J. Account. Public Policy 40 (2021) 106881

Contents lists available at ScienceDirect

J. Account. Public Policy

journal homepage: www.elsevier.com/locate/jaccpubpol

Blockchain and other distributed ledger technologies: Where is the accounting?

Miles Gietzmann, Francesco Grossetti *

Department of Accounting and Bocconi Institute for Data Science and Analytics (BIDSA), Bocconi University, via Röntgen 1, 20136 Milano, Italy

ARTICLE INFO

Article history:
Available online 8 August 2021

Keywords:
Blockchain
Distributed ledger
Regulatory compliance
Smart contracts
Asset provenance

ABSTRACT

In a recent survey of academic research, Fintech related topics, broadly classified as crypto-currency studies, were by far the most researched topics in the social sciences. However, we have observed that, perhaps surprisingly, even though crypto-currencies rely on a distributed accounting ledger technology, relatively few of those studies were conducted by accounting academics. While some of the features of a system like Bitcoin do not necessarily rely on a traditional accounting knowledge, this knowledge is key in designing effective real-world distributed systems. Building on a foundational framework developed by Risius and Spohrer (2017), we provide support for their hypothesis that to date, research in this area has been predominantly of a somewhat narrow focus (i.e., based upon exploiting existing programming solutions without adequately considering the fundamental needs of users). This is particularly reflected by the abundance of Bitcoin-like crypto-currency code-bases with little or no place for business applications. We suggest that this may severely limit an appreciation of the relevance and applicability of decentralized systems, and how they may support value creation and improved governance. We provide supporting arguments for this statement by considering four applied classes of problems where a blockchain/distributed ledger can add value without requiring a crypto-currency to be an integral part of the functioning system. We note that each class of problem has been viewed previously as part of accounting issues within the legacy centralized ledger systems paradigm. We show how accounting knowledge is still relevant in the shift from centralized to decentralized ledger systems. We advance the debate on the development of (crypto-currency free) value-creating distributed ledger systems by showing that applying accounting knowledge in this area has potentially a much wider impact than that currently being applied in areas limited to auditing and operations management. We develop a typology for general distributed ledger design which assists potential users to understand the wide range of choices when developing such systems.

1. Introduction

In a recent research article in Science, Falco et al. (2019) argued that one possibility why research into cyber-risk was challenging was because the field requires an interdisciplinary perspective in order to address the broad set of issues that

* Corresponding author.
E-mail addresses: miles.gietzmann@unibocconi.it (M. Gietzmann), francesco.grossetti@unibocconi.it (F. Grossetti).

https://doi.org/10.1016/j.jaccpubpol.2021.106881
0278-4254/

arise with the new technological possibilities. Included in the team calling for an inter disciplinary approach is a Professor of Accounting.

In a survey on the research topics of papers submitted to ssrn.com (SSRN), the managing director concludes that "Fintech related pre-prints and early stage research papers received over 550,000 paper downloads in the past two years from the platform. This far outstrips other emerging research topics on the platform like big data (under 160,000 papers downloaded) and the controversial subject of *fake news* (under 50,000 papers downloaded) in the same time-period" (SSRN, 2019). When the mix of authors is investigated, it is difficult to conclude that, on the basis of numerical representation, accounting academics are active participants in leading interdisciplinary teams studying these issues. This paper addresses this puzzle of non-representation in a number of ways.

It is suggested that since Bitcoin is the most well known application of blockchain, the relatively sparse use of specific accounting knowledge required to run such a system has perhaps lead some accountants to infer that distributed ledgers are outside the main realm of accounting. We argue that, while it is correct to conclude that Bitcoin relies more heavily upon cryptography and computer science design issues than on accounting ones, for other types of Distributed Ledger Technologies (DLTs) the irrelevance of accounting knowledge may no longer hold. Another concern arises because of the increasing use of (centrally controlled) distributed computing; one may question what a discussion of "new" distributed ledger systems like Bitcoin adds to the debate on distributed computing.

In summary, we argue that distributed ledger systems are significantly different from traditional centrally controlled distributed computing systems Meunier (2016), and that the blockchain-powered Bitcoin architecture illustrates a special type of distributed ledger which may not be appropriate to use in all settings. To address these two issues, we have developed a taxonomy for distributed ledger design. Our motivation is that, equipped with this new taxonomy, users will start to ask questions such as "what consensus system (to confirm transactions) do we envisage?" Or "should permission to write to the ledger be restricted (permissioned)?" This approach contrasts the assumption that support for a crypto-currency is required. We show that, indeed, in some applications of distributed ledger systems a crypto-currency may not be necessary. Instead, we argue that the contribution of accounting knowledge is a critical *development* requirement. We structure our argument by identifying four pivotal accounting-based constructs that need to be considered when choosing which form of DLT to adopt: ledger access transparency rules, regulatory compliance requirements, contracting variable design and provenance traceability.

In our discussion below, we begin by reviewing what a blockchain is and how it differs from (centrally controlled) distributed computing. We then define, in a more general sense, what a distributed ledger is. We adopt this introductory ordering because blockchains used to support the delivery of a crypto-currency are the most well known (and historically widely used) type of DLT. We then provide a definition of a DLT and show how a blockchain is a special case. From Section 3 we widen our discussion to other forms of distributed ledger thus generally referring to DLTs rather than limiting our scope to the special case of a blockchain system.

This paper is organized as follows: Section 2 characterizes blockchain technology and shows how a blockchain is a subclass of the set of all possible DLT's. Next, we discuss the critique of Risius and Spohrer (2017) that to date, outside the narrow area of creating a crypto-currency, the literature has provided few papers demonstrating the business case for implementing a blockchain. We address this concern by recommending that for a wide class of business settings, organizations should start by considering the full class of possible DLTs rather than restricting attention to the special case of a blockchain. We characterize four generic representative choices that organizations would need to make when deciding on how to design an applicable value creating DLTs. We do not claim that the four choice variables provide an exhaustive list but suggest that they are amongst the most important and can provide a useful starting point.[1] These four representative choices are discussed in the following sections: Section 3 discusses distributed ledger access and transparency choices, Section 4 discusses regulatory compliance issues, Section 5 discusses smart contracting designs and issues, Section 6 discusses issues related to traceability in a distributed system and in Section 7 we conclude the paper.

2. Fundamentals on blockchain and DLT

The first appearance of a blockchain dates back to 2008 when Satoshi Nakamoto introduced the original Bitcoin protocol (Nakamoto et al. 2008). There are four fundamental pillars which are the constituents of a blockchain system: the distributed nature of the network of users, the permissionless access, the consensus mechanism, and the cryptographic algorithms which form the security layer of the system.

The key construct behind a blockchain stems from the theory of distributed ledgers (Froystad and Holm, 2016). A blockchain is a distributed system which maintains a continuously growing list of ordered data records (i.e., transactions) organized in blocks with their integrity being verified through the use of digital signatures and cryptographic algorithms. In this context, integrity is a specific aspect of a software system with three major components: data integrity, behavioral integrity, and security (Boritz, 2005). In other words, the system has to ensure that the data used and maintained by it are complete, correct and that there are no contradictions, that it behaves without any logical error, and that it dispatches data access and

[1] We acknowledge that in the management literature, other research topics have also been explored. For instance, see the following papers: Mendling et al. (2018), Treiblmaier (2018), Hughes et al. (2019).

permissions to authorized users only. Here the concept of system integrity plays a major role since the main goal of a blockchain is to achieve and maintain its integrity over time. There are number of ways in which a blockchain is able to do this. In particular, such systems make extensive use of strong cryptographic algorithms which make the alteration of a previously recorded data impossible.[2] The key innovation of blockchain was to introduce a methodology for ensuring the accurate recording and tracking of transactions in a zero trust public system (Goldreich and Oren, 1994). The methodology uses cryptography and specific game-theory computational methods to record, verify and synchronize data across the system. The set of computational methods used to ensure data consistency and integrity goes under the name of *distributed consensus algorithms*. We refer the reader to the review papers by Yli-Huumo et al. (2016) and Liu et al. (2019) about blockchain technology and game-theory based applications. We now develop our definition of a blockchain in terms of the above four pillars.

