6

Property Rights, Knowledge Commons, and Blockchain Governance

*Darcy W. E. Allen, Chris Berg, Sinclair Davidson, and Jason Potts**

Shared knowledge about ownership – that is, knowledge about property rights – facilitates market exchange and economic coordination. Knowledge about property rights is necessary because participants must believe that the seller owns the goods or services being offered for sale in order to complete the contract. Without that belief, they put themselves at risk of fraudulent sales: this is the infamous 'I have a bridge to sell you' problem, named after the con man George C. Parker, who repeatedly sold the Brooklyn Bridge to gullible buyers at the turn of the twentieth century (Cohen 2005). But while this knowledge about property rights has value, and supports markets and economic coordination, it needs to be produced. Today we rely on a complex mix of formal and informal norms, institutions, and practices to ensure knowledge about property rights is both shared and trusted. In this chapter we ask how this shared knowledge resource is governed through shared infrastructures; how those rules shift due to advances in distributed ledgers; and the implications of this for our broader understanding of robust political economy.

The institutions supporting our knowledge of property rights are a core part of the 'scaffolding' that sustains the market economy (see North 2005: xi). This scaffolding is a mix of norms, technologies, and both formal and informal institutions. Hodgson (2015) has identified the elision in the property rights literature about the distinction between 'economic' property rights – those derived from possession and the ability to use or dispose goods one possesses – and 'legal' property rights – the ability to have those rights recognised in law, which allows for property holders to make complex financial exchanges (such as a business loan) using their legal property as collateral (De Soto 2000; Hoffmann 2013). Further, different classes of property are governed

* Darcy W. E. Allen, Chris Berg, Sinclair Davidson, and Jason Potts are with the RMIT Blockchain Innovation Hub. We would particularly like to thank Mark Miller and Bill Tulloh of Agoric for discussions that have led our research in this direction, and Erwin Dekker and Pavel Kuchař for their comments on an earlier version.

https://doi.org/10.1017/9781108692915.008 Published online by Cambridge University Press

through different institutional mechanisms. Berg et al. (2019) argue that small-value property tends to be governed through a complex process of heuristics and norms, where the observation that an individual (or firm) possesses a good is taken as a proxy for legal ownership. Opportunism in this environment is mitigated by market institutional mechanisms such as reputation and discounting. Where the costs of opportunistic behaviour are larger, the task of governing shared facts about property tends to be assumed by hierarchical organisations, particularly the state – consider, for instance, registers of property ownership in land, or intellectual property (such as patents). Even further, the institutional mechanisms by which property rights are governed are themselves the subject of entrepreneurial innovation – they change as new technologies emerge, including institutional technologies (Allen et al. 2020).

Our approach in this chapter is to view knowledge about property rights as a knowledge commons, with various types of infrastructure to govern that commons. We aim to begin to understand this institutional complexity through Elinor Ostrom's Institutional Analysis and Development (IAD) framework (Ostrom 1990), and to place these understandings within a broader political economy context. Since Ostrom's pioneering work on natural resource commons, the IAD framework has been modified and extended in the context of knowledge, information, and cultural commons (see Hess and Ostrom 2005; Ostrom 2005; Madison et al. 2009; Frischmann et al. 2014). Knowledge commons need to be understood differently to physical comments because they "usually must create a governance structure within which participants not only share existing resources but also engage in producing those resources and, indeed, in determining their character" (Frischmann et al. 2014: 16). For property rights, that knowledge commons involves creating ledgers that are updated accurately and enforced – whether by government, hierarchy, norms, or other technologies.

Our aim is not to analyse a particular 'action arena' where participants interact – we leave that for later empirical work – but rather provide theoretical attention to several parts of the IAD framework, including the underlying social dilemma at the heart of property rights knowledge commons, the nature of the shared infrastructure that helps to sustain the commons, and the 'rules in use' that maintain property rights knowledge commons. Indeed, the challenge of property rights knowledge commons is underpinned by a range of social dilemmas. That knowledge not only needs to align with community norms (so that those property rights are enforced), but that knowledge needs to be trusted by market participants to be accurate, and it must do so in the context of potentially hostile actors seeking to undermine property rights (see Potts 2019: ch 9). The research programme that we reveal focuses on the evolution in the 'rules in use' in the property rights knowledge commons, particularly as new technologies such as distributed ledger technology (including blockchain) change how the property rights knowledge commons are created, protected, and enforced.

We also seek to understand these contributions from a broader political economy perspective. To do so we connect two previously unconnected areas: robust political economy and Byzantine consensus. Economists have developed a framework – robust political economy – for thinking about policy design in the context of information and incentive constraints for government planning (Boettke and Leeson 2004; Leeson and Subrick 2006; Pennington 2011). A system is (more) robust if its institutions deal comparatively well with incentive and information problems – that is, by better coordinating distributed information or by better aligning the incentives of people within a complex system. Computer science also provides a new way of looking at the economic problem of coordinating distributed knowledge about property rights. Byzantine political economy takes its name and inspiration from the Byzantine Generals' Problem in computer science (Lamport et al. 1982). Solutions to this problem are said to be Byzantine fault tolerant, describing a class of methods of maintaining consensus over shared facts in the presence of possible miscommunication – in other words, scaffolding for the maintenance and consensus over knowledge either non-hierarchically, or where hierarchy is incomplete. Our understanding of information about property rights as a knowledge commons – and how such commons governance shifts in the context of technological disruption can be understood through this 'Byzantine political economy' connection: the study of the coordination of shared knowledge by distributed consensus mechanisms at the intersection of economics and computer science.

