

This PDF includes a chapter from the following book:

Distributed Ledgers

Design and Regulation of Financial Infrastructure and Payment Systems

© 2020 Massachusetts Institute of Technology

License Terms:

Made available under a Creative Commons
Attribution-NonCommercial-NoDerivatives 4.0
International Public License

<https://creativecommons.org/licenses/by-nc-nd/4.0/>

OA Funding Provided By:

The open access edition of this book was made possible
by generous funding from Arcadia—a charitable fund of
Lisbet Rausing and Peter Baldwin.

The title-level DOI for this work is:

[doi:10.7551/mitpress/13382.001.0001](https://doi.org/10.7551/mitpress/13382.001.0001)

1

Introduction

Distributed ledger technology (DLT) or, better put, its various features in isolation and in combination, has the potential to be transformative. Nevertheless, this subject has engendered controversy and sharp debate as well as a lack of clarity in the terminology researchers use to discuss it.

The first part of this introduction outlines the debate, and the second part outlines the point of view of this book.

Of note, Bitcoin is thought of as having created the blockchain as part of its validation system, so some people consider Bitcoin, blockchain, and distributed ledger technology to be synonymous. They are not. Some distributed ledger technologies exist without the blockchain technology and without coins. Indeed, much of the technology for distributed ledgers existed before Bitcoin and blockchain.

So this book proceeds in reverse. It starts with distributed ledgers, works backward to blockchain, and defers a more in-depth discussion of cryptocurrencies to the end. Concepts, definitions, applications, and impact are discussed at each turn. The term “distributed ledger” is sometimes used synonymously (if incorrectly) with the term decentralization. Computer science and data science are needed to clarify the distinction, and we compare and contrast with the meaning of decentralization in economics.

1.1 General Motivation: A View from All Sides

We begin with selected quotes from policymakers and academics (not practitioners or fintechs) in support of the premise that technology is fundamental.

DLT refers to the processes and related technologies that enable nodes in a network (or arrangement) to securely propose, validate, and record state changes (or updates) to a synchronized ledger that is distributed across the network's nodes. In the context of payment, clearing, and settlement, DLT enables entities, through the use of established procedures and protocols, to carry out transactions without necessarily relying on a central authority to maintain a single “golden copy” of the ledger.

DLT may radically change how assets are maintained and stored, obligations are discharged, contracts are enforced, and risks are managed. Proponents of the technology highlight its ability to transform financial services and markets by: (i) reducing complexity; (ii) improving end-to-end processing speed and thus availability of assets and funds; (iii) decreasing the need for reconciliation across multiple record-keeping infrastructures; (iv) increasing transparency and immutability in transaction record keeping; (v) improving network resilience through distributed data management; and (vi) reducing operational and financial risks [Mills 2016]. DLT may also enhance market transparency if information contained on the ledger is shared broadly with participants, authorities and other stakeholders.

—Bank for International Settlements (BIS 2017a, 2, 1)

Contracts, transactions, and the records of them are among the defining structures in our economic, legal, and political systems. They protect assets and set organizational boundaries. They establish and verify identities and chronicle events. They govern interactions among nations, organizations, communities, and individuals. They guide managerial and social action. And yet these critical tools and the bureaucracies formed to manage them have not kept up with the economy's digital transformation. They're like a rush-hour gridlock trapping a Formula 1 race car... With blockchain, we can imagine a world in which contracts are embedded in digital

code and stored in transparent, shared databases, where they are protected from deletion, tampering, and revision. In this world every agreement, every process, every task, and every payment would have a digital record and signature that could be identified, validated, stored, and shared.