Definition 1. A blockchain is an open and permissionless distributed peer-to-peer system of ledgers that utilizes a software algorithm, called distributed consensus protocol, to validate and permanently (immutability) store every transaction in a timely ordered chain of blocks. Each block contains a set of valid transactions connected together by cryptographic algorithms such that the system maintains its integrity over time.

Meunier (2016) clarifies this point by stressing that the difference between traditional (centrally managed) decentralized systems, such as distributed SQL, is that those decentralized systems are controlled by one central authority (the enterprise), whereas blockchains are built on the initial presumption that there is no central control of any sort. For illustration, we reproduce an adapted version of the schematic as in Meunier (2016) in the Appendix. That discussion reinforces the point that a key differentiating feature of blockchain is that the control over the distributed ledger is not with a single entity but rather with a subset, if not all, the nodes in the network. This is a fundamental difference between a blockchain and other technological infrastructures such as cloud computing or data replication which are commonly used in existing shared ledgers (Lamehamedi et al., 2002; Nordin et al., 2006; Milani and Navimipour, 2016). This absence of centrality ensures that a single entity cannot approve the addition of a given transaction to the system. Since trust between nodes is unknown, system's integrity is achieved through a set of rules that every peer has to comply with in order to identify, authenticate and then authorize a given data record for it to be appended. Hence, in a blockchain, the appending of new data and its validation process is a joint effort and it is strictly regulated by a consensus mechanism specified by design. A consensus mechanism is necessary to establish whether a particular transaction is legitimate or not through a set of predefined cryptographic validation methods. There are several consensus mechanisms available. For example, the Bitcoin protocol uses the so-called *Proof-of-Work* (PoW) and the data validation process is called mining (Nakamoto et al., 2008). As we shall see later in Section 6 about provenance, the decentralized nature of governance on a blockchain is a desirable feature when users do not necessarily trust a government or an enterprise to centrally record all features of a transaction with full accuracy. That is, the blockchain was initially developed when trust between parties and their recording of transactions was an issue. Cryptography is used to establish immutability. In a blockchain, on each new transaction record a cryptographic hash function is applied to the original data. In summary, the intuition is that a hash function is able to take any input data and compute a digital fingerprint that cannot be changed unless the input data itself is changed (Menezes et al., 1996).[3] Moreover, the output of a hash function cannot be reverted to its original input.[4]

The new paradigm introduced by the blockchain technology certainly comes with both advantages and disadvantages. For instance, among the advantages we highlight disintermediation, a high level of integrity, immutability and transparency (although with caveats), faster transactions with lower associated costs, and a much improved traceability. In fact, a blockchain is a distributed system which implies that intermediaries are not a requirement.[5] In addition, the distributed nature of the system makes it more resilient with respect to more standard infrastructures since there is not a single point of potential failure (Dooley, 2001; Lynch, 2009). Compared to any other network system, a blockchain offers the highest level of integrity ensuring that all the data stored in the chain will always be consistent and robust to tampering. Moreover, the immutability feature plays another important role since it preserves the history of every transaction record in the system which is an additional guarantee of high integrity. Also, in its original design (e.g., supporting Bitcoin), a blockchain is fully transparent. If anyone in the system tries to alter any record, all the other users would notice.[6] The system also offers faster transactions compared to the ones managed by traditional centralized systems such as banks. This becomes more evident when we think about transactions involving money transfers overseas. This process can take several days to align the whole system whereas in a blockchain every transaction occurs in a matter of seconds. In addition to increased speed, a blockchain performs transactions at a lower cost because there are no multiple intermediary steps with associated fees along the way. Transactions do not come for free of

[2] A cryptographic algorithm is a computer protocol whose objective is to protect certain data or to make them inaccessible to anybody except authorized users only. A well known algorithm which introduced modern cryptographic paradigms is contained in the work by Diffie and Hellman (1976a,b).

[3] One of the most common hash functions is SHA-256 which is used by the Bitcoin protocol and was created by the National Security Agency (NSA) in 2001 (Dworkin, 2015).

[4] We use the word *cannot* whereas the correct statement would be computationally unfeasible. In other words, the computational problem of reverting a hash function is so complex that it would take an infinite amount of time to achieve the goal.

[5] This has further consequences for the whole system's trust. By removing all the intermediaries, the system reduces costs and concerns about potential malicious actors going undetected. This still has lots of implications for both research and applications since it is a specific characteristic of a blockchain framework (Hawlitschek et al., 2018; Werbach, 2018; Casey and Vigna, 2018).

[6] There are different blockchain designs with specific properties including transparency. We discuss some of these designs in sub-Section 2.1, but even in the most private design, there is shared ledger information that anyone can see at any given time.

course, but because the system has no intermediaries, the final cost is much less than in traditional means. As a last advantage, blockchain has an improved traceability. Each transaction past and present can be traced in a transparent way so that every node involved can be quickly and easily identified.

As anticipated, blockchain does also come with some disadvantages. Among others, we highlight integration and privacy concerns, large energy consumption, and uncertain regulations. If we think about large enterprises with complex and well-consolidated legacy systems, the integration process is not fully functional yet. For example, many of the existing blockchains simply do not have the capability to work with some legacy systems. This has huge implications for the enterprises which see themselves in the position of having to choose between the legacy system and the new one. In addition to the potential replacement of the whole system, privacy is another key aspect especially for enterprises. This is one of the reasons why the original public blockchain design was not attractive to some potential users. Next, we note that at the heart of a blockchain powered crypto-currency is the consensus protocol which guarantees the validity of each transaction. However, the original PoW protocol is very energy inefficient since it requires a huge amount of (computational) effort from every node (O'Dwyer and Malone, 2014). Each node has to communicate back and forth with everyone else in the network to finalize the validation. In more recent years, we started seeing the development of more energy efficient protocols (Ismail and Materwala, 2019; Li et al., 2020). To conclude our list, we want to highlight another very important issue: regulation. Blockchain is by definition a world-wide network in which nodes are located all around the globe. There are no boundaries and there are no particular legal constraints. This gives rise to an important set of new concerns. For instance, it is often very difficult to correctly establish which jurisdictions' laws and regulations apply to a given transaction. In other words, it could be the case that a given transaction performed by an organization could fall under every jurisdiction in which a node in the blockchain is located. There are a number of other legal issues and concerns. For example, the difficulties of applying the existing regulatory regime can be seen clearly when it comes to the use of crypto assets (Kaal and Calcaterra, 2018). Other instances are given by conflict and dispute resolution which need particular attention in a system when there is no centralized party that takes responsibility for the provision of services (Swan, 2015; O'Shields, 2017).

To summarize, we can characterize a blockchain as a set of codified blocks containing the transactions occurring in the distributed system. Such system can be parametrized by a tuple $(\zeta, \rho, \xi, \gamma, \alpha)$ where:

- ζ is the central administrator variable with $\zeta(0)$ denoting no central administrator.
- ρ defines the set of access rights (e.g., read and write) in the blockchain. $\rho(O)$ denotes an open (permissionless) system.
- ξ is the defined consensus mechanism. $\xi(PoW)$ denotes use PoW consensus.
- γ is the set of cryptographic security methods. $\gamma(\tau)$ identifies a specific set of such methods τ.
- α is the set of transactional assets. $\alpha(C)$ denotes use of crypto-currency asset C.

The above parameters indicate a specific set of choices that need to be made when implementing a blockchain system.[7]

More generally, a range of other choices is possible. For instance, a different consensus mechanism besides PoW could be specified. Relaxing the specific tuple arguments required for a blockchain allows us to introduce a more general concept that is often referred to as a DLT (Distributed Ledger Technology). DLTs are a broader and much more comprehensive set of tools which incorporate the classic blockchain technology as a special case.[8] As in Nakamoto et al. (2008), the original blockchain has the central feature of being a $\rho(O)$ system that is completely open, public, and permissionless, $\zeta(0)$. In addition, the validity of a transaction is checked by a PoW consensus protocol ($\xi(PoW)$) to support a crypto-currency $\alpha(C)$. In DLTs though, we can relax some of the initial requirements, such as a completely open system or the need for a distributed consensus protocol. Thus we are now in a position to introduce a definition of a DLT.

Definition 2. A DLT is a distributed peer-to-peer system of ledgers for which choices over user permissions, consensus protocols and cryptographic security to maintain immutability and achieve integrity need to be made.