6.1 BLOCKCHAINS, PROPERTY RIGHTS, AND KNOWLEDGE COMMONS

Our understanding of knowledge about property rights is usually analysed in relation to search costs (Stigler 1961), but there is also a cost of identifying who owns what. Does the counterparty who offers a good for sale actually own that good? Market participants require knowledge about property rights – who owns what and who can trade it – to make exchanges and to coordinate. This is in *addition* to other knowledge, such as knowledge about the rules of contract (see Epstein 2009). A transaction consisting of goods that are not lawfully owned is a risky one – opening the risk of later disputes with the rightful owner, or legal liabilities around trafficking stolen goods. In jurisdictions with low-quality legal infrastructure those who possess land, without holding formal title to that land (i.e., a formal record of ownership), are unable to fully exploit the value of their ownership (De Soto 2000). Therefore, we can identify a kind of 'social dilemma' underpinning property rights knowledge commons – that knowledge about property rights might be underproduced or might be of low quality.

Ideally, knowledge about property rights is easily accessible (to facilitate widespread coordination) and also trusted (mapping to the 'general ledger' – a concept that we introduce below). Blockchains are a subcategory of distributed ledgers that

https://doi.org/10.1017/9781108692915.008 Published online by Cambridge University Press

were first developed by Nakamoto (2008a) for the cryptocurrency Bitcoin. While other distributed ledger technologies exist (the word 'blockchain' describes a data structure rather than a consensus mechanism), blockchains are currently the dominant technology, and here we use the word blockchain as a stand-in for all distributed ledger technologies (for a more general survey see Rauchs et al. 2018). As Berg et al. (2019) argue, blockchains are an institutional technology that allows groups to record, and come to consensus over, property rights (see also Davidson et al. 2018) – acting as additional 'rules in use' for the property rights knowledge commons that will shift commons governance at the margin. Importantly, we do not suggest that blockchains will act as a complete substitute to existing infrastructure to maintain the commons, but rather that blockchain represents a technological disruption of those rules at some margins. Precisely on what margins blockchains will improve the governance of the knowledge commons is an entrepreneurial question that will be revealed through time, and one that will manifest differently across jurisdictions and in relation to the quality of existing commons infrastructure. In the same way that there is no regulatory or institutional panacea for other commons, the property rights knowledge commons are maintained through a complex range of nested institutional 'rules in use'. How those commons are maintained and updated – such as through the transfer of property rights – will depend on the particular 'action arena' and the participants in it.

What are the underlying 'resource characteristics' of knowledge about property rights? Understanding knowledge as a commons sits alongside a range of other 'new commons' that examine underlying resources such as culture, science, or intellectual commons (e.g., see Boyle 2003; Hess and Ostrom 2005; Bollier and Helfrich 2014). These 'new commons' differ from early natural resource commons because the resource at hand does not have distinct biophysical characteristics. Rather than natural resource commons, the resource in the property rights knowledge commons has different characteristics. The problem is less that the underlying resource will be depleted through overuse, such as in a fishery. But rather that individuals will underprovide the resource – that is, by free riding – or that there will be a lack of coordination, because the effective production of the knowledge resource requires joint participation by many parties. The knowledge about property rights necessary to exchange and coordinate is similarly a commons where having widespread, trusted and enforced knowledge is a critical input. Our focus below is on the institutional governance arrangements that help to maintain the knowledge about property rights – that is, the infrastructure that is a "form of community management or governance of a shared resource" (Madison et al. 2018: 2). These knowledge commons are governed through a complex set of informal norms and more formalised hierarchical governance structures (e.g., legal enforcement, public registries).

What are the 'rules in use' by which we manage the property rights knowledge commons? One common way is to use possession as a heuristic for ownership (see Berg et al. 2018, 2019). That is, in most exchanges, the fact that a counterparty holds

a good (or money) works well enough as an indication that they are the rightful owner of the item. Another common mechanism is through legal documentation that identifies a person or entity as the owner (see De Soto 2000; Hodgson 2015). This approach corresponds to a ledger, whether real or imaginary, that maps relationships between property and individuals (or any other property-owning entity). From an Ostromian perspective – utilising the IAD framework – we can understand these institutional mechanisms as 'rules in use' that shape various action situations of economic, social, and political exchange and coordination.

We can understand each of these governance mechanisms as attempts to govern the underlying property rights knowledge commons, and we can understand the effectiveness of those mechanisms in relation to what Berg et al. (2018, 2019) call a *general ledger*: a hypothetical map of all owned items and property relationships. Possession is a heuristic for the navigation of this general ledger, but possession is also coexistent with formal property ledgers – paper titles to (for example) vehicles that are mapped in centralised (usually government) vehicle registrations databases, and centralised ledgers of titles to land. To understand potential discrepancies between different infrastructures for managing property rights knowledge commons we can think of a *perfect ledger*. From this perspective we can see that better 'rules in use' – or institutional infrastructure – for managing property rights knowledge commons will tend us towards more widespread and more trusted information, facilitating market exchange and making the knowledge commons more robust.