—Marco Iansiti and Karim R. Lakhani, *Harvard Business Review*, 2017 (119–120)

But, of course, there are concerns and qualifications. We note some of these immediately, from the same sources. The Bank for International Settlements (BIS 2017a) lists risks associated with using DLT for payments that include potential uncertainty about operational and security issues arising from the technology; the lack of interoperability with existing processes and infrastructures; ambiguity relating to settlement finality; questions regarding the soundness of the legal underpinning for DLT implementations; absence of an effective and robust governance framework; and issues related to data integrity, immutability, and privacy. The Committee on Payments and Market Infrastructures (CPMI) chair, Benoît Cœuré, writes,

Central banks have traditionally played an important catalyst role in payments and settlements. This report will help central banks, other authorities, and the public to identify the risks as well as the benefits associated with the emerging technology, which could be the basis for next-generation systems (BIS 2017b).¹

Iansiti and Lakhani (2017) focus on the difficulty of adoption of transformative technologies. They distinguish between novelty and complexity, laying out various historical examples of innovation and current, ongoing experiments in DLT, in the end dividing innovations into four categories along the lines of high/low novelty and high/low complexity. This allows them to make predictions about not only where innovations are likely to succeed first, but also to identify those that could take considerable time, possibly decades, if they happen at all.

An implicit point: There is a distinction between invention of something new and its actual innovation and implementation. Lags in adoption are a murky criterion to use in the evaluation of the merits of inventions. This conflates the search for something “new” in DLT. Relatedly, innovation depends on context. In some settings, innovation is on the margin with much of the technology already in place. This may make the value of innovation marginal, potentially not worth the cost. But even if the gains from innovation could be incrementally large, vested interests with legacy systems can block change. In contrast, innovation can happen in settings where there is little if anything already in place on the ground, in which case implementation of key components singly, or in combination, can make a huge difference, even for innovations that are mundane and already adopted in other contexts.

An example of a low novelty–low complexity innovation given by Iansiti and Lakhani (2017) is Bitcoin. Their argument is that Bitcoin is another object like money for the transfer of value—hence, nothing novel. This, however, belies both Bitcoin’s creative algorithm and the controversy around it. To some, Bitcoin is singularly innovative. This has a lot to do with differences between computer scientists’ and economists’ perspectives, which we seek to clarify in this book. For others, Bitcoin is extremely problematic (we will return to this debate shortly). In any event, the bulk of innovations and experiments in new technologies occur under what Iansiti and Lakhani refer to as “localization.” That is, they introduce highly innovative uses and products but with a limited number of users. The list of these types of localized technologies is growing in length, some moving beyond commitments to experimentation and actual implementation.

One DLT use case that immediately reveals what DLT can do involves land-title projects such as those in Georgia, Sweden, and the Ukraine (Reese 2017). To buy property in these

locations, the lawful owner must have a secure title to sign over to the purchaser. DLT uses hashes to record every real-estate transaction and make them immutable, publicly available, and searchable so that titles can be transferred quickly, without costly title searches. Propy.com is an example of a proprietary company innovating in this space, with a distributed ledger in active use. The same idea underlies the emergence of digital assets to facilitate ownership and transfer.

In practice, in many markets, there are gaps and pauses in transaction time lines even for the most obvious transactions. A key example: In financial markets, trade, clearing, and settlement are separated in time. An agreement to trade between two parties can happen quickly, but it is then recorded into the private and proprietary legacy systems of each party, hence requiring reconciliation later. Trades in equity on a central stock exchange can take two days or more to settle, and, in part, this is not a matter of choice as there is no immutable synchronized record on which all parties can rely. Digital Asset is a company that has entered into an agreement with the Australian stock exchange to allow trade and confirmation in equities in real time, which was scheduled to be in operation by 2020.

TReDS, in India, is a platform for the discounting and sale of trade receivables. There are two other competing platforms currently operating in India. Since 2017, these three platforms have been operating a common distributed ledger for the recording of submitted buyer-seller receivable transactions. Each transaction has a unique ID number, so there can be no duplicates and thus no fraud. The Hong Kong Monetary Authority is also implementing a DLT system to avoid double invoicing.

Stellar is a not-for-profit entity that Iansiti and Lakhani would place in the category of an innovation with low novelty and high coordination needs. Stellar focuses on banking,

micropayments, and remittances for people without access to the formal financial sector or those who have access but at a high cost. Stellar has been operating since late 2014 and has a current market valuation, at the time of writing, of \$2 billion. Ripple is a for-profit entity with an even larger valuation, \$13.5 billion, founded in 2012. Stellar emerged from Ripple.