Similarly to what we have done for a classic blockchain system, we characterize a DLT by a tuple $(\zeta(.), \rho(.), \xi(.), \gamma(.), \alpha(.))$ with a range of possible choices over the five defining variables. For instance, whereas for a Blockchain the access was characterized by open (permissionless) access $\rho(O)$, for a DLT we allow for the possibility of restricted (permissioned) access to only certain agents. This can be denoted as $\rho(\Omega)$ where Ω is the set of permissioned agents with full access rights. In order to understand why some organizations may prefer to make specific DLT choices regarding the components identified by the tuple, we now investigate some of the design choices.

2.1. Permissionless versus permissioned

In light of what has been discussed above, we now focus on how some types of DLT differ from blockchains in terms of their control over read and write rights. In general, we have two main families of such systems:

[7] To give an example, a Bitcoin-like system would have the following parameters: $\zeta(0), \rho(O), \xi(PoW), \gamma(\tau = \text{SHA} - 256), \alpha(C)$.

[8] In this paper, the concept of classic blockchain refers to the original system describe by Nakamoto et al. (2008).

Permissionless (open) - $\rho(O)$:
in a permissionless system every user can join or leave the network at will without being pre-approved or vetted by any entity. The only requirement to participate is a computer with the necessary software installed. Each participant has free access to the full transaction history and can decide to actively participate in the consensus mechanism implemented in the architecture. The most notable permissionless blockchains are Bitcoin (Nakamoto et al., 2008) and Ethereum (Bogner et al., 2016; Tapscott and Tapscott, 2016; Hirai, 2017) which have accumulated around $200 billions in market capitalization (Coin Market, 2020). To summarize, what follows are some of the most relevant features of permissionless blockchains: they are decentralized and distributed so that there is no single entity (i.e., either an enterprise/organization and/or government) that can bring the network down and censor parts of its transaction history. Transparency is paramount in this architecture. For example, users need transparency into the ordering of transactions and how they are grouped into blocks and chained together as monetary incentives for running nodes are based on these details. Moreover, these incentives are fundamentals to support participants (and miners) to run and trust the network itself. This is the working principle of the Bitcoin protocol (Nakamoto et al., 2008). Another feature is anonymity. A permissionless blockchain, to a certain extent, guarantees that participants can stay anonymous. Recent work by cryptographers has shown that there exist limits on achievable anonymity, specifically by showing how a barrier given by pseudonyms can be overcome in order to unveil the true identity of the node (Meiklejohn et al., 2013).

Permissioned (restricted) - $\rho(\Omega)$:
a permissioned DLT is a closed ecosystem in which each node and its role within the network are well defined and formalized. Members are pre-selected by a given authority (e.g., an owner or an administrator of the ledger) who controls network access and sets the rules of the ledger. Usually, this type of architecture is built to allow an organization to efficiently exchange and store information in a "secure" way. This partially solves for a number of concerns governments and regulators have about permissionless distributed ledgers such as identity verification of network members, whom to license and regulate, and legal ownership of the ledger. The given authority bears the responsibility to ensure that the participants in the network are reliable and trustworthy, something that is in complete contrast with permissionless systems. Besides privacy concerns, another appealing feature of this type of architecture is that typically, it does not require a computational intensive PoW to verify transactions but relies on different algorithmic rules to establish consensus among members. In permissioned DLTs, any node can propose an addition of a transaction, which is then replicated to other nodes, potentially even without any consensus mechanism. The set of features we described for permissionless systems also apply in this context but with an important difference. In fact, we see a dedicated open source community focusing on the development of a suite of frameworks, tools and libraries for enterprise-level applications which has been started by the Linux Foundation in 2015 (Linux Foundation, 2015).[9] Among others, early commitment to this project came from IBM which has contributed in the development of IBM's Hyperledger Fabric Project (Cachin, 2016).[10],[11] A direct competitor of Hyperledger is the platform developed by the R3 consortium called R3 Corda (Brown et al., 2016; Mohanty, 2019).[12] Although there are common features to both the systems, there are also differences. We refer the reader to the work by Valenta and Sandner (2017) for a detailed review of their features.

In reality, there is not just a binary categorization of architectures based purely on access rights, but rather a spectrum of possible configurations. Each set of design choices has a specific degree of openness and decentralization for a given the distributed ledger system. At the boundaries of the spectrum, we recognize a fully open, permissionless blockchains, such as Bitcoin, whereas at the other end are permissioned DLTs hosted by private entities, such as Hyperledger Fabric or R3 Corda. Between the two extremes, there is a taxonomy of features which, if combined together, give rise to the rich ecosystem of configurations we observe today.

In the following discussion, we refer almost exclusively to DLTs rather than blockchains. We do this to highlight that we are considering distributed systems other than highly tailored blockchains used to support Bitcoin transactions. Indeed, we go even further and argue that going forward, DLTs that do not necessarily have a built-in Bitcoin (crypto-currency) support are perhaps the most promising candidates for widespread use in private and public sector organizations. Before considering these wider issues, we complete our brief historical overview by noting that the initial application of blockchain (i.e., Bitcoin), came with some limitations.

2.2. Smart contracts

The original blockchain design introduced by Nakamoto et al. (2008) is well suited to make payments but falls short in its ability to execute any instruction organized into codes. This is one of the main reasons why, in 2015, a different network

[9] https://www.linuxfoundation.org/about/.
[10] The full list of organizations who committed to the initial development include: Accenture, ANZ Bank, Cisco, CLS, Credits, Deutsche Börse, Digital Asset Holdings, DTCC, Fujitsu Limited, IC3, IBM, Intel, J.P. Morgan, London Stock Exchange Group, Mitsubishi UFJ Financial Group (MUFG), R3, State Street, SWIFT, VMware and Wells Fargo.
[11] https://www.hyperledger.org.
[12] https://www.r3.com.

called Ethereum was launched after an Initial Coin Offering (ICO) that raised $16 M in bitcoins. These two systems, Bitcoin and Ethereum, have several key technical differences and one of them fundamentally changed the way we interact with the latter. For instance, in Bitcoin in order to prove that one can spend *x* bitcoins from a given wallet, it needs to be shown that there is an *Unspent Transaction* of at least *x* bitcoins that was sent and stored to the wallet beforehand. In other words, we cannot transfer any money if we do not prove we already own that quantity of money in our wallet. Conversely, in Ethereum in order to prove that one can spend *x* ether from a given wallet, it needs to be shown that the balance of the account is above *x* ether. The account balance is updated after every transaction so that the proof is almost always real-time. This is a much more intuitive approach than the one found in Bitcoin. We move from an *Unspent-Transaction-Output* (UTXO) system of the Bitcoin network to an *Account Base* one of the Ethereum network. In addition to its less complex approach, the Account Base system has a couple of advantages. First, it allows one to associate a *memory state* with every account and second, it allows one to have non-human accounts capable of holding funds. This last point is incredibly powerful since it enables machines to execute codes running on the nodes of the Ethereum network. These codes are generically referred to as *smart contracts*. A smart contract is an agreement between two or more parties that binds them to the future resolution of a given set of conditions. A smart contract is the digital version of a standard contract where the conditions and the agreement itself are evaluated and executed by a computer. In addition, smart contracts ensure a *trustless execution*, meaning that the contracting parties are not forced to rely on external entities (such as notaries) to execute the conditions. Following prior literature (Tapscott and Tapscott, 2016; Savelyev, 2017), we define a smart contract as follows.

Definition 3. A smart contract is a trustless transaction protocol which enables the automatic verification, execution and resolution of all the terms as specified in a contract.

This new way of resolving contracts is potentially disruptive. Besides the inherent ability to eliminate unnecessary parties in the process, smart contracts provide a well-defined and robust language which is not subject to human interpretation and does not suffer from the presence of unintuitive legal vocabulary. We discuss potential applications and contexts in which these structures could be relevant in Section 5. Having now outlined the historical development of blockchain (to support Bitcoin) and smart contracts, we now consider how the field has developed.

2.3. The Risius and Spohrer Blockchain research framework

In a comprehensive multidisciplinary review of the blockchain literature, Risius and Spohrer (2017) adapt the Aral et al. (2013) social media research agenda so that it can be applied to research on this topic. This framework allows one to conceptualize the various approaches in terms of two top level features: activities and units of analysis. The three identified groups of activities are: Design and Features; Measurement and Value; Management and Organization. The four units of analysis are: Users and Society; Intermediaries; Firms and Industries; Platforms. In Table 2 of their paper, they classify 70 leading papers in terms of the above four features and conclude that "extant publications still focus primarily on technological and business related topics and are often confined to the disciplines of computer science and information systems rather than addressing the broader societal, political or judicative questions". Furthermore, they argue that the literature review provides further support for the view "that blockchain is an innovative technology in search of use cases" (Glaser, 2017).

