"Distributed Ledgers: Innovation and Regulation in Financial Infrastructure and Payment Systems"

Robert M. Townsend*
Elizabeth & James Killian Professor of Economics
Massachusetts Institute of Technology

Prepared for Presentation at
Sveriges Riksbank, De Nederlandsche Bank, and Deutsche Bundesbank
Annual Macroprudential Conference
Stockholm, Sweden
June 16, 2018

Revised April 18, 2019

Abstract

Distributed ledgers have the potential to transform economic organization and financial structure. Yet the subject is embroiled in controversy, hype, and terminological inconsistencies. Rather than get waylaid by alternative possible definitions of distributed ledgers, or decentralized ledgers, we focus more broadly on an economic analysis of what distributed ledgers can do. We proceed by analyzing key individual components. We also compare and contrast the economics framework with the frameworks of computer science and data management disciplines, to clarify the technology.

The familiar but key component parts of distributed ledgers are: ledgers as financial accounts, e-messages and e-value transfers, cryptography, and contracts including multi-party mechanisms. Each component is discussed, evaluated, and illustrated through the context of historical and contemporary economies, with featured applications in both developed economies and emerging market countries. These use cases are a

*Townsend gratefully acknowledges research support from the Eunice Kennedy Shriver National Institute of Child Health and Human Development (NICHD) (grant number R01 HD027638), the Centre for Economic Policy Research (CEPR) and the Department for International Development (DFID) under grant MRG002_1255; and funding assistance for the continuation of the field surveys by the Thailand Research Fund, the Bank of Thailand and the University of Thai Chamber of Commerce. I also am grateful for the collaborations with the Federal Reserve Banks of New York, Boston, and Chicago; Velo for decentralized settlement banking network; and EvryNet for decentralized custodian banking network in Thailand. The views expressed are my own. Thank you to Ernst-Ludwig von Thadden, James McAndrews, discussants of this paper, and other participants at the Sveriges Riksbank Annual Macroprudential Conference in June, 2017, participants at the "Kellogg Conference on Development Economics" conference at Northwestern University in September, 2018, participants at the "Cryptocurrencies and Blockchains" conference at the Becker Friedman Institute in November, 2018, participants at the First New York Fed Conference on FinTech, March 2019, participants at the Penn State Conference in Celebration of Neil Wallace's Contribution to Economics, April, 2019, Daniel Aronoff , Neha Narula, Neil Wallace, Joshua Gans and Jesus Fernandez-Villaverde, for very helpful comments, and Jennifer Roche and Deborah Jamiol for wonderful editing. Thank you also for comments on preliminary drafts to: Marios Angeletos, Pablo Aznar, Zach Chao, Co-Pierre Georg, Michael Lee, Rhys Lindmark, Jacky Mallet, Antoine Martin and Nish Patel.

hallmark of the paper. A recurrent focus is the general equilibrium impact of innovations and welfare gains from innovations featuring key components. This does not require that all components be introduced at the same time.

Contract theory is used to derive optimal arrangements, constrained only by obstacles to trade, featuring how the various aspects of ledgers can deepen infrastructure. Mechanism design and monetary theory are used to study public versus partitioned ledgers and improvements in payment systems. Prudential regulation, rather than being a barrier to innovation, can be improved with the use of distributed ledger technologies.

The goal is to provide blueprints for the ex ante optimal design and regulation of financial systems, including not only choices at the end points of the spectrum, of centralized versus decentralized systems, as in the hype, but the choice of hybrid forms in between. Overall, the paper provides a vision for where we are heading, being clear about obstacles along the way.

Table of Contents

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 lack of clarity in terminology

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

Of note, Bitcoin created the blockchain as part of its validation system, and so some refer to Bitcoin, blockchain and distributed ledger technology as being 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 prior to bitcoin and blockchain.

So this paper proceeds in reverse. It starts with distributed ledgers, works backward to blockchain, and defers a discussion cryptocurrencies to an appendix. 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 for clarification, compared and contrasted 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 nor fintechs) in support of the premise that technology is fundamental.

> "Distributed Ledger Technologies (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 synchronised 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…" (page 2)

> "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."

– Committee on Payments and Market Infrastructures, Bank for International Settlements (BIS), 2017a (page 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

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, which include the 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; the 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 transformation technologies. They distinguish between novelty and complexity, laying out various historical examples of innovation and current, ongoing experiments in DLT, in the end dividing them into four

[1] See also Mills et al (2016).

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, 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 costs, 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 which are mundane and already adopted in other contexts.