Distributed ledgers, including blockchains, have the potential to more closely align the property rights knowledge commons to the perfect ledger. Blockchains are a technology for the distribution, maintenance, and verification of social facts. This is the basic feature of their design. At their most essential they are databases (ledgers) to which participants propose modifications (make/announce transactions) that are accepted, or not accepted, by other participants (the transaction is included in a new block adopted by the majority of the network). Once those modifications have been accepted by all participants, they are *in effect* rendered immutable – a permanent addition to the shared ledger. What makes (some) blockchains unique is that they achieve this consensus without a single, central authority to manage that process. Acceptance and synchronisation of new transactions occur without the use of a trusted third party to validate those transactions. For a full public blockchain (such as Bitcoin or Ethereum) no one needs permission from a given authority to access and make transactions on the chain. In their ideal or maximalist type, blockchains are thus open-access systems where openness is written into the protocol, rather than norms or laws.

The social facts that blockchains record are property rights (Berg et al. 2018). More specifically, they record the outcome of contractual agreements for the creation and transfer of property rights. In this way blockchains can be understood as a new infrastructure for maintaining the knowledge commons of information about property rights. Blockchain tokens can be inherently (i.e., intersubjectively)

valuable on their own right, such as Bitcoin or other cryptocurrencies designed to be money substitutes. Alternatively, tokens can be markers pointing to other digital or material property – such as intellectual property in the case of blockchain applications that manage and track information about cultural or patentable goods, or real-world items like goods travelling across a supply chain. The use, disposal, and exchange of property rights is managed through a public key/private key infrastructure that requires publicly visible transactions on a public key to be signed by a corresponding, secret, private key (Rauchs et al. 2018). The network treats the use of the private key as the endorsement that the owner of the relevant tokens has initiated the transaction. This places a lot of reliance on individual token holders protecting their private keys from theft or loss, particularly given the fact that transactions are irreversible.

Blockchains offer an infrastructure for the social verification of information and access to knowledge about property rights. Blockchains are not the only technology required for secure exchanges in all circumstances – for example, distributed-ledger-managed supply chains rely on significant Internet of Things (IoT) infrastructure and alternative technologies of trust (such as government enforcement of rights). Nonetheless, they offer a publicly accessible record of property ownership with a built-in payment (i.e., transaction) system for exchange. Alongside other centralised or informal ledgers and enforcement mechanisms, the immutability and accessibility of blockchain infrastructure might act as additional 'rules in use' for the underlying knowledge commons.

There are several ways that blockchain infrastructure might augment existing infrastructure. By being freely accessible, distributed ledgers reduce the cost of consulting the ledger representation of the knowledge commons and reduce the risk that information of the ledger has been tampered with by the authority that manages it. Blockchains also offer the potential for the 'tokenisation' of more real-world assets and the association of those assets with ledger entries, allowing for significantly lower transaction costs in the exchange of those assets. Blockchains not only contribute to the scaffolding or infrastructure of the property rights knowledge commons, but also to the contracts underpinning the transfer of those resources. Even further, a decentralised sharing economy might organise more than just houses and drivers, but also ledger entries of lower value assets like hardware and moveable furniture (Catalini and Gans 2016; Munger 2018). Such knowledge about property rights facilitates exchanges on/with those goods in the same way that De Soto (2000) identified the opportunities for capital markets on top of more secure property titling in land.

There are limitations of blockchain-based property rights. This reality raises important questions about their interaction with other mechanisms – whether informal or centralised mechanisms – that also facilitate the knowledge commons. Blockchains sit alongside a range of other institutional mechanisms to maintain knowledge commons. Blockchains add to the complex ecosystem of different ways

we manage the property rights knowledge commons, and should not be thought of as a complete substitute for existing institutions of property rights. Rather, different 'action arenas' wherein market participants interact will variously rely on different sources of underlying knowledge, and rights will be exchanged according to different rules. For instance, the potential for blockchain to augment property rights knowledge commons will differ depending on the nature of those rights. For some property rights they are native digital items. Digital tokens (such as Bitcoin) are enforced endogenously by the blockchain itself (Ishmaev 2017) and are less reliant on other institutions to maintain the commons. On the other hand, tokens that refer to property outside the blockchain – physical property whose ownership is governed by a distributed ledger – has no such endogenous enforcement mechanism. Ensuring that a digital exchange is reflected by a change in ownership in the real world faces the same challenges as any other property transfer; that is, it relies on the institutional mechanisms (such as legal rules, policing, and norms) that govern other property rights systems (Djankov et al. 2003). The extent to which blockchains are entangled and reliant on existing structures is a function of several factors, including whether those rights come into disputes and how those disputes are resolved (see Allen et al. 2020). The norms of disputes around blockchain-created and blockchain-enforced property rights are still developing, including where, for instance, there have been malicious hacks of blockchain tokens, and the community governing the blockchain has reversed those transactions.

'Onboarding' existing property rights frameworks onto a distributed ledger is a non-trivial problem. The real-world property rights systems studied by De Soto feature contested rights in environments where owners do not trust a central authority to recognise those rights. However, to realise the benefits of a blockchain-based property system, those contested rights need to be adjudicated by an authority – so that the definitive ownership can be recorded on the ledger. This need for a trusted authority to facilitate the move to a 'trustless' ledger is a common theme of bootstrapping decentralised systems (Berg et al. 2019; Saadatmand et al. 2019). Further, customary and informal property rights systems are not always amendable to the sort of formalisation that a digital ledger requires – a challenge faced by all property formalisation movements from the English enclosure movement (McCloskey 1972) to contemporary Papua New Guinea (Stead 2017). In this sense, blockchain offers a complement, not a replacement, for the existing institutional framework around property rights (Allen et al. 2019). Identifying where blockchain is most valuable at the margin for the governance of shared knowledge is the entrepreneurial task that has been ongoing since Nakamoto (2008a). Blockchains are a complement to existing governance structures of the property rights knowledge commons; they are governance frameworks that can coexist – and interoperate – with other knowledge commons about property rights. Blockchain infrastructure also serves as an input to processes and exchanges that might otherwise not occur. It expands the feasible number of trades that can or could occur

within an economy. The open nature of blockchain as a distributed ledger ensures that they constitute shared goods or information that facilitates trade.

