Governance of Blockchain and Distributed Ledger Technology Projects: a Common-Pool Resource View

Bronwyn E. Howell, School of Management, Victoria University of Wellington, bronwyn.howell@vuw.ac.nz

Petrus H. Potgieter, Dept of Decision Sciences, University of South Africa, potgiph@unisa.ac.za

11th June 2019

Prepared for delivery at the Workshop on the Ostrom Workshop (WOW6) conference, Indiana University Bloomington, June 19–21, 2019 in the themed panel *Blockchain: Innovation in Distributed Governance in the Cryptosphere*. Copyright 2019 by author(s).

Abstract

In this paper, we utilise Elinor Ostrom's Institutional Analysis for Development framework (E. Ostrom 2010) to explore the economic, technological, political, social and psychological contexts in which distributed ledger systems operate in order to understand their broad governance arrangements as polycentric systems operating simultaneously at many different levels of interaction. Somewhat ironically given the purported motivations of decentralisation and revolutionary changes purported for the blockchain technology compared to established firms and governments, it appears that most DLS applications substitute centralised control by one set of stakeholders with centralised control by a different set of stakeholders. We describe the features of DLs that render them candidates as common-pool resources and use the IAD framework to explore their governance arrangements and discusses the implications of this analysis for the viability of some specific DLS applications. We provide a comprehensive overview of the use of distributed ledgers (referring mainly to Bitcoin and Ethereum) and examine the rights roles and incentives of stake-holders, including miners, ordinary coin holders as well as other application users. This includes the role of payments as incentives in proof-of-work and proof-of-style systems.

Keywords

Blockchain, distributed ledger, polycentric governance, club governance, distributed consensus

1 Introduction

> "That money can be kept separate from the state and from the political process leading to the formation of our governments and their policies is a dangerous illusion."[1]

Blockchains are the first, and most well-known example of a distributed ledger technology (DLT). The first known commercial use of a blockchain was for the cryptocurrency Bitcoin. The Bitcoin network was established by a pseudonymous identity within the open-source computer programming community, from which the code underlying distributed ledger (DL) systems was developed and continues to be made available for common use (Nakamoto 2008). A fundamental objective of Bitcoin was to provide a decentralised, de-nationalised and apolitical currency competing with its very antithesis: centralised and politically-influenced national fiat currencies overseen by central banks and intricately interwoven with national political agendas considered by many to be captured by elite interests (Varoufakis and Moe 2018). A further defining feature of Bitcoin was its ability to enable individuals to transact anonymously on the platform, and thereby outside of other centralised,

[1] Varoufakis and Moe (2018) p 159

government-controlled regulatory regimes. The anonymity of those creating it served to reinforce the perception that the Bitcoin system was "owned and controlled by no-one" and operated solely to support the interests of those using it.

Due to their departure from traditional top-down hierarchical control of data, distributed ledger systems (DLSs) – including blockchains – have been heralded as revolutionary tools capable of transforming both commercial and other digital interaction. DLS are distinguished from classical database systems by the absence of a centralised entity 'owning' a single authoritative copy of the ledger and controlling access to and updating of its content. Rather, ledger data is considered to be in the public domain, as it can be viewed without permission and is updated without the need to channel through a single centralised portal. By removing centralised intermediaries, DLS are held to eliminate the exertion of market power by centralised controllers and have been heralded as a new paradigm of management potentially capable of disrupting traditional forms of governance (Pilkington 2016), in part due to their potential to support 'smart contracts' – a computerised transaction protocol that executes the terms of a contract (Szabo 1997; Gans 2019). Whereas Davidson, Filippi, and Potts (2018) propose blockchains are an institutional technology, thereby offering an alternative to firms and markets as a means of organising transactions, A. Berg, Berg, and Novak (2018) go so far as to suggest that DLS function as constitutional catallaxies governing interactions by numerous disparate 'citizens' in a democratic community of activity.

In the view of A. Berg, Berg, and Novak (2018), "blockchain co-ordination changes and adapts not only to the technological limitations of the available protocols, but to mutual expectations and influence of interacting stakeholders." However, they conceptualise the blockchain protocols and software associated with the DLS as setting the boundaries of the constitutional arrangements of interest. They do not extend their analysis to take account of how the behaviours and expectations of the interacting stakeholders in regard to the DLS may be influenced by elements external to it. Yet if DLS operate as "democratic" constitutional catallaxies, then the community of "citizens" must necessarily be governed by an overarching set of constitutional arrangements defining both the transactional relationships between them and the way in which the community itself will be governed. While DL software may govern the former, it begs the question of how the latter are determined and how they will operate over the lifetime of the DL and its community. If an inefficient 'tragedy of the commons' for resources with no owners is to be avoided (Hardin 1968), some form of governance control of the community as well as the resource is essential (Blomquist and Ostrom 1985).

In this paper, we utilise Elinor Ostrom's Institutional Analysis for Development (IAD) framework (E. Ostrom 2010) to explore the economic, technological, political, social and psychological contexts in which DLSs operate in order to understand their broad governance arrangements as polycentric systems operating simultaneously at many different levels of interaction. We draw on prior conceptualisatons of the internet as a common-pool resource (CPR) (e.g. Bernbom 2000), and analyses of community governance of the digital commons (e.g. Schweik 2007; Dulong de Rosnay and Le Crosnier 2012) to propose that a DLS is a common-pool resource. However, the examples considered here mostly exhibit governance arrangements that fail to meet E. Ostrom (1990)'s principles of effective CPR governance (derived from Blomquist and Ostrom 1985). In particular, while proposing that they are predicated on distributed ledgers managed by peers operating in an egalitarian fashion using democratic consensus processes, a rather different picture emerges of different classes of system participants exercising very different rights of system governance. While this is not unusual in polycentric systems, we find that one of the primary beneficiaries of value extraction – coin holders – typically have no formal ability to influence system decision-making, while those controlling most management rights – system software developers – need not have any material interest in the system (i.e. they do not need to be coin holders). This leads to questions about the sustainability of resource management in the long run.

Furthermore, somewhat ironically given the purported motivations of decentralisation and revolutionary changes purported for the blockchain technology compared to established firms and governments, it appears that most DLS applications substitute centralised control by one set of stakeholders (e.g. nation state governments in the case of currencies; or the owners and controllers of centralised databases and portals to them) with centralised control by a different set of stakeholders (e.g. those designing the DLS and the rules governing how and when changes can be made to the algorithms and software governing the systems). The extent to which stakeholders at various levels can interact with each other in the governance of the DLS as an entity in its own right (separate and distinct from the ownership and management of the records held on it) will ultimately determine its viability in the long run. We suggest that DLSs giving consideration to Ostrom's eight principles will prove more resilient and sustainable, especially in support of commercial activities. Already, we observe DLSs exhibiting more of these principles emerging to support commercial endeavours (Howell, Potgieter, and Sadowski 2019).

The paper proceeds as follows. The subsequent section defines distributed ledgers. Section 3 summarises the key features of common-pool resources and the IAD framework whereas Section 4 identifies the features of DLs that render them candidates as common-pool resources. Section 5 uses the IAD framework to explore their governance arrangements and Section 6 discusses the implications of this analysis for the viability of some specific DLS applications. Section 7 concludes.

2 Distributed Ledger: a definition

A distributed edger (DL) is a database (or file) spread across several nodes or computing devices. Each node in a network has access to (and probably saves) an identical copy of the entire ledger. Unlike classic databases, the ledger is not maintained by any central authority. The integrity of the ledger is maintained automatically by an algorithmic consensus process whereby nodes vote and/or agree upon the authoritative version, which is then updated and saved independently by each node. In effect, the consensus algorithm operates in the manner of a decision-making process within a governance system. DLs underpin the burgeoning number of cryptocurrencies, but the technology is also anticipated to facilitate a much wider range of applications, principally because of the capacity of the software implementing them to embed instructions to be executed automatically in the future when certain criteria are met – so-called "smart contracts."

The concepts of an electronic *ledger* and *distributed ledger* are very well described by Anta et al. (2018). By ledger we mean, as in usual parlance, an object which starts empty and grows through records being appended to it. Note that this does not necessarily imply that the records are appended in a linear fashion, so the ledger itself need not appear strictly linear. It can, for example, be a tree. In traditional accounting (based on physical books) the objects are generally linear and of course the data objects that constitute different kinds of ledgers are also linearly representable so the distinction is actually unimportant as long as the records being appended can themselves be represented in a linear fashion. The linear order can, trivially, be implied by the time of addition of an item (which may included a description of its relation to preceding items).

The old-fashioned ledger is necessarily centralised – there is an authority that keeps the ledger and it may or may not be open to inspection but is not generally available for copying. By a distributed ledger, we generally imply not only that there are numerous copies of the ledger but also that there is no centralised authority that broadcasts records to be appended to the ledger. This directly implies a consensus decision mechanism for adding records that the ledger copy owners agree on. This need not necessarily imply a convention that resides in part in the normal social realm, if only on how to refer to (i.e. name) the specific common ledger. It also implies in fact that the common data object is or has an underlying ledger in the usual sense of the word. For, in order to agree on how to change the common data object, distributed ledger owners (or, keepers) need to receive a sequence of stepped changes. This sequence of stepped changes can be seen as the underlying ledger even if the visible data object is something like a database, in which the underlying ledger is just a journal.

Blockchains are a specific kind of distributed ledger where the common rules for changing the data object are of a particular kind. Every blockchain is a distributed ledger but not vice versa. A prime example of a distributed ledger which is not a blockchain, is Holochain. The technological characteristics of blockchain systems are well documented (Narayanan et al. 2016). Considerable faith has been placed in the technology as a means of revolutionising digital transacting (Mulligan et al. 2018; Crosby et al. 2015; Czepluch, Lollike, and Malone 2015; Swan 2015), due to their promise of transparent, tamper-proof and secure systems enabling novel business solutions (Andoni et al. 2019). However to date, outside of the arena of highly-publicised cryptocurrencies such as Bitcoin, few examples exist of the use of the technology to support significant economic activities. Nonetheless, Ethereum (as well as offering the cryptocurrency ether) supports 'smart contracts' and plans for many other blockchain systems have been announced – for example Sovrin for identity management, and Halo for supply chain management.

3 Common-Pool Resources and the IAD Framework

E. Ostrom (2000) and E. Ostrom (2009), drawing on a long heritage of governance studies including but not limited to Samuelson (1954), James M Buchanan (1965), James M. Buchanan (1967), Olson (1965) and Hardin (1968), defines *the commons* as a resource shared by a group of people. She draws a clear distinction between the commons as a resource or resource system and the commons as a property rights regime. Common-pool resources are defined as economic goods independent of any property rights associated with them (e.g. fish in the sea), whereas common property defines a legal regime governing a jointly-owned set (or bundle) of rights associated with the resource. Whereas Samuelson (1954) distinguishes between private and public goods on the basis of the nature of rivalry in use (one person's use excludes another for private goods), and the ability (or high cost) of excluding others from using them and James M Buchanan (1965) delineates a non-rival good where exclusion is easy as a club good, whereas one where exclusion is difficult is a classic public good, E. Ostrom and Ostrom (1977) demarcate common-pool and public goods on the basis of subtractability.

With classic public goods, there is no limit as to how many people can utilise the item for which it is costly to exclude individuals from using it. A piece of information or the sunset are public goods, because utilisation of them is limited only by demand and excluding individuals from using them is costly. On the other hand, while it is costly or difficult to

		Subtractability	
		Low	High
Exclusion	Difficult	Public goods Useful knowledge Sunsets	Common-pool resources Libraries Irrigation systems
	Easy	Toll or club goods Journal subscriptions Day-care centres	Private goods Personal computers Doughnuts

Table 1: Types of goods (Source: E. Ostrom and Ostrom 1977)

exclude someone from using a common-pool resource, one person's use does subtract from the benefits from use available to others. For example, the water in a river is a common-pool resource, because when one user extracts water for a given use (e.g. irrigation), that water is not available for use by someone else downstream. This leads to the four-quadrant classification of Table 1.