Consistent with the framework above, the focus of this paper is to move beyond the sometimes grandiose unsubstantiated claims that blockchain is a "game changing disruptive technology" by introducing an element of realism into the debate based upon actual accounting systems that have been previously used in practice. That is, we build on the suggestion from Risius and Spohrer (2017) to "take on the challenge and achieve contributions that advance the general knowledge on blockchain systems, particularly regarding value creation and management". We do this by identifying four classes of problem in which when accounting knowledge is appropriately combined with computer science knowledge it can advance our understanding of the limits of achievements with DLTs. We hope that the proposed benefits of looking at the problems through an accounting research lens are further clarified by us focusing on settings where a crypto-currency is not a necessary requirement. To an extent, we propose that some researchers have been too quick to reuse what they have learned from the Bitcoin experiment and have recycled it claiming general applicability. We argue that, potentially, the most valuable DLT development for businesses may be crypto-currency-free permissioned systems.

3. Ledger transparency access rules

In the extreme case of Bitcoin, the blockchain is visible to all because it is a permissionless, zero trust system (Davidson et al., 2016; Xu et al., 2017). However, for a wider class of DLTs, the decision of who to give permission "to see details of transactions" is a key variable.

Many companies have chosen to opt for permissioned DLTs using frameworks such as IBM's Hyperledger. This is one of the most used open source environments which offers advanced facilities for robust and general development. Interestingly, some of these permissioned systems have purposely limited features for crypto-currency development. For permissioned DLTs, one of the first design criteria to resolve is what level of transparency to grant member agents (nodes). In some settings, this may be relatively straight forward, but in others this may result in a more complex task. For instance, if a DLT is designed for a supply-chain application, in order to improve coordination and planning, does a lower level supplier get to see

the blocks appended by a higher level one? Moreover, does the lower level supplier get to see appended blocks from same level peers in order to encourage competition? These sort of design issues are important because when coordination is achieved via a traditional centralized system, suppliers are typically contracted with through some form of competitive tendering (auction). In such a setting, an approved supplier working in isolation (and privacy) from other approved suppliers simply has to bid the lowest price to win a supply contract. To illustrate these issues, we now focus upon the automotive industry to better explain the basic logic of building decentralized blocks of information. We then discuss how these principles can be generalized to wider industry settings.

In the automotive sector, the competitive tendering systems of some western Original Equipment Manufacturers (OEMs) are referred to as *arms length contracting*. This is in sharp contrast to the forms of relationship-based contracting used in the Japanese automotive industry (Sako et al., 1992). A key element of early relationship contracting systems in the Japanese automotive industry was the sharing of information between suppliers and the final assembler. With a DLT, this feature can be readily built into the information system itself. That is, whereas for those early systems, suppliers needed to rely on the OEM to (strategically) distribute information, with DLTs such features are available by design without the need for routing via a centralized OEM. Through careful choices of inter-nodal access rules (i.e., conceptually this is the parameter $\rho(.)$), an OEM can more effectively facilitate information sharing between a network of suppliers. It is thus interesting to speculate whether accounting research on arms length versus relational contracting could have a bearing on the implications for those firms considering the introduction of different forms of DLT-based supplier contracting. In previous works by Gietzmann (1996),Mouritsen (1999), and Hoyt and Huq (2000), it is argued that the choice of contracting style (relational versus arms length) should depend on the extent to which the final assembler desires suppliers to make relation specific investments. That is, we are not arguing for a universal DLT applicability. Instead, we suggest an understanding of the need for relation specific investments in contractual relationships mediates the demand. This motivates us to identify the first class of research settings in which accounting knowledge can meaningfully contribute to DLT design:

RS1a: Mediating Exchange via careful choice of ledger access rules ($\rho(.)$) for a distributed ledger system is more likely to be successful when a final assembler wants suppliers in a permissioned DLT to make relation specific investment.

To summarize, it would be interesting to test whether implementing a distributed ledger system is more likely to be successful in a setting in which a final assembler wants suppliers to make relation specific investments (Wagner and Bode, 2014; Wu et al., 2017). Conversely, if such investments are not required, a distributed ledger system with variable ledger access rules is less likely to be helpful. More specifically, distributed systems are unlikely to have universal appeal in all industries. These considerations generate a demand for a careful analysis of the high level requirements needed before starting off any implementation. We suggest that this has long been within the the purvey of accounting.

Keeping within the automotive sector, a second application of DLT arises. When a supplier wants to raise private debt, a concern of potential funders is that the financial viability of the supplier may critically depend on the dominant relationship with a single specific OEM. To some lenders, such as banks, this is seen as a significant risk which triggers a red flag. In response to this, a number of economic agents have tried to introduce supply-chain financing methods, but have been hampered by the fact that while the supplier will share information, the OEM may not (Hofmann et al., 2017). We argue that through a thoughtful design, permissioned DLTs could more easily grant the needed supply-chain information to lenders in order to better manage risk exposures.

The above examples, drawn from the automotive industry, illustrate the sort of design issues that management may face when deciding on whether to adopt a DLT system approach. Choices regarding which nodes and data fields to give an agent access to are not neutral decisions. They are at the heart of effective organizational design and have relevant consequences. Historically, design choices were previously limited by the processing ability of a centralized agent "owning" the ledger database (Neil, 2019). We emphasize how these are issues that are all central in the study of accounting. Interestingly, these issues have been pushed to the forefront by the COVID-19 pandemic (Ting et al., 2020). In designing contact tracing apps, fundamental decisions over whether a centralized versus a distributed system should be adopted have been debated (Lomas, 2020). In addition, who has direct access to the data and whether a centralized coordinating authority can be trusted has also been questioned. This has been further probed by Ian Grigg with his development of a triple-entry system where a third entry may be made to give public transparency on a blockchain to certain critical transactions (Grigg, 2004; Cai, 2021).

A third area where the interface between distributed data transparency and access rules is relevant is public services provision. Some have gone as far to say researchers need to consider a new field of study: algorithmic government (Engin and Treleaven, 2019). An example better shows this point. In many countries, vehicle drivers applying for a driving license need to provide information on their relevant medical history. Typically, this is self-certified without any independent and formal check since licensing authorities cannot access to personal medical records. Alternatively, at additional cost the driver can choose to undergo an independent assessment by an approved medical practitioner. If the data was stored on a DLT, restricted medical data would be easily accessible and verifiable by nodes thus making querying processes possible. Obviously, issues of confidentiality and privacy arise. Once again, careful design choices ensure that information contained in

medical records are released with precision (e.g., for specific medical conditions only) and consistency (e.g., with a high level of reliability). Moreover, immutability and transaction history tracking would reduce the likelihood of fraud attempts.[13]

The reason why DLTs make possible such transformation in public services delivery lies in the change of paradigm regarding how individuals are treated and view by these systems. For instance, an interpretation is that each individual is no longer viewed as a pensioner who pays taxes, who receives benefits, and who has medical insurance and a driving license. Instead, it is a block of data fields that can be accessed by the various public service departments before any license (or payment or any other type of service) is made or requested. If, on the one hand, the technical procedures to code and link the decentralized blocks may be a major task, on the other hand the very same procedures, which ensure the validity of the different transactions, are well within the realms of accounting.[14] While some early DLT developments have been funded because they had enthusiastic internal technology champions, hopefully future developments will be based on a better appreciation of the critical costs and benefits involved. Hence, we believe that an important area for future research in public sector accounting will be in the evaluation and documentation of such costs and benefits of DLT implementations. This motivates us to identify a second class of DLT design choices in which accounting knowledge can meaningfully contribute to economic activity.

RS1b: Public sector record integration via careful choice of ledger access rules ($\rho(.)$) for a distributed ledger system is more likely to be cost effective when different public services share a higher number of data fields.