There is also innovation in nonfinancial markets. In 2017, Maersk and IBM implemented a distributed ledger technology for freight shipping, both for tracking and for improved logistics, sharing information and documentation among connecting nodes: port and terminal operators, customs authorities, customs brokers, transportation companies, and cargo owners. They project a substantial reduction in shipping costs.² Walmart has partnered with IBM to develop a system to track the supply chain of leafy vegetables from farms to stores so that in case of contamination, Walmart can quickly pinpoint and pull suspect produce. There are many such projects at the prototype stage—for example, pharmaceutical blockchain for reliable drugs.

Regarding smart contracts, Iansiti and Lakhani (2017) assign them to the most innovative yet hardest-to-implement category. Ethereum, R3's Corda, and Hyperledger are examples of smart-contract technologies running on distributed ledgers. As an example, Universal Market Access (UMA) allows contracts in financial derivatives that pay off as a function of the price of underlying assets.

One should be cognizant of the hype in the field and the difficulty of getting accurate, up-to-date information. There is continued discussion of improvement of the trade, clearing, and settlement systems at the Depository Trust and Clearing Corporation (DTCC) for repurchase agreements (repos) in the New York financial market—critical to the execution of the Federal Reserve's monetary policy. The reconciliation process takes up to two hours every trading day, creating an obvious friction. Yet the announced agreement between DTTC and

Digital Asset to install DLT did not move forward. One view is that a consensus in syndicates with conflicting interests is difficult to achieve, especially with existing infrastructure in place. A second view is that the proposed system added infrastructure on top of the old, adding to complexity and costs. In contrast to these negative experiences, DTCC and 15 leading global banks, in collaboration with IBM and Axoni, are implementing a re-platformed version of its Trade Information Warehouse for credit derivatives using synchronized distributed ledgers.

The Maersk platform is criticized by some as being proprietary to Maersk and apparently has had difficulty attracting other major shipping companies, arguably for that reason. This too is an increasingly typical experience, an obstacle recognized by participants in the industry. Surfing the web, a non-trivial set of initiatives seems to pitch blockchain, and billions of dollars are being spent on development, yet DLT may be pushed where there may be little need (Columbus 2019).

Controversy seems intrinsic to the technology or, better put, intrinsic to the way it is sometimes pitched. Observers draw an analogy between DLT today and the distributed computer networking technology known as TCP/IP (transmission control protocol/internet protocol), which is the communication protocol that laid the groundwork for the development of the internet. Following Iansiti and Lakhani (2017) closely, before TCP/IP, bilateral connections between two parties or machines had to be preestablished and sustained throughout an exchange, which was achieved though billions of dedicated communication lines. In contrast, TCP/IP transmitted information by digitizing it, breaking it up into very small packets, releasing it into the network, and finally, with smart receiving nodes, reassembling the packets and interpreting the encoded data. TCP/IP created an open, shared public network without any central authority or single party responsible for its maintenance and improvement.

This idyllic vision is maintained by many in the computer science research community and in industry and permeates the websites and white papers describing new releases. There are, by now, hundreds of cryptocurrencies and various leading platforms for the exchange of digital assets.³ Nevertheless, there are concerns and qualifications worthy of emphasis here.

First, the computer science community recognizes that there are trade-offs in the design of communication, computation, and decision-making systems: limited capacity in communication, latency (time lags) in transmission, and especially interpretation of information received (Mallett 2009). These trade-offs are not necessarily taken into account in some discussions of validation systems, but they are an intrinsic part of such systems. The synchronization of so-called decentralized ledgers actually requires centralization or coordination across nodes, and this is costly. Arguably, the validation systems of Ripple and Stellar arose to deal with some of these problems in Bitcoin, which is slow and done by blocks, not only to economize on costly proof of work but to deal with network latency (Hinzen, John, and Saleh 2019).

Likewise, hierarchical top-down systems do have some virtues, as in military command-and-control systems. Often, an optimizing choice would be a hybrid in between, which is hard to describe as either decentralized or hierarchical. The point is that there are trade-offs and choices that depend on context and goals.⁴ This is the interesting challenge of design. More generally, the language used to describe DLT, as if decentralized, is misleading. The term “disintermediation,” in its most favorable light, means, presumably, the elimination of the profits of financial intermediaries and market makers. Yet the financial platforms of fintechs and liquid high-velocity financial assets accomplish financial intermediation by almost any reasonable definition economists could use.