Common-pool resources are comprised of resource systems and a flow of resource units or benefits from these systems (Blomquist and Ostrom 1985). The resource system (or alternatively, the stock or the facility) is what generates a flow of resource units or benefits over time (Lueck 1995). Examples include dams, forests, rivers and "common" fields. Facilities constructed for joint use, such as mainframe computers and the internet may also be common-pool resources. The resource units or benefits from a common-pool resource include (for example) water, timber, medicinal plants, fish, fodder, central processing units, and connection time. Successful management of a common-pool resource requires rules that limit access to the resource system and rules that limit the amount, timing and technology used to withdraw diverse resource units from the resource system (Ostrom and Hess 2010).

Hardin's "tragedy of the commons" holds that the high costs of excluding individuals from using common-pool resources creates incentives for over-exploitation that can be managed (second-best) efficiently by government assuming ownership and control of the resource and using its legislative and regulatory powers to constrain its exploitation (including, for example, the creation of a limited number of tradeable usage rights that effectively privatise the resource – e.g. fishing quotas or spectrum bundles). Ostrom and Hess (2010), however, contend this is not the only possible solution. Citing a number of studies (e.g. Robert M. Netting 1976; Robert McC. Netting 1981; Ellickson 1993; Sengupta 1991; Nugent and Sanchez 1993) they provide many examples (e.g. Swiss mountainside grazing land; colonial territories; irrigation systems; Sudanese grazing land) where decentralised groups of stakeholders in a common-pool resource can develop co-operative management strategies enabling durable, robust systems that mitigate the perverse incentives to over-exploit without the need for overarching hierarchical government control.

These self-governing commons exhibit collective action in decision-making based on a high degree of social capital derived from strong social networks and shared norms of reciprocity, which, depending upon the contexts in which they are applied, offer more efficient governance arrangements than exogenously-imposed, centralised alternatives. These arrangements are characterised by user-created boundary rules for determining who could use the resource, choice rules related to the allocation of the flow of resource units, and active forms of monitoring and local sanctioning of rule breakers (E. Ostrom, Gardner, and Walker 1994).

3.1 The IAD Framework

The IAD framework was developed as a diagnostic tool to investigate situations where humans repeatedly interact within rules and norms that guide their choice of strategies and behaviours. It seeks explanations of how fallible humans come together, create communities and organisations, and make decisions and rules in order to sustain a resource or achieve a desired outcome. It can be used to analyse static situations crafted by existing rules and relating to an unchanging physical world and relevant community, and equally to analyse dynamic situations where individuals develop new norms, new rules, and new physical technologies. E. Ostrom and Hess (2007) observe its particular suitability for analysing resources where new information technologies are developing at an extremely rapid pace and have redefined knowledge communities, juggled the traditional world of information users and information providers, rendered obsolete many existing norms, rules, and laws and have led to unpredicted outcomes.

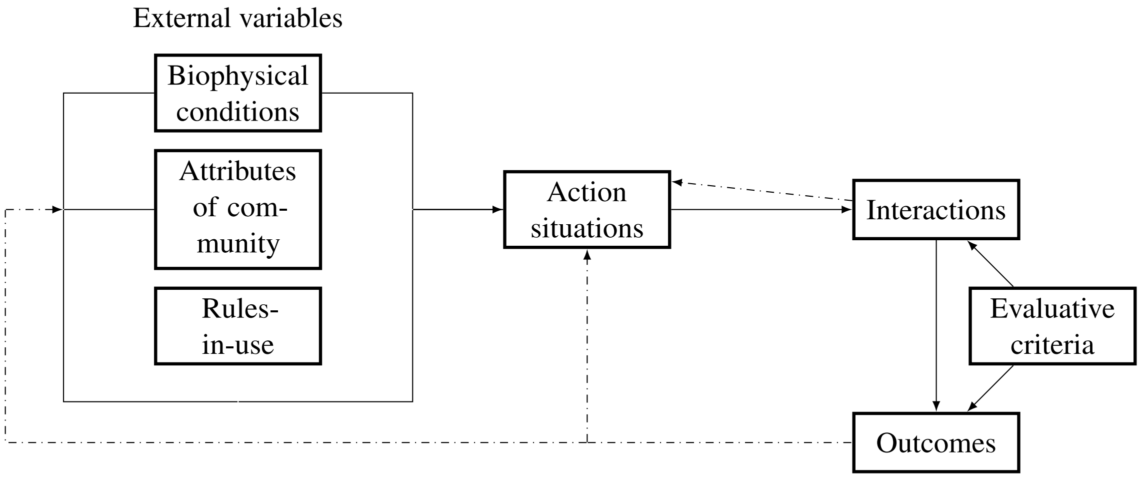

Figure 1: The IAD Framework (Source: E. Ostrom 2010)

The IAD framework (Figure 1) embodies a general set of variables that an institutional analyst may want to use to examine an institutional setting, and is sufficiently broad enough to encompass human interactions within markets, private firms, families, community organisations, legislatures and government agencies. It can be 'unpacked' multiple times to reflect complex interactions.

At any particular time, external variables affecting an action situation are:

1. Biophysical conditions, which can be simplified in some analyses to be one of the four goods in Table 1;
2. Attributes of a community, which may include history of prior interactions, internal heterogeneity or homogeneity of key attributes and knowledge and social capital of participants; and
3. Rules-in-use, specifying common understandings related to who must, must not or may take which actions affecting others (subject to sanctions). They may evolve over time as interaction occurs, or be self-consciously changed via a collective choice or constitutional-choice process.

External variables interact via action situations (which can be modelled using a variety of different tools, such as, but not limited to, game theory and agent-based modelling) leading to outcomes, the success of which can be determined objectively via a range of evaluative criteria. These evaluations can then contribute to changes in the patterns of (inter)action and/or changes to the variables themselves. The IAD framework has been used to explore many systems involving common-pool resources. For example, it is used in E. Ostrom and Hess (2007) to analyse the knowledge commons, which provides an exemplar for this paper drawing together elements of the shared data resource held on DLS, created maintained and shared using software derived from the open-source community.

3.2 Bundles of property rights related to common-pool resources

Common-pool resources can be owned and managed as government property, private property, community property or owned by no-one. They can be conceptualised as a bundle of property rights that individuals using a common-pool resource may cumulatively have. E. Ostrom (2010) categorises these as:

1. Access: The right to enter a defined physical area and enjoy nonsubtractive benefits (for example, hike, canoe, sit in the sun);
2. Withdrawal: The right to obtain resource units or products of a resource system (for example, catch fish, divert water);
3. Management: The right to regulate internal use patterns and transform the resource by making improvements;
4. Exclusion: The right to determine who will have access rights and withdrawal rights, and how those rights may be transferred; and
5. Alienation: The right to sell or lease management and exclusion rights.

The allocation of these rights as they relate to the biophysical characteristics of the institutional arrangement, across the various stakeholders and other attributes of the system community, and how they are managed under the rules-in-use is fundamental to understanding any institutional arrangement. They define actions: insofar as if an individual has a right, then someone else has a duty to observe that right.

3.3 Adaptive governance in complex systems

Rapid change in the environment and in a community is a major challenge to any governance system. For any governance system to be resilient and capable of adaptation over time, it requires arrangements at all levels (operational, collective and constitutional) covering the interactions of all stakeholders (users, providers, policy-makers) that align interests and enable modifications to be made when circumstances (either exogenous or endogenous) change. This requires good trustworthy information available where and when it is needed, effective inducement of rule compliance, means of dealing with conflict and an openness to the need for future change. E. Ostrom (2010) observes: "fixed rules are likely to fail because they place too much confidence in the current state of knowledge, while systems that guard against the low probability, high-consequence possibilities and allow for change may be suboptimal in the short run but prove wiser in the long run."

Based on a very large number of cases, Blomquist and Ostrom (1985) identify a core of underlying features characterising long-sustained common-pool resource regimes that were not present in systems that failed. These relate to the following.

1. Clearly-defined boundaries: User Boundaries between legitimate users and nonusers are present; Resource Boundaries clearly separate a specific common-pool resource from a larger social-ecological system.
2. Congruence of rules regarding the appropriation and provision of resources that are adapted to local conditions: appropriation and provision rules are congruent with local social and environmental conditions; appropriation rules are congruent with provision rules; the distribution of costs is proportional to the distribution of benefits.
3. Collective choice arrangements: most individuals affected by a resource regime are authorized to participate in making and modifying its rules.
4. Effective monitoring (of the condition of the resource and its usage): by monitors who are part of or accountable to the appropriators.
5. Graduated sanctions: sanctions for rule violations start very low but become stronger if a user repeatedly violates a rule.
6. Conflict resolution mechanisms: rapid, low cost, local forums exist for resolving conflicts among users or with officials.
7. Minimal recognition of rights: the rights of local users to make their own rules are recognized by the government.
8. Nested enterprises: when a common-pool resource is closely connected to a larger social-ecological system, governance activities are organized in multiple nested layers.

3.4 For analysis

The IAD framework provides a means of analysing the system in which a DLS operates. Rather than focusing solely on its technological features, it allows consideration of the entire context in which interaction occurs between its stakeholders across multiple levels. The analysis requires an understanding of the DLS as an institutional system in its own right, in addition to the particular resources it governs. This requires an understanding of how the rights to the system and the property within it are governed. If a DL system is indeed a common-pool resource, then the extent to which we observe the characteristics associated with sustainable management systems will inform us of the potential for such systems to operate sustainably and beneficially in the long-run.

4 DLS as a common-pool resource

DLSs offer an alternative to centralised control as a means of addressing the classic economic dichotomy between (physical) resource scarcity and informational resource reproducibility.

4.1 Tension: cost reduction, free-riding and information creation incentives

Information (knowledge) is the classic non-rival and non-excludable public good, under-produced because of its high costs of initial creation and difficulty in preventing others from free-riding on its use once created (Arrow 1962). To encourage its creation when copying and communicating costs are low, government-managed mechanisms such as copyright and patent legislation have placed formalised constraints on free-riding, at least insofar as commercial use is concerned. Digitisation of information has, however, created new challenges. The costs of copying, transportation and storage have been reduced dramatically, along with the ability to observe such activities, leading to the risks of free-riding increasing.

The approach most commonly adopted to address the free-riding risk has been to 'enclose' digital information within encrypted and password-protected repositories accessible only under strict centrally-controlled authorisation protocols. On the one hand, these arrangements provide some protection against free-riding, at least in relation to the original reason for creating it. On the other hand, they prevent the information from being used by non-authorised parties for other unrelated value-creating activities. For example, information on an insured individual's medical care collected by the insurance company is not available to medical researchers, without explicit authorisation first being obtained. The need to acquire authorisation necessitates a centralised gate-keeper who can be held accountable for any special access rights granted (e.g. ensuring personal identifying information is not revealed), but brings with it the attendant potential for monopoly pricing of (or other non-price restrictions on) those rights, and therefore inefficient under-utilisation of the digital information. A single point of access also creates a vulnerability to unauthorised data manipulation, both due to weak access authorisation protocols and having a single point of attack.