4. Regulatory compliance requirements

For many economic transactions (e.g., trade finance transfers and freight movement across international borders), there can exist complex and lengthy compliance procedures which must be completed before a transaction can be finalized (Sathye, 2008). In such circumstances, a DLT system can outperform traditional ones in terms of reliability and speed. For instance, proof of recorded and immutable transactions related to compliances can be provided independently, asynchronously and in advance. This is something that traditional compliance systems cannot do. In addition to mentioned benefits, several others arise. Yeoh (2017) and Maull et al. (2017) argue that in general, DLT can reduce transaction costs. Finken and Finkemeyer (2019) discuss instead applications in the banking industry in detail with Blidholm and Johnson (2018) giving a thorough explanation on how transaction cost efficiencies can occur in trade finance operations. They illustrate how slow and cumbersome traditional centralized processing of trade credit are. In response, they argue that:

> *"Today's traditional process of issuing trade and export finance products requires manual labour and physical handling of paper in several of the necessary steps (*Ramachandran et al., 2017b*). The application for a guarantee is commonly sent to banks via email or by traditional postal service. The application is reviewed by an employee at the trade finance department manually (*McWaters et al., 2016*). If the application is correctly signed and the necessary documents are attached, such as a copy of the purchase agreement, the process of issuing the product can start. An up-to-date sanctions lists and Anti Money Laundering (AML) registries must then be inspected. If there are no matches in these lists and the Know Your Customer (KYC) documentation is correctly filed, the issuing can proceed (*Ramachandran et al., 2017b*). Commonly, the majority of processing of sanctions, AML and KYC is manual (*McWaters et al., 2016*). Further, the bank must ensure that the client wishing to issue the trade finance product has enough credit. Then, the terms of the contract regarding the product are agreed upon and the actual contract for the product can be written. This is also commonly done manually, and the document may need to be inspected by several colleagues depending on the routines of the bank.*"

Based on survey evidence, they found that one of the major advantages of using a DLT system in such a setting is the reduction of compliance costs. This is possible because there is no longer a need for a constant production/updating of reports aimed at regulators if they are given direct access to the distributed ledger system. Treat and Brodersen (2017) estimate that this reduces compliance costs by as much as 30–50%. Interestingly, much of the estimated cost savings are in the form of disintermediation of intermediaries such as notaries and external compliance officers in addition to procedures. Rather than assuming a transaction can be overseen by an intermediary, the original blockchain system was specifically designed with no trust at all in intermediaries and other economic agents — hence the importance of cryptographic controls on potential tampering of records and the use of decentralized consensus without the need for an intermediary. In their review of disintermediation facilitated by a blockchain, Zamani and Giaglis (2018) point out that, in reality, trust still needs to be managed. This leads them to conclude the following: "However, we don't consider full disintermediation to be a possible scenario, as this would require businesses to operate their own DLT. Rather, it is likely that this will lead to new roles for intermediaries. Presently, DLT is an attractive technological solution when the requirements include proof of ownership, trade ability, and trust among the participating actors with the aim to achieve real time transactions, increased reliability and resilience to external threats. As a result, we consider that intermediaries may still have a role to play within a blockchain-enabled environment." At a more applied level, Smits and Hulstijn (2020) present three case-studies of blockchain applications in the diamond industry explaining how issues arise when a trusted but (costly) notary is replaced by

[13] This is particularly relevant and has significant potential in the area of public benefit payments (Hyvärinen et al., 2017; Cheng et al., 2018).

[14] Accounting knowledge is key here since that for those costs to be worth incurring, the public services need to model what potential service delivery benefits and cost savings are achievable. In other words, the final decision makers should be an accounting professional.

a complex software system. How trust is endangered and made accountable is of primary importance in the application of permissioned DLT systems.[15] A further example of these issues is provided by Charles et al. (2019) who discuss how different DLT platforms result in different possibilities for the management of medical records. All of the above motivates the identification of the second class of problems in which accounting knowledge can contribute.

RS2: When regulatory compliance procedures specify cross-validation with an increasing numbers of external data providers, the costs associated with contacting, verifying and updating a large set of nodes (economic agents) within a centralized ledger system can increase dramatically. There may be significant compliance cost reduction possibilities switching to a distributed ledger system.

Within a research setting, it may be possible to implement a classic difference-in-differences design (Card and Krueger, 1993; Card and Krueger, 2000; Abadie, 2005) in which some departments introduced a DLT and others did not. Managing and tracking the costs of compliance procedures is again a classic (cost) accounting based control issue for which accounting knowledge is required in order to correctly evaluate and confirm any potential advantage.

5. Smart contracting variables

In addition to the focus on transaction verification, recording and ex-post compliance procedures, DLTs can also be used to exchange terms of a contract. Some users started to formalize sets of instructions in specific structures called smart contracts to be put on the system. One of the advantages is the self-execution which can be triggered when pre-defined conditions are met. One way to refer to these self-executing contracts goes under the tag of *robot contracts* (Crosby et al., 2016). However, as Frantz and Nowostawski (2016) put it "when interacting in an open environment, contractual specifications should be accessible to all engaging entities, whether artificial or human. While the DLTs assure deterministic execution and consistent state representations, the codified contracts as *de facto* coordination protocols, still require careful design and implementation, and cannot guard against badly written or insecure contracts, an aspect recently witnessed in the massive theft of funds from the Ethereum's most successful Decentralized Autonomous Institution (Finley, 2016)". The *language* developed in businesses to codify the terms and performance of complex contracts is a mixture of accounting and legal practice (Savelyev, 2017). Understanding how the measurement of potential contracting variables such as profit (and even revenue) are susceptible to strategic or illegal acts assists in contract design.[16]

Another aspect that differs from traditional transactional systems is the *tokenization* of assets. Put simply, a *token* is the digital version of a physical asset belonging to the real-world. In this context, an asset can be anything from a car, a diamond, or a building. Tokenization activates interesting options which are simply not available in the corresponding physical world. To illustrate this point, we consider a common market such as the Real Estate Industry (REI) and we then proceed to compare an investment approach based on tokenization with a more traditional one. In the traditional REI, there are a number of problems which restrict access. For instance, some entry point barriers are related to citizenship, credit score, intermediary fees and cash requirements. Intermediary fees may vary dramatically when buying a new asset either for investment or private residency. Also, there can be many different forms of fees such as exchange fees, notary fees, taxes, investment fees and others. An even more common barrier is related to liquidity requirements. The combination of these barriers makes any type of initial investment complicated. As we discuss below, DLTs can bring fresh air into this business model by enabling micro-investments on digital properties created through a tokenization process. This could make possible a significant enlargement of the contracting space.

Keeping the focus on REI, assets such as entire buildings can be *tokenized* (i.e., decomposed and/or disaggregated into smaller parts) and registered on a blockchain as multi-class assets. Ideally, there is no lower bound in terms of how aggressive the tokenization could be. For example, even a single room in the building can be considered as an asset to buy or sell (Chang, 2020). Following the standard protocol of any DLT/blockchain, each tokenized asset (i.e., a room) is then assigned with a unique address within the network to ensure immediate identification and traceability. This step effectively transforms any room in the building into a *smart asset* (Hargrave et al., 2018), the digital version of the physical asset being tokenized and transferred into a blockchain.

With a centralized system, having a "single room asset" was also possible but before any trade could be performed, a financial intermediary had to check whether a claim on one asset could in any way affect claims on other assets. In other words, the building documentation needed to be formally evaluated in its entirety by an certified professional. In contrast, with a DLT, by design, separable parts of the building are fully specified by their address in the system. As Walport (2016) put it, DLT can "be applied in a broad variety of areas, particularly in smart contracts and asset registration. By registering assets on a distributed ledger, all property could effectively become smart assets, providing a robust and trustworthy proof-of-record for a broad variety of services that currently cost Small-Medium size Enterprises (SMEs) time and money. Examples include registering Intellectual Property (IP) and patents, wills, notary services, health data and SIPPs/pensions. DLTs offer a

[15] DLT systems are very complex and have several components, including trust. The relevance of accounting knowledge in this context lies in the ability to identify critical liabilities and address them in a timely manner.

[16] An interesting question is to what extent Solidity, the programming language underlying Ethereum, (Dannen, 2017b) is capable of capturing the complexities of contracting in an efficient and reliable way (Rubinstein, 1996).

new way to coordinate these types of services." A DLT greatly enhances the possibility for the implementation of subsidiary ledger structures in which very detailed and disaggregated information is held for monitoring and tracking purposes. Tasks like asset registry management, which were perhaps viewed as somewhat mundane in the past, are now critical strategic variables in a world that is moving towards the digitalization of assets. Once again, we highlight how asset registry management is central to the study of accounting. In addition, how performance metrics will be designed to motivate economic agents in such settings is also of considerable interest.