Blockchain innovation and diffusion is driven by complex dynamics that are not readily subject to the usual replicator dynamics and adoption-diffusion models that we are used to seeing with production technologies (Allen et al. 2020). Blockchains are an institutional technology – a competing technology in the mainline institutional schema of markets (Smith 1976), firms (Coase 1937), governments (Buchanan and Tullock 1962), clubs (Buchanan 1965), and the commons (Ostrom 1990) – which is subject to the evolutionary pace of production and communication technologies. They are network technologies, subject to network effects and adoption curves, but in which not only the object of exchange (the technology) is rapidly changing but also the institutional framework in which the exchanges occur.

Our aim here is not to describe the complexity of blockchain governance – itself a rapidly evolving field (Allen and Berg 2020). Indeed, the IAD framework has already been applied specifically to distributed ledgers to understand blockchains themselves as common pool resources (see Howell and Potgieter 2019). Blockchains deal with challenges of consensus over property rights in different ways – trading off various features such as privacy and permission in recording and enforcing property rights. The way that any given blockchain governs the knowledge shared within it can also be understood as evolutions in its 'rules in use'. An open permissionless public blockchain, for instance, has murkier resource boundaries than a permissioned one, where read and write access to the distributed ledger is restricted. The 'rules in use' differ across blockchains – depending on the code and the relationship between exogenous and endogenous governance. The question of whether blockchains are private, commons, or club-like depends on how that protocol is designed and the nature of the applications built on top of the blockchain infrastructure. Decentralisation and distribution occur on many margins in blockchain infrastructure. The validation of the network might be decentralised (i.e., there are many miners in Bitcoin) yet the governance of the networks (decisions about how the underlying protocol is modified) might be highly centralised. Private or permissioned blockchains such as IBM's Hyperledger Fabric trade-off the open and public features of decentralised systems for speed or privacy. It is for this reason that we have elsewhere described blockchains as a 'universal Turing institution' – a meta-institution that can be structured to simulate any other institution, just as a universal Turing machine can simulate any other Turing machine (Berg et al. 2019). Nevertheless, on the margin blockchains augment the broader understanding of property rights as knowledge commons.

Why did blockchain-based property rights emerge? Blockchain innovation can be understood as a form of self-governance (Leeson 2009; Skarbek 2011; Stringham 2014) – where they are forms of governance that are developed from the bottom-up with the intention of enabling better coordination. We can therefore understand the

process of formalising and developing mechanisms of property rights in a similar way to Demsetz (1967), who argued that the emergence of property rights occurred where the marginal benefits exceeded the marginal costs – that is, where the costs of developing the institution were lower than the potential benefits from trade. The parallel here for knowledge about property rights is that this is a process of discovery over the most effective way to maintain that knowledge commons – whether through government registries, informal norms, private protection, or 'ledgerisation' using blockchains. New mechanisms of governance for the knowledge commons will emerge where the benefits of those new rules in use exceed the costs to maintain them – or, put another way, where demand for knowledge about property rights increases, whether in market or other forms of non-market coordination.

The extent to which blockchain governance is effective is a comparative question, relating to how protocol design solves coordination problems – in this case the recording and verifying of knowledge about property rights – compared to other governance structures. In a public permissionless blockchain those rules are social and communal – just as individual transactions are subject to consensus, so too is the protocol. This goes beyond the voluntary choice to participate or withdraw from participation on the network. Users (and developers) can at any time choose to 'fork' the code or even the history of the chain itself in order to vary its rules. These are complex events within 'action arenas' that are deserving of empirical studies. Proposals for upgrades to the protocol (to deal with, for instance, security problems, or to increase transaction throughput) require the ascension of a (usually) super-majority of users – in the Ostrom schema these are more "collective choice" rules. Bargaining power around these upgrades is unevenly distributed between classes of users (such as 'core' developers, miners, application developers and token holders – see Allen and Berg (2020)), each of whom claim a degree of legitimacy and influence over the choices to upgrade. This dynamic has brought about the crisis in blockchain governance, where these self-organised communities have disputed how decisions about protocol changes (i.e., constitutional level changes) should be made (De Filippi and Loveluck 2016).

What we have seen in this section is both that knowledge about property rights can usefully be thought of as a commons, and that there is a complex range of infrastructural rules that seek to maintain that commons. We have also seen that advances in blockchain technology – as a technology for recording and transferring property rights – are new 'rules in use' for those commons. Now we turn to the implications of this for our broader political economy understanding of knowledge commons and market infrastructure, combining the dual lenses of 'robust political economy' from economics and the 'Byzantine Generals' Problem' in computer science into what we call Byzantine political economy. The upshot from our analysis is that to the extent blockchains augment and complement existing hierarchical and norms-based governance of property rights, they will contribute a more robust political economy.