Bernbom (2000) identifies the tension between the apparently limitless ability for individuals to access and copy digital information over the internet and the physical constraints of congestion on the infrastructures that such use engenders. He discusses the various communities of interest that have emerged to enable these tensions to be addressed. Success is evident in the burgeoning use of digital information over an ever-growing infrastructure where incentives are provided to minimise losses arising from congestion. Indeed, the degree of formalisation of some of these collaborative arrangements over time, with associated reductions in the costs of exclusion (albeit in part influenced by government policy and regulation), has led Raymond (2012) to conclude that the internet infrastructure at least is now no longer appropriately characterised as a common-pool resource but rather more closely resembles a club good. Dulong de Rosnay and Le Crosnier (2012) similarly illustrate how classical copyright measures have imposed constraints on the use of digital copies of creative works not encountered in the physical world, constraining the ability for further creation relative to the physical world, but the creation of the digital commons embodying a different set of norms controlling the use of digital works has mitigated some of these inefficiencies.

A specific DLS can be considered, like the internet, as a technological common-pool resource in respect of apparently limitless access to the information held on the ledger, and limitless ability to copy and amend the software managing it. The algorithms determining when and how information can be added act in a similar manner to congestion or a centralised editorial function, constraining the ability to duplicate or amend that information, thereby providing the incentives to use the ledger to store information.

4.2 DL technologies and the information tension

DLs dispense with a centralised gatekeeper 'problem' entirely by enabling every stakeholder in the system (nodes) to hold a copy of the ledger. The software used to interpret and manipulate the ledger is created and distributed according to the principles of the open-source community. Open-source software itself is a common-pool resource established and maintained by a collaborative community as an alternative to the commercial licensing model of proprietary software limiting the number of installations on which the software can be installed and precluding the license-holder from examining and/or modifying it (Schweik 2007). In open-source arrangements, actors (typically computer programmers and users of the software) with defined rights can access, use, redistribute and modify the software with few limits (e.g. "Copyleft", with minimal restrictions, and more restrictive open-source licensing[2] which requires users to adhere to a moral code designed to give credit to original creators and to protect them and their outputs from risks of reputational harm from derivative works).

As every node in a DLS has a copy of the ledger and software, there appear to be few limits to copying or altering either, suggesting the potential for inefficient over-copying. However, the integrity of the DL content is maintained by the use of a consensus algorithm embedded in the software. Nodes, communicating with each other using peer-to-peer protocols, agree on

[2] https://opensource.org/

what constitutes the “correct” information. Transactions that will change the ledger content are proposed by another user class (transactors), which may or may not also operate as nodes. As the ledger is strictly time-bound, the only changes permitted are additions (in a blockchain system, additions occur when new links are added to the chain). Two different protocols are typically used to ensure the legitimacy of the data added to the ledger: “proof of work” and “proof of stake.”

4.2.1 “Proof of work” and the information tension

In “proof of work” (POW) systems, the nodes solve complex but meaningless unique cryptographic puzzles generated by the transactions, and then compete to propose the next update. The ‘winner’ proposing the most popular ‘solution’ becomes the accepted update (that is, proof that work has been undertaken is provided by the provision of an answer). Nodes are discouraged from knowingly proposing a chain that will be rejected because it differs from the others, as considerable resources are expended for negligible chance of a return being obtained.

Providing there are sufficiently many nodes, it is extremely costly (though not infeasible) for nodes to collude to propose a fraudulent block (e.g. “double-spending” the same coins on a cryptocurrency blockchain, or altering the content of historic blocks containing smart contract instructions in a blockchain enabling such a feature). To propose a “double spend”, the fraudulent transactor would have to ensure that the node winning the competition contained the fraudulent information – something statistically feasible only when the perpetrator controls close to 50% of the node processing capacity (hence such activity is called a “51% attack”).

In a similar fashion, changes to the software controlling material transaction and consensus processes can also be achieved only when a significantly large-enough number of node operators agree to adopt the software incorporating these changes. But whereas the processes for achieving consensus on ledger content are automated using the blockchain protocol, the decision to implement a different form of processing software on any particular node is typically a human-mediated activity. When two (or more) different forms of software compete to be used by the nodes, a “fork” can occur, with the community dividing into two (or more)sub-communities. The success of each sub-community depends upon a sufficiently large-enough number of nodes supporting each variant. If too few nodes support a variant, the alternative ledger content generated might lack integrity and/or credibility. Such a variant will be unable to attract transactors, so will cease operating. The larger is the DL network, the higher will be the costs of gaining agreement of sufficient nodes to change to the alternative software. These costs therefore counteract the risk of ledger fragmentation for sufficiently large systems. However, in extreme cases (akin to constitutional changes in a nation state), the change may not be well-supported and division (secession) ensues.

The most high-profile case to date of a POW fork being orchestrated concerned Ethereum. In 2016, a smart contract-based application using the Ethereum platform was hacked and ether worth around $50 million was stolen. While the attack was made possible by a flaw in the application and not the Ethereum platform itself, the reputational consequences for the Ethereum system were considerable. Following discussion amongst the Ethereum community (self-identified stakeholders expressing an interest via either the “official” Ethereum forum https://forum.ethereum.org/ or on social media such as https://www.reddit.com/r/ethereum) a “hard fork” was executed: a change was made to the software code to move the stolen funds to a new smart contract allowing the defrauded owners to withdraw their funds. Node operators were presented with a choice of adopting the new software or not. A significant minority disagreed with the action because altering the code was viewed to be contrary to the principles that blockchain systems remain secure, anonymous, tamper-proof and unchangeable and hence unchanged, even in the event of an undesiraable event.

The result was that the community split in two, with those node operators preferring no change continuing to transact using the old software on the Ethereum Classic blockchain, while the Ethereum chain continued to be operated by nodes using the amended software. Thus, both chains are identical up to the point of forking; from that point onwards, the two different chains reflect the different software protocols. This incident well illustrates the effect of real-world arrangements on technocratic DLS.

4.2.2 “Proof of stake” and the information tension

“Proof of stake” (POS) systems aim to achieve a similar degree of assurance as POW systems without the need to expend large sums in computing resources (notably electricity) to solve the cryptographic puzzles. These systems work on the assumption that those with the greatest stake in the system have a vested interest in ensuring the integrity of the ledger is maintained. This is achieved by requiring the node proposing the next update to post a financial stake (bond) as an assurance of the correctness of the information contained in it. If it is subsequently demonstrated that the information contained in the proposed update is erroneous, the proposer (known as a ‘forger’) loses the stake and the right to participate in future forging activities. The stake associated with a given update can be held in the equivalent of an escrow for a considerable period (several months is not unusual). The larger is the stake the forger stands to lose, the greater is the reliance that can be placed on the validity of the information in the ledger update proposed by that forger.

4.2.3 A note on coins

Coins (tokens) created and issued by a POS DLS are integral to its operation, regardless of whether or not its primary purpose is as a cryptocurrency (see below for the more on cryptocurrency tokens). Node operators must first acquire coins, either at the outset from an “initial coin offering” reminiscent of the initial public offering of shares in a listed company, or subsequent purchases in a market operated as one of the functions of the DLS. The larger and more diverse is the base of node operators, the more broadly will the allocation of the right to forge the next block be spread. The more valuable are the coins (measured by the price in the market), the more the forger risks losing if lodging an incorrect update. To this end, much rests on the allocation of coins in the initial offering, and how the coin market develops over time.

Typically, a fixed number of coins is issued at a low price (they may even be given away), with a maximum number allowed per node to guard against concentrated ownership (at least at the node level – anonymity of node owners makes it difficult to guard against one party or coalition controlling many nodes). The initial node operators may not pay much for their coins, but they face strong incentives to eschew unethical behaviour which will detract from the reputation of their DLS which is reflected in the value of the coins in the market. If the number of coins is fixed, the risk that coin value will vary due to factors other than confidence in the integrity of the DLS is low. Over time, the higher (lower) is the coin value, the more (less) the node operators risk losing by lodging untrue information. A single node operator can dilute individual risk by lowering their stake or selling part of it to a new operator. To prevent over-much dilution, typically a minimum number of coins is required for each stake.

4.2.4 Selecting the ‘winning’ node

Selecting the node providing the next update typically relies on either or a combination of randomised selection and coin age-based selection. The general principle of randomised block selection is that the node operator with the largest stake has the highest chance of producing the next update. As the size of the stakes is public knowledge, each node is able to predict which proposer will be selected to forge (i.e. cast, as in blacksmithing) the next block. However, these systems favour those with the largest coin holdings, who can post the largest stakes. Coin age-based selection spreads block formation more broadly. In this method, the system selects the next forger based on the ‘coin age’ of the stake. This is calculated by multiplying the number of days the coins have been held as a stake by the number of coins that are being staked. Users who have staked older and larger sets of coins have a greater chance of being assigned to forge the next block.

In one example, coins must have been held for a minimum of 30 days before they can compete for a block. Once a node operator has forged a block, their coin age is reset to zero and then they must wait at least 30 days again before they can sign another block. To prevent users with very old and large stakes from dominating, all users are assigned to forge their next block within a maximum period of 90 days. Because a forger’s chance of success goes up the longer they fail to create a block, forgers can expect to create blocks more regularly in a coin age-based system than one based solely on randomised block selection.

As “proof of stake” systems do not require any special computing equipment or large energy expenditures to solve cryptographic puzzles, they are argued to have lower entry costs and therefore be more likely to attract large numbers of users/node operators than “proof of work” systems. This may hold at the outset, when the coin price is low. As the coin price rises, buying a stake may become costly (a barrier to entry, as coins are required to participate), unless there are provisions for the minimum stake to vary inversely with coin price. However, the lower is the coin price, the higher is the risk of a “nothing at stake” problem occurring, where nodes have nothing to lose from voting for multiple ledger variants and preventing consensus from being achieved. Lowering prices also increase the risk of consolidated node holdings and associated co-ordinated activity occurring.

4.3 Tokens and transaction-processing incentives

While they exhibit considerable technological and computational sophistication in the digital environment, if DLS are to be successful, they must support real-world activities that are valued by (at least) some human participants at more than their costs of operation – including a return on capital, including the opportunity cost of human capital deployed in (notably) open source software communities (Gans 2019). It is important now to distinguish between the costs of operating the DL infrastructure and the costs of developing and maintaining the meaningful applications it supports. This parallels the costs of operating the internet: the appropriate infrastructure will not be provided unless telecommunications operators can make at least a fair return on the large amounts of capital deployed, while application providers must be able to meet the long-run costs of their outgoings to survive.

In the earliest days of the internet, the provision of both infrastructure and applications were typically bundled together (e.g. ISPs provided both access to infrastructure and services such as email), but over time the two have become more separate (though we are currently witnessing increased incidence of mergers between content distribution networks and ISPs, and Google has experimented with owning a fibre network in some locations). Similarly, current DLS typically bundle together infrastructure and applications (e.g. the ledger record updating mechanisms and the currencies they support). It is helpful to use Hansmann (1996)'s classification of costs as ownership and market contracting costs when considering how the various stakeholders in a DL system are remunerated. Transactions incur costs for node operators (digital activity costs) separate and distinct from the costs incurred by the transacting parties in striking their agreement (real-world costs) and monitoring and verifying performance (which may or may not be lower using blockchain applications than other alternatives). As both node operators and transactors must have coin holdings to participate, they face (risks) associated with fluctuations in coin values, separate from the costs of transacting, which can also be uncertain. On the other hand, node operator costs per transaction are more predictable.

Necessarily, the evaluation of incentives to own infrastructure and offer services will depend on exactly what services are offered (as with the internet, demand for DL services is a derived demand, dependent upon demand for the applications using them and not for the technology per se). So far, two main application models have emerged: cryptocurrencies (as per Bitcoin) and the provision of a DL platform on which independent applications can be hosted (as per Ethereum). It also depends on the choice of consensus algorithm.

