost of verification and (with the addition of bitcoin) of the cost of networking. Athey et al. (2017) discuss the privacy paradox, whereby cryptocurrencies offer people the chance to escape government surveillance, but do so by making transactions themselves public on a “blockchain”. The effect of small incentives may explain the privacy paradox, where people say they care about privacy but are willing to relinquish private data quite easily. Biais et al. (2018) undertake a full, detailed theoretical dynamic analysis of the equilibria in a blockchain. Our work is distinct from the above in our emphasis on how emerging governance solutions deal with existing problems

and pointing the way to future governance, and in comparing the robustness to incentive conflicts of centralized and decentralized networks.

Our work builds on two recent streams of literature in organizational economics on coordination and adaptation (e.g. Dessein et al. 2010) and on relational contracts (Levin, 2003). We also rely on the law and economics literature on agency law and the birth of the corporation (e.g., Hansmann and Kraakman 2000, Harris 2000, Arruñada 2010) to highlight that comparisons of current blockchain governance against centralized governance is unfair to blockchain: centralized governance applies ready-made solutions provided by a centuries-old evolution of institutional and organizational solutions, while blockchain is just starting a similar process of evolutionary discovery.

Readers familiar with these topics may be advised to jump now to Section 2, skipping the rest of this Introduction (Section 1.2), in which we describe blockchain technology, its main platforms (Bitcoin and Ethereum), its application to smart contracts, and one of its key problems—splits caused by hard forks.

1.2. Blockchain platforms: a primer

We conceive a “platform” as the combination of software and hardware resources enabling the functioning of an exchange network. In turn, a “network” is the community of individuals (including computer nodes) exchanging goods or services through a given platform.

Blockchain is a type of (distributed ledger) technology in which related transactions are bundled together into blocks. These blocks are made public to all the network after being linked to one another using “hashes” which take an input string of any length and convert it into a fixed-length string.

For our purposes, the key economic contribution of blockchain technology is that it allows person-to-person (P2P) transactions to take place safely. It therefore holds the potential to decentralize and disintermediate all sorts of markets and activities.

This is most visible in the first application of blockchain: the creation of decentralized currency without any intermediary.

For a currency to be a store of value and medium of exchange, it must fulfill three conditions: (1) it must be easy to verify its validity or authenticity; (2) it must be impossible to spend it twice; and (3) it must be robust to opportunistic minting by its creator or its owner.

Up to blockchain, the solution to these problems was to centralize the issuance of currency in a trusted intermediary—normally the state. Bitcoin and all other cryptocurrencies based on blockchain provide a decentralized alternative. Instead of relying on an intermediary, they rely on a peer-to-peer network, in which each node replicates and encrypts the particular transaction, as well as the full history of previous transactions and sets them on a “chain” linked in a permanent and immutable way, which it then makes public. Moreover,

validation of a transaction requires consensus: any transaction with more (implicit) acceptances becomes the “official” next block in the chain.

In addition, players have explicit monetary incentives to choose the transaction that will engender more consensus, an implicit “truth telling” device: for them to benefit from adding a block to the chain, subsequent blocks must be built on top of their block.

In particular, the initial key players in the bitcoin blockchain “community” were:

- (1) **Core Developers:** Engineers who launched the initiative and then often retain privileges to maintain and suggest changes to the core protocol.
- (2) **Miners:** Owners of computer nodes who do the calculations to validate transactions in exchange for a block reward (initially only) and (later, increasingly) transaction fees.¹ Mining equipment can easily switch to mine on different blockchains, making miners unlikely candidates for hold up.
- (3) **Investors:** Funders of blockchain ventures who directly fund development and, initially mainly, mining equipment.
- (4) **Users:** Purchasers of the cryptocurrency to pay for transactions, to store value or for speculation.

The emergence of “decentralized Apps” on Bitcoin and Altcoins

Bitcoin was created as a currency but has unintendedly become a platform for blockchain Apps: once it enabled secure payments between strangers, this removed a key obstacle to the development of P2P trade and encouraged developing techniques that, building on top of the bitcoin’s blockchain, execute simple, “smart”, contracts with low transactions costs.

These smart contracts apply blockchain technology to allow for the recording and verification of transactions over all types of assets. Being built directly into a blockchain, they are enforced automatically without intermediaries, therefore applying the “code is law” principle defined by Lessig (1999). Essentially, claims on the asset, like a bitcoin, circulate through the peer-to-peer network through an encrypted chain of blocks in which the full history of transactions is recorded and preserved.

Smart contracts are the basic units of more complex decentralized applications (in blockchain jargon, “DApps”), which, instead of running on a central server, as usual Apps, rely on a decentralized computer network to provide some specific valuable service (for example, a system for sharing apartments or renting cars).

¹ Miners add blocks of transactions to the chain in the form of the public distributed ledger of transactions, the “blockchain”. To add a block, they must first compile the past transactions and then solve a hard mathematical puzzle. The first miner to solve the puzzle adds the block and is rewarded with some cryptocurrency, such as bitcoin, once the proof of work is verified by other nodes.

This proliferation of DApps ended up creating “overlay” networks (Monegro 2014) built up on top of Bitcoin’s blockchain, using it to provide incentives, time-stamp transactions and validation, and benefitting from its liquidity, reliability and network effects. In this way, Bitcoin eased the entry of a myriad of applications that did not even need to bootstrap their own cryptocurrency and/or start a new blockchain.

Such ad-hoc development of DApps was, however, plagued with compatibility problems similar to those suffered by conventional Internet networks (Kasireddy 2017). Each DApp built on the bitcoin chain stands alone, and in order to communicate with others it must adapt to the proprietary standards set by other applications in their “application programming interfaces” or APIs, the software that sets specific methods of communication between different software elements.

The emergence of Ethereum and other blockchain platforms

These shortcomings of Bitcoin-based DApps led to the development of blockchain platforms such as Ethereum, which is now the second most important network, well above all others. Its purpose is to facilitate the development of DApps and to integrate them in a much more consistent ecosystem than those based on Bitcoin (Narayanan et al., 2016:263–70).

Specifically, blockchain platforms create value by setting up standards and other inputs essential for DApps’ development and operation. This involves both operating standards, including money (“ether” in Ethereum); a central reputation for the whole ecosystem, which is partly shared across many DApps; and, most importantly, direct network effects through a communication protocol that allows value to be created by connecting applications.

In particular, Ethereum defines itself as a computer platform running all sorts of smart contracts. All Ethereum’s DApps are run on the same virtual machine, use the same language and the same “primitives” (e.g., smart contracts, accounts, addresses, etc.), rely on the same blockchain, validate their transactions in the same way, pay the same use fees, etc.

Compared to a situation in which each DApp would start independently from zero, blockchain platforms lower entry barriers by sharing fixed costs, which are incurred by the core. However, this also means that this investment in the core must be funded and DApp developers become more dependent with respect to the core, to the extent that it is costly or even impossible to switch platforms. In particular, the core (potentially also, some infrastructure apps) may hold up the application developer and appropriate the value of their investment. Such accumulation of market power by the development core could reduce the incentives to invest in DApps’ development and in their quality.

Moreover, the fact that multiple independent DApps share a common core means that difficulties may arise when changes in this core need to be introduced.

Thus incentives must be structured ex ante to contain both conflicts with a governance structure, providing both the core and DApps’ developers adequate incentives to invest and adopt changes efficiently.

Coordinating change in blockchain: “hard forks” and splits

Such governance structure must consider two sets of constraints, one present in all types of platforms. Another specific to those based on blockchain.

The general constraint is the familiar issue that solving the contracting problem *ex ante* is well-nigh impossible mainly because intense uncertainty and the sequential arrival of applications mean that contingencies cannot be anticipated and proper incentives cannot be designed *ex ante*. The contract must therefore be relational, as platforms and applications will be distributing their gains *ex post*.

The specific constraint of blockchain governance is due to the fact that the decision to change the rules—what is called a “fork”—is intrinsically decentralized. It can be initiated by anyone who proposes an upgrade in the protocol but it fully succeeds only if the whole network adopts the upgrade.

There are two types of forks. A *hard fork* is any change in protocol that involves forward incompatibilities with the previous protocol, so that transaction that are accepted by the new protocol are invalid in the old one and will not be accepted by non-upgraded nodes. Conversely, a *soft fork* restricts the validation rule, so that any transaction that is valid in the new protocol would still be valid in the old one.

Hard forks may pursue different objectives, from eliminating security hazards in the code to implementing new functions or even reversing transactions. The latter are the most interesting from a contract theory perspective: blockchain enthusiasts see them as anathema, as they deny the immutability principle. However, as we will see, by providing a means for efficient contract breach in extreme circumstances, they may serve to efficiently complete smart contracts that—at least by now—remain necessarily incomplete.

With every hard fork, there is a risk of a split in two halves, with the owner of each currency unit, or “coin”, receiving two new coins. Both currencies then start functioning as separate entities, with their own different sets of rules (see, e.g., Narayanan et al., 2016:171–72).

Before the well-publicized splits following hard forks that took place in Ethereum and Bitcoin blockchains in 2016 and 2017, respectively, a few precedents had happened in what then were small blockchains.² This included Bitcoin itself, which forked in 2010 after someone minted billions of bitcoins. However, given that the network was still small, it was easily handled without much difficulty. In 2014, the MintPal exchange suffered a hack that led to two million USD in VeriCoin tokens being stolen. Subsequently, developers reclaimed the funds by what is said to be the first hard fork. Also in 2014, after Nxt had suffered a 1.75 million USD theft, developers also proposed a hard fork, but it was rejected. Most of the funds were recovered through negotiations but only after paying ransom to the hacker. It has been alleged that the different outcomes were aligned with the different causes of the hacking and, consequently, the merits of the cases.