Second, the phrase “absence of a central authority” naturally creates controversy among policymakers. Specifically, as stated in the Bitcoin white paper written under the alias Satoshi Nakamoto (2008, 1): “What is needed is an electronic payment system based on cryptographic proof instead of trust, allowing any two willing parties to transact directly with each other without the need for a trusted third party.” Or, to put this crudely and more provocatively, the aim is to create a payment system that eliminates the need for central banks in the provision of money. Denison, Lee, and Martin (2016) make the point that, with exceptions, people do trust third parties: Both central banks that provide currency and reserves and (derivative) payment systems run by the named and trusted institutions that maintain the ledgers and operating systems.⁵

In what follows we will highlight hybrid systems where some of the features of distributed ledgers allow useful innovation, while other parts of the same systems rely on trusted third parties. It is the view of this book that neither side of the debate should dismiss these hybrids on the grounds they do not qualify as DLTs under some overly stringent definition.

1.2 Methods and Philosophy

To lay out the point of view of this book more specifically, we start, first, with the premise: Technological improvements in the design of mediation/intermediation systems could potentially, if executed properly, allow economies to be more connected in a positive way. Connectedness can come from new forms of mediation, though, as is already evident, not necessarily traditional intermediation through existing formal-sector financial institutions. Rather, the idea is (or should be) to create needed missing markets and institutions to fill in gaps in financial access and reduce inefficiencies, some of which are

large. Technological change is not new per se; we have witnessed various historical episodes of innovation. Communication systems have evolved from oral assignment to paper recording and written messages to e-messages and electronic registries (Townsend 1990; 1987). These episodes are instructive, as some of the basics of these episodes are the same as those observed in the current wave of DLT advanced communication and recording systems. These episodes of innovation also serve the purpose of allowing us to step back from the hype and controversy of DLT in order to highlight its key components and welfare gains.

Again, this book distinguishes invention from innovation. Was Bitcoin, with its blockchain and distributed ledgers, a sharply defined invention so that we can imagine innovations are now possible that were not possible before? Or was Bitcoin an innovation of previous inventions, a part of a longer, slow-moving process? The latter, for sure. Bitcoin with DLT was incremental. Nick Szabo described a decentralized digital currency, bit gold, with a public ledger and cryptographic puzzles a decade earlier (Moskov 2018; Narayanan et al. 2016). Szabo (1998) is also thought of as the originator of the smart contract implemented on distributed ledgers. Indeed, some argue that Szabo is Nakamoto, though he denies this.

Also relevant here are advances in database management and distributed computing, dating from years earlier. Secure multiparty computation goes back to the late 1970s, not simply to conceal content but also to conceal partial information about data while computing with data from many sources to produce publicly correct output. At least two of the key components of DLT, ledgers and cryptography, are so familiar to us that they seem mundane. Furthermore, some of these components have very deep historical roots dating back hundreds or thousands of years.

Nevertheless, there are recent, important, and contemporaneous innovations using these familiar components that have seen large welfare gains.⁶ Likewise, there remain gaps to fill. This book seeks to make these potential welfare gains and gaps transparent by being explicit about economic frictions and technological capabilities. It gives blueprints for the design of markets and institutions using suitably implemented components of distributed ledger technology, which includes formalizations of the limits of communication and database management systems widely discussed in the computer science literature.

On the negative side are bugs and troubling episodes. Consider the disturbing hacking event involving the Decentralized Autonomous Organization (DAO), a smart contract on Ethereum Classic. A hacker found a loophole in the coding and drained the equivalent of \$70 million in Ether cryptocurrency in the first few hours of the attack (Falkon 2017). The subsequent fork between Ethereum and Ethereum Classic also illustrates that smart contracts are not necessarily immutable after all, begging larger questions about consensus and legal frameworks. There has been apparent fraud in cryptocurrency exchanges. Largely unregulated at first, the Securities and Exchange Commission in the United States has argued that tokens are securities, though there is no regulatory consensus. On the other hand, it is not clear that traditional regulatory frameworks are appropriate either, revealing a gap in understanding that this book also tries to fill.