To summarize, for smart contracts and tokenization to function effectively, a language that gives complete clarity on claims and liabilities has to be implemented and used. In addition, digital transaction transfers need to be recorded with sufficient detail. This gives rise to our third class of problems in which optimal organizational design depends in part on accounting knowledge.

RS3: The relative success of smart contracting, together with tokenization, depends upon the ability of contract designers to create asset registers and transfer protocols that account for transactions in an unambiguous manner and which are not open to manipulation or legal contention.

The most prominent example of an open source blockchain devoted to smart contracting is Ethereum. While considerable effort has been put into coding smart contracts, one cannot ignore the fact that if a contract is to be automated, an ad-hoc terminology is required to ensure that terms and conditions are presented in a clear and unambiguous way. The dedicated language Solidity serves this purpose (Dannen, 2017a).[17] Interestingly, since the details of Ethereum contracts are reported publicly, a potential area of research would be to study how prominent the use of standard accounting terminology in contract specification is.[18] It would be then interesting to see if variation in accounting intensity in contracts has any performance implications.

6. Provenance traceability

In areas such as antiques, classic cars and fine wine, the concept of provenance is well developed. Clearly, a DLT with its immutable transaction history can contribute in the development of specific systems for the sole purpose of tracking these fine objects. However, what is particularly interesting is the relatively low operating cost of a DLT, making objects as mundane as bunches of spinach traceable. After various food scandals, the provenance and traceability of food along the supply-chain has become an important issue. Kamath (2018),Yiannas (2018) and Zhao et al. (2019) document how leading retailers have implemented provenance tracking systems using DLT systems. That is, provenance is no longer only achievable for expensive items when traceability is a key to sales performance and compliance. Moreover, some valuable objects such as designer clothes, which are often subject to the attentions of counterfeiters, can benefit from this improved tracking. DLT systems such as IBM Hyperledger have allowed fashion retailers and other manufacturers to make forgery more difficult (Liang et al., 2017; RRamachandran et al., 2017a; Ramya et al., 2018).

It is helpful here to clarify that to implement a provenance tracking system other long established technologies exist. For instance, Feng et al. (2020) point out that "Traditional Internet of Things (IoT) traceability systems can monitor and store the specific information in all stages of production, processing, distribution and consumption by using Radio Frequency Identification (RFID), Wireless Sensor Network (WSN), Near Field Communication (NFC) technology, etc. It can provide valuable information for the food quality monitoring and traceability. However, it is based on the centralized server-client paradigm, the stakeholders and consumers have to rely on a single information point to store, transmit, and share the traceability information. As a result, most consumers have difficulty in acquiring full transaction information and tracking the origins of products". We argue that what is required is instead a system that allows a consumer to confirm provenance in a decentralized way (Aung and Chang, 2014). Put simply, if a key feature of a product is its place of origin, some consumers do not put absolute trust in a logistics tracking system that is 100% controlled (i.e., recorded and supervised) by the end-supplier retailer. DLT systems that allow other external observers to confirm and/or assess provenance may have value. The fact that independent observers have permission to read and/or write to a supply-chain DLT may increase the value of the final product and in part address concerns about the trustworthiness of data.

More generally, and as discussed in Section 3, the design of a DLT allows end-users or purchasers to have selective decentralized access to associated earlier ledger transactions such as point of manufacture or collection and other key performance variables such as tracking route and tracking time. Taking all of the above, we define our fourth class of problems below.

RS4: The benefits of implementing a DLT provenance system depend upon multiple features such as ease and frequency of forgery, chance of medical harm and ability to prevent tampering. For some (ethical and ESG) products, consumers may value the fact that a DLT is not wholly controlled by the (end) retailer.

Returning to the original call by Risius and Spohrer (2017) for researchers to clearly identify how the use of blockchains creates value, the issue of how a blockchain (DLT) could support data analytics and decision making is a key question.

[17] Solidity is completely open source and its development code can be found here: https://github.com/ethereum/solidity/.
[18] All the smart contracts coded in Ethereum are available here: https://etherscan.io/contractsVerified.

Dillenberger et al. (2019) argues that DLTs "provides a rich warehouse of information for analytics and Artificial Intelligence (AI). Currently, data scientists spend 80% of their time collecting, preparing, cleaning, and organizing data for analysis (Press, 2016). For an enterprise DLT, the data has already been identified, collected, prepared in a common format, and organized". The objective of allowing a DLT to support data analytics implies that common tasks such as data visualization, predictive modeling, together with provenance and compliance queries would be based on the data structures defined in the DLT. Dillenberger et al. (2019) explains that if a DLT with data analytics capabilities is combined with external data, it can be used to develop an AI model with trusted DLT data.[19] To make this point, they argue that typically AI models "may have been trained with biased data or poisoned data that would cause the models to create predictions that are detrimental to the company using the very same models". The data analytics could be carried out directly on the DLT (see for instance *BlockBench* (Dinh et al., 2017)) or alternatively using a paradigm called *off-chain analytics*. Dillenberger et al. (2019) propose a DLT-based AI system which supports the sort of queries and predictive analysis data analytics and accounting professionals are familiar with — the big difference being the data is not coming from a legacy SQL system. The development of these sorts of "federated learning" systems needs input from accounting professionals if they are to be considered relevant in the decision making process.

Going forward, the way distributed ledger systems are developed to support data analytics and AI is a key issue if an application is likely to support ongoing value creation. It is perhaps surprising that it is relatively late in the day that candidate designs are being proposed and these issues are being actively debated. We suggest that accounting professionals need to become actively involved in the conversation and take action rather than sitting in the back seat.

7. Conclusion

This paper builds on the Risius and Spohrer (2017) call for blockchain research to be more broadly based. We have identified four research settings in which we argue that accounting knowledge is relevant and critical to the design choices and the selection of features of a decentralized ledger system. These are: (1) choices over inter nodal transparency; (2) the means of achieving cost effective regulatory compliance; (3) designing effective means of disaggregating asset registers and designing smart contracts; (4) insuring transactions can be effectively recorded, tracked and analyzed to enable proper and effective provenance analysis.

While much initial attention was almost exclusively focused on how blockchain could facilitate the trading of cryptocurrencies, we suggest that such, almost exclusive, focus is no longer warranted. Going forward, public and private sector entities need to consider and evaluate best strategies to decentralize their information systems. In carrying out this task, a lot can be learned from combining an appreciation of existing centralized ledger systems with the new range of possibilities that arise with decentralized ledger technology. Building the management system of the future requires a clear understanding of how to combine features, like the immutable data tracking and recording system, together with AI and predictive analytics tools. This is a fundamental step in order to achieve the level of integration that is necessary for decision-making support systems to create value. The future of DLTs needs an interdisciplinary approach that includes computer science, cyber-security, cryptography, statistics as well as accounting. In conclusion, this research argues that specific topics in modern accounting research are highly relevant to the development of state of the art value creating distributed ledger systems. We operationalize this view by proposing that ledger design can be conceptualized in terms of the DLT design tuple $(\zeta, \rho, \xi, \gamma, \alpha)$ with accounting knowledge being used to address how best choices can be made to create value.

Acknowledgements

We would like to thank the editor and two anonymous reviewers for their comments which considerably improved the paper. We would also like to thank Dan Amiran, Daniel Rabetti and Maarten Sies for their useful comments.

Appendix A. Schematic illustration

Fig. A.1

[19] They characterize these systems as trusted because typically AI systems just work with public or corporate data as provided. If it is collected from a DLT, it is reasonably coming from an immutable tamper proof source and hence with a higher level of trust.