² Arruñada (2018) provides further analysis and sources on these hard forks.

The possibility of splits in decentralized platforms drastically changes the incentives with respect to those present in traditional, centralized ones (such as Apple's iTunes or Google's Play Store). In centralized platforms, the architect unilaterally can make design decisions and "complete" the contracts whenever contingencies arise for which no specific action was foreseen in the contract. For instance, Apple often adds new features or software updates that instantly renders previous investment obsolete.³ This may lead to an incentive conflict in the form of hold up, as the platform architect may ignore the costs incurred by others.

In blockchain, the platform architect may play a limited role, as in Ethereum, or may even play practically no role (in Bitcoin there is, in fact, no architect as such). Changes require the independent actions of a wide range of players and thus a large level of decentralization is inevitable: the nature of the network is that all nodes can unilaterally determine which protocol they run, and whether they update it or not is their decision. The larger cost here is thus coordinating changes and avoiding both damaging network splits and inefficient delays as features that are welfare improving are not adopted.

The challenge for blockchain networks lies in developing efficient mechanisms producing a necessarily "soft" form of governance that takes as given the ultimately decentralized decision-making (via potential splits); but, by promoting the good "equilibria", ensures efficient nodes' coordination when adapting to new circumstances.

2. Coordination and hold up in centralized networks

We consider throughout a network with many nodes that must adapt to change. It confronts two problems (1) change has asymmetric benefits, for instance, some users benefit from it while others do not. (2) The costs of change are partly private.

2.1. Spot contracting in a centralized network leads to hold-up by the architect of the nodes investments

To fix ideas, consider a network with a given set of nodes, each of which enjoys an operational benefit $b(n)$ from operating, depending on the size of the network n . The network has a central player, such as for instance Apple in the mobile phone application

³ Apple is quite clear about it when it says that "the proposed changes are massively source-breaking for Swift code, and will require a migrator to translate Swift 2 code into Swift 3 code." <https://www.infoworld.com/article/3120171/open-source-tools/apple-swift-3-forward-looking-but-not-backward-compatible.html>

platform, or PayPal in payments. We shall call this player the “network architect.”⁴ The player controls and determines the adoption protocol—by pushing a button she installs the new protocol on all network nodes.

The network architect appropriates all network value net of the compensation w to the individual nodes. Thus her profitability is given by:

$$\Pi = n(b(n) - w) \quad (1)$$

where $b'(n) > 0, b(1) = 0$. When all possible nodes join the network, its size is N .

For each node to willingly participate, his utility must be larger than the outside utility, which we normalize to 0, plus the sunk cost of developing the application specific to the platform, C . Thus the network architect must offer ongoing value to the developer at least equal to the cost of participation-given competition, $w = C$. The utility attained by the developer is

$$U = w - C = 0 \quad (2)$$

Suppose also that, for the sake of adaptation, the architect is given the right to make changes by compensating them for the verifiable costs. Imagine now that an unforeseen opportunity (an unforeseen contingency in the sense of Kreps, 1990) for change or adaptation appears after the development cost has been incurred. This opportunity requires updating the network protocol. However, not all benefit equally from it: a (random) subset of $m \leq N$ nodes obtain a benefit from the change of Δ_h per node if it is implemented, whereas the rest of the nodes, the *losers*, get a benefit $\Delta_l < \Delta_h$. Any node who implements the change incurs a private cost $k \in R^+$. However, only a share γ of these costs can be verified in court. We assume large network externalities— $b(N)$ is large, $b(N) \gg k, b(N) \gg \Delta_h$.

First best implementation rule: At any period t , it is efficient for the new protocol to be adopted if:

(1) The value of the change is larger than its cost:

$$\begin{aligned} m\Delta_h + (N - m)\Delta_l &> Nk \\ \alpha_h\Delta_h + \alpha_l\Delta_l &> k \end{aligned} \quad (3)$$

Where α_h is the share of winners, $\alpha_h = m/N$ and α_l the share of losers $\alpha_l = (N - m)/N$.

(2) A split of the network, where only winners implement, is inefficient:

⁴ We are simplifying: even in a centralized network there may be more than one architect, like in the “Wintel” case where two architects must coordinate upgrades with one another (see Casadesus and Yoffie, 2007).

$$m\Delta_h + (N - m)\Delta_l + Nb(n) - Nk > m\Delta_h + mb(m) - mk + (N - m)b(N - m)$$

or equivalently:

$$\alpha_l\Delta_l + b(N) - (\alpha_h b(m) + \alpha_l b(N - m)) > \alpha_l k \quad (4)$$

The expression is straightforward: the losers must adopt the change optimally if the gain to them, Δ_l , plus the gain from a unified network is larger than their adoption costs.

We assume for now that the network effects are sufficiently high that it is not efficient to split the network.

Assumption 1. Splits are inefficient. Network economies are sufficiently high that it is not efficient to split the network—i.e., $\alpha_l\Delta_l + b(N) - (\alpha_h b(m) + \alpha_l b(N - m)) > \alpha_l k$.

In a centralized network, the architect has full bargaining power and will therefore adjust the compensation of the nodes (e.g., App developers) to extract the additional profitability. Moreover, given that some of the costs of transitioning to a new protocol are unverifiable they may be expropriated by the architect. The compensation of the developers after all of them incur the cost k to implement the new protocol (note that in this one-period world $w = C$) will be:

$$w' = w - (1 - \gamma)k = C - (1 - \gamma)k \quad (5)$$

Resulting in ex-post utility for the developer

$$U = C - (1 - \gamma)k - C = -(1 - \gamma)k < 0 \quad (6)$$

However, the sunk cost of development is incurred already. Thus the node (an app developer) implements the protocol change (rather than shutting down the node) as long as total ex post compensation is positive, that is

$$C - (1 - \gamma)k \geq 0$$

$$C \geq (1 - \gamma)k$$

The argument is standard (Williamson, 1975; Grossman and Hart, 1986): as long as the sunk cost and consequent quasi-rents are significant enough, hold up results and the investment is partially expropriated by the network architect. The node would have been better off if she had not invested in development. The fear by the developer that she may be expropriated by the network architect may lead her to ignore any promises he may make and not invest in developing apps.

Result 1 [Hold up]. *Nodes in a centralized network will accept to undertake changes that would have had an ex ante negative value. This hold up risk will lead to underinvestment by application developers.*

Consider now the decision by the network architect to introduce the upgrade. Let $J=1$ if the protocol change is made. Her profits are

$$\Pi = Nb(N) - Nw' = N(b(N) - C) + J (\alpha_h \Delta_h + \alpha_l \Delta_l - \gamma Nk) \quad (7)$$

Where the architect appropriates all of the profit from the change plus part of the agents quasirents (the non-verifiable share of cost). The architect's objective is to maximize the value of the network, net of compensation to nodes. Because only a share γ of the costs is verifiable, he ignores the rest $(1-\gamma)$ of the cost. As a result, the architect imposes changes even when they lead to large losses to the nodes:

$$\text{Decision by network architect : } \begin{cases} \text{Impose protocol change if: } \alpha_h \Delta_h + \alpha_l \Delta_l \geq \gamma k \\ \text{Do not impose change if: } \alpha_h \Delta_h + \alpha_l \Delta_l < \gamma k \end{cases}$$

Result 2 [A centralized network imposes too much investment and implements too many changes]. *Since a central architect ignores the unobservable costs of nodes that are committed to the network, a centralized network implements too many changes.*

Figure 1: The inefficiency of centralization: inefficient change is imposed by architect




Figure 1 illustrates this inefficiency. In a centralized network the architect chooses to impose changes as long as the benefit is larger than a fraction of the cost, which leads to excessive inefficiency. This inefficiency is due to the switching costs incurred by the nodes together with the sunk investments they have made, specific to this particular network. As long as sunk costs specific to the system are large enough, $(C > (1 - \gamma)k)$, they will implement the change. This may lead to welfare reducing changes.

Examples. Traditional web and mobile platforms are centralized, as we explained in the introduction. Although (for reasons that will become clear momentarily) these are the exceptions, the popular press contains accounts of expropriation of the investments made by developers. For instance, Apple often adds new features or software updates that, being backwards incompatible, may render developers' previous investments obsolete: "Swift 3 is a great opportunity to start fresh with the language, but if you have existing Swift codebases, be prepared to rewrite—or dump—them" (Yegulalp 2016). Something similar happens when Apple updates the operating system of its iPhones, in order to enhance their functionality and ensure their security, which often makes old apps obsolete and requires updating their code. And updates do not only cover software but also apps' marketing standards: for instance, the promotional texts and keywords of the apps.

On the other hand, the overall success, including millions of App developers willing to invest in these networks, suggests this type of opportunistic behavior is less prevalent than a short term analysis suggests. The reason, of course, is that the future matters a lot, and allows for a reputation to develop that protects specific investments.

2.2. The solution to hold up in centralized network: Relational capital as a base for relational contracts

The network architect has a lot at stake. If she expropriates the App developers, they may not invest, and the network will not take off. The loss of such future surplus may allow the network to exist in the presence of temptations by the architect to expropriate developers' investments.

Since unverifiable changes are however observable to network participants, a possible solution is that, when the network gets set up, players implicitly agree on a relational contract: a self-enforcing agreement that specifies a rule for efficiently updating the network J_t (an indicator function that is 1 if the change is implemented), a generalized multi-period version of the wage, w_t , as above and a bonus payment β_t to all the network participants. We denote the contract by $\mathbb{C} = (J_t, w_t, \beta_t)_{t=0}^{\infty}$. For simplicity, we shall refer to the contract as \mathbb{C} .