An example of 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 the creative algorithm and the controversy. To some, Bitcoin is singularly innovative. This has a lot to do with a difference between computer science and economics perspectives, which we seek to clarify in this paper. For others, Bitcoin is extremely problematic (we will return to this debate in this introduction, shortly). In any event the bulk of innovations and experiments in the 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 technologies is growing in length, some moving beyond commitments to experimentation and actual implementation.

One DLT use case that reveals quickly what DLT could do are land title projects such as those in Georgia, Sweden, and the Ukraine (Reese 2017). To buy property in these locations, there must be a secure title that the lawful owner can sign over to the purchaser. DLT uses hashes to identify every real estate transaction and make it 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.

In practice, in many markets, there are gaps and pauses in timelines 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 onto private and proprietary legacy systems, 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 should 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 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, that is, no double issue, and thus no fraud.

Stellar is a not-for-profit entity, which Iansiti and Lakhani would put in the category of having low novelty and high coordination needs. Stellar focuses on banking, micropayments, and remittances for people without access to the formal financial sector or 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. There are also various other, indeed many other cryptocurrencies in use.

There is also innovation in non-financial markets. In 2017, Maersk and IBM implemented a distributed ledger technology for freight shipping, both for tracking and for better logistics, sharing information and documentation among connecting nodes: port and terminal operators, customs authorities, customs brokers, transportation companies, and cargo owners. They are projecting 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 on ledgers. Universal Market Access (UMA) allows contracts in financial derivatives that payoff as a function of the price of underlying non-held assets.

One should be cognizant of hype in the field and the difficulty of getting accurate and 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 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, an obvious friction. Yet the announced agreement with 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, DTCC and 15 leading global banks, in collaboration with IBM and Axioni, are implementing a re-platformed version of its Trade Information Warehouse for credit derivatives using synchronized distributed ledgers.

The Maersk platform, in contrast, is criticized by some as being proprietary to Maersk and apparently has had difficulty attracting other shipping companies, arguably for that reason. This too is an increasing typical experience, an obstacle recognized by participants in the industry. Surfing the web, a nontrivial set of initiatives seem to pitch blockchain, and billions 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. An analogy is drawn between DLT today and distributed computer networking

[2] Real-time episodes are featured on their websites, see https://www.ibm.com/blockchain/hyperledger.

technology of TCP/IP (transmission control protocol/internet protocol), which 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 pre-established and sustained throughout an exchange, achieved though billions of dedicated lines. In contrast, TCP/IP protocol 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.

Though the bulk of this analogy between TCP/IP and DLT is about the technology itself, and one can see clearly from the way it is presented as an innovation why the technology took hold, there are two qualifications we add here.

First, the computer science community recognizes that there are tradeoffs 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 (Mallet 2009). Yet this is not necessarily taken into account in some discussions of validation systems, though this is an intrinsic part of such systems; more generally, tradeoffs are typically not part of the language used to describe distributed ledgers. The synchronization of so-called decentralized ledgers actually requires centralization or coordination across nodes. Further, 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 hierarchal. The point is that there are tradeoffs and choices that depend on context and goals.[3] This is the interesting challenge of design. The language used to describe DLT as if decentralized is misleading.

Second, the phase "absence of a central authority" naturally creates controversy among policymakers. Specifically, as stated in the Bitcoin white paper by Satoshi Nakamoto (2008), a pen name or alias: "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" (page 1). Or, to put this crudely and more provocatively, created to eliminate the need for central banks in the provision of money.[4] Similarly, cryptocurrencies and DLT are applauded by those in industry, featuring as they do the

[3] A stark example is the tradeoff between smart contacts on public ledgers and the industry collusion this allows, see Cong and He (Forth.).

[4] According to the BIS (2018, p. 91), "Decentralised technology of cryptocurrencies, however sophisticated, is a poor substitute for the solid institutional backing of money." In other words, the cryptocurrencies require independent and accountable central banks. Shin (2018) makes this case based on the economics of user fees for Bitcoin. The fees are high due to deliberately engineered congestion, but that, in turn, limits the use of cryptocurrency as a 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) payments systems run by the named and trusted institutions that maintain the ledgers and operating systems. Trustless, expensive decentralized systems such as Bitcoin are not needed, they argue.

words “disruption” and “disintermediation” which sometimes do not have an unambiguous meaning, and which do not appeal to regulators trying to protect the public.

1.2 Methods and philosophy of the paper

In the second part of the introduction we lay out more specifically the point of view of this paper.