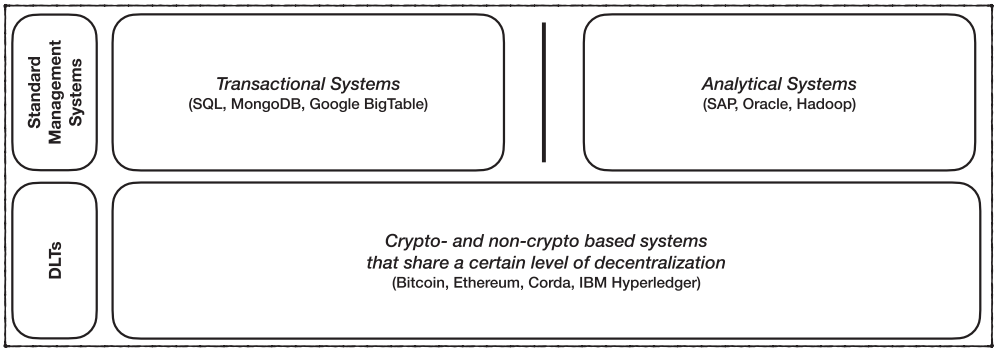

Fig. A.1. Schematic of industry-level management systems adapted and simplified from Meunier (2016).

References

Abadie, A., 2005. Semiparametric difference-in-differences estimators. Rev. Econ. Stud. 72 (1), 1–19.
Aral, S., Dellarocas, C., Godes, D., 2013. Introduction to the special issue—social media and business transformation: a framework for research. Inform. Syst. Res. 24 (1), 3–13.
Aung, M.M., Chang, Y.S., 2014. Traceability in a food supply chain: Safety and quality perspectives. Food Control 39, 172–184.
Blidholm, G., Johnson, M., 2018. The adoption of distributed ledger technology in trade and export finance operations of Swedish banks.
Bogner, A., Chanson, M., Meeuw, A., 2016. A decentralised sharing app running a smart contract on the ethereum blockchain. In: Proceedings of the 6th International Conference on the Internet of Things. ACM, pp. 177–178.
Boritz, J.E., 2005. Is practitioners' views on core concepts of information integrity. Int. J. Acc. Inform. Syst. 6 (4), 260–279.
Brown, R.G., Carlyle, J., Grigg, I., Hearn, M., 2016. Corda: an introduction. R3 CEV, August, 1:15.
Cachin, C., 2016. Architecture of the hyperledger blockchain fabric. In: Workshop on distributed cryptocurrencies and consensus ledgers, vol. 310, pp. 4.
Cai, C.W., 2021. Triple-entry accounting with blockchain: How far have we come? Acc. Financ. 61 (1), 71–93.
Card, D., Krueger, A.B., 1993. Minimum wages and employment: A case study of the fast food industry in New Jersey and Pennsylvania. Technical report. National Bureau of Economic Research.
Card, D., Krueger, A.B., 2000. Minimum wages and employment: a case study of the fast-food industry in new jersey and pennsylvania: reply. Am. Econ. Rev. 90 (5), 1397–1420.
Casey, M.J., Vigna, P., 2018. In blockchain we trust. MIT Technol. Rev. 121 (3), 10–16.
Chang, C., 2020. from securitization to tokenization. In: 5. Building the New Economy. MIT Press.
Charles, W., Marler, N., Long, L., Manion, S., 2019. Blockchain compliance by design: Regulatory considerations for blockchain in clinical research. Front. Blockchain 2, 18.
Cheng, J.-C., Lee, N.-Y., Chi, C., Chen, Y.-H., 2018. Blockchain and smart contract for digital certificate. In: 2018 IEEE international conference on applied system invention (ICASI). IEEE, pp. 1046–1051.
Coin Market, 2020. Cryptocurrency market capitalizations. Retrieved on May.
Crosby, M., Pattanayak, P., Verma, S., Kalyanaraman, V., et al, 2016. Blockchain technology: Beyond bitcoin. Applied. Innovation 2 (6–10), 71.
Dannen, C., 2017a. Introducing Ethereum and solidity, vol. 318. Springer.
Dannen, C., 2017b. Solidity programming. In: Introducing Ethereum and Solidity. Springer, pp. 69–88.
Davidson, S., De Filippi, P., Potts, J., 2016. Economics of blockchain. Available at SSRN 2744751.
Diffie, W., Hellman, M., 1976a. New directions in cryptography. IEEE Trans. Inform. Theory 22 (6), 644–654.
Diffie, W., Hellman, M.E., 1976b. Multiuser cryptographic techniques. In: Proceedings of the June 7–10, 1976, national computer conference and exposition, pp. 109–112.
Dillenberger, D., Novotny, P., Zhang, Q., Jayachandran, P., Gupta, H., Hans, S., Verma, D., Chakraborty, S., Thomas, J., Walli, M., et al, 2019. Blockchain analytics and artificial intelligence. IBM J. Res. Dev. 63 (2/3), 1–5.
Dinh, T.T.A., Wang, J., Chen, G., Liu, R., Ooi, B.C., Tan, K.-L., 2017. Blockbench: A framework for analyzing private blockchains. In: Proceedings of the 2017 ACM International Conference on Management of Data, pp. 1085–1100.
Dooley, K., 2001. Designing Large Scale Lans: Help for Network Designers. O'Reilly Media Inc..
Dworkin, M.J., 2015. Sha-3 standard: Permutation-based hash and extendable-output functions. Technical report.
Engin, Z., Treleaven, P., 2019. Algorithmic government: Automating public services and supporting civil servants in using data science technologies. Comput. J. 62 (3), 448–460.
Falco, G., Eling, M., Jablanski, D., Weber, M., Miller, V., Gordon, L.A., Wang, S.S., Schmit, J., Thomas, R., Elvedi, M., et al, 2019. Cyber risk research impeded by disciplinary barriers. Science 366 (6469), 1066–1069.
Feng, H., Wang, X., Duan, Y., Zhang, J., Zhang, X., 2020. Applying blockchain technology to improve agri-food traceability: A review of development methods, benefits and challenges. J. Cleaner Prod., 121031
Finken, S., Finkemeyer, D., 2019. The impact of blockchain on the transaction banking business model. J. Digital Bank. 4 (1), 19–34.
Finley, K., 2016. A $50 million hack just showed that the dao was all too human. Wired https://www.wired.com/2016/06/50-million-hack-just-showed-dao-human/ (June 2016).
Frantz, C.K., Nowostawski, M., 2016. From institutions to code: Towards automated generation of smart contracts. In: 2016 IEEE 1st International Workshops on Foundations and Applications of Self* Systems (FAS* W). IEEE, pp. 210–215.
Froystad, P., Holm, J., 2016. Blockchain: powering the internet of value. Evry Labs.
Gietzmann, M.B., 1996. Incomplete contracts and the make or buy decision: governance design and attainable flexibility. Acc. Organ. Soc. 21 (6), 611–626.
Glaser, F., 2017. Pervasive decentralisation of digital infrastructures: a framework for blockchain enabled system and use case analysis. In: Proceedings of the 50th Hawaii international conference on system sciences.
Goldreich, O., Oren, Y., 1994. Definitions and properties of zero-knowledge proof systems. J. Cryptol. 7 (1), 1–32.
Grigg, I., 2004. The ricardian contract. In: Proceedings. First IEEE International Workshop on Electronic Contracting, 2004. IEEE, pp. 25–31.
Hargrave, J., Sahdev, N., Feldmeier, O., et al., 2018. How value is created in tokenized assets. In: Blockchain Economics: Implications Of Distributed Ledgers-Markets, Communications Networks, And Algorithmic Reality. World Scientific.
Hawlitschek, F., Notheisen, B., Teubner, T., 2018. The limits of trust-free systems: A literature review on blockchain technology and trust in the sharing economy. Electron. Commerce Res. Appl. 29, 50–63.