Given these assumptions, the payoffs to the entire network (the combined surplus available to all of the players) in a multiple-period zero-discount setup is:

$$s = N \sum_{t=0}^{\infty} ((b_t(N) - C_t) + J_t(\alpha_h \Delta_h + \alpha_l \Delta_l > k)) \quad (8)$$

The first term is the operational profit. The second term $\alpha_h \Delta_h + \alpha_l \Delta_l > k$ captures the additional value of efficient change to the members of the network.

Given our assumptions, Theorem 3 of Levin's (2003) general model holds, and we can find the first best stationary contract by simply searching for contracts that implement first best investment rules and ensuring that the relational capital is large enough to enforce them. Letting the surplus that the parties enjoy outside of the contract \bar{s} , and the highest payment that the nodes can require $\bar{\beta}$, we can restate this result in terms of our model as:

Lemma 1. A rule J_t that generates first best surplus can be implemented with a stationary contract if and only if there is a payment schedule β_t such that $J_t = 1$ when (3) holds in every period (first best adoption of changes) and the dynamic enforcement constraint holds, meaning, at all times, the surplus of continuing in the network is

$$RC = \frac{\delta}{1-\delta} (s - \bar{s}) \geq \sum_{i=1}^N \bar{\beta} \quad (9)$$

Intuitively, we need to ensure that there is sufficient relational capital to pay for the largest possible deviation by the architect. Essentially we need to find the most costly deviation in each case, and then make sure that the network architect owns sufficient relational capital to, even in that case, make the relationship persist.

Through what follows, we focus on what amount of relational capital would be necessary to sustain first best. That is we assume the architect wants to implement the first best decision

$$\text{First best decision} : \begin{cases} J = 1 \text{ if } \alpha_h \Delta_h + \alpha_l \Delta_l > k \\ J = 0 \text{ if } \alpha_h \Delta_h + \alpha_l \Delta_l \leq k \end{cases}$$

(since we are looking at stationary contracts, we suppress the time notation). Since the conflict concerns the network architect, who partially ignores the private costs of the nodes, we need enough relational capital to provide the right incentives to the network architect. Otherwise, the threat of expropriation will lead to underinvestment: nodes (e.g., application developers) will not enter into a centralized network if they know only a share γ of ex post investment will be compensated.

The total bonus that the network architect must promise in this case is $\beta = (1 - \gamma)k$ per node. This bonus is highest when the network architect is indifferent between changing the network protocol and not changing it, thus when:

$$\alpha_h \Delta_h + \alpha_l \Delta_l = k$$

And thus the total relational capital required, is:

$$RC \geq N(1 - \gamma)k = N(1 - \gamma)(\alpha_h \Delta_h + \alpha_l \Delta_l)$$

To the extent that relational capital is larger than that level, the centralized network may attain efficient adaptation.

Result 3 [Relational capital in a centralized network]. *A centralized network with sufficient relational capital (larger than $RC = N(1 - \gamma)(\alpha_h\Delta_h + \alpha_l\Delta_l)$) may attain first best. Thus the necessary relational capital is larger the more important the unverifiable cost share $(1-\gamma)$, and the larger the value of the changes.*

Thus relational capital required is larger the larger the unobservable share of costs, the larger the cost of the changes and the larger the size of the network.

Remark (Reducing costly relational capital): Governance in a centralized network aiming to economize relational capital must reduce the network architect's temptation to expropriate the nodes.

2.3. Evidence on centralized governance with large relational capital: Google and Apple

Apple is the network architect of a closed network. It offers its iOS operating system along with the Apple App Store (iTunes). Google is the network architect for an open standard,⁵ the less centralized Android operating system, along with its Google Play Store.⁶ While hold up fears periodically flare, particularly with Apple, both Apple and Google have enough reputation to (normally) restrain their ability to hold up members of the network.

Both Apple and Google subsidize in important ways adaptation and programming investments by their developers. First, Apple and Google vertically integrate their "application suites" containing core apps packaged with the operating system and including by default not only the phone software but that for text messaging, email, calendar, etc. Moreover, they provide not only a distribution channel to apps but many essential services, including development tools, quality control and policing of free riders. Both help developers with free software, more so the more centralized Apple platform (Bessarabova 2017). Indeed, Apple provides a specific programming language (Swift) which makes programming Apps easier than at Android (for which developers must use Java and have to write more code), which has lead Google to support the Kotlin language.

⁵ Within Android, even something analogous to a fork takes place when, for instance, a partly competitive platform such as Amazon develops the software for its Kindle by adding its own proprietary features to the open version of Android.

⁶ They are the clear market leaders: Google Play had 2.8 million apps and Apple Store 2.2 million in March 2017, with the others far behind: Windows, 0.669; Amazon, 0.60 and Blackberry 0.2345 million ("Number of apps available in leading app stores as of March 2017," <https://www.statista.com/statistics/276623/number-of-apps-available-in-leading-app-stores/>).

Both platforms also provide integrated environments: Apple's XCode points out issues rapidly thanks to its built-in background compilation, making it superior to Android's Studio, which requires an explicit building stage. Both also provide help and learning materials.

Even if both platforms charge fixed fees to App developers (an annual fee of \$99 at the Apple Store and a one-time \$25 charge at Google Play), these fixed fees are probably much lower than the value for most developers of the implicit services they receive from the platform.

Apple is more centralized, and thus would seem to have greater incentives to hold up developers. However, consistently with the relational-capital interpretation, Apple is also more active than Google in filtering bad Apps out and controlling App imperfections, leading to more rejections. At Apple, a dedicated team reviews App applications according to a detailed set of standards (covering from children protection to reliability) and allows communication in case of rejection, while Google Play runs only automated tests, does not allow such communication and was slow to reverse its policy of free App placement (so long as the App is not offensive or harmful) and start punishing poorly-made Apps (Dutta 2017), a policy that had led to proliferation of useless and insecure Apps, even scams (Bell 2014). Apple also provides a users' review and rating service (to which developers can also respond) and punishes or even terminates free riders (for instance, spammers or those who manipulate the reviews and ratings, steal users' data or incur in plagiarism, which are removed from the Apple Store and expelled from its Developer Program). Both platforms also encourage success by, e.g., getting good Apps "featured" at the store or even elected "App of the Year" or "Free App of the Week".

Platforms' integration of core Apps, low fixed prices and their policing of App's quality is consistent with the incentives of developers to underinvest. First, vertical integration ensures maximum control of investment. Second, the above-mentioned difference between the value of free services and the low fixed fees charged to developers allegedly subsidizes them, containing their incentives to underinvest in the face of hold-up and adaptation risks. The difference between both platforms' fixed fees is also in line with the extent of the services they provide.

Lastly, something similar can be said about policing, as Apple's stricter review is also consistent with it pursuing greater reliability and higher quality. This elicits complaints from developers about rule ambiguity, subjectiveness, slow speed, poor quality (e.g., Dave 2015), as well as relying on little evidence and even favoring its business partners (Yang, Xinning and Waters 2017). Obviously, these complaints can be interpreted either as a sign that Apple is in a position to exercise holdup or as a sign that it is performing well its policing function to avoid underinvestment and preclude free riding. Developers agree ex ante to assign unilateral review and termination-at-will rights to Apple, which might well be optimal to both parties, given Apple's central position, incentives and knowledge, as well as the collective action dilemma faced by developers. Other early features of the review process point out that Apple takes this review task seriously: for instance, it made available an internal appeal process. After all, Apple Store generates more revenue for developers than competitive platforms (Sims 2015). Moreover, it would be probably suboptimal to rely on verifiable evidence and objective standards, as critics demand (Yang, Xinning and

Waters 2017, Dave 2015). Thus, Apple has a point when asserting that “We will reject apps for any content or behavior that we believe is over the line. What line, you ask? Well, as a Supreme Court Justice once said, ‘I’ll know it when I see it’. And we think that you will also know it when you cross it” (Apple Inc. 2017).

3. Blockchain networks: Coordination difficulties protect developers from hold-up

The technology of blockchain, as discussed in Section 1, alters substantially the contractual possibilities and the tradeoff between coordination and hold up.

3.1. Spot contracting: the possibility of splits protects specific investments

Consider now a decentralized network where in each period a node must decide whether or not to accept a proposed change. Each agent must decide on her own whether to adopt the change, according to her own preferences. Crucially, the $(N - m)$ losers may split and create another network if they are forced to implement individually-damaging updates. Suppose for now (we remove this assumption below) that such losers can act in a coordinated way. Their network will have a size $(N - m)$ and thus generate an outside value for each of them of $b(N-m)$. This provides them with an outside utility which works as their reservation value. On the one hand, it implies that many profitable changes will not be implemented. On the other hand, it entirely eliminates the threat of hold up threat, as we will make clear immediately.

Assuming by now perfect coordination among losers, the change takes place if losers prefer to obtain the benefits of the larger network, $b(N)$, plus their small gain from the change, Δ_l , to staying with the current technology, which yields the benefits of the smaller network, $b(N - m)$, and avoids the cost of adaptation k :

$$\Delta_l + b(N) > b(N - m) + k \tag{10}$$

That is, in a spot decision by a fully decentralized network, the nodes with the least to gain dictate to the rest. Only win-win changes are implemented.

When compared to the efficient condition to preclude a split, which we can rewrite, from assumption 1, as


$$\alpha_l(\Delta_l + b(N) - b(N - m) - k) + \alpha_h[b(N) - b(m)] > 0$$

the node’s individual decision does not consider the additional value of the network effect to the winning nodes. To restate, in a decentralized network, a coordinated change that is welfare improving only takes place if all the losers prefer it, which is a highly demanding threshold.

Potentially, a decentralized network could be structured to decide adopting the change by a majority rule. The problem in blockchain (and this is what makes it decentralized) is that individual nodes are always able to renege on any commitment to abide by the voting—splits are always possible because in blockchain the individual adoption of a given piece of software by a node is not verifiable.