The conceptual framework adopted here is that of general equilibrium, and the welfare metric for deciding if something is good or bad, as well as how it should be regulated, is the Pareto criterion for a given economy. The recommended way of proceeding with this artillery is to assess what can be accomplished in a given economy relative to what is there now and, more specifically (at least in some contexts), exactly how to innovate.

Indeed, we can distinguish three possible metrics for this assessment of what can be accomplished, which, as a warning, can be easily confounded with one other. One would be to reduce obstacles to exchange and mitigate frictions. A second is to place value on products/systems that are commercially viable and actually potentially profitable. The third is again our basic standard: to allow allocations that are Pareto improvements relative to previous outcomes, though losers may need to be compensated. These metrics are not always equivalent with each other. One reason for failures of equivalence comes from the potential failure of the first fundamental welfare theorem in economics. The theorem states that under certain assumptions any competitive equilibrium, decentralized through a price system, must be Pareto optimal. But under some frictions, competitive equilibrium allocations are not necessarily Pareto optimal, and potential failures are intimately associated with some of the properties of e-money as money more generally. Another reason for failure is the political economy of reform. There can be losers from removing an obstacle, especially if there is no compensation, as noted. This may also explain slow adoption or failure to innovate even when the technology is well understood and potentially Pareto improving.

A broader view also comes naturally with mechanism design, where the distinction between public and private ownership has no real meaning. Agents enter into social agreements, subject to information, resources, and other constraints. It is as if a “planner” were acting on behalf of agents as a collective group. But the word “planner” is a misnomer, especially in this book as we sort through language issues. A planner would refer to a highly centralized system, juxtaposed with autonomous decentralized markets. Here, the planner of mechanism design theory acts through a secure multiparty computation framework in which underlying states are not required to be revealed. We will revisit this issue later, in particular how to compute and

implement optimized solutions to multiparty smart contracts and the limitations in doing so.

A counterexample to a forced distinction between private and public ownership is the case of private clearinghouses as a consortium of financial institutions—an industry association with tight rules for membership, collateral, and operations. Clearinghouses were not always connected to central bank accounts for settlement (Tucker 2014) and yet were public institutions in many ways. Campbell-Kelly (2010) describes the Bankers Clearing House in Britain and, with modifications, the Clearing House of New York, as algorithms for netting and settlement of checks among bankers, implemented by humans rotating around tables rather than computers yet sharing much with the code and liquidity issues of contemporaneous e-payment systems.⁷

The book proceeds, then, as outlined here. This introduction serves as an executive summary and the final chapter reviews the context again while also attempting to draw some conclusions from our analysis.

Chapter 2 describes the concept of what we mean by an economy: the underlying commodity space, general enough to include time, uncertainty, and geography. All the examples in this book fit into this general framework. The welfare criterion is made clear. When additional information and other constraints are appended onto programming problems for the determination of the class of optimal allocations, we refer to the solutions as *constrained optimal*. Measurement is also featured, ideally with integrated financial accounts, if available. We present the Townsend Thai project as an example of an economy, and it appears repeatedly through the chapters as a source of examples, chosen because many of the ingredients we wish to discuss come together there.

Chapters 3 through 6 describe four key components of distributed ledgers: ledgers as accounts, e-messages and e-value transfers, cryptography, and contracts including multiparty

mechanisms. We discuss, evaluate, and illustrate through the context of historical and contemporary economies each component, with featured applications in both developed economies and emerging-market countries. A recurrent focus is the general equilibrium impact of innovations and welfare gains from innovations featuring these key components, which does not require that all components be introduced at the same time.

Specifically, chapter 3 introduces ledgers in the context of various emerging markets and advanced economies. The ledgers are linked to statements of currency flows in Thailand as a first example, showing conceptually how common yet distributed accounts could be created from a database of transactions at the level of households and small and medium-sized enterprises (SMEs). One gain from a common database is that discrepancies in entries across agents can be detected and, in principle, corrected as they occur. There is an analogy with how financial accounts and double-entry bookkeeping allow for greater accuracy at the individual level. Illustrative applications of the use of financial accounts indicate new and important uses for the application of distributed ledgers.