Hirai, Y., 2017. Defining the ethereum virtual machine for interactive theorem provers. In: International Conference on Financial Cryptography and Data Security. Springer, pp. 520–535.
Hofmann, E., Strewe, U.M., Bosia, N., 2017. Supply chain finance and blockchain technology: the case of reverse securitisation. Springer.
Hoyt, J., Huq, F., 2000. From arms-length to collaborative relationships in the supply chain: An evolutionary process. Int. J. Phys. Distrib. Logist. Manage. 30 (9), 750–764.
Hughes, L., Dwivedi, Y.K., Misra, S.K., Rana, N.P., Raghavan, V., Akella, V., 2019. Blockchain research, practice and policy: Applications, benefits, limitations, emerging research themes and research agenda. Int. J. Inf. Manage. 49, 114–129.
Hyvärinen, H., Risius, M., Friis, G., 2017. A blockchain-based approach towards overcoming financial fraud in public sector services. Bus. Inform. Syst. Eng. 59 (6), 441–456.
Ismail, L., Materwala, H., 2019. A review of blockchain architecture and consensus protocols: Use cases, challenges, and solutions. Symmetry 11 (10), 1198.
Kaal, W.A., Calcaterra, C., 2018. Crypto transaction dispute resolution. Bus. Law. 73, 1–37.
Kamath, R., 2018. Food traceability on blockchain: Walmart's pork and mango pilots with ibm. J. Brit. Blockchain Assoc. 1 (1), 3712.
Lamehamedi, H., Szymanski, B., Shentu, Z., Deelman, E., 2002. Data replication strategies in grid environments. In: Fifth International Conference on Algorithms and Architectures for Parallel Processing, 2002. Proceedings. IEEE, pp. 378–383.
Li, A., Wei, X., He, Z., 2020. Robust proof of stake: A new consensus protocol for sustainable blockchain systems. Sustainability 12 (7), 2824.
Liang, X., Shetty, S., Tosh, D., Kamhoua, C., Kwiat, K., Njilla, L., 2017. Provchain: A blockchain-based data provenance architecture in cloud environment with enhanced privacy and availability. In: 2017 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID). IEEE, pp. 468–477.
Linux Foundation, 2015. Linux Foundation Unites Industry Leaders to Advance Blockchain Technology.
Liu, Z., Luong, N.C., Wang, W., Niyato, D., Wang, P., Liang, Y.-C., Kim, D.I., 2019. A survey on applications of game theory in blockchain. arXiv preprint arXiv:1902.10865.
Lomas, N., 2020. UK gives up on centralized coronavirus contacts-tracing app — will 'likely' switch to model backed by Apple and Google. Tech Crunch.
Lynch, G.S., 2009. Single point of failure. Wiley Online Library.
Maull, R., Godsiff, P., Mulligan, C., Brown, A., Kewell, B., 2017. Distributed ledger technology: Applications and implications. Strateg. Change 26 (5), 481–489.
McWaters, R.J., Bruno, G., Galaski, R., Chaterjee, S., 2016. The future of financial infrastructure: An ambitious look at how blockchain can reshape financial services. World Econ. Forum 49.
Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G.M., Savage, S., 2013. A fistful of bitcoins: characterizing payments among men with no names. In: Proceedings of the 2013 conference on Internet measurement conference. ACM, pp. 127–140.
Mendling, J., Weber, I., Aalst, W.V.D., Brocke, J.V., Cabanillas, C., Daniel, F., Debois, S., Ciccio, C.D., Dumas, M., Dustdar, S., et al, 2018. Blockchains for business process management-challenges and opportunities. ACM Trans. Manage. Inform. Syst. (TMIS) 9 (1), 1–16.
Menezes, A.J., Katz, J., Van Oorschot, P.C., Vanstone, S.A., 1996. Handbook of applied cryptography. CRC Press.
Meunier, S., 2016. Blockchain technology—a very special kind of distributed database. Saatavissa: https://medium.com/@sbmeunier/blockchain-technology-a-very-special-kind-of-distributed-database-e63d00781118. Hakupäivä, 4:2018.
Milani, B.A., Navimipour, N.J., 2016. A comprehensive review of the data replication techniques in the cloud environments: Major trends and future directions. J. Netw. Comput. Appl. 64, 229–238.
Mohanty, D., 2019. Corda architecture. In: R3 Corda for Architects and Developers. Springer, pp. 49–60.
Mouritsen, J., 1999. The flexible firm: strategies for a subcontractor's management control. Acc. Organiz. Soc. 24 (1), 31–55.
Nakamoto, S. et al., 2008. Bitcoin: A peer-to-peer electronic cash system.
Neil, S., 2019. Here's how blockchain will change the oem business model.
Nordin, R.J., Bauer, A.L., Krishnan, S., Lazar, G.W., 2006. Data replication facility for distributed computing environments. US Patent 7,054,910.
O'Dwyer, K.J., Malone, D., 2014. Bitcoin mining and its energy footprint.
O'Shields, R., 2017. Smart contracts: Legal agreements for the blockchain. NC Banking Inst. 21, 177.
Press, G., 2016. Cleaning big data: cost time-consuming, least enjoyable data science task, survey says. Forbes.
Ramachandran, A., Kantarcioglu, D., et al., 2017a. Using blockchain and smart contracts for secure data provenance management. arXiv preprint arXiv:1709.10000.
Ramachandran, S., Porter, J., Kort, R., Hanspal, R., Garg, H., 2017b. Digital innovation in trade finance.
Ramya, U., Sindhuja, P., Atsaya, R., Dharani, B.B., Golla, S.M.V., 2018. Reducing forgery in land registry system using blockchain technology. In: International Conference on Advanced Informatics for Computing Research. Springer, pp. 725–734.
Risius, M., Spohrer, K., 2017. A blockchain research framework. Bus. Inform. Syst. Eng. 59 (6), 385–409.
Rubinstein, A., 1996. Why are certain properties of binary relations relatively more common in natural language? Econometrica: J. Econometric Soc., 343–355
Sako, M. et al, 1992. Price, quality and trust: Inter-firm relations in Britain and Japan, Number 18. Cambridge University Press.
Sathye, M., 2008. Estimating the cost of compliance of amlctf for financial institutions in australia. J. Financ. Crime 15 (4), 347–363.
Savelyev, A., 2017. Contract law 2.0:'smart'contracts as the beginning of the end of classic contract law. Inform. Commun. Technol. Law 26 (2), 116–134.
Smits, M., Hulstijn, J., 2020. Blockchain applications and institutional trust. Front. Blockchain 3, 5.
Social Science Research Network (SSRN), 2019.
Swan, M., 2015. Blockchain: Blueprint for a new economy. O'Reilly Media Inc..
Tapscott, D., Tapscott, A., 2016. Blockchain revolution: how the technology behind bitcoin is changing money, business, and the world. Penguin.
Ting, D.S.W., Carin, L., Dzau, V., Wong, T.Y., 2020. Digital technology and covid-19. Nat. Med. 26 (4), 459–461.
Treat, D., Brodersen, C., 2017. Using distributed ledgers: Blockchain moves to early adoption.
Treiblmaier, H., 2018. The impact of the blockchain on the supply chain: a theory-based research framework and a call for action. Supply Chain Manage. Int. J.
Valenta, M., Sandner, P., 2017. Comparison of ethereum, hyperledger fabric and Corda. no. June, pp. 1–8.
Wagner, S.M., Bode, C., 2014. Supplier relationship-specific investments and the role of safeguards for supplier innovation sharing. J. Oper. Manage. 32 (3), 65–78.
Walport, M., 2016. Distributed ledger technology: beyond block chain (a report by the uk government chief scientific adviser). UK Government.
Werbach, K., 2018. Trust, but verify: Why the blockchain needs the law. Berkeley Tech. LJ 33, 487.
Wu, A., Wang, Z., Chen, S., 2017. Impact of specific investments, governance mechanisms and behaviors on the performance of cooperative innovation projects. Int. J. Project Manage. 35 (3), 504–515.
Xu, X., Weber, I., Staples, M., Zhu, L., Bosch, J., Bass, L., Pautasso, C., Rimba, P., 2017. A taxonomy of blockchain-based systems for architecture design. In: 2017 IEEE International Conference on Software Architecture (ICSA). IEEE, pp. 243–252.
Yeoh, P., 2017. Regulatory issues in blockchain technology. J. Financ. Regulation Compliance 25 (2), 196–208.
Yiannas, F., 2018. A new era of food transparency powered by blockchain. Innov. Technol. Govern. Globaliz. 12 (1–2), 46–56.
Yli-Huumo, J., Ko, D., Choi, S., Park, S., Smolander, K., 2016. Where is current research on blockchain technology?—a systematic review. PloS One 11 (10), e0163477.
Zamani, E.D., Giaglis, G.M., 2018. With a little help from the miners: distributed ledger technology and market disintermediation. Ind. Manage. Data Syst.
Zhao, G., Liu, S., Lopez, C., Lu, H., Elgueta, S., Chen, H., Boshkoska, B.M., 2019. Blockchain technology in agri-food value chain management: A synthesis of applications, challenges and future research directions. Comput. Ind. 109, 83–99.