Result 4 [Decentralized networks only implement win-win changes]. *Win-lose changes are hard to implement in a decentralized network. Losers may block the change, leading to a large loss of value due to loss of coordination gains. These losses from coordination are larger the larger the number of losers.*

Figure 2: The inefficiency of decentralization: Too much inertia



3.2. Multiple equilibrium and expectations management

In the previous subsection, we assumed for simplicity that the losers can act together in adopting or not the proposed changes. More realistically, in a fully decentralized setting, such coordination among losers is not possible: both the losers and the winners must act according to what they expect all other players will do.

Although a fully dynamic analysis of the possible equilibria is possible (see, for the blockchain case, Biais et al. 2018), a simple, static, analysis allows us to capture the main features of this problem.

Result 5. [Multiple Nash Equilibria]. *Suppose Assumption 1 holds, so that it is efficient to have one chain and implement the change. Then there exist values of k such that there exist three Nash equilibria:*

1. *Successful fork: All nodes accept the change, a single network of size N ensues.*
2. *Split: Losers stay on the old network, now of size $(N-m)$, winners create a new network of size m with the new protocol*
3. *No change: All nodes stay on the current network of size n , no change is adopted.*

Proof.

1. Consider the problem of a single losing node, who knows that all other nodes accept the new protocol. She decides to accept as well if:

$$\Delta_l + b(N) - b(1) \geq k$$

Which is always true, since network effects are “large” by assumption, $b(1) = 0$ and $b(N) \gg k$. Network effects ensure that being the only hold out is never optimal. Thus there is always an “all change” equilibrium. (Note that, a fortiori, this is even more true for a winning node).

2. To see that there may exist equilibria with splits, suppose that the m winning nodes split and adopt the change while the $(N-m)$ losing nodes stay on the old protocol. This is an equilibrium as well as long as a node who is a loser does not prefer to adopt the change,⁷

$$k \geq \Delta_l + b(m) - b(N - m),$$

and a winning node does not prefer to abandon the change,

$$\Delta_h + b(m) - b(N - m) \geq k$$

Both may be true for intermediate k as long as the splitting group is neither too large nor too small:

$$\Delta_h + b(m) - b(N - m) \geq k \geq \Delta_l + b(m) - b(N - m)$$

To see that this condition holds for some parameter values, suppose $m=N-m$. Then $\Delta_h \geq k \geq \Delta_l$.

3. To see that “no protocol change” may also be an equilibrium, note that if all nodes stay in the old equilibrium, the payoff for the single changer is positive if:

$$\Delta_h + b(1) > b(N) + k$$

$$\Delta_h > b(N) + k$$

⁷ Note that to simplify slightly we assume a node is small so adding it to the network does not change the payoff $b(m + 1) \approx b(m)$, $b(N - m + 1) \approx b(N - m)$.

Again, with “large” network effects, $b(N) \gg \Delta_h$ this condition never holds and “staying” is always a Nash equilibrium.

In sum, in a decentralized network, we generally find multiple equilibria. When change is efficient and network effects are important, at least always two (all change and no one changes) and potentially three, whenever a sufficiently sizeable group of losers exist.

3.3. Evidence on splits and inertia in blockchain

Examples of the two most valuable networks, Bitcoin and Ethereum, show the two sides of the coin we have just studied. On Bitcoin, difficulties in coordinating adaptation. On Ethereum, the protection of the interests and investments of network members (mainly developers, as miners’ investments are not chain-specific but easily redeployable) through the possibility of their undertaking splits.

Bitcoin Cash Split. At Bitcoin, decisions take place through the interaction between the main constituencies (core developers, miners, users), whose proposals, mining and investment decisions determine, via a hard or soft fork, any change in the rules and prices. Before launching important proposals there is usually some bargaining and in key cases some agreements have even been formalized. However, participants are not committed to these agreements at the time of adopting the change in the protocol.⁸

In the summer of 2017, the Bitcoin community was trying to reach a consensus to solve the technical, economic and ideological conflict between miners, who wanted bigger block sizes, and code developers, who stressed security. The choice had serious consequences for the different participants, including blockchain applications with different business models: “Many [blockchain startups] have business models that would be affected by how the block-size problem is solved. Blockstream, a firm that employs some Bitcoin Core developers, builds ‘sidechains,’ the sort of secondary system that would be more in demand if bitcoin itself doesn’t start accepting more transactions. On the other hand, there’s BitPay, which has sold merchants the idea of bitcoin as a low-fee retail payment system, and for whom the strangled state of the bitcoin blockchain has been a serious headache.” (Morris 2017).

This heterogeneity of preferences among users meant that some but not all nodes incorporated the hard fork upgrade launched by a small and well coordinated coalition, and a split took place, with the creation of another coin (named “Bitcoin Cash”), leading to two incompatible networks. The solution was therefore determined in the market and led to two different designs: A few days after the fork, it had mined the first 8BM block.

Interestingly, at the time of the split, there were fears that this would damage the total value of the system but, these fears did not materialize. (Total value fell later but this later

⁸ For instance, the “New York Agreement” related to the Segwit2x update (Dinkins 2017).

fall cannot be imputed to the split). Even though Bitcoin prices had soared, Bitcoin Cash became the fourth cryptocurrency by market capitalization, fluctuating in the following two months around 10% of that of Bitcoin, with a price around 12% of that of Bitcoin; data which admittedly was probably not informative given its relative lack of liquidity.⁹ This suggests that the split may have indeed protected the interests of the minority.

The Ethereum Classic split in 2016. The DAO (an acronym that stood for “Decentralized Autonomous Organization”) was a sort of venture capital fund within Ethereum to which any investor could contribute ether (Ethereum’s cryptocurrency), thus purchasing shares and voting rights, which they then used on the projects they decided to support. In June 2016, after it had raised up to \$250 million from thousands of backers, it emerged that someone had used a bug in its code to “siphon” from its original owners about \$60 million worth of ether. After using similar tactics to fight a so-called DAO war for weeks, the Ethereum team decided to implement a hard fork that, once the changes proposed by the Ethereum team were adopted by miners, by simply upgrading their software, would effectively delete the allegedly fraudulent transactions and refund the money to its previous owners.

However, this denied the immutability that was predicated of smart contracts, which were supposed to make enforcement automatic and dispute resolution unnecessary. The Ethereum team was accused of conflict of interests and, in particular, of supporting the conclusiveness of transactions only when it served their interests. Under this malevolent interpretation, they were willing to endanger immutability to protect their investment. Alternatively, a benevolent interpretation is that, by endangering a bit of immutability they not only protected their investment but also the broader long-term interest of the whole community. In the law-and-economics terms mentioned above, they were aiming for “efficient contractual breach.”

At last, the community was split and some important miners and exchanges started backing an alternative currency, called “Ethereum Classic” (ETC), which relies on the original blockchain.¹⁰ Those who held ether on it retained their rights, but for the funds stolen in the DAO attack. One of the beauties of hard-fork splits, is that the evolution of both coins in the market, in terms mainly of price and volume, provides an imperfect but independent measure of the relative value of the two sets of rules.¹¹

⁹ Calculated on October 2, 2017, with data obtained from <https://coinmarketcap.com/>.

¹⁰ Later incidents have triggered somewhat similar controversies and the issues remain far from solved (e.g. O’Leary, 2018).

¹¹ In the DAO case, the evolution of market prices may offer some support to an efficient “contractual breach” hypothesis. Certainly, Ethereum Classic survived but ten months later, its market capitalization was only 4.88% of Ethereum. Even if its price had increased between those two dates by a multiple close to 13, this was much less than Ethereum’s 24.4. Given that, at that point, the main difference between the two coins was the original conflict, the market (and, crucially, the exchanges, as Classic was only traded by a few of

Inertia in protocol changes: Segwit2x. Multiple instances of the inertia of blockchain networks can be reported. Consider for instance the update Segwit2x, which would increase block size and allow the Bitcoin network to lower fees and become more reliable was ready for November 16, 2017. The debate on the need to allow the network to scale up has been going on for three years.¹² The community was divided about its convenience. Promoters finally abandoned the update, fearing another split, after observing prices in the new futures market for cryptocurrencies as well as other signals hinting lack of consensus.

3.4. Improving blockchain governance: pre-mining and soft governance

Blockchain networks follow several strategies to tackle the main conflicts that the previous model highlights. They are designed to minimize the number of losers and display a variety of governance devices to encourage nodes to reach good equilibria and avoid bad equilibria—in short, they aim to facilitate the coordination of winners and hinder that of losers.¹³

Pre-mining

Practically every new blockchain platform has been launched together with a crypto currency through an “Initial Coin Offering” (ICO). Usually in exchange for a established coin such as bitcoin or ether, founders issue what is called “tokens” which are a digital assets that can be used as a means of payment in the platform, to buy or sell products or services.

them) was apparently not very appreciative of the conservativeness of Ethereum Classic with respect to immutability. Moreover, Ethereum Classic’s claims of code-as-law were somehow diluted, as, at least for fraud cases, it relies on standard legal recourse—what could also be understood as a form of third-party contract completion—and blockchain integrity is dissociated from self-enforcement. Based on data from <https://coinmarketcap.com/> (August 13, 2017).

¹² “Bitcoin’s scaling debate has been going on and the need to seize this opportunity of increasing capacity beyond what is provided by Segregated Witness (SegWit).” <https://www.forbes.com/sites/ktorpey/2017/10/31/heres-why-bitcoin-businesses-are-pushing-for-a-protocol-change-without-clear-consensus/#33a1eca27b40>.

¹³ Given the two-sided nature of its markets, centralized platforms also face multiple-equilibrium problems, and they must manage them, for instance, by relying on coordination bias to decide between subsidizing buyers or sellers (Halaburda and Yehezkel 2016). The problem for blockchain platforms is much more complex: there are multiple levels of platforms, with additional key players (mainly, miners) and, most essential, design decisions (e.g., whom is subsidized by whom) are decentralized.