This section on ledgers concludes with an important discussion from the computer science literature on the advantages and disadvantages of traditional database management versus the decentralized database management of distributed ledgers. With decentralization, in the presence of latency, impossibility theorems arise with regard to consistency (ledgers the same), accuracy (up-to-date and without error), and partition tolerance (in the presence of a partition tolerance, one has to choose between consistency and availability). Furthermore, even when the system is running normally, there is a tension between consistency and latency. From the distributed computing literature there is a theorem that with asynchronous systems, consensus is impossible. Yet with synchronous systems each node must be connected to every other node and

thus, with communication costs, this raises the issue of scaling up to large systems. Trusted third parties solve this problem, but this centralization not only requires trust but raises the issue of data integrity, as those with the correct access can also (accidentally) destroy or corrupt data, and there is data security/cyber-risk. In practice, choices are made and hybrids emerge.

Roughly speaking, database management systems have not paid much attention to incentives among parties with conflicting interests. Distributed ledgers for business applications err toward keeping everything secret, not toward solving a design problem. In this book we thus point an arrow toward where we could go. An example from economics with transactions costs from linking provides an illustration of an optimal hybrid system.

Chapter 4 features the second component: e-messages and e-payments. We compare and contrast examples of e-money, looking at Thailand, with its dominant use of paper currency and little e-value transfer, and at Sweden, where the use of e-payments now dominates and the use of currency has fallen off tremendously. This sets the stage for understanding an e-money innovation with large welfare gains: the case of M-Pesa through the mobile company Safaricom in Kenya. Many low-income and developing countries could experience such welfare gains. The Kenyan system uses a trusted third party and so what is to some a defining characteristic of DLT, no trusted third party, might lead one to dismiss this innovation. Alternatively, the gains from the e-transfer component are large with trusted parties. However, there are several caveats. Trust is less obvious when one takes into account the larger financial system. In addition, there are some infrastructure issues and the provision of liquidity that deserve attention in these contexts.

Chapter 5 deals with cryptography, validation, and consensus. A description of how contemporaneous e-systems work

without a universal trusted third party does beg the issue of consensus. However, there is no one single way to achieve consensus. Among the various consensus systems are Bitcoin cryptography with proof of work, Byzantine fault tolerant systems with proof of stake and voting rights based on coin ownership, and federated Byzantine protocols with layers of trust. There are trade-offs across these systems that involve fault tolerance, safety/consistency, liveness, latency, and transaction speed. Bitcoin and consensus algorithms have attracted much of the academic and industry interest, and there are interesting issues being revealed as computer scientists, economists, and their literatures interact. In addition, chapter 5 explores the validation systems and economics of Ripple, Stellar, Algorand, and HotStuff, the basis for Libra.

Chapter 6 presents the fourth component, contracts and multi-agent arrangements that are implemented as smart contracts. This chapter includes a discussion of how contract and mechanism design theory delineate various distinct concepts of trust, thus helping to clarify the debate concerning trusted third parties and what is needed, or not. Smart contracts operate on distributed ledgers and overcome obstacles—namely, there is commitment in entering into an agreement and in carrying it out (the immutability of terms). The language is similar to that used in financial accounting, with states for the balance sheet and flows as in cash flow, executed with commands. There are hybrid smart-contact systems with lots of possibilities and flexibility: unique and nonunique consensus; single or multiple trusting or nontrusted notaries; public versus private nodes; oracles for public information; and broadcast versus selectively private communication. The contracts and mechanisms that economists envision to deal with specified obstacles and frictions in an environment now have DLT as a natural implementation technology.

Basics include messages sent and recording on the ledger; truth telling and hence no further verification; past messages recorded as immutable history; enduring relationships under multi-period contacts; promised utilities as the key summary of past history; incentives to take appropriate actions; utility threats for lying about past history; costly state verification with no messages over a range of states; full commitment versus limited commitment distinguishing one meaning of trust; renegeing and thus a restriction to time consistency or the building of an internal scoring reputation mechanism; collusion and remedies for implementation; and commitment to limit sequential play to solve hold-up and bargaining problems.