6.2 TOWARDS MORE ROBUST KNOWLEDGE COMMONS

Thomas Hobbes described governments as *automata* – machines that move itself, like a watch – and the metaphor of government as machine has been a deep thread throughout political thought (Agar 2003). When Walter Bagehot offered his vision of government as divided into the efficient (bureaucracy and administration) and the dignified (the monarch), he did so through an elaborate machine metaphor: the dignified part provided the power, regulator, and safety valves (as required by a steam engine), for the efficient part which 'works and rules' (Bagehot 1872: 4). As the potential and power of computing became evident, Oskar Lange (1967) declared that the socialist calculation debate was now an anachronism. The state's management of the socialist economy could be now delegated to the electronic computer, which would perform better than the market had done. Planning was a matter of extremely complex (although still linear) calculations. In the industrial age, 'machinery' was deployed as a metaphor; in the digital age it could now be a practicality.

What connects these visions together is government as a single apparatus that chooses and acts. The government is seen as *a* single machine – it is a single hierarchical authoritative mind that produces, allocates, and moves resources. To the extent that it has diverse centres of implementation – multiple machines, ministries, or domains of action – they are organised hierarchically (i.e., top-down, according to a command and control system). Orders are passed through the system and implemented by fiat. This approach to government is subject to the Hayekian (Hayek 1945) and public choice critiques (Buchanan and Tullock 1962) – governments do not work like that; in fact, they may be better conceived as orders (see Eusepi and Wagner 2011). Decision makers and bureaucrats do not smoothly, efficiently, and selflessly direct and implement policies in full knowledge of relevant trade-offs. But it is also the case that *machines* do not work like that either. Machines, especially those that interface with hardware might break, erode, deteriorate, fail to communicate, experience noisy communication channels, or even be hacked for malicious or malign purpose.

The Byzantine Generals' Problem is just one of a group of problems around the reliability of communications in decentralised systems. The question facing engineers is how to build reliable computation systems out of components that have non-zero failure rates. Improving the reliability of a component was subject to diminishing returns: eventually it was uneconomical or just impossible given technological limitations to improve reliability further. Thus, engineers accommodated the non-zero failures rates of components by ensuring a degree of redundancy in systems. If one part failed, another could take its place (Pierce 1964; Short 1968).

Such fault tolerance was particularly sensitive in the fields of aviation and space exploration, where it was proposed in the post-war years that an increasing amount of the control of avionics be handed off to computers and where the consequences

of failure were, and remain, severe. Two separate research projects into fault-tolerant systems was supported by NASA, Software Implemented Fault Tolerance (SIFT) at SRI International and Fault Tolerant Multiprocessor (FTMP) at C. S. Draper Laboratory (Lee and Anderson 2012). Each approach relied on providing at least three times redundancy in components. In SIFT, for example, processes were run in parallel on three separate components and an executive function reconciled the outputs based on a majority vote (Wensley 1972; Wensley et al. 1978). Establishing how such 'voting' schemes would function and ascertaining the minimum required level of redundancy led to the conceptualisation and construction of the Byzantine Generals' Problem (Lamport 1978; Pease et al. 1980). This is a computer science equivalent to the knowledge problem that we have alluded to before: How does a buyer know that the seller actually owns the good being offered for sale?

The description of the Byzantine Generals' Problem in Lamport et al. (1982) is as follows. The Byzantine army surrounds an enemy castle. The army is divided into divisions, each of which is led by a general. The generals need to come to a consensus over a plan to attack the castle. But not all generals are loyal – a certain number of the generals are traitors, who want either to prevent consensus or for the generals to come to consensus around a bad plan. The messengers that pass information between the generals could be killed – and information therefore does not get through – or they could be late – which we describe as asynchronous messaging. In the simplest set-up of the problem, consider an army with three generals and a binary decision about whether to attack or retreat. Each general sends a message to the other two about their proposed action, and listens for messages from the other. If the traitorous general shares contradictory messages (i.e., recommending to one that they attack, and the other that they retreat) then coordination will fail, risking defeat on the battleground. To deal with the fact that a 'bad plan' is hard to formalise, they designated one general as a commanding general who formulates the plan and distributes orders to the others. Any general, commander, or lieutenant can be a traitor. In business, any seller could be selling unowned goods.

The Byzantine Generals' Problem is thus that: (1) all loyal generals need to obey the same order and (2) if the commanding general is loyal, then all loyal generals obey the order of the commanding general. Metaphorically, traitorous generals are unreliable components that report inconsistently to other distributed components, impeding consensus. Lamport et al. (1982) wanted to know what proportion of loyal/traitorous generals a distributed system could handle. Their finding was that, contrary to prior belief, under the simplest set-up three generals were not tolerant to one of their members being traitorous. This finding has generalised to the statistical expectation that decentralised Byzantine fault-tolerant systems cannot sustain more than one-third of nodes failing or acting maliciously.

The Byzantine Generals' Problem, therefore, describes a class of engineering failures where a system needs to come to consensus but fails to do so because of inconsistencies in communication (a Byzantine fault). The original problem

statement gives an anthropomorphic implication to the problem – where generals betray their comrades – but the motive for misreporting is not the key. Whether the fault is the result of faulty reporting, faulty communications links, or capture by hostile actors, Byzantine faults are endemic to distributed systems (Driscoll et al. 2003, 2004).