The ICO often involves pre-mining and also selling tokens at a discounted price or even granting valuable options to network participants, in order to facilitate the formation of a community, getting liquidity in the system, facilitating trading and participation for programmers and service providers. Some tokens are even sold before the ICO, in highly discounted “token pre-sales.”

There are complementarities between the cryptocurrency and the platform.¹⁴ First, and most obviously, launching a platform increases the viability of the currency, which becomes a medium of exchange in that platform, so that transaction demand anchors to some extent the value of the platform.

Second, the cryptocurrency also benefits the platform. Pre-mining the cryptocurrency allows founders and those allowed to pre-mine to capitalize on the future value of the platform and thus aligns to some extent their investment incentives with those of the network. By providing the promoters of the new platform with some coins, and distributing them strategically so that potential losers hold coins, therefore weakening the constraint that $\Delta_l + b(N) \geq b(N - m) + k$ to:

$$\Delta_l + b(N) + (v' - v) > k + b(N - m)$$

where $v' - v$ is the increase in the value of the coin “stash” in the hands of such loser. To the extent that $v' - v$ is large, we can get closer to efficient adoption, particularly since many critical changes and investments are implemented in the early stages of the network.

For instance, Ethereum counts with a governing non-profit foundation (Stiftung Ethereum) which has “the purpose of managing the funds that were raised from the ether sale in order to best serve the Ethereum and decentralized technology ecosystem” (Ethereum 2017). The foundation holds a substantial amount of ether that was premined before the ICO aligns its incentives with the aim of maximizing the value of the Ethereum ecosystem.

However, even if this initial allocation of coins was initially right, it cannot provide a long term solution to the coordination problem. First, it is unclear ex ante who the losers from future changes will be. Second, its effectiveness decreases with the lapse of time. Indeed, newcomers (such as future developers) do not hold any of the pre-mined coins. Perhaps more importantly, founders and early holders are tempted to sell, as the vesting and sale restrictions typical of IPOs are not easy to implement.

Some platforms have introduced explicit formal mechanisms to overcome these limitations. For instance, the above mentioned Ethereum-based DAO issued a DAO token that allowed token holders to vote on projects and receive a share on the profits from them. However, such efforts have had limited success, not least because the SEC has ruled that such tokens

¹⁴ This is not entirely new: large centralized platforms such as Facebook and Amazon have introduced, private digital currencies (see Gans and Halaburda, 2015).

are securities and thus their sale needs to be regulated as such (that is, like an IPO of stock).¹⁵

Remark [Life Cycle]. *A crypto currency or any explicit profit participation economizes on relational capital by allowing networks to reward members even in the absence of cash flow. We expect them to be most important early on in the relationship.*

Soft governance

The multiplicity of equilibria opens the door to using “soft governance” to facilitate decision making. The idea is not to “order” or “decide” as in standard centralized governance, often not even to lead but to merely nudge players to coordinate on the good equilibria.

The effectiveness of these mechanisms of “decentralized governance” is intrinsically limited, because of the simple fact that in a blockchain network (even with on-chain decision mechanisms) splits are always possible, subject only to the condition of them achieving enough support in the community. Since individual nodes keep the right to run or not updates, all these soft governance devices can do is to coordinate the expectations of all participants on the equilibrium selected, reinforcing the expectation that everyone else will chose accordingly. Therefore, instead of allocating decision rights across participants, soft governance merely allocates nudging rights among them.

Soft anarchy: Lack of governance at Bitcoin. At Bitcoin, there is practically no ex post soft (or hard, for that matter) governance as such. It has not even had a leader, as its supposed founder (who remains anonymous) retired soon after its creation. Decisions take place through the interaction between the main constituencies (core developers, miners, users), whose proposals, mining and investment decisions determine, via a hard or soft fork any change in the rules and prices. Usually, before launching important proposals there is bargaining and in key cases some agreements have even been formalized.

¹⁵ Securities and Exchange Commission, Securities Exchange Act of 1934, Release No. 81207, July 25, 2017: “This Report reiterates these fundamental principles of the U.S. federal securities laws and describes their applicability to a new paradigm—virtual organizations or capital raising entities that use distributed ledger or blockchain technology to facilitate capital raising and/or investment and the related offer and sale of securities. The automation of certain functions through this technology, “smart contracts,” or computer code, does not remove conduct from the purview of the U.S. federal securities laws”.

However, participants are not committed to these agreements at the time of adopting the change in the protocol.¹⁶

Moreover, the lack of centralized governance in decentralized networks such as Bitcoin makes it difficult to properly finance protocol developers when they are funded externally (Wirtdum 2016). Initially, Bitcoin was created by volunteers. This became unsustainable when founders lost interest and the demands for protocol updates kept increasing. The core foundation has been financed with donations and sponsorships, following the steps of open-source software (Evans, Hagi and Schmalensee, 2006:403–408), but this has also been criticized for allowing for free riding by non-contributors and exacerbating potential conflicts of interest (Wiecko 2017).¹⁷

Soft autocracy: Leadership at Ethereum. Unlike at bitcoin, and in spite of the lack of centralization, Ethereum relies strongly on the existence of a “benevolent leader”, its founder Vitalik Buterin, who crowdfunded issuing a new cryptocurrency (“ether”) in August 2014 when he was 19-year old. He has described his functions as “Constantly keep an eye out as to what users are looking for, and make sure that we are satisfying people's concerns. Be transparent about what we're doing. Have frequent developer calls between the various client developer teams and researchers, and publish the minutes and audio as much as possible. If there are controversial choices or hard tradeoffs to be made, present the tradeoff and do our best to give the community an input. For anything truly controversial, try to gauge community consensus via carbonvote or similar tools as well as other polls”.¹⁸ Moreover, Buterin has aimed to instill an ethos in the network, a “strong culture of mutual respect and tolerance”.¹⁹

The visible and active presence of Buterin in the Ethereum community is a testimony to the importance he attaches to the importance of coordinating expectations: He is a frequent participant in Reddit discussions, and often posts longer and substantive articles on his website, discussing the main strategic decisions to be made by Ethereum. Relational capital

¹⁶ For instance, the “New York Agreement” related to the Segwith2x update (Dinkins 2017).

¹⁷ At Ethereum and other networks by contrast. At Ethereum and at Dash, for instance, 10% reward goes to developers.

¹⁸ For an online reddit discussion of Ethereum versus Bitcoin governance, see “How Is Ethereum Governance Different from Flawed Bitcoin Governance?” *Ethereum Subreddit*, https://www.reddit.com/r/ethereum/comments/5ybocd/how_is_ethereum_governance_different_from_flawed/.

¹⁹ E.g., “How does the Ethereum community feel that it can overcome/is overcoming the governance problem that is killing Bitcoin?” *Ethereum Subreddit*, https://www.reddit.com/r/ethereum/comments/61tpul/how_does_the_ethereum_community_feel_that_it_can/.

seems to be present here in at least two directions. First, the sharing valuable knowledge increases Buterin's reputation. Second, his reputation in turn reinforces his authority and power as the leader of the community.

Soft aristocracy: Contributors and masters. The massive split caused by The DAO in 2016 took place in spite of the strong leadership of Buterin. This has led to a variety of proposals.

For instance, the Yellow Paper Committee, proposed by Gavin Wood, a cofounder and one of its three top Ethereum developers, argues for "a blockchain-based organization [structured as a DAO] to manage the evolution of the Ethereum protocol specification" (Wood 2016), concentrating decision rights on the technical elite: The wider community of "middleware developers, application developers, infrastructure projects, commercial endeavors powered by Ethereum, substantial holders of ether and affiliate organizations" would form an Ethereum General Assembly (EGA) as "a forum to help air, clarify and formalize the views of the community." Then, a highly selective group of contributors (those who had contributed an independent implementation of the protocol in the previous six months) would form the Ethereum Implementers Group (EIG), who would ratify changes to the constitution and the specification. No single person or entity would be allowed to influence more than 25% of the total membership. Each EIG member would have a vote and a three-month veto right at the EIG. The whole EGA would have a right to vote but would not enjoy veto rights at the EIG. The EIG chair, would submit potential changes (so-called EIPs, for Ethereum Improvement Proposals describing the standards for the platform) to the EIG through a two-step budget-allocation process for code development, which would then be decided by the EIG's voting. Revealingly, the chair might also propose modest expenditures for accommodation and travel, but not subsistence, to EIG meetings.

Another aristocratic solution based on creating "master nodes", has been implemented by Dash, which defines itself as a "privacy-centric digital currency with instant transactions" based on Bitcoin software. It also aims to introduce leadership by a "technical elite", a sort of aristocratic democracy of nodes for governance and budgeting, in which it attempts to collocate decision rights with long term incentives.²⁰ Specifically, master nodes are invested in the future of the currency. In particular, it is they who allocate funds among proposed development changes and decide on protocol changes. To make decisions (establishing the budget to pay for the core team, contractors and other costs; increasing block size, etc.), each master node operator has one vote. Master nodes are also essential to achieve anonymization. Anyone can become a master node by ownership of 1,000 Dash

²⁰ Sources: Miller (2016), Wiecko (2017), Dash White Paper (<https://dashpay.atlassian.net/wiki/spaces/DOC/pages/1867864/What+is+a+masternode>).
<https://dashpay.atlassian.net/wiki/spaces/DOC/pages/1867864/What+is+a+masternode>

and are paid by the network with 45% of the block reward that used to go to miners (another 10% go for the budget). It has successfully handled a number of hard forks without splits.

Soft democracy. Several networks are also trying to directly coordinate the entire community with the help of smart contracts.