In chapter 6 there is also a discussion of the similarities and differences between economics and computer science, including some concluding comments on integration. In much of computer science, nodes are trusted or not, and designs center around having a sufficient number of trusted nodes, as in fault tolerance. Yet smart contracts for mechanism design problems, if entirely implemented on Ethereum with its proof of work, including validation of code, can be prohibitively costly. Messages internal to the contract might also be done on-chain, but this is largely unnecessary if we respect internal incentives. Likewise, databases and documents can be secured off-chain. There is, however, encouraging common ground. Costly and imperfect messaging can be incorporated into the mechanism design, or even no messages at all, yet versions of the revelation principle still apply in some contexts. Though there are impossibility results regarding consensus and common knowledge in the economics literature, one can build on this and draw a distinction between following naïve or simple communication protocols versus incentives and strategic behavior in well-defined economic games. Optimal design of communication systems thus becomes key. Recent contributions establish the

effectiveness of multiple repeat messages and how iterations of decentralized validation can be truncated to achieve coordination. These are the nuts and bolts we need going forward.

Chapter 7 builds on smart contracts, going on to explore issues in “decentralization.” In some environments it is necessary to partition ledgers so that messages are sent but not seen by everyone, which improves incentives. Likewise, ultimate outcomes are randomized to facilitate concealment, which improves constrained risk-sharing. There is also a discussion of the need in some environments to have a preprogrammed, third-party custodian making portfolio decisions. Another consideration is how the burden of validation can be lessened if there are portable, concealable tokens that can be carried about and displayed voluntarily on request. A system that records histories of trade—which DLT provides—can improve welfare relative to a decentralized, partitioned system in which information is lost. Indeed, in a hybrid decentralized system, tokens can play a distinct role in implementing the centralized mechanism design outcome by conveying histories of trades—a kind of communication system, but one that avoids problems of scale. This can provide insurance and smoothing over time, even if tokens, or line items on ledgers, could in principle be concealed. Incentives for revelation of tokens or revelation of private accounts take care of that problem. The messages are endogenous, not forced or required as under centralized systems, but the messages are fully revealing. Tokens can also be used to track trades over multiple commodities, such as when there are preference shocks, but here, multiple colored coins may be needed if there are multiple dimensions to keep track of, for example, if over time there are shocks associated with preference reversals. This has a parallel in cryptography: Coins are not fungible, in the sense that coins have publicly verified histories, to trace ownership.

Several chapters deal with specific applications. The Thai setting is featured in chapter 8, making the point that context matters. Large gaps in services exist for credit, savings, payments, and insurance (Asian Development Bank 2017). The Townsend Thai project, on the ground with 20 years of data, shows that informal, local risk-sharing is good, achieved through credit chains and networks. But there are shortfalls in cash management. There are also shortfalls in county-level risk-sharing. The risk premium is low for idiosyncratic risk because of good pooling. Yet the risk premium is high for aggregate risk, even though aggregates differ across villages and could be pooled and better insured. Interventions have helped, but more are needed. A government village fund intervention allowed increased consumption overall by alleviating borrowing constraints and a cashing in of buffer stocks; better intermediation, especially for the lower-wealth households, through a costly state verification regime, with lower costs of verification for kin; and profits and increases in assets for high-productivity SMEs that received funds. But loans were more readily available for village committee members and those with connections to them; some kinship pathways mitigated distortions, but evidently these rely on preexisting trust. High-productivity SMEs without kin could have benefited more. Smart contracts on distributed ledgers can help overcome some of these trust issues, allowing trade among strangers.

The point is that innovations that make use of distributed ledgers have great potential. There are gains from individual contracts and services, such as escrow services with non-banks, savings products for automated deposit and portfolio management, and securitized waterfall payments along the path of supply chains from buyer to seller to employee loans. There are also gains from competition, because a common platform can provide structure for competition in contracts

with open access to providers writing smart contracts, as with free entry in general equilibrium models with an intermediary broker sector. A featured example of innovation is EvryNet, an intelligent financial automation operation system that provides open-source banking services and financial contracts to unbanked and underbanked populations.