Some parallels between consensus requirements in the presence of Byzantine faults and the economic coordination problem should be already apparent. Hayek (1945) describes the economic problem as the coordination of activity in an environment where information is "dispersed . . . incomplete and frequently contradictory" (519). Williamson (1985) argues that economic institutions are structured around the contractual problems of opportunism ('self-seeking with guile') and bounded rationality (leading to risks around investment in specific assets). Boettke and Leeson (2004), Leeson and Subrick (2006), and Pennington (2011) all describe the economic problem as being a dual challenge of structuring social institutions to handle information problems, and motive problems in public and private choices, leading to the agenda of robust political economy. These approaches, which can be loosely grouped around the transaction cost and Austrian schools, coalesce around the same problem of social organisation in the presence of disagreement and inconsistent, incomplete, and unreliable communication. In that sense, Byzantine fault tolerance is a parallel conceptualisation of this same class of problem. Nakamoto (2008a) designed the blockchain to operate under these very conditions – creating incentive structures to promote cooperative behaviour in the presence of unreliable trading partners and incomplete information.

The relevance of Byzantine consensus is more apparent when we consider this engineering problem in the context of blockchains – and particularly in relation to the governance of property rights. As we have seen, blockchains offer a shared database that can be updated – recording individual transactions and account states and allowing for complex economic organisation. Blockchains achieve Byzantine fault tolerance through their consensus mechanisms. Gramoli (2017) frames the Byzantine consensus problem in blockchains as "(i) agreement: no two correct processes decided different blocks; (ii) validity: the decided block is a block that was proposed by one process; (iii) termination: all correct processes eventually decide". Nakamoto's (2008a) innovation was to use a group of pre-existing technologies – peer-to-peer networking, asymmetric cryptography, and proof-of-work hashing algorithms – to come to distributed agreement over the state of a shared ledger (Narayanan and Clark 2017). The possible Byzantine faults range from the dramatic (hostile users of the ledger who seek to disrupt agreement and thereby allow for cryptocurrency tokens to be 'double spent') to the seemingly mundane (devices disconnecting from the network, disagreement about the order in which transactions are logged on the ledger due to geographic distances and a lack of synchronised clocks). Unlike in the example of the Byzantine Generals' Problem, there is no

'commander'. Each participant in the network that wants to keep a copy of the ledger and register new transactions must both listen to orders and announce orders.

While the Bitcoin white paper (Nakamoto 2008a) did not make explicit reference to Byzantine consensus, Satoshi Nakamoto's later forum posts argue that the consensus mechanism, part of which involves 'miners' solving a computationally difficult proof-of-work puzzle, gives Bitcoin characteristics which solve the Byzantine Generals' Problem (Nakamoto 2008b). Blockchains order their transactions through a majority agreement mechanism. Each 'miner' (in the case of the classic proof-of-work system in Bitcoin) wants to ensure that their mined block is accepted by the rest of the network, so they only extend the longest chain of blocks that they can observe on the network. Proof-of-work blockchains are vulnerable to '51 per cent' attacks, which allows a single attacker to interrupt honest behaviour if that attacker gains control of more than 50 per cent of the 'hash power' (computational resources directed at the proof-of-work puzzle) of the network (Budish 2018).

Where the anthropomorphic structure of the original Byzantine consensus set-up was somewhat misleading – components do not choose to be loyal or disloyal – Nakamoto consensus is distinct by adding incentives to the consensus mechanism. In Lamport et al. (1982) and the subsequent literature on consensus between unreliable components, reliability is exogenous. Whether a general is honest or dishonest is not affected by the operation of the consensus mechanism. In a real-world system of components this abstraction cannot be entirely true – the probability of component fault increases with the frequency of use due to aging. Nonetheless, generals can choose to betray their comrades, but components do not 'choose' to be faulty. Under Nakamoto consensus by contrast, each general faces distinctly economic incentives as to whether to be honest. Bitcoin is structured to reward miners for acting in the interests of the network (i.e., updating the ledger according to predetermined rules) and to punish other miners forwarding invalid transactions by orphaning chains that do not meet majority agreement. To our knowledge, Narayanan and Clark (2017) were the first to point out that this incentive mechanism made Bitcoin distinct from prior approaches to Byzantine consensus. It is this distinction that brings the study of Byzantine faults squarely into the realm of social science.

6.3 CONCLUSION

In this chapter we have made three interrelated points. First that a knowledge of property rights is a shared resource. Importantly, however, that this is a shared resource that must be produced in the economy. Second that 'rules in use' are an important component of those shared resources and those rules and the governance of those rules are subject to technological disruptions. Finally, we point to distributed ledger technology – blockchain – as an innovation that can better

provide a robust political economy solution to existing knowledge commons governance problems.

To the extent that blockchains are adopted as an alternative social organisation for knowledge, how those blockchains are governed and how they interact with the existing set of institutions will become increasingly important. Their creation by Nakamoto (2008a) has opened up a large field of investigation, underlining the role that centralised ledgers play in the economy, and the long-run possibilities brought about by more and more decentralised ledgers. While we have suggested a research programme that examines blockchains as an innovation in how we govern the property rights knowledge commons – using the IAD framework for instance – these studies should also be laid within a broader conception of Byzantine political economy. That is, how can new systems of decentralised governance of knowledge about property rights – and their decentralised enforcement – be understood in the context of the robustness of the knowledge commons that fundamentally underpin our market system?

In this chapter we have laid out the principles at stake – what we have not provided is any empirical validation of these principles. This is largely due to the empirical applications currently being the subject of entrepreneurial discovery. Unlike previous innovative disruptions to the economy, we are seeing the disruption occurring in real time and not over the course of decades. Governance structures and institutions that have evolved over decades, if not centuries, to inform rules in use and resolve commons problems could be disrupted in a matter of years.