An example is Buterin’s proposal to replace the current “difficulty time bomb” for “a more general ‘governance gadget’ for hard forks”. The “Difficulty Time Bomb” is an exponential increase in difficulty implemented in Ethereum on September 7, 2015, to incentivize everybody to switch to the new blockchain once the hard-fork moving it from Proof-of-Work to Proof-of-Stake is implemented, thus avoiding the risk that the chain is duplicated. (Given the greater difficulty, miners could not cope and the blockchain would freeze, hence the “Ice Age” name).²¹ Instead, under Buterin’s alternative “Bomb 2.0” proposal, participants would vote by sending ether on a hard fork and, once enough people approved it, the hard fork takes place automatically. It would still be “a bicameral model: a fork requires consent from both holders and devs because devs have to find the fork acceptable enough to be willing to write up the code” (Buterin 2016).

Similarly, Tezos defines itself as “the self-amending cryptographic ledger”. This startup, which raised \$232M in July 2017, features on-chain governance to achieve a smooth upgrading of its protocol instead on relying on hard forks (Xie 2017). The process works by having developers submit upgrade proposals, which include a compensation for developers’ work. This compensation should provide incentives and avoid conflicts of interests. Holders of the application coins (“tezzies”) then decide by delegated proof of stake, which should be cheaper, as it does not waste electricity; provide better incentives, including penalties; and allow transfer of voting rights to those with time and knowledge. (As with corporate “empty voting” by means of financial derivatives [Hu and Black 2007], it is doubtful if this delegation of votes—now popular in many innovative governance structures in blockchain—could exacerbate conflicts of interests, as it separates economic and political rights). Funds raised in the crowdsale will be managed by a Swiss foundation with can veto proposals in the first year but has not other control over the submission and upgrade process. Moreover, its continuity will be decided by the on-chain governance system.²²

A similar solution is the EVM Dfinity, a sister network of Ethereum, which would provide for its apps an automatic and decentralized decision system: a “Blockchain Nervous System with privileged control over token ownership”, which would be used instead of the

²¹ Based on “What is the Ethereum Ice Age?” *CryptoCompare*, July 5, 2017, <https://www.cryptocompare.com/coins/guides/what-is-the-ethereum-ice-age/>.

²² Interestingly, in spite of this effort to improve governance, Tezos suffered in early 2018 a nasty governance crisis (Dale 2018).

automatic “code is law” of Ethereum (Williams 2017a). It is designed to prevent hacking and aims to freeze “miscreant smart contracts that harm the interests of those using the platform” (Williams 2017b). It would rely on quasi proof of stake (“mining clients connect to the network using a ‘mining identity’ that is acquired by making a security deposit in tokens”).

4. Comparing optimal centralized and decentralized networks

When will a centralized or a decentralized network be preferred? The choice depends on their ability to resist opportunism—what one could call their contractual “robustness”. This hinges, in turn, on the relation between the future value of staying in the network (the relational capital) and the reneging temptation to key players or coalitions of players.

A centralized network, which we studied in Section 2, confronts a relatively simple problem: precluding the reneging temptation of the network architect, who may otherwise expropriate the specific investments of the nodes.

A decentralized network confronts a more complex problem, as analyzed in Section 3: preventing the temptation of splits from “coalitions of losers”. Moreover, in a centralized network relational capital is easier to visualize in terms of reputation and future gains from trade because most of it is owned (as the model assumed) by the network architect—it is *centralized* itself. In contrast, relational capital is dispersed across the whole community in a decentralized network.

We proceed next to compare, optimally-designed centralized and decentralized networks—meaning by optimal design that they are able to preclude their essential problems of, respectively, hold ups and splits. Obviously, the terms of this comparison are exceedingly generous to both structures, since we assume, first, that the task of optimally compensating losers is itself costless. Secondly, in the centralized networks, there is no single individual architect who centralizes all relational capital, but a whole set of hierarchical layers of principal and agents with their own incentive problems. After making the comparison between these idealized optimal types, we will examine more deeply in Section 4.3 how difficult is for each type of networks to reach such optimality in the real world.

4.1. An optimal decentralized network

The glue holding together the blockchain is its future value. At any stage, a coalition of members may try to split and form a separate network, by refusing to implement a protocol change. In a decentralized network, losers who implement a change lose each $k - \Delta_l - (b(N) - b(N - m))$. For them not to be tempted to split, the network relational capital must be larger than the worst case scenario. That is the point where the

network as a whole faces the largest reneging temptation. The worst case scenario (and thus the highest reneging temptation) is where the cost of adapting to the new technology is highest, and makes adoption marginally profitable, that is when

$$\alpha_h \Delta_h + \alpha_l \Delta_l = k \quad (11)$$

Let RC_D , the relational capital of a decentralized network, be the net present value of the future gains from the network. For this relational capital to be larger than the worst reneging temptation of the $N-m$ losers we must have.²³

$$RC_D \geq (N - m) \left(\alpha_h (\Delta_h - \Delta_l) - (b(N) - b(N - m)) \right) \quad (12)$$

Result 6 [Relational capital in decentralized network]. *The relational capital needed in a decentralized network is larger the higher the heterogeneity in benefits among network members, and the lower the larger the network effect.*²⁴

4.2. Which of the optimal networks requires less relational capital?

Thus, an optimally-designed decentralized network will be preferred if it needs less relational capital:

$$RC_D < RC_C$$

$$\alpha_l \left(\alpha_h (\Delta_h - \Delta_l) - (b(N) - b(N - m)) \right) < (1 - \gamma) (\alpha_h \Delta_h + \alpha_l \Delta_l)$$

²³ We are not being explicit about how the payments are actually made. We shall come back to this point later. For now, suppose there is a timeline that includes first a vote on a change, then a realization of winners and losers, then a bonus for losers paid by winners. The earnings of the winners from the change, in this worst case scenario, are 0:

$$\Delta + b - k - \frac{\bar{w}(n-m)}{m} = \Delta + b - \frac{m}{n} \Delta - b - \frac{(n-m)\Delta}{n} = \Delta \left(1 - \frac{m}{n} - \frac{(n-m)}{n} \right) = 0$$

²⁴ When splits are possible, there are two relevant cases: inefficient split and inefficient status quo. Given that k is random, relational capital is needed in both types of situation but always to compensate losers in order to reach the assumed unanimous-change first best. In both cases, RC is the same as in the no-split case. Therefore, contemplating splits does not change anything if the relational mechanism is effective in precluding inefficient change.

Or equivalently

$$(\alpha_h \Delta_h + \alpha_l \Delta_l)(\alpha_1 - (1 - \gamma)) < \alpha_1(b(N) - b(N - m) + \Delta_1) \quad (13)$$

Proposition 1. Network choice. *An optimally-designed decentralized network will be preferred iff network effects are large, if the gains for the losers are large (so that changes are win win). Centralized networks are preferred is the risk of hold up by the network architect is small ($(1 - \gamma)$ is small), that is when adoption costs are easy to verify.*

4.3. Comparing real networks

Comparing ideal economic optima is fruitful only to a certain point (Coase 1960). To explore how likely centralized and decentralized networks are to reach such ideals, we must compare them from a static and an evolutionary perspective.

Avoiding holdup and optimizing change not only requires that the networks own enough relational capital to be able to compensate losers. The model assumes that compensating losers is equally costless in the two types of networks. This assumption seems to depart from reality particularly in the case of decentralized networks, since it implies that it is costless to make optimal payments between potentially large numbers of decentralized winners and losers (footnote 23). The empirical analysis of Sections 2 and 3 indeed confirms that decentralized networks experience serious difficulties to structure such payments. They therefore suffer an additional disadvantage in this dimension, as may have been suspected from not having a single central payer and repository of reputation or reputational capital, more generally.

But centralized networks also suffer from incentive conflicts which remain hidden in our model. This is simply because we modeled both network architects as simple “black boxes” whose internal organization costs we also assumed to be zero. However, in reality the two types of architect could hardly be more different. The architect of a centralized network is a huge firm which integrates many functions while decentralized network architects are at most tiny teams. Centralized networks economize on the costs of organizing the network by increasing the costs of organizing internally, the firms-versus-markets tradeoff discovered by Coase (1937). Therefore, by applying the same black-box assumption to the two architects, our model grants centralized networks an unrealistic advantage.

For instance, centralized architects exhibit multiple agency levels, which reduces the effectiveness of any given amount of relational capital, as agents’ interests are imperfectly aligned with those of the firm and may be tempted to cheat on nodes for their personal gain. Centralized architects therefore either spend capital trying to align agents’ interests and have to deploy additional relational capital to preclude hold up. Such multilevel principal-agent problems are quite prevalent, as an enormous recent literature has shown.

Lastly, from an evolutionary perspective, centralized networks likely receive more institutional support and benefit more from accumulated managerial knowledge, which are by now less developed and therefore less adapted to the new decentralization P2P possibilities that blockchain has just opened up.²⁵ Note that it took centuries to develop the institutions and organizational solutions applied for the effective governance of centralized relational capital—that is, a corporation with multiple levels of agents acting as the architect and then contracting effectively as a single unit with a myriad of network members which may in turn be big firms themselves.²⁶ Neither currently available institutions nor organizational techniques suit the demands of decentralized governance, which is in itself a wholly new phenomenon.

Therefore, we should not be too harsh when judging the current state of decentralized governance in blockchain networks. Instead, we should pay more attention to its ability to improve over time and evolve solutions to its current problems. Considering the evidence in section 3.4, the industry seems to be starting a *Hayekian* discovery process (Hayek 1982). Time will say if and how fast are participants able to find governance solutions that best economize on relational capital by distributing both capital and decision rights in a way that encourages good equilibria.