4.4 Payment and incentives in "proof of work" systems

The real costs of node operation in POW systems consist of the costs of obtaining the necessary computer capacity and the electrical power resources required to solve the cryptographic puzzles associated with transactions. Assuming that all nodes are equally likely to propose the 'winning' solution, most DLSs reward the winner with a fee associated with the transactions reflected in the update. This fee can be paid explicitly by the transactors, or embodied in the fabric of the ledger itself, using the system-generated and governed currency. Some systems use a combination of the two methods.

To the extent that the DLS provides services that are valuable to transactors, that value can be attached to system-specific tokens (coins) which can then be traded just like any other commodity. This arrangement forms the basis of cryptocurrencies. In essence, all DLS operate as cryptocurrencies in the first instance – either explicitly as per Bitcoin, implicitly as a means to other transaction ends, or both (e.g. Ethereum).

4.4.1 Cryptocurrencies: e.g. Bitcoin

Cryptocurrency DLS record the transfer of claims to the tokens among the various members of a 'community' as they transact amongst each other (e.g. buying and selling real-world commodities paid for with system tokens). It has been argued that cryptocurrencies may be more properly described as payment systems rather than currencies, as few merchants cite prices in bitcoin (for example), and the extent to which bitcoin can be considered a reliable store of value is debatable, especially following the 'spike' in prices in late 2017[3]. The first cryptocurrency DLS, Bitcoin, was established as a decentralised, and allegedly anarchistic, competitive alternative to fiat currencies operated by central banks. The originating white paper states "The root problem with conventional currencies is all the trust that's required to make it work. The central bank must be

[3] Cryptocurrencies are considered by some to be inferior to fiat currencies as a store of value because the inability to alter the number of coins in circulation in response to specific circumstances (as occurs when central banks intervene in currency markets) renders these systems especially prone to volatility (Varoufakis and Moe (2018)).

trusted not to debase the currency, but the history of fiat currencies is full of breaches of that trust" (Nakamoto 2008). To this end, the decentralisation of the Bitcoin ledger was in large part to eliminate the need for trust in a centralised entity for a currency system, but in effect Bitcoin also eliminates the need for a (trusted) central clearing process in a payment system, e.g. Visa, Paypal, etc. We consider that cryptocurrencies meet both requirements, as for the purposes of token management, they act in the manner of a central bank issuing coins, but for the purposes of transaction fulfilment, they act in the manner of online retail banking systems, enabling account holders to assign the property rights to funds in an account in their control to another account (which is most likely controlled by another actor). The Bitcoin system is calibrated so that one new 'block' is added to the chain around every ten minutes, so the transfers can be considered to be "batched" in a similar manner to retail banks' clearing house processes, rather than being confirmed in near real-time (as in centralised systems such as Paypal).

The property rights associated with tokens are effectively assigned to members of the DLS community (transactors and node operators) and recorded in the ledger (linked to the 'owner') by way of unique codes (addresses) part of which is known only to that owner (the private key) and part of which constitutes the address (the public key). The ability for users to generate addresses through corresponding keys relies entirely on two things:

1. a presumption that large good pseudo-random numbers can be generated; and
2. the mathematical theory of public-key cryptography that was discovered in the 1970s.

As the ledger records are unique and indelible, token rights are rival and excludable. Hence, they constitute a classic private good with a value determined by trading them in a market e.g. as cigarettes were used in World War 2 prisoner of war camps (Varoufakis and Moe 2018). The transactor proves ownership of the rights by submitting a message generated using the private key associated with the address of the source address (identified by a public key, another mathematical intricacy!), and can transfer the rights to another transactor by linking them in a new transaction to an address specified by the new owner via a public key associated to the recipient's private key. The value of the tokens alters as a proportion of the expected value of the ledger services to transactors in aggregate, denominated by the number of tokens issued. The value of tokens can, just as with any other currency, be influenced by altering the number of tokens in circulation. The difference between a DL token currency and a fiat currency is that the number of tokens and the conditions of their issue can be and generally are controlled by the software and not a central agency such as a central bank. Confidence in the performance of the ledger, rather than confidence in the ability of a sovereign government to fulfil its obligations, underpins token value.

By way of illustration, the Bitcoin POW blockchain coins are called bitcoin. The number of bitcoin is capped at 21 million; the 17 millionth was reportedly issued on April 26 2019.[4] In the first instance (i.e. until all 21 million bitcoin are in circulation), new bitcoin are issued and paid to the 'winning node' each time blocks are added to the chain. The assumed costs to each node of performing the calculations is known as "proof of work." The number of bitcoins issued each time a block is added decreases as a function of the number of coins already issued. In the early days of the system, to encourage both the addition of nodes (known as miners) and transactors, miners were rewarded fully by the issue of new bitcoin.

However, as the mining reward decreased, miners increasingly needed to extract payment from transactors. Thus, over time, the fee has included a larger proportion from an explicit fee paid (in bitcoin) by the transactor initiating the transaction, and a smaller proportion from winning the mining race. When all 21 million bitcoins are issued, rewards will be derived solely from user-paid transaction fees. However, given the parameters of the algorithm, this is not anticipated to occur within the next 122 years. Incidentally, it is possible that many bitcoin have been lost. If bitcoin are owned by an address and the entity controlling that address loses the attendant private keys, it becomes practically (we think) impossible to spend the bitcoin 'in' that address. Similarly, if someone accidentally or on purpose sends bitcoin to an address for which no-one has the private keys, that bitcoin is allocated in the DL but become unspendable.

Miners select the transaction problems to be solved according to the magnitude of the combined system and transactor fee reward on offer. Both the system and transaction fees are independent of the bitcoin value associated with the transaction, (a set 'prize' is paid for 'winning'). Transactors can increase the probability of having their transaction added to the ledger chain (the equivalent of 'clearing' in bank processes) sooner by proposing a high transaction fee. Transactions with low fees may sit around unfulfilled for long periods. At 1.19am GMT on April 29, 2019, a transactor wishing to have a transaction in the next block added (that is, 'cleared') within 10 minutes needed to pay bitcoin to the value of at least US$ 0.69, within three blocks (30 minutes) US$ 0.69 and within 6 blocks (one hour) US$ 0.26.[5]

Average transaction fees in US dollars have varied over time due to exchange rate fluctuations and transaction numbers. The number of transactions per day was initially small, but grew rapidly to a peak of around 490,000 in late December 2017.

[4] https://www.cnbc.com/2018/04/26/there-are-now-17-million-bitcoins-in-existence\T1\textendashonly-4-million-left-to-mine.html

[5] https://bitcoinfees.info/

Numbers fell off dramatically in January 2018 (to around 140,000) but have climbed steadily since. The average for April 2019 is around 360,000[6]. The average value per transaction per day also varies. The 2019 average from 1 January to 30 April is around US$ 22,000[7]. The median transaction value is close to US$ 250. Thus, Bitcoin appears to underpin a comparatively small number of very large transactions, rather than being used for a large number of small everyday transactions.

Each block added to the chain generates a 'winner-takes-all' payment. On average, if all miners face identical costs and transactions, the 'random' nature of the cryptographic puzzles means all miners have an equal chance of winning. Hence, in the long run, miners will recover their mining costs so long as the combined transaction payment exceeds these costs. However, miners with better (faster) equipment have a higher chance of winning as they can generate more "work" in the same time as ones with less-capable equipment. Combined with the downward price pressure applied by the release of new coins and transactors seeking to pay the least sum for each transaction, the ability to alter the odds of winning provides strong incentives for node operators to to keep improving the efficiency of their computing equipment and find cheaper sources of electricity. For example, specialised equipment using application-specific integrated circuits (ASIC), which are chips designed to perform the single function of block mining, have been developed specifically for the Bitcoin mining market. However, as with any market, absolute cost advantage can lead to concentration of market power. One firm claims a nearly 75% share of the Bitcoin mining equipment market, and via its proprietary mining pools accounted for over 50% of the bitcoin hash rate (work) in September 2018 (Ferreira, Li, and Nikolowa 2019). Furthermore, the number of reachable nodes (both sending to and receiving) on the Bitcoin network fell by approximately 19% in 2018, suggesting industry consolidation is occurring as mining costs and technology options change.[8] On May 1 2019, Bitnodes reported 9476 reachable Bitcoin nodes.[9]

4.4.2 Application platform: Ethereum

Unlike Bitcoin, Ethereum is not designed primarily to support a cryptocurrency-based trading platform. Rather, it has been established as an open platform enabling an entire panoply of applications written using its own programming language, Solidity. These applications include – but are not limited to – other cryptocurrencies and trading platforms. A particular feature, drawing on the Turing-complete features of its programming language, is its ability to implement 'smart contracts' – instructions that will self-execute at a future date when specified preconditions are met. 'Smart contract' code operates on the subsidiary Ethereum Virtual Machine (EVM) runtime environment, which also hosts the platform's fundamental consensus mechanism. It has been proposed that by using 'smart contracts', entire autonomous self-governing entities can be created and operate on Ethereum. This feature is both the platform's strength and its weakness. By supporting a very wide range of applications, its value to users is potentially much greater. Whereas Bitcoin has been likened to a single web-based application, Ethereum has been likened to the internet itself as a platform on which many different applications can operate. However, as an open platform, there are few controls on the quality of the applications hosted on it. Weaknesses in those applications can impinge negatively on Ethereum's reputation and value, as illustrated by the circumstances leading to the fork in 2016. It also allows the platform to be used to host smart contracts executing illegal transactions e.g. Ponzi schemes and investment fraud (Bartoletti et al. 2017).

The Ethereum DL system operates using a modified version of the POW arrangements developed for Bitcoin. Ethereum's digital currency is ether. There is no apparent limit on the number of ether that will be issued, although there are limits on the number issued each year. Ether are issued each time a block is added, with the number per block varying inversely with average transaction volume. The number in circulation has grown steadily over time, and linearly since around the beginning of June 2017 at around 20,000 per day. At 10:16pm May 1 2019 GMT, there were 105,892,613.41, valued at US$16,992,592,487.33 in circulation.[10] From inception in 2015, the number of transactions per day rose to a peak of 13,349,890 on January 4 2018, but fell away to around 550,000 for the latter quarter of 2018. The number in the first 4 months of 2019 has risen steadily, reaching 707,771 on April 30 2019.[11] On May 1 2019, there were 8856 nodes on the Ethereum system[12] – 6,043 in the top ten countries for node hosting,[13] most of them running some version of the Linux operating system.

[6] https://bitinfocharts.com/comparison/bitcoin-transactions.html

[7] https://bitinfocharts.com/comparison/transactionvalue-btc.html

[8] https://www.ccn.com/number-of-reachable-bitcoin-nodes-fell-19-in-2018, reporting on Bitnodes data https://bitnodes.earn.com/dashboard/?days=365

[9] https://bitnodes.earn.com/

[10] https://etherscan.io/chart/ethersupplygrowth

[11] https://etherscan.io/chart/tx

[12] https://www.ethernodes.org/network/1

[13] https://etherscan.io/nodetracker

Because Ethereum processes operations which may generate very different transaction volumes (e.g. a payment generates a single action, but a single smart contract may generate a stream of payments and therefore actions), the payments to node operators is more complex than for Bitcoin. Ethereum gets around this with “gas” – a measure of the amount of computational effort an operation initiated on the Ethereum network generates that decouples the market value of ether from the units measuring computational effort. Each operation costs some discrete amount of gas depending on the complexity of the instructions to be executed and the amount of data that has to be stored in the EVM. Operation initiators specify the gas price they are willing to pay and the maximum amount of gas they are prepared to commit. Gas is ‘purchased’ with ether.