Chapter 9 turns to systems for actual or pending payments and describes two innovations on distributed ledgers. One was run as an experiment by a central bank for commercial banks in an advanced country, Project Jasper in Canada, and is replete with sophisticated algorithms implemented as smart, multiparty contracts for queuing and clearing. The other application embraces the full set of possibilities for payment systems designed to achieve constrained-optimal trade, credit, and insurance in the context of cross-border payments among money transfer operators (MTOs) in Southeast Asia. Lightnet features an optimized liquidity management layer that efficiently searches for offsetting cross-country fiat balances for expedient clearing. MTOs can be viewed as agents with varying income and balance sheets hit by shocks, with the need for trade for their customers, hence themselves needing credit and insurance. Velo Labs acts in conjunction with a digital reserve bank.

Chapter 10 deals with regulation as an integral part of financial system design using distributed ledgers. Each section in this chapter is separate from the others, with important ideas for actual uses. First, DLT can improve on the current technology used by banks and markets in order to mitigate, if not eliminate, resulting runs. The idea is to exploit the time-stamped and immutable nature of ledgers to keep track of history and to condition current outcomes. Second, competition among providers can fail to complete the financial system because of unexploited complementarities and lack of coordination. Traditional regulation by sector and product can exacerbate this problem. The time line of when to allow competition also

matters. Not all forms of competition are good. Ex ante competition in rights to provide services and contracts can be fine, but there needs to be exclusivity and restrictions on contract execution and ex post spot trade. DLT systems can provide these, in principle.

The last section of chapter 10 shows that, in some instances, public information on ledgers is necessary for coordination and prudential regulation. To achieve an optimum, one has to know where the system is headed, the remaining options in the future to get there, and hence what trades have been accomplished in the past to establish what trades are needed now. Classic work established the generic impossibility of efficient yet decentralized monetary exchange. Knowledge of identities of agents, histories of trade and payments, and initial excess demands are needed for implementation, not simply pairwise knowledge of those contemporaneously matched but also information from others with whom contemporaneously matched payment parties have not been matched previously, so they could not know the history in that way. Likewise, there are potential crashes with cryptocurrencies and digital assets. As with circulating private debt as a medium of exchange, mismatch is likely as too much or too little debt is issued. Clearly, there could be problems in using distributed ledgers to keep track of all the information. An encouraging aspect of the featured examples is that there are only key instances that require public information on trade, not entire histories of everyone.

Chapter 11 considers cryptocurrency with an expanded discussion of the role and value of tokens in economies with distributed ledger systems. There are various types of money in existence, some fiat, some backed, as well as credit objects. The standard definitions measure “moneyness” via velocity and frequency of use in payments. In practice, in any given economy, one notes the prevalence of multiple media of exchange. Mechanism design theory pinpoints the special role of tokens

relative to fiat money; an unrestricted use of fiat money along the time line of contract implementation can hurt the ability to achieve a constrained optimum. From monetary theory, the value of money can be endogenous, as displayed in various distinct models, depending on who meets whom when. In these models the fundamental welfare theorems fail. Competitive equilibria without money are not Pareto optimal, but money as an intrinsically useless device has value. There is an associated empirical literature that examines whether these bubbles are large enough. Typically, however, there are multiple equilibria, a kind of indeterminacy in value, intrinsic to monetary theory. In other environments money is not needed when market structure is closer to being complete. Still, money can have a value because of its required use to pay taxes or legal stipulations—that is, media of exchange by fiat. Likewise, tokens in hybrid systems can play a role and have value on top of the fiat structures. Indeterminacy of token values in these contexts has remedies in the same roots of monetary theory: real interest and use requirements as with utility coins, both implemented with smart contracts. Tokens can have value via backing, potentially, and we critically review various types of stable coins. An algorithmic digital reserve system with commitment can implement optimal activist token policies armed with transactions data from the distributed ledger.

Chapter 12 offers a partial summary while focusing on major conclusions. An online appendix (which can be found at <http://www.robertmtownsend.net/research/online-appendix>) provides easy access to some of the key papers in the economics literature and thus details of the models or empirical work, with clickable links.