5. Conclusions: Centralization trade-offs and the deficit of governance in decentralized networks

Contrary to what some early blockchain participants believed (applying the “code is the law” principle in Lessig 1999), even smart contracts are and will remain incomplete. More so when the contract refers not to a simple exchange of goods and services but to such

²⁵ The fact that decentralized networks are less regulated (for instance, the ICOs) may favor its finances in the short run but does not help and may even hinder their governance. On this regard, lack of enabling regulation may be a serious handicap, as it was at the inception of the modern corporation (Harris 2000, Arruñada 2010).

²⁶ With respect to legal solutions, firms rely heavily in ready-made solutions that we now give for granted but took thousands of years to develop, such as the protection *in rem* of good-faith third-party acquirers, agency law enabling agents to commit their principals, legal personality to handle conflicts between personal and corporate creditors without endangering legal entities, and, quite recently, corporate governance to enable separation of ownership and control (e.g., Hansmann and Kraakman 2000, Arruñada 2012). Similarly, with respect to organizational techniques, little could firms achieve without, e.g., double-entry accounting, divisionalization, transfer pricing or the asymmetric allocation of decision rights characteristic of franchising (e.g., Chandler 1977, Arruñada, Garicano and Vázquez 2001).

complex and conflictive interactions as those we have sketched here (mainly, the protocol governing the blockchain).

However, unlike in centralized networks, in a decentralized network nodes' investments are automatically protected from expropriation by the fact that any change requires their consent.

In particular, by achieving enough support from other nodes, they are always theoretically able to force a split and thus unanimity is required to change network protocols. However, this need for consensus comes at the cost of hindering adaptation and coordinated change. Given that different equilibria are possible, blockchain networks may stall in an outdated solution or repeatedly split as different subgroups of agents, developers and miners, having different objectives, seek to impose, through sheer force of commitment, their preferred solution on the entire network. This could fatally weaken the blockchain model. Therefore, in decentralized networks, the "subeconomy" (Holmstrom, 1999), all the governance tools available to organizations, should be brought to bear on ensuring coordinated adaptation and change. A key element here is that this will not be standard governance but "soft" decentralized governance.

Given these tradeoffs, we expect decentralized networks to be preferred when the risks of expropriation are high, as their decentralization automatically protects member investors from expropriation. Conversely, when heterogeneity is high, centralized platforms, which easily adopt changes, are more likely to support efficient contracting.

After exploring how emerging solutions, such as pre-mining coins and other soft governance tools help to deal with the incentive problems present in these networks, we have concluded that current decentralized governance is very far from solving them adequately. Compared to the advantages that, for instance, stock ownership and options confer to the emerging network that is a new corporation, ex post alignment is quite primitive. Coins or tokens, for instance, are yet woefully inadequate—no-one can be forced to hoard coins in order to align his incentives, anyone can sell them as needed, or worse, founders have incentives to do so in search of diversification.

But there is merit in decentralization, and we are in the presence of a very young, competitive and well-endowed evolutionary process searching for better solutions. Moreover, there is an important technical reason to be hopeful: blockchain networks are unlikely to fall into the traps of inferior equilibria characteristic of network industries, with their lock in and tipping features. The reason is that blockchain offers a new set of tools that, effectively used, hold the potential to avoid these bad equilibria, as we discussed. The key difference is that, while each game with network externalities is a one-off game, in blockchain we observe games with multiple stages (many protocol updates), where the potential for competition with other networks provides incentives for coordinating on good equilibria, and for evolving those decentralized governance tools that are better at selecting good equilibria, as we discussed in Section 3.4.

We submit that the platform that best solves these incentive problems will have a large advantage in this race. Right now the platform that has travelled furthest in this road is Ethereum, where the Ethereum Foundation is tasked with coordinating decisions in the network. However, as the split of Ethereum Classic showed in 2016, even Ethereum is

unable to deal effectively with consensual change. Other emerging platforms are competing, as we discuss in the main text, on all of these dimensions.

In this paper we have focused on the organizational difficulties that blockchain faces to achieve optimal adaptation. But blockchain also raises related bounded-rationality and social welfare considerations. It is particularly worth exploring if individuals would really want to live in a world of smart contracts, in which they alone are expected to take the full responsibility such contracts entail. (Note that, without intermediaries, the whole value of assets and transactions hinges only on individuals preserving and properly using their cryptographic keys). When freedom comes with responsibility, individuals often prefer to give up freedom, and true P2P trading through blockchain requires maximum individual responsibility.

This bounded rationality limitation might therefore decide where blockchain ends up enjoying greater comparative advantage: likely more for small-value transactions (e.g., related to the Internet of Things) than for those dealing with large-value assets such as real estate. Moreover, it might also define the supposed capacity of blockchain to defuse the current threat that Internet giants pose. If and once blockchain networks achieve effective decentralized governance, they will still face a steep slope—in terms of individuals' bounded rationality—to fulfill the promise of outcompeting those giants. But this is a problem for future work.

6. References

- Apple Inc. , 2017, “App Store Review Guidelines,” <https://developer.apple.com/app-store/review/guidelines/>.
- Arruñada, Benito, 2010, “Institutional Support of the Firm: A Theory of Business Registries,” *Journal of Legal Analysis*, 2(2): 525–76.
- Arruñada, Benito, 2012, *Institutional Foundations of Impersonal Exchange: Theory and Policy of Contractual Registries*, Chicago: University of Chicago Press.
- Arruñada, Benito, 2018, “Blockchain’s Struggle to Deliver Impersonal Exchange,” *Minnesota Journal of Law, Science & Technology*, 19: 55–105.
- Arruñada, Benito, Luis Garicano and Luis Vázquez, 2001, “Contractual Allocation of Decision Rights and Incentives: The Case of Automobile Distribution,” *Journal of Law, Economics, and Organization*, 17(1): 256–83.
- Arthur, Charles, and Samuel Gibbs, 2014, “The Hidden Costs of Building an Android Device,” *The Guardian*, 23 January, <https://www.theguardian.com/technology/2014/jan/23/how-google-controls-androids-open-source>.
- Athey, Susan, Christian Catalini, and Catherine E. Tucker, 2017, “The Digital Privacy Paradox: Small Money, Small Costs, Small Talk,” MIT Sloan Research Paper No. 5196-17, September 27.

- Barabas, Chelsea, Neha Narula, and Ethan Zuckerman, 2017, "Defending Internet Freedom through Decentralization: Back to the Future?" The Center for Civic Media & The Digital Currency Initiative, MIT Media Lab, August 2017, http://dci.mit.edu/assets/papers/decentralized_web.pdf.
- Bell, Killian, 2014 "Android Dev Scams More than 10,000 Customers with Fake Antivirus App," *Cult of Android*, April 14, <https://www.cultofandroid.com/60926/android-dev-scammed-10000-customers-fake-antivirus-app/>.
- Bessarabova, Elina, 2017, "6 Differences Between iOS and Android App Development," February 21, <https://themindstudios.com/blog/6-differences-between-ios-and-android-app-development/>.
- Biais, Bruno, Christophe Bisiere, Matthieu Bouvard and Catherine Casamatta. 2018. "The Blockchain Folk Theorem," *Swiss Finance Institute Research Paper* No. 17-75 (January 5), <http://dx.doi.org/10.2139/ssrn.3108601>.
- Böhme, Rainer, Nicolas Christin, Benjamin Edelman, and Tyler Moore, 2015, "Bitcoin: Economics, Technology, and Governance," *Journal of Economic Perspectives*, 29(2): 213–38.
- Bresnahan, Timothy, and Shane Greenstein, 2014, "Mobile Computing: The Next Platform Rivalry," *American Economic Review*, 104(5): 475–80.
- Buchholz, Nicholas, 2016, "Spatial Equilibrium, Search Frictions and Efficient Regulation in the Taxi Industry," *Working Paper v.* presented at CEMFI, https://www.cemfi.es/ftp/pdf/papers/wshop/taxi_draft.pdf.
- Buterin, Vitalik, 2016, "Proposal: Bomb 2.0," *Ethereum Improvement Proposal* 138, August 6, <https://github.com/ethereum/EIPs/issues/138>.
- Buterin, Vitalik, 2017, "Notes on Blockchain Governance," Vitalik Buterin's website, December 17, <https://vitalik.ca/general/2017/12/17/voting.html>.
- Casadesus-Masanell, Ramon, and David B. Yoffie, 2007, "Wintel: Cooperation and Conflict," *Management Science*, 53(4): 584–98.
- Catalini, Christian, and Joshua S. Gans, 2016, "Some Simple Economics of the Blockchain," NBER Working Paper 22952.
- Chandler, Alfred D., 1977, *The Visible Hand: The Managerial Revolution in American Business*, Cambridge, MA: Belknap Press.
- Coase, Ronald H., 1937, "The Nature of the Firm," *Economica*, 4(16): 386–405.
- Coase, Ronald H., 1960, "The Problem of Social Cost," *Journal of Law and Economics* 3(1): 1–44.
- Dale, Brady. 2018. "Tezos Has a New Board, But What About the Money?," *CoinDesk*, February 2, <https://www.coindesk.com/tezos-new-board-mean-money>.