Miners are rewarded via a combination of rewards for winning the block competition and fees associated with the transactions in the block. The fee is determined by a combination of the gas price and gas limit specified by the operation proposer. Miners choose transactions to work on – usually those with the highest gas price. Different operations have different gas prices; miners stop executing an operation when the gas runs out; and the value of any gas ‘left over’ when the transaction is completed is returned in ether to the operation generator. However, the desirability of the operations to miners is influenced by the gas limit (which varies each block and is determined by consensus amongst the miners – consistently averaging around 8 million from December 2017). Miners can add transactions up to be less than or equal to the block gas limit. Their returns are optimised across the gas price paid and the gas limit specified for all operations added to a block. Operations with too low a gas limit will not be selected. Operations with low fees even with sufficient gas specified will also not be appealing to miners as the miners forfeit the potential to include more transactions in the block and thereby gaining more transaction fees. Hence the most desirable transactions are those with high gas prices and a gas limit exactly equal to (or just a very little bit higher than) the actual amount of work required. Recommended gas prices (in Gwei: 0.000000001 or one billionth of an ether) for transactions completed fast (less than 2 minutes), standard (less than 5 minutes) and safe low (less than 30 minutes) at 27:51 on May 10 2019 (GMT) were 0.018, 0.011 and 0.004 Gwei respectively according to https://ethgasstation.info.

Ethereum has been compared favourably to Bitcoin because assurances of transaction processing are obtained more quickly (around 15 seconds between blocks). Also, as with Bitcoin, there is evidence of node consolidation occurring. The number of nodes on Bitcoin was recorded to exceed the number on Ethereum for the first time since at least 2016 on January 9 2019 (10,266 vs 10,078[14]). However, by May 1, the numbers of both had fallen further (Bitcoin 9476 and Ethereum 8856), despite the amount of gas used per day averaging around 39 billion (its all-time high being 41 billion) and apparently increasing, albeit slowly.[15] This is despite some commentators holding the view that Ethereum provides better incentives for node operators. Others however, disagree, with one node operator worrying about “centralization risks over time as running a node is not about just having some free computing resources, but also being willing to dedicate time to operations work (along with knowing how to do that operations work.”[16]

Less controversially, Ethereum appears to have a larger developer team working on it that Bitcoin (measured from software lodgements in the open source repository[17] arguably because of the greater scope afforded to develop new applications on Ethereum (and thereby derive some return for the human capital deployed in software development). However, concerns have been expressed about the incentives for developers of Ethereum-based applications, as the focus on developing open-source software is arguably coming at a cost of attention to understanding exactly what end-users want in their applications.[18]

This is attributed in part to the lack of a clear business model that includes the clients in the development process or building a competitive product attracting users. That is, “similar to government work, there isn’t a strong incentive to build something that actively incites people to use it, you just have to build something that is good enough to pass muster with the people funding you” (ibid). If those funding projects are not actually those who will use the completed application, the potential for misaligned incentives appears considerable. This mimics the early days of the internet, when a huge range of applications were developed, but only a very small proportion went on to become successful commercially and sustainable in the long-run.

4.5 Payment and incentives in “proof of stake” systems

In POS systems, as there is no mining reward for forming the block, all remuneration from the transaction fees associated with the block forged are awarded to the selected forger. They are more cost-efficient, as there are no computation competitions and no need for specialised equipment, so the remuneration for node operators is lower. The lack of need for specialised equipment also guards against the centralisation of node operation activity observed in Bitcoin and Ethereum. However, as

[14] https://www.trustnodes.com/2019/01/09/bitcoin-overtakes-ethereum-in-node-numbers

[15] https://etherscan.io/chart/gasused

[16] https://ethresear.ch/t/incentives-for-running-full-ethereum-nodes/1239

[17] https://medium.com/@ElectricCapital/dev-report-476df4ff1fd2

[18] https://ethresear.ch/t/incentives-for-running-full-ethereum-nodes/1239

there are few system mechanisms other than striving for a high coin price and the risk of losing a stake to align the interests of node operators, the probability of forking is greater, especially when the system is new and the coin price low. There are very few examples of POS applications, although it has been proposed that Ethereum could transition from being a POW system to a POS system in order to reduce transaction costs and increase scalability as growth in the number of transactions is leading to the need for higher levels of computational resource to be applied and exacerbating the degree of centralisation observed. However, it is unclear how even in a POS system the aggregation of a limited number of tokens by concentrated interests can be avoided when the identity of those controlling the nodes is not known.

5 Applying the IAD framework to a DLS community

DLS commons such as the Bitcoin and Ethereum systems can be described using Ostrom's IAD Framework as applied to the knowledge commons in E. Ostrom and Hess (2007) (Figure 2).

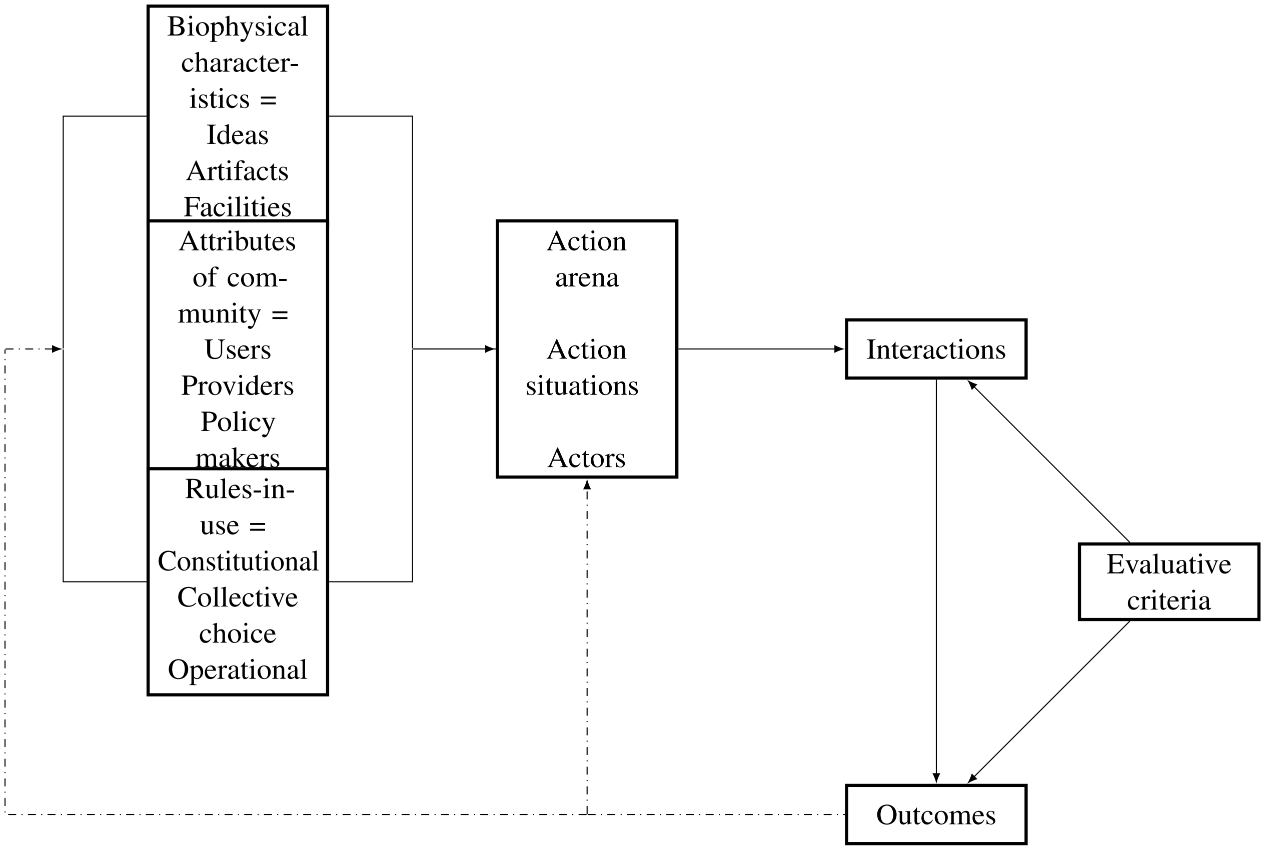
Biophysical characteristics = Ideas Artifacts Facilities
Attributes of community = Users Providers Policy makers
Rules-in-use = Constitutional Collective choice Operational
Action arena
Action situations
Actors
Interactions
Evaluative criteria
Outcomes

Figure 2: The IAD Knowledge Commons (Source: E. Ostrom and Hess 2007)

5.1 Biophysical characteristics

The biophysical characteristics of the DLS commons consist of:

- the system tokens (coins);
- software for operating the DLS;
- applications hosted on the DLS (for smart contract systems such as Ethereum);
- the ledger content (i.e. the blockchain records);
- transactions using system tokens;
- transactions within applications hosted on the DLS;
- information created from analysing the DLS and its community; and

- the electronic (computing) and energy resources required to operate the system.

While the system tokens, electronic and energy resources are rival and excludable, single manifestations of the DL software and applications hosted on a specific DLS can be utilised by multiple users at the same time. As the software content is open-source and the ledger content is widely distributed, both can be appropriated without limit to support new DLS communities (i.e. "forking" or "secession"). The software, ledger content and information created from the DLS and its community are stocks, while the transactions using system tokens and transactions within applications hosted on the DLS constitute *flows*. These flows may correspond with activities in the physical world (e.g. a bitcoin payment for the sale and purchase of a real-world item or the recording of information relating to the fulfilment of a smart contract, such as a particular item being recorded as passing a checkpoint in the real world triggering a transfer of funds within the DLS) or be generated within the DLS (e.g. payment of mining rewards).

5.2 Attributes of the community

DLS user communities consist of:

- external application users;
- coin owners, consisting of
  - coin application users;
  - other application users;
  - DLS node operators (including their co-ordinated activity within node pools); and
  - application operators, including currency exchanges;
- software developers, creating and amending applications utilising the DLS;
- software developers, creating and amending DLS code; and
- DLS custodians.

Using Schlager and Ostrom (1992)'s categorisation of the property rights most relevant to common-pool resources cited in E. Ostrom (2010) (access, withdrawal, management, exclusion and alienation), we observe that coin owners exercise all of these rights in relation to coins that they hold the keys to. However, with regard to the Bitcoin and Ethereum DLSs, none of the community members possesses an alienation right – that is, the right to sell or lease management rights or exclusion rights – because, by very definition, the DLS is owned by no-one. The subsequent analysis focuses on the rights associated with the DLS – noting that the arrangements regarding the coins are embedded within this nexus.

Coin owners are the users of the primary application offered on the Bitcoin blockchain and the ether currency component of Ethereum: the transfer of their property rights to the system coins, recorded in the ledger, to other application users. Coin ownership is a necessary condition to participate in DLS updating activities. Coin owners who participate only for the ability to trade coin ownership (cryptocurrency customers) are termed "coin application users." Bitcoin wallet owners fall into this category. Those users with coin holdings in order to participate in applications hosted on the DLS are termed "other application users." Many Ethereum wallet owners fall into this category. By definition, DLS node operators and application operators are also coin owners, as coin ownership is fundamental to their existence. Coin owners exercise rights of access to the DLS and the withdrawal of the value generated by it, as manifested in their ability to trade coins. In comparison to a forest common-pool resource, they resemble users who can walk in the forest and harvest the trees.

However, some applications hosted on a DLS can serve the interests of individuals who do not have a coin holding. For example, a user of a supply chain management system recording the movement of goods, who provides information for a smart contract managed on the system, may be remunerated using a currency other than the DLS one. Similarly, a researcher or analyst has access to all the information held on the ledger in order to produce reports without needing a coin holding, because all information is in the public domain. These users are termed "external application users." They possess rights to access the DLS, but not to withdraw value directly from it. In comparison to a forest common-pool resource, they resemble users who can walk in a forest and derive value from that activity, but not harvest the trees.