REFERENCES

Agar, Jon. 2003. *The Government Machine: A Revolutionary History of the Computer*. Cambridge, MA: MIT Press.

Allen, Darcy W. E., and Chris Berg. 2020. "Blockchain Governance: What We Can Learn from the Economics of Corporate Governance". *Journal of the British Blockchain Association* 3 (1): 1–10.

Allen, Darcy W. E., Chris Berg, Brendan Markey-Towler, Mikayla Novak, and Jason Potts. 2020. "Blockchain and the Evolution of Institutional Technologies: Implications for Innovation Policy". *Research Policy* 49 (1): 1038–1065.

Allen, Darcy W. E., Aaron M. Lane, and Marta Poblet. 2019. "The Governance of Blockchain Dispute Resolution". *Harvard Negotiation Law Review* 25: 75–101.

Bagehot, Walter. 1872. *The English Constitution*. London: Henry S. King & Company.

Berg, Chris, Sinclair Davidson, and Jason Potts. 2018. "Ledgers". *SSRN*. https://ssrn.com/abstract=3157421

2019. *Understanding the Blockchain Economy: An Introduction to Institutional Cryptoeconomics*. Cheltenham: Edward Elgar.

Boettke, Peter J., and Peter T. Leeson. 2004. "Liberalism, Socialism, and Robust Political Economy". *Journal of Markets & Morality* 7: 99–111.

Bollier, David, and Silke Helfrich, Eds. 2014. *The Wealth of the Commons: A World beyond Market and State*. Amherst, MA: Levellers Press.

Boyle, James. 2003. "The Second Enclosure Movement and the Construction of the Public Domain". *Law and Contemporary Problems* 66(1–2): 33–74.
Buchanan, James M. 1965. "An Economic Theory of Clubs". *Economica* 32: 1–14.
Buchanan, James M., and Gordon Tullock. 1962. *The Calculus of Consent, Logical Foundations of Constitutional Democracy*. Ann Arbor: University of Michigan Press.
Budish, Eric. 2018. "The Economic Limits of Bitcoin and the Blockchain". Working Paper 24717. Working Paper Series. National Bureau of Economic Research. https://doi.org/10.3386/w24717.
Catalini, Christian, and Joshua S. Gans. 2016. "Some Simple Economics of the Blockchain". Working Paper 22952. Working Paper Series. National Bureau of Economic Research. https://doi.org/10.3386/w22952.
Coase, Ronald H. 1937. "The Nature of the Firm". *Economica* 4: 386–405.
Cohen, Gabriel. 2005. "For You, Half Price". *The New York Times*, 27 November.
Davidson, Sinclair, Primavera De Filippi, and Jason Potts. 2018. "Blockchains and the Economic Institutions of Capitalism". *Journal of Institutional Economics* 14 (4): 639–658.
De Filippi, Primavera, and Benjamin Loveluck. 2016. "The Invisible Politics of Bitcoin: Governance Crisis of a Decentralized Infrastructure". *Internet Policy Review* 5 (4): 1–28.
De Soto, Hernando. 2000. *The Mystery of Capital: Why Capitalism Triumphs in the West and Fails Everywhere Else*. New York: Basic Books.
Demsetz, Harold. 1967. "Toward a Theory of Property Rights". *The American Economic Review* 57 (2): 347–359.
Djankov, Simeon, Edward Glaeser, Rafael La Porta, Florencio Lopez-de-Silanes, and Andrei Shleifer. 2003. "The New Comparative Economics". *Journal of Comparative Economics* 31: 595–619.
Driscoll, Kevin, Brendan Hall, Håkan Sivencrona, and Phil Zumsteg. 2003. "Byzantine Fault Ttolerance, from Theory to Reality". In International Conference on Computer Safety, Reliability, and Security, 235–248. Springer.
Driscoll, Kevin, Brendan Hall, Michael Paulitsch, Phil Zumsteg, and Håkan Sivencrona. 2004. "The Real Byzantine Generals". In The 23rd Digital Avionics Systems Conference. IEEE Cat. No. 04CH37576, 6. D. 4-61. IEEE.
Epstein, Richard A. 2009. *Simple Rules for a Complex World*. Cambridge, MA: Harvard University Press.
Eusepi, Giuseppe, and Richard E. Wagner. 2011 "States as Ecologies of Political Enterprises". *Review of Political Economy* 23 (4): 573–585.
Frischmann, Brett M., Michael J. Madison, and Katherine J. Strandburg, Eds. 2014. *Governing Knowledge Commons*. Oxford: Oxford University Press.
Gramoli, Vincent. 2020. "From Blockchain Consensus Back to Byzantine Consensus". *Future Generation Computer Systems* 107 (June): 760–769. https://doi.org/10.1016/j.future.2017.09.023.
Hayek, Friedrich A. 1945. "The Use of Knowledge in Society". *The American Economic Review* 35 (4): 519–530.
Hess, Charlotte, and Elinor Ostrom. 2005. "A Framework for Analyzing the Knowledge Commons". In *Understanding Knowledge as a Commons: From Theory to Practice*, edited by Charlotte Hess and Elinor Ostrom, 41–82 (Chapter 3). Cambridge, MA: MIT Press.
Hodgson, Geoffrey M. 2015. "Much of the 'Economics of Property Rights' Devalues Property and Legal Rights". *Journal of Institutional Economics* 11: 683–709.
Hoffmann, Sabine. 2013. "Property, Possession and Natural Resource Management: Towards a Conceptual Clarification". *Journal of Institutional Economics* 9: 39–60.