- Dave, Kushal, 2015, "Apple's App Store review process is hurting users, but we're not allowed to talk about it," *Medium*, April 10, <https://medium.com/@krave/apple-s-app-store-review-process-is-hurting-users-but-we-re-not-allowed-to-talk-about-it-55d791451b>.
- Dessein, Wouter, Luis Garicano, and Robert Gertner, 2010, "Organizing for synergies," *American Economic Journal: Microeconomics* 2(4): 77-114.
- Dinkins, David, 2017, "SegWit2x Could be in Trouble as Bitwala Becomes First Signatory to Renege on New York Agreement," *The Coin Telegraph*, August 24 <https://cointelegraph.com/news/segwit2x-could-be-in-trouble-as-bitwala-becomes-first-signatory-to-renege-on-new-york-agreement>.
- Dutta, Tushar Subhra, 2017 "Now Google Will Punish Poorly-Made Apps In The Play Store," *Tech Viral*, May 22, <https://techviral.net/google-punish-play-store-apps/>.
- Ethereum, 2016, *White Paper: A Next-Generation Smart Contract and Decentralized Application Platform*, GITHUB <https://github.com/ethereum/wiki/wiki/White-Paper>.
- Ethereum, 2017, *Ethereum Homestead Documentation, Release 0.1*, Ethereum, community, March 1, <https://media.readthedocs.org/pdf/ethereum-homestead/latest/ethereum-homestead.pdf>.
- Evans, David S., Andrei Hagiu and Richard Schmalensee, 2006, *Invisible Engines: How Software Platforms Drive Innovation and Transform Industries*, Cambridge, MA: MIT Press.
- Farrell, Joseph, and Garth Saloner, 1985, "Standardization, Compatibility, and Innovation," *RAND Journal of Economics*, 16(1): 70-83.
- Finley, Klint, 2016, "EU Accuses Google of Antitrust Violations Over Android," *Wired*, April 20, <https://www.wired.com/2016/04/eu-accuses-google-antitrust-violations-android/>.
- Gans, Joshua S. , and Hanna Halaburda, 2015, "Some Economics of Private Digital Currency," in Avi Goldfarb, Shane M. Greenstein, and Catherine E. Tucker, eds., *Economic Analysis of the Digital Economy*, pp. 257–76, NBER and University of Chicago Press: Chicago.
- Gartenberg, Chaim, 2017, "Google Will Brand Certified Android Device Boxes With a Google Play Protect Logo," *The Verge*, August 28, <https://www.theverge.com/circuitbreaker/2017/8/28/16215530/google-certified-android-device-play-store-apps-protect-logo-boxes>.
- Grossman, Sanford J., and Oliver Hart, 1986, "The Costs and Benefits of Ownership: A Theory of Lateral and Vertical Integration," *Journal of Political Economy*, 94(4): 691–719.
- Halaburda, Hanna, and Yaron Yehezkel. 2016. "The Role of Coordination Bias in Platform Competition," *Journal of Economics & Management Strategy*, 25: 274–312.
- Hansmann, Henry, and Reinier Kraakman. 2000. "The Essential Role of Organizational Law," *Yale Law Journal*, 110: 387–440.

- Harris, Ron, 2000, *Industrializing English Law: Entrepreneurship and Business Organization, 1720-1844*, Cambridge: Cambridge University Press.
- Hayek, Friedrich A. von, 1982, *Law, Legislation, and Liberty: A New Statement of the Liberal Principles of Justice and Political Economy*, London: Routledge & Kegan Paul (original ed. Chicago: University of Chicago Press, 1973-1979).
- Holmstrom, Bengt, 1999, "The Firm as a Subeconomy," *Journal of Law, Economics, and Organization*, 15(1): 74–102.
- Hu, Henry T.C. , and Bernard Black, 2007, "Hedge Funds, Insiders, and the Decoupling of Economic and Voting Ownership: Empty Voting and Hidden (Morphable) Ownership," *Journal of Corporate Finance* 13(2-3): 343–67.
- Jentzsch, Christoph, 2016, "Control Protocol Parameters with On-Chain Governance," *Ethereum Improvement Proposal* 182, December 6, <https://github.com/ethereum/EIPs/issues/182> .
- Jentzsch, Christoph, 2017, "Ethereum Governance Mechanism: What to Govern?," Address to EDCOM (European Ethereum Development Conference), Paris, February 17-18, [https://edcon.io/ppt/two/Christoph%20Jentzsch%20Ethereum%20Governance%20Mechanism EDCON.pdf](https://edcon.io/ppt/two/Christoph%20Jentzsch%20Ethereum%20Governance%20Mechanism%20EDCON.pdf).
- Kasireddy, Preethi, 2017, "The Synergies Gained from Building on Ethereum's Decentralized App Ecosystem," *Medium*, <https://medium.com/swlh/the-synergies-gained-from-building-on-ethereums-decentralized-app-ecosystem-22a709a675d2>.
- Katz, Michael L., and Carl Shapiro, 1985, "Network Externalities, Competition, and Compatibility," *American Economic Review*, 75(3): 424-40.
- Katz, Michael L., and Carl Shapiro, 1986, "Technology adoption in the presence of network externalities," *Journal of Political Economy*, 94, (4): 822-41.
- Kreps, David M, 1990, "Corporate Culture and Economic Theory," in James E. Alt y Kenneth A. Shepsle, eds., *Perspectives on Positive Political Economy*, pp. 90–143, New York: Cambridge University Press.
- Lessig, Lawrence, 1999, *Code, and Other Laws of Cyberspace*, Basic Books: New York.
- Levin, Jonathan, 2003, "Relational Incentive Contracts," *American Economic Review* 93(3): 835–57.
- Mender, Bedeho, 2016, "Why Bitcoin Apps Are Hard, and Are Ethereum Apps Easier?" *Medium*, September 25, <https://medium.com/@bedeho/why-bitcoin-apps-are-hard-and-are-ethereum-apps-easier-7bbeccab4136>.
- Miller Ira, 2016, "Not Too Late for Humans to Save Ethereum," *Medium*, June 18, <https://medium.com/@guld/not-too-late-for-humans-to-save-ethereum-2f42f5fdb75>.
- Monegro, Joel, 2014, "The Shared Data Layer of The Blockchain Application Stack," Union Square Ventures, December 9, <http://joel.mn/post/104755282493/the-shared-data-layer-of-the-blockchain>.

- Morris, David Z., 2017, "Bitcoin's King Solomon Moment," *Slate*, June 6, http://www.slate.com/articles/technology/future_tense/2017/06/internal_conflict_could_split_bitcoin_in_half.html.
- Mow, Samson, 2017, "The Bitcoin Cash Fork Was a Dangerous Trick," *Fortune*, August 7, <http://fortune.com/2017/08/07/bitcoin-cash-bch-hard-fork-blockchain-usd-coinbase/>.
- Murck, Patrick, 2017, "Who Controls the Blockchain," *Harvard Business Review*, April 19, <https://hbr.org/2017/04/who-controls-the-blockchain>.
- Narayanan, Arvind, Joseph Bonneau, Edward Felten, Andrew Miller, and Steven Goldfeder, 2016, *Bitcoin and Cryptocurrency Technologies: A Comprehensive Introduction*, Princeton University Press: Princeton.
- O'Leary, Rachel Rose, 2018, "Ethereum Users Are Losing Money and Devs Don't Quite Know What to Do," *CoinDesk*, February 7, <https://www.coindesk.com/ethereum-users-losing-money-devs-dont-quite-know/>.
- Roberts, Jeff John, 2016, "Facebook Can Use Controversial Law to Punish Spammy Startup, Court Rules," *Fortune*, July 13, <http://fortune.com/2016/07/12/facebook-cfaa/>.
- Sims, Gary, 2015, "Google Play Store vs the Apple App Store: by the Numbers," April 20, <http://www.androidauthority.com/google-play-store-vs-the-apple-app-store-601836/>.
- Smit, Jorijn, and Jacko Smit, 2016, "Carbonvote: Third-Party Interference or Community Tool?" *Medium*, June 17, 2016, <https://medium.com/@jorijnsmit/carbonvote-third-party-interference-or-community-tool-5476daac031a>.
- Wenger, Albert, 2017a, "Attention on Digital Monopolies," *Blog Post*, July 25, http://continuations.com/post/163405066045/attention-on-digital-monopolies#disqus_thread.
- Wenger, Albert, 2017b, "ICOs and Governance," *Blog Post*, August 15, http://continuations.com/post/164232486650/icos-and-governance#disqus_thread.
- Wiecko, Robert, 2017, "Self-sustainable Decentralized Governance by Blockchain (Budget System)," *DASH*, <https://dashpay.atlassian.net/wiki/spaces/DOC/pages/8585240/Self-sustainable+Decentralized+Governance+by+Blockchain+Budget+System>
- Williams, Dominic, 2017a, "The DFINITY "Blockchain Nervous System," *Medium*, January 4, <https://medium.com/dfinity-network-blog/the-dfinity-blockchain-nervous-system-a5dd1783288e>.
- Williams, Dominic, 2017b, "DFINITY In A Nutshell — A Non Technical Introduction," *Medium*, April 5, <https://medium.com/dfinity-network-blog/dfinity-in-a-nutshell-a-non-technical-introduction-ec45ec5967c1>.
- Williamson, Oliver E, 1975, *Markets and Hierarchies: Analysis and Antitrust Implications*, New York: Free Press.

- Wirdum, Aaron van, 2016, "Who Funds Bitcoin Core Development? How the Industry Supports Bitcoin's 'Reference Client,'" *Bitcoin Magazine*, April 6, <https://bitcoinmagazine.com/articles/who-funds-bitcoin-core-development-how-the-industry-supports-bitcoin-s-reference-client-1459967859/>.
- Wood, Gavin, 2016, "Yellow Paper Committee," April 10, <https://github.com/gavofyork/curly-engine>.
- Xie, Linda, "A Beginner's Guide to Tezos," *Medium*, August 3, <https://medium.com/@linda.xie/a-beginners-guide-to-tezos-c9618240183f>.
- Yang, Yuan, Liu Xinning, and Richard Waters, 2017, "Chinese App Developers Lodge Apple Antitrust Complaint," *Financial Times*, August 9, <https://www.ft.com/content/39efe43e-7cc6-11e7-9108-edda0bcbc928>.
- Yegulalp, Serdar, 2016, "Apple Swift 3: Forward-Looking, but not Backward-Compatible," *Info World*, September 14, <https://www.infoworld.com/article/3120171/open-source-tools/apple-swift-3-forward-looking-but-not-backward-compatible.html>.