DLS node operators and application operators enjoy rights of access and withdrawal as coin owners. In addition, they hold some management rights, because their activities directly influence the volume of coins in circulation (in POW systems) and their value (in both POS and POW systems). They can choose to exercise their rights in a manner consistent with the sustainable long-term operation of the common-pool resource (e.g. posting only valid transactions, via their choices of allegiance when forks are proposed, taking care with the development of applications hosted on the DLS) or they can choose to act otherwise. However, they do not enjoy a full set of management controls of the DLS in the manner of the managers of forests, as they cannot bring about any changes to the structure or operation of the DLS, in the same way as the managers of a forest can. While the forest managers can make material changes to the forest over time – for example changing the species of tree or expanding the acreage planted – DLS node and application operators have no ability to make changes to the DLS algorithms within the existing setup.

As illustrated with Ethereum, changes to the algorithms or chain can be made only by way of a "fork" – that is, the creation of a new derivative DLS. This resembles the establishment of new communities within tribal arrangements, when a subset of an original group secedes (e.g. following conflict between an old leader and a younger challenger) and sets up its own community under new leadership (i.e. the 'young pretender'). The remainder stays under the old leadership. Either, both or neither may survive, but from the point of the secession, they are separate communities operating under (at least slightly) different rules and leadership. In DL terms, short of "forking", node and application operators can exercise management rights in relation to regulating internal use patterns of the original DLS, but are severely constrained in their ability to make improvements, because its very design prevents changes being instigated within it.

It is noted that application operators exercise full management control over the design of their their own applications – including derivative cryptocurrency applications operating on the DLS host's system e.g. on Ethereum's EVM. In this case, a similar 'nesting' of management rights applies – that is, coin owners associated with a currency or other application managed by an operator hosting the application on the original DLS exert no management rights over either the rival and excludable 'token' created for the application or the DLS itself unless they hold a separate role in that DLS. However, depending upon the arrangements of the application, they may exert management rights over it. However, these should not be confused with management rights over the underlying DLS, even though transactions generated by the application can influence the value of that currency.

This raises the question of who does hold the management rights regarding design and improvement of the DLS. If no change is ever permitted to be made to the DLS design (as advocated by the minority of Ethereum stakeholders at the 2016 fork), then these management rights have been effectively extinguished. Even if it is indicated that a change should be made to the DLS design to ensure it is better managed, it cannot be instigated. This would be the equivalent of a forest being established under a perpetual trust arrangement where no change in any element of its design or operation could ever be contemplated, even if new information or opportunities came to light after it was first established (e.g. a new more productive tree variety being developed which could replace existing plantings).

We contend that such an arrangement would be contemplated only by individuals motivated by ideological reasons. Given the bounded rationality of human beings and their inability to perfectly foresee all future contingencies, most constitutional arrangements (for nation states and corporations, as well as non-profit trusts and other informal entities) make allowances for the allocation of some reserve decision-making rights that can be exercised if unexpected circumstances arise. These may be found in formal documents, or exist informally within the cultures and norms of the groups concerned (for example, even in tribal situations, reserve powers may be granted to a subset of community members to exercise a transfer or sharing of power to prevent its fragmentation in the face of a leadership challenge).

We propose that the management rights concerning the design, implementation and future changes to the DLS structure and software are held by a category of stakeholder called DLS Custodians. DLS Custodians hold both change management rights and exclusion rights, as they determine if not exactly who can have access and withdrawal rights, the processes via which those rights can be obtained and transferred. But who are they? And how are their interests aligned with those of the coin holders who rely upon the efficient design and management of the DLS to maintain and increase coin value? This brings us to consideration of the remaining stakeholders: software developers creating and amending applications utilising the DLS and those creating and amending the DLS code.

Software developers creating and amending applications utilising the DLS and those creating and amending the DLS code exercise considerable influence over in the first instance, the content of the application software operating on the DLS and in the second instance, the content of the DLS software itself. Applications software developers design the software executing the limited management rights of applications operators, but have no special rights in or influence over the design of the DLS itself. In effect, they act as agents of the applications operators when programming the relevant code). By corollary, DLS software developers act as agents of the DLS Custodians when designing, implementing and amending the DLS code. Software developers' incentives as agents may conflict with those of their principals, but can be aligned to some extent if they too are exposed to some of the principal's risks. This can be achieved if the software developers are also coin holders, but it is not axiomatic that they will be. In the absence of any other clear specification, DLS software developers likely exercise the role of de facto DLS Custodians, even though they need not also be coin holders. This resembles some governance arrangements – notably government control of common-pool resources – where disinterested third-party administrators exercise these management rights. In notable exceptions (e.g. the Sovrin identity DLS), separate custodians may be explicitly identified (Howell, Potgieter, and Sadowski 2019).

DLS Custodians control the initial design and development of the DLS. They act as principals to the software developers converting those designs into code. To date, in most cases the original DLS developers are software developers since quite sophisticated knowledge of code and progamming is required to develop the systems in the first instance. This was certainly the case for both Bitcoin and Ethereum, and arguably also for Sovrin. In both Bitcoin and Ethereum, the original custodians determined that the systems would be fully open to any individuals who wanted to participate, via coin ownership in both cases, and via application hosting on Ethereum. They also determined that participation in both could be anonymous.

The originating Bitcoin DLS custodians have chosen to remain anonymous behind the Nakamoto pseudonym. The originating Ethereum DLS Custodian, Vitalik Buterin, revealed himself, and has been a prominent and vocal trustee-advocate for the DLS, albeit posessing no formal power. To the extent that the originating custodians may have intended no further changes would be made to the DLS, no formal arrangements were specified about the exercise of these rights once the system was implemented. However, such an arrangement is optimal only to the extent that no circumstances could ever be imagined where the need to exercise reserve management rights concerning DLS design would ever be needed. Yet no constitutional arrangement can be sustained without attention being given to the exercise of reserve decision-making rights in exceptional circumstances. When Ethereum did face an unexpected crisis in 2016, without explicit arrangements about how to deal with conflicts of views between stakeholders, only a stark choice between "accept the new software" or "secession" was available. There was no provision to broker any alternative arrangements ("amendments"), as allowed under formally-instituted constitutions.

In effect, the Bitcoin and Ethereum arrangements presuming that the DLS could be created and set operating in a fully autonomous manner without any subsequent change or any process facilitating one does not appear to follow sound constitutional practice. In the forest analogy, it is the equivalent of failing to provide any long-term oversight by DLS custodians, and presuming instead that when a crisis arose, the community could be left to organise itself to work out how to deal with the situation. This is hardly an efficient arrangement, even though it may appear "democratic", to the extent that node and application operators (and by extension, affiliated coin holders) got to "vote with their feet" by choosing to go to Ethereum Classic or remain with the reformed Ethereum. In the absence of any formal arrangements, the costs of informing and co-ordinating the decision were extremely high, and arguably left a significant number of community members effectively disenfranchised.

While multiple unofficial forums exist for interested parties to interact and discuss Ethereum software, applications and other matters, these are not formally aligned with Ethereum itself. These forums, participated in mainly by software developers, enable information to be shared but confer no special rights or obligations with regard to DLS governance. As with other voluntary systems, participation is likely strongly contingent upon the amount the participants have "at risk" – those with most invested in the system (as a proportion of their total wealth, including reputational and human capital invested in blockchain software) face the strongest incentives to participate. In the absence of any formal arrangements, in the event of a crisis one would expect a strong bias towards those with the most at stake (in terms of anticipated future returns from the system) being most active in forums seeking to influence the outcome of its resolution. Those with small stakes, even though collectively commanding a majority of stakeholders by number and collective value (e.g. coin holders), are unlikely to participate as the expected return from participation is less than the cost. Thus, without any other arrangements specified, it is quite likely that, in the initial stages at least, DLS Custodianship is almost surely effectively exercised by the DLS software developers implementing it, in just the same manner as any other new venture is effectively governed by its originators. These are aligned with long-term interests of coin holders only so long as the gains come from coin ownership. If the developers are paid by some other means (e.g. salaries) then alignment of interests cannot be assured.

Even though management decision rights may not have been formally assigned, it seems that for Ethereum and Bitcoin at least, software developers have appropriated them in what is in fact a form of a "governance tragedy of the commons" arising because the rights were not clearly specified in the first place. Over time, if coins become widely distributed, software developers' control becomes disproportionate with their stake (in contrast to firms, where shares carry with them voting rights as they transfer from owner to owner, clubs where membership brings voting rights or in nation states where decision-making control is ultimately vested in citizens). It is not, however, axiomatic that all DLS be constituted in the manner of Bitcoin and Ethereum, with no formal assignment of DLS Custodian rights. The Sovrin DLS developed for identity verification applications has an explicit constitution that allocates DLS Custodian rights between a group of stakeholders known as Stewards who also participate as applications operators, and members of a governing trust instrument controlling some superior decision-making rights. This suggests a much more "democratic" distribution of residual management control rights, even though the form adopted has much in common with the corporate or 'establishment' form that some DLS developers claim to eschew with their 'anarchic' instrument (Howell, Potgieter, and Sadowski 2019). This is explored further with 'rules-in-use'.

5.3 Rules-in-use

The rules-in-use operate at three levels – operational, collective choice and constitutional. Operational rules-in-use articulating how individuals interact in the day-to-day operation of the CPR are specified at the outset and hard-coded into the DLS. They consist of the consensus mechanism(s) via which the ledger content is agreed, and give effective effect to the limited management rights conferred on node and application operators in the system design. They also enforce the originally-agreed access, withdrawal and exclusion rights.

A. Berg, Berg, and Novak (2018) contend that these rules-in-use alone comprise a complete constitutional catallaxy. We differ by contending that they constitute only the expression of the lower-level operational rules, and not the governance of the processes under which they are determined in the first place, and subsequently modified – that is, the collective choice (or policy) rules-in-use, via which the operational rules-in-use are made. Neither do they provide constitutional rules-in-use governing the processes of deciding who may participate in making collective choices. The operational rules-in-use utilise artefacts of democratic processes – that is a 'voting process' for reaching consensus about ledger content – but this is separate and distinct from the collective choice and constitutional rules-in-use governing the DLS over its lifetime, from origination to extinction. By analogy to the internet, they confer no ability on community members to participate in material decisions about technical characteristics of the system, in the same way that collective choice rules governing the internet common-pool resource outline how consensus is obtained, for example, on new technical standards or management of the DNS system governing the allocation of internet addresses (Bernbom 2000).

Indeed, the Bitcoin and Ethereum DLSs are conspicuous for the lack of any clearly-articulated collective choice or formal constitutional rules-in-use. To the extent that any rules of this type have emerged over time, they appear to be derived from the cultural norms of the open-source software communities in which the DLS developers and application system developers participate. These communities too lack explicitly-articulated constitutional and collective choice rules-in-use, relying instead on voluntary participation and self-selection of members into sub-communities with similar views on how the community should operate (Dulong de Rosnay and Le Crosnier 2012). They too are characterised by "forking" when potential conflicts emerge between different interests. For example, the community favouring open-source licensing "forked" away from the original open-source community favouring the complete freedom of anonymous modification embodied in the original "copyleft" principles (Schweik 2007).