Howell, Bronwyn E., and Potgieter, Petrus H. 2019. "Governance of Blockchain and Distributed Ledger Technology Projects: A Common-Pool Resource View". Workshop on the Ostrom Workshop (WOW6) Conference, Indiana University Bloomington, June 19–21, 2019.

Ishmaev, Georgy. 2017. "Blockchain Technology as an Institution of Property". *Metaphilosophy* 48: 666–686.

Lamport, Leslie. 1978. "The Implementation of Reliable Distributed Multiprocess Systems". *Computer Networks* 2: 95–114.

Lamport, Leslie, Robert Shostak, and Marshall Pease. 1982. "The Byzantine Generals Problem". *ACM Transactions on Programming Languages and Systems* 4: 382–401.

Lange, Oskar R. 1967. "The Computer and the Market". In *Capitalism, Socialism and Economic Growth: Essays Presented to Maurice*, edited by C. Feinstein, 158–161. Cambridge: Cambridge University Press.

Lee, Peter A., and Thomas Anderson. 2012. *Fault Tolerance: Principles and Practice*. New York: Springer.

Leeson, Peter T. 2009. *The Invisible Hook: The Hidden Economics of Pirates*. Princeton, NJ: Princeton University Press.

Leeson, Peter T., and J. Robert Subrick. 2006. "Robust Political Economy". *The Review of Austrian Economics* 19: 107–111.

Madison, Michael J., Brett M. Frischmann, and Katherine J. Strandburg. 2009. "Constructing Commons in the Cultural Environment". *Cornell Law Review* 95: 657–709.

Madison, Michael J., Katherine J. Strandburg, and Brett M. Frischmann. 2018. "Knowledge Commons". In *Routledge Handbook of the Study of the Commons*, edited by Blake Hudson, Jonathan Rosenbloom, and Dan Cole. Abingdon: Routledge.

McCloskey, Donald N. 1972. "The Enclosure of Open Fields: Preface to a Study of Its Impact on the Efficiency of English Agriculture in the Eighteenth Century". *The Journal of Economic History* 32: 15–35.

Munger, Michael C. 2018. *Tomorrow 3.0*. Cambridge: Cambridge University Press.

Nakamoto, Satoshi. 2008a. "Bitcoin: A Peer-to-Peer Electronic Cash System". www.bitcoin.org.

2008b. "Re: Bitcoin P2P e-cash Paper". In *The Book of Satoshi: The Collected Writings of Bitcoin Creator Satoshi Nakamoto*, edited by Phil Campagne. E53 Publishing LLC.

Narayanan, Arvind, and Jeremy Clark. 2017. "Bitcoin's Academic Pedigree". *Communications of the ACM* 60: 36–45.

North, D. C. 2005. *Understanding the Process of Economic Change*. Princeton, NJ: Princeton University Press.

Ostrom, Elinor. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press.

2005. *Understanding Institutional Diversity*. Princeton, NJ: Princeton Uuniversity Press.

Pease, Marshall, Robert Shostak, and Leslie Lamport. 1980. "Reaching Agreement in the Presence of Faults". *Journal of the ACM* 27: 228–234.

Pennington, Mark. 2011. *Robust Political Economy: Classical Liberalism and the Future of Public Policy*. Northampton, MA: Edward Elgar.

Pierce, William H. 1964. "Redundancy in Computers". *Scientific American* 210: 103–113.

Potts, Jason. 2019. *Innovation Commons: The Origin of Economic Growth*. Oxford: Oxford University Press.

Rauchs, Michel, Andrew Glidden, Brian Gordon et al. 2018. "Distributed Ledger Technology Systems: A Conceptual Framework". *SSRN*. https://ssrn.com/abstract=3230013.

Saadatmand, Fatemeh, Rikard Lindgren, and Ulrike Schultze. 2019. "Configurations of Platform Organizations: Implications for Complementor Engagement". *Research Policy* 48: 1037–1070.

Short, Robert A. 1968. "The Attainment of Reliable Digital Systems through the Use of Redundancy – A Survey". *IEEE Computer Group News* 2: 2–17.

Skarbek, David. 2011. "Governance and Prison Gangs". *American Political Science Review* 105: 702–716.

Smith, Adam. 1976. *An Inquiry into the Nature and Causes of the Wealth of Nations*. 2 Vols., edited by R. H. Campbell, A. S. Skinner, and W. B. Todd. Oxford: Oxford University Press.

Stead, Victoria. 2017. "Landownership as Exclusion". In *Kastom, Property and Ideology: Land Transformations in Melanesia*, edited by Siobhan McDonnell, Matthew G. Allen, and Colin Filer. Canberra: ANU Press.

Stigler, George J. 1961. "The Economics of Information". *Journal of Political Economy* 69: 213–225.

Stringham, Edward. 2014. *Private Governance: Creating Order in Economic and Social Life*. Oxford: Oxford University Press.

Wensley, John H. 1972. "SIFT: Software Implemented Fault Tolerance". In *Proceedings of the December 5–7, 1972, Fall Joint Computer Conference*. Part I, 243–253. New York: ACM.

Wensley, John H., Leslie Lamport, Jack Goldberg et al. 1978. "SIFT: Design and Analysis of a Fault-Tolerant Computer for Aircraft Control". *Proceedings of the IEEE* 66: 1240–1255.

Williamson, Oliver E. 1985. *The Economic Institutions of Capitalism*. New York: Free Press.