Consistent with the analysis in the previous section, both the constitutional and collective-choice rules-in-use of Bitcoin and Ethereum have been determined by the DLS developers to reflect their principles of openness (anyone can participate) and anonymity (at the operational level at least). The processes for reaching agreement on the development of and changes to collective-choice rules are similarly governed by principles of voluntary participation in forums by any interested party, regardless of whether they have a coin ownership stake. But while participation is open, the discussions are predominantly technical in nature, so are accessible only by individuals with the requisite software knowledge – that is, DLS and application system developers. Thus, ordinary coin holders are effectively unable to participate – that is, disenfranchised from - the collective choice and constitutional rule-making processes because they lack the specialist human capital. They can chose to patronise or not, but have no effective participation in DLS governance (Howell, Potgieter, and Sadowski 2019).

The informal nature of governance arrangements exhibited by early platforms such as Ethereum, combined with the comparative indelibility of management rules hard-coded into the DLS software poses risks to the potential for their use in commercial applications where it is not clear in the first instance that the prevailing operational rules-in-use will be optimal for all time. Normally, where large capital sums are involved for application development, investors will want to be able to exercise some formal control over constitutional and collective-choice matters, even if they do not have a formal ownership stake. Howell, Potgieter, and Sadowski (2019) suggest that more-inclusive formal arrangements – such as observed in clubs – specifying how the rights to participate in policy and constitutional decision-making are both created in the first instance, and are transferred over time – may be required if such developments are to proceed. This will likely lead to some compromises to the principles of openness and anonymity governing DLSs such as Bitcoin and Ethereum. They illustrate how Sovrin uses both a formal trust arrangement specifying constitutional rules-in-use and the roles of a special member class – Stewards – in the development and operation of the policy and operational rules-in-use to engender greater confidence in the operation of the DLS.

5.4 Action arena

These charateristics have largely been detailed in the preceding discussion, but are summarised here for completeness.

Actors

The actors in a DLS consist of:

- external application users;
- coin owners, consisting of
  - coin application users;
  - other application users;
  - DLS node operators (including their co-ordinated activity within node pools); and
  - application operators, including currency exchanges;
- software developers, creating and amending applications utilising the DLS;
- software developers, creating and amending DLS code; and
- DLS custodians.

Action situations

The action situations include:

- operational actions:
  - coin issuing
  - transaction generation
  - transaction processing
- governance actions:
  - policy change processes (agreeing changes to the DLS software design)
  - constitutional processes (agreeing the boundaries of the DLS; facilitating forking processes if indicated),

5.5 Patterns of interactions

Stakeholders interact as per their rules outlined above, in either the day-to-day operation of the DLS or in its governance.

5.6 Evaluative criteria

A wide range of information is available to assess DLS performance, because the ledger is publicly available. Just as with financial markets, specialist entities have created applications reporting regularly on blockchain performance. These too tend to be in the public domain. For example, Etherscan^[https://etherscan.io/charts[] provides information on Ethereum performance and and Bitcoin.com does likewise for Bitcoin.[19]

The primary evaluative criteria relate to coin number (in a POW system) and value (both POW and POS systems). If the DLS is operating in a manner aligned with its members' interests, then coin value will be at least stable, and at best, increasing as a reflection of the value of the platform to its users. Specific applications will have other measures (e.g. their own subsidiary coin or application-specific measures). Other relevant information is transaction price. Again, stability in transaction price suggests a stable system.

[19] https://charts.bitcoin.com/bch/

The number of active nodes provides some indication of platform strength, but like concentration levels in a market must be interpreted with caution. A growing node number can indicate a new platform or one where anticipated coin value increases or transaction price changes are stimulating new entry. A falling node number may indicate either a loss of confidence in the coin or the consolidation of node operation in a more mature platform as a consequence of scale economies in operation being leveraged. Information about node location (both geographically and by associated internet addresses) provides information about the concentration of node operation, as does information about the distribution of blocks mined across mining pools.

Other system statistics also yield relevant performance information. These include the number of blocks added, block size, the amount of work (hash rate) required to complete a block (in POW systems), the average time between block lodgements.

5.7 Outcomes

The desired outcome for a DLS is measured by its survivability. Successful systems will be able to navigate through requisite changes whilst retaining its original identity. Unsuccessful ones will either cease to operate through lack of support for their business or be superseded by more successful forks.

6 Implications

The IAD framework has been used to identify the governance arrangements of selected DLSs. In terms of Table 1, there are grounds to consider them as common-pool resources. But do the governance arrangements of the examples considered above – Bitcoin and Ethereum, conform to Blomquist and Ostrom (1985)'s governing principles for long-enduring resources, as reinterpreted in E. Ostrom (1990) and applied in E. Ostrom and Hess (2007)'s eight design principles for long-enduring resources?

1. Clearly defined boundaries (effective exclusion of external un-entitled parties)

Dulong de Rosnay and Le Crosnier (2012) notes that for the internet, and E. Ostrom and Hess (2007) for knowledge resources in general, it is the absence of limits that matters for qualification of information-related CPRs. Enclosure, in the form of walled gardens or other limits, parallels the privatisation of the resource as per Hardin (1968). As anyone can participate in the Bitcoin and Ethereum DLSs, they are apparently boundaryless. Anyone capable of obtaining the relevant operating software and a copy of the ledger can participate without having to gain permission. This equates to the open access available for individuals wishing to obtain any information held in the public domain on the internet. Likewise, the software code is also freely available for modification. Yet, implicit limits are imposed by the fact that only those capable of accessing the internet are in a position to access the resources in the first place. So in fact, a boundary – albeit a very extensive one – exists, at least in respect of those capable of accessing the resource (the least-rigorous property right).

A tighter boundary could be said to exist in regard to those able to exercise higher-level property rights (non-passive engagement) as as coin ownership is required to participate as either a miner or a transactor (for Ethereum, application operators generally need at least "gas"). However, no-one can be stopped from mining, in general so the the boundaries are not clearly defined. Also, no-one can be ejected from the network exactly because everyone in fact anonymous.

Yet effective boundaries are placed around participation in the governance activities of the systems, as only those with the necessary software development knowledge are able to actively participate; however no boundaries exist provided these criteria are met, so long as the individual has an interest. This is both a strength and a weakness: individuals without a coin stake can participate, and it is unclear that their interests will be aligned with those of other system participants. This enables open appropriation of the software and the ability to set up a rival network, which can harm the original system. At the very least, it suggests an inherent instability. It is noted Sovrin avoids this risk by allowing only Stewards to participate at this level.

2. Rules regarding the appropriation and provision of common resources that are adapted to local conditions

This principle looks for congruence between appropriation and provision rules and local conditions. In the case of digital resources which are global in nature, local conditions can be understood as the structure and the features of the resource, its evolution in time, the culture and the custom of the community. To the extent that the Bitcoin and Ethereum DLSs have emerged from within the open-source software development community, by their very design they reflect the originators' adherence to openness to information and software, anonymity and and a degree of anti-establishment ethos. However, as with the internet, which began as an open platform developed for scientific endeavours, but required some modification to include some commercial values in its provision as the resources of commercial telecommunications operators were required to expand its scope and networks and avoid congestion (Bernbom 2000), it is likely that changes will be observed in DLSs as they embrace more commercial applications.

This has been achieved for the internet without any major compromises to the originating principles of access to all those meeting the entry criteria of being able to afford a connection in the first place. However, it has led to some restrictions in participation at the infrastructure layer. Arguably, this parallels the restrictions in software developer participation loosely-controlled in Ethereum and Bitcooin, bit more explicitly controlled in Sovrin. Furthermore, there is nothing to stop anyone from setting up a DLS which has rules that are adapted to specific local conditions. The rules for such a DLS can reflect any required local conditions, in principle, although the authors are not aware of many geographically local applications.

3. Collective-choice arrangements that allow most resource appropriators to participate in the decision-making process

Arguably, it is the inability of Bitcoin and Ethereum to satisfy this principle that potentially stands as a threat to their long-term survival and effectiveness. Neither platform provides an effective means for most resource appropriators (coin holders) to participate in system decision-making. As most rules are determined before these participants even join the system, and they are effectively excluded from any subsequent system-related decision-making by their inability to understand the technical elements attending changes to the software elements. While the same can be said for the internet, the nexus of commercial agreements between internet access seekers and the ISP connecting them to it provides at least a chain of representation into the the higher-level governance arrangements (participation in internet exchanges and international standards forums etc.). The anonymity of those stakeholders participating at the higher levels of the system and the lack of any meaningful commercial agreement leaves ordinary coin-holders relying on the goodwill of individuals they cannot identify to ensure their interests are considered. This differs substantially from other CPRs where the identify of higher-level agents is known and other links between them (e.g. access payments) can be utilised to align interests.

The lack of clear links of this type in the Bitcoin and Sovrin DLSs is what leads us to conclude that under current arrangements, coin-holders are effectively disenfranchised. While this may matter less for Bitcoin, as it is a very simple, one-application ledger, it may prove to be more significant for Ethereum, where there are many different applications, and hence stakeholders with very different interests but still no clear way of aligning their interests back into the broader governance arrangements. Again, by contrast, Sovrin addresses this by placing clear observability and accountability rules in place. Identity system participants must have a commercial arrangement with one of the Stewards in the first instance to actively use the platform. Even though there may be no explicit links provided in the platform itself, these wider engagements in the Sovrin ecosystem afford all users a direct link to key system decision-making.

4. Effective monitoring by monitors who are part of or accountable to the appropriators

Again, the separation of coin value appropriators from the design and operation of the system-controlled monitoring systems leaves this element of DLS operation falling short of effective CPR governance arrangements for Bitcoin and Ethereum, although they are addressed to a much greater extent by Sovrin.

5. A scale of graduated sanctions for resource appropriators who violate community rules

There are no clear sanctions for rule violation, short of system forking within Bitcoin and Ethereum, despite the fact that the operations of some stakeholders can have profound effects on coin value (as evidenced by the 2016 Ethereum fork). Once again, Sovrin provides a clearer set of arrangements, in that the Stewards are identified and as they have other commercial relationships with coin holders, increasingly stronger sanctions are likely to be executed if repeated transgressions occur.

6. Mechanisms of conflict resolution that are cheap and of easy access

By very design, DLSs are immutable, so it is likely to be very costly to compensate coin holders for any errors attributable to Bitcoin or Ethereum system design. As there are a large number of stakeholders and their identities are unknown, conflict resolution (if it is indeed possible) will be extremely costly to bring about. While the Sovrin system is as immutable as any DLS, it does ensure lower-cost dispute resolution processes because the relevant parties are identifiable. It also offers further assurances by way of the processes outlined in its constitution immutability.

7. Self-determination of the community recognized by higher-level authorities

As both Bitcoin and Ethereum are global entities without any specific formal agreements regarding participation, it is not clear what higher authorities they could be held to account by (other than the recognition of any licensing arrangements the open-source software code is made available under – see Dulong de Rosnay and Le Crosnier (2012)). Indeed, these DLSs had their origins in their ability to provide alternative means to those governed by states for specific applications. Again, by contrast, the arrangements under which Sovrin is established are recorded in its incorporation documentation lodged in the State of Utah (Howell, Potgieter, and Sadowski 2019).

An interesting issue pertains to the recognition and enforceability of the terms of "smart contracts" lodged on Ethereum and similar DLSs. Patel et al. (2018) acknowledges that they "cannot be enforced in some of the existing judicial frameworks." They propose an arrangement whereby smart contracts can be made legally enforceable in some national legal frameworks by incorporating crypto primitives like digital signatures. However, much depends on exactly what the nature of the instructions executed by the code are. This matter will be explored in another paper.

8. In the case of larger common-pool resources, organization in the form of multiple layers of nested enterprises, with small local CPRs at the base level

As inherently distributed structures, DLSs do not exhibit a nesting structure (except by way of the communication tools via which the nodes peer with each other – which is an operational and not a governance arrangement). As Ostrom's formulation of these principles predated much recent thinking on the governance of 'flat' institutional structures, it is not clear that this is necessarily a requirement for efficient CPR governance. Notwithstanding, it is not inconceivable that some DLS arrangements may be predicated upon hierarchical interactions – for example, where groups of coin-holders may be required to channel all their transactions only through specific applications. However, this is not the case for either Bitcoin or Ethereum.

7 Conclusion

We have reviewed the Ostrom IAD framework and have looked in detail at the various stakeholders in distributed ledger systems (DLSs), specifically Bitcoin and Ethereum. After careful consideration of the rights exercised by different members of the community of DLS users, we propose that the management rights concerning the design, implementation and future changes to the DLS structure and software are held by a category of stakeholder called DLS Custodians. DLS Custodians hold both change management rights and exclusion rights, as they determine if not exactly who can have access and withdrawal rights, the processes via which those rights can be obtained and transferred.

The originating Bitcoin DLS custodians have chosen to remain anonymous behind the Nakamoto pseudonym. The originating Ethereum DLS Custodian, Vitalik Buterin, has however revealed himself. Bitcoin and Ethereum arrangements presuming that the DLS could be created and set operating in a fully autonomous manner without any subsequent change or any process facilitating one does not appear to have followed sound constitutional practice and this is borne out by the way in which several forks have played out. In our view, DLS Custodianship is currently almost surely effectively exercised by the DLS software developers implementing it, in just the same manner as any other new venture is effectively governed by its originators. Their interests might not necessarily be fully aligned with those of the coin holders.

The IAD framework has been used to identify the governance arrangements of selected DLSs. There are grounds to consider them as common-pool resources and we considered how the governance arrangements of Bitcoin and Ethereum conform to Blomquist and Ostrom (1985)'s governing principles for long-enduring resources as applied in E. Ostrom and Hess (2007)'s eight design principles for long-enduring resources. Most strikingly, clear collective-choice arrangements are not in place for the two DLS systems but this leaves the way open to argue that a minimalistic, vague and somewhat mathematical arrangement is perhaps superior to conventional arrangements. On the whole DLS adherence to the eight principles is poor for the case of Bitcoin and Ethereum and much less so for POS blockchain Sovrin. Nevertheless, once cannot exclude the possibility that stable arrangements outside the IAD framework are possible.

References

Andoni, Merlinda, Valentin Robu, David Flynn, Simone Abram, Dale Geach, David P. Jenkins, Peter McCallum, and Andrew Peacock. 2019. "Blockchain Technology in the Energy Sector: A Systematic Review of Challenges and Opportunities." *Renewable and Sustainable Energy Reviews* 100 (February). Elsevier Limited: 143–74. doi:10.1016/j.rser.2018.10.014.

Anta, Antonio Fernández, Kishori Konwar, Chryssis Georgiou, and Nicolas Nicolaou. 2018. "Formalizing and Implementing Distributed Ledger Objects." *ACM SIGACT News* 49 (2). Association for Computing Machinery (ACM): 58–76. doi:10.1145/3232679.3232691.

Arrow, Kenneth. 1962. "The Economic Implications of Learning by Doing." *Review of Economic Studies* 29 (3): 155–73. https://EconPapers.repec.org/RePEc:oup:restud:v:29:y:1962:i:3:p:155-173.

Bartoletti, Massimo, Salvatore Carta, Tiziana Cimoli, and Roberto Saia. 2017. "Dissecting Ponzi Schemes on Ethereum: Identification, Analysis, and Impact." In.

Berg, Alastair, Chris Berg, and Mikayla Novak. 2018. "Blockchains and Constitutional Catallaxy." *Available at SSRN 3295477*.

Bernbom, Gerald. 2000. "Analyzing the Internet as a Common Pool Resource: The Problem of Network Congestion," January.

Blomquist, William, and Elinor Ostrom. 1985. "Institutional Capacity and the Resolution of a Commons Dilemma." *Review of Policy Research* 5 (2): 383–94. doi:10.1111/j.1541-1338.1985.tb00364.x.

Buchanan, James M. 1965. "An Economic Theory of Clubs." *Economica* 32 (125). JSTOR: 1–14.

Buchanan, James M. 1967. "Cooperation and Conflict in Public-Goods Interaction." *Economic Inquiry* 5 (2): 109–21. doi:10.1111/j.1465-7295.1967.tb01944.x.

Crosby, Michael, Nachiappan, Pradhan Pattanayak, Sanjeev Verma, and Vignesh Kalyanaraman. 2015. "Blockchain Technology: Beyond Bitcoin." *Sutardja Center for Entrepreneurship & Technology*. http://scet.berkeley.edu/wp-content/uploads/BlockchainPaper.pdf.

Czepluch, Jacob Stenum, Nikolaj Zangenberg Lollike, and Simon Oliver Malone. 2015. "The Use of Block Chain Technology in Different Application Domains." *The IT University of Copenhagen, Copenhagen*.

Davidson, Sinclair, Primavera De Filippi, and Jason Potts. 2018. "Blockchains and the Economic Institutions of Capitalism." *Journal of Institutional Economics* 14 (4). Cambridge University Press (CUP): 639–58. doi:10.1017/s1744137417000200.

Dulong de Rosnay, Melanie, and Hervé Le Crosnier. 2012. "An Introduction to the Digital Commons: From Common-Pool Resources to Community Governance," September.

Ellickson, Robert C. 1993. "Property in Land." *Yale Law Journal* 102: 1315–44.

Ferreira, Daniel, Jin Li, and Radoslawa Nikolowa. 2019. "Corporate Capture of Blockchain Governance." Working Papers 880. Queen Mary University of London, School of Economics; Finance. https://ideas.repec.org/p/qmw/qmwecw/880.html.

Gans, Joshua S. 2019. "The Fine Print in Smart Contracts." Working Paper 25443. Working Paper Series. National Bureau of Economic Research. doi:10.3386/w25443.

Hansmann, Henry. 1996. *The Ownership of Enterprise / Henry Hansmann*. Book. The Belknap Press of Harvard University Press Cambridge, Mass.

Hardin, Garrett. 1968. "The Tragedy of the Commons." *Science* 162 (3859). American Association for the Advancement of Science: 1243–8. doi:10.1126/science.162.3859.1243.

Howell, Bronwyn E., Petrus H. Potgieter, and Bert M. Sadowski. 2019. "Governance of Blockchain and Distributed Ledger Technology Projects." *SSRN Electronic Journal*. doi:10.2139/ssrn.3365519.

Lueck, Dean L. 1995. "Property Rights and the Economic Logic of Wildlife Institutions." *Natural Resources Journal* 35 (3): 625–70.

Mulligan, CJ, Z Scott, S Warren, and JP Rangaswami. 2018. "Blockchain the Hype." In *World Economic Forum. Http://Www3.*

Weforum. Org/Docs/48423_Whether_Blockchain_WP. Pdf. Accessed. Vol. 2.

Nakamoto, Satoshi. 2008. “Bitcoin: A Peer-to-Peer Electronic Cash System,” Http://Bitcoin.org/Bitcoin.pdf.”

Narayanan, Arvind, Joseph Bonneau, Edward Felten, Andrew Miller, and Steven Goldfeder. 2016. *Bitcoin and Cryptocurrency Technologies: A Comprehensive Introduction*. Princeton University Press.

Netting, Robert M. 1976. “What Alpine Peasants Have in Common: Observations on Communal Tenure in a Swiss Village.” *Human Ecology* 4 (2): 135–46. doi:10.1007/BF01531217.

Netting, Robert McC. 1981. *Balancing on an Alp : Ecological Change and Continuity in a Swiss Mountain Community / Robert Mcc. Netting*. Cambridge University Press Cambridge [Eng.] ; New York.

Nugent, Jeffrey B, and Nicolas Sanchez. 1993. “Tribes, Chiefs, and Transhumance: A Comparative Institutional Analysis.” *Economic Development and Cultural Change* 42 (1): 87–113. https://ideas.repec.org/a/ucp/ecdecc/v42y1993i1p87-113.html.

Olson, Mancur. 1965. *The Logic of Collective Action*. Cambridge, Mass.: Harvard University Press.

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

———. 2000. “Collective Action and the Evolution of Social Norms.” *Journal of Economic Perspectives* 14 (3): 137–58. doi:10.1257/jep.14.3.137.

———. 2009. “A General Framework for Analyzing Sustainability of Social-Ecological Systems.” *Science* 325 (5939). American Association for the Advancement of Science: 419–22. doi:10.1126/science.1172133.

———. 2010. “Beyond Markets and States: Polycentric Governance of Complex Economic Systems.” *American Economic Review* 100 (3): 641–72.

Ostrom, Elinor, and Charlotte Hess. 2007. “A Framework for Analyzing the Knowledge Commons.” In *Property Law and Economics*, edited by Charlotte Hess and Elinor Ostrom. Chapters. MIT Press. https://www.jstor.org/stable/j.ctt5hhdf6.

Ostrom, Elinor, and Vincent Ostrom. 1977. “Public Economy Organization and Service Delivery.” In *Financing the Regional City, Project Meeting of the Metropolitan Fund*. Vol. 1.

Ostrom, Elinor, 1947- Gardner Roy, and 1950- Walker James. 1994. *Rules, Games, and Common-Pool Resources*. Book; Book/Illustrated. Ann Arbor : University of Michigan Press.

Ostrom, and Charlotte Hess. 2010. “Private and Common Property Rights.” In *Property Law and Economics*. Edward Elgar Publishing. https://EconPapers.repec.org/RePEc:elg:eechap:12900_4.

Patel, Dhiren, Keivan Shah, Sanket Shanbhag, and Vasu Mistry. 2018. “Towards Legally Enforceable Smart Contracts.” In *Blockchain – Icbc 2018*, edited by Shiping Chen, Harry Wang, and Liang-Jie Zhang, 153–65. Cham: Springer International Publishing.

Pilkington, Marc. 2016. “Blockchain Technology: Principles and Applications.” In *Research Handbook on Digital Transformations*, edited by F. Xavier Olleros and Majlinda Zhegu, 225–53. London: Edward Elgar. doi:10.4337/9781784717766.

Raymond, Mark. 2012. “Puncturing the Myth of the Internet as a Commons.” In.

Samuelson, Paul A. 1954. “The Pure Theory of Public Expenditure.” *The Review of Economics and Statistics* 36 (4). The MIT Press: 387–89. http://www.jstor.org/stable/1925895.

Schlager, Edella, and Elinor Ostrom. 1992. “Property-Rights Regimes and Natural Resources: A Conceptual Analysis.” *Land Economics* 68 (3). [Board of Regents of the University of Wisconsin System, University of Wisconsin Press]: 249–62. http://www.jstor.org/stable/3146375.

Schweik, Charles M. 2007. “Free/Open-Source Software as a Framework for Establishing Commons in Science.” In *Understanding Knowledge as a Commons: From Theory to Practice*, 277–309. Cambridge, MA: MIT Press; MIT Press.

Sengupta, Nirmal. 1991. *Managing Common Property : Irrigation in India and the Philippines / Nirmal Sengupta*. Book.

Sage Publications New Delhi ; Newbury Park, Calif. http://www.loc.gov/catdir/enhancements/fy0657/90049970-t.html.

Swan, Melanie. 2015. *Blockchain: Blueprint for a New Economy*. O'Reilly Media, Inc.

Szabo, Nick. 1997. "Formalizing and Securing Relationships on Public Networks." *First Monday* 2 (9).

Varoufakis, Y., and J. Moe. 2018. *Talking to My Daughter About the Economy: Or, How Capitalism Works–and How It Fails*. Farrar, Straus; Giroux. https://books.google.co.nz/books?id=hi86DwAAQBAJ.