First, the premise: Technological improvements in the design of mediation/intermediation systems could potentially, if done properly, allow economies to be more connected in a good way. Connectedness can come from new forms of mediation, though, as is already evident, not necessarily traditional intermediation through formal sector financial institutions, with agents connected to a common third party. The term “disintermediation,” presumably meaning the elimination of the profits of financial intermediaries and market makers, clouds the picture. 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 run from oral assignment to paper recording and written messages to e-messages and electronic registries.[5] These episodes are instructive, as some of the basics are the same as in the current wave of 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.

This paper distinguishes invention from innovation. Was Bitcoin with its blockchain and distributed ledgers a sharply defined invention, so that we can imagine new things are now possible that were not possible before? Or was Bitcoin an innovation, a part of a longer, slow moving process?

Arguably, it is the latter. 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 is also thought of as the originator of the smart contract implemented on distributed ledgers (Szabo1998). Indeed, some argue that Szabo is Nakamoto, though he denies this.

Related are advances in database management and distributed computing, dating from years earlier. Secure multi-party computation goes back to the late 1970s, not simply concealing content but also concealing partial information about data while computing with data from many sources to produce publically correct output. Many key components of DLT, ledgers and cryptography, are so familiar to us they seem mundane. Further, some of these components have very deep historical roots dating back hundreds of years.

Nevertheless, there are recent important contemporaneous innovations using these familiar components that have had large welfare gains.[6] Likewise, there remain gaps to be filled. This paper seeks to make these potential welfare gains and gaps transparent by being explicit about

[5] See Townsend (1990; 1987).
[6] A related, though opposite, example of slow moving innovations, pre-existing invention waits until there is a shock and then the system jumps to a new configuration, see Crouzet, Gupta and Mezzanotti (2019).

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 that are discussed in the computer science literature.

On the negative side are bugs and troubling episodes. The DAO hacking event is disturbing. Dao was a Decentralized Autonomous Organization (DAO), a smart contract on Ethereum Classic. A hacker found a loophole in the coding that allowed him to drain in the first few hours of the attack the equivalent of $70 million in Ether (Falkon 2017). The subsequent fork between Etherium and Etherium classic also illustrates that smart contracts are not necessarily immutable after all, begging larger questions about consensus and legal frameworks. Largely unregulated at first, the SEC in the US has recently 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, a gap this paper also tries to fill.

The conceptual framework adopted in this paper 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, how to do it.

Indeed, three possible metrics for this assessment can be distinguished, which, as a warning, can be easily confounded with each other. One, reduce obstacles to exchange and mitigate frictions. Two, have value in products/systems as commercially viable, actually be profitable. Three, again our basic standard: 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 failure of equivalence comes from 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; 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 not compensation, as noted. This may also explain slow adoption or failure to innovate even when the technology is well understood.

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, resource 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 so in this paper as we sort through language issues. A planner would be 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 come back to this below and revisit the issue in the paper.

A counter-example to a forced distinction between private and public ownership is the case of private clearinghouses as consortia 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 being algorithms for netting and settlement of checks among bankers, implemented by humans rotating around tables rather than computers, yet sharing much with codes and liquidity issues of contemporaneous e-payment systems.[7]

The language of the "planner" begs some key issues at the heart of DLT. How do we know that the code that implements the optimal social design is functioning correctly? This is where we enter the realm of distributed computing, running the code several times and independently. Computer science provides guidance from the standpoint of fault tolerance levels. Likewise, implementation and execution of a multi-party smart contract solution program in real time as a computer algorithm needs to take into account communication across nodes, the issues of capacity and latency, and the possibility of hacking. These problems can be mitigated by having one or several trusted or competing notaries, taking a step in the direction of trusted third parties, but with potentially decentralized features that take into account other lessons from mechanism design (Cao, Cong and Yang 2018; Janin, Gervais and Mamageishvili 2019).

The term "economy" here does not refer exclusively to nation-states. It means villages or towns as communities, regions within a country, and international cross-country economies. We are not shy about starting in villages, as small open economies, in a developing country, then moving to inter-regional and international money transfers, for example. Of great interest to this conversation are the economies of developing countries, which typically have limited financial infrastructure. This is illustrated with the data from the Townsend Thai project (Townsend 2016). This is a useful context in which to think about innovation and regulation, in part because so much is known from research about the underlying environment with its current institutions and markets, quite limited on the formal side, and in part because this is a useful step-back context in which to think about the advantages of various possible versions of DLT for those more familiar with heavily banked economies. There is much that is needed to be done in developing countries, where the potential for welfare gains is large and significant. However, obviously, innovation is going on in developed economies where gaps also remain, and that is also featured.

1.3 Outline of the Paper

The first part of the paper describes four key component parts of distributed ledgers: ledgers as accounts, e-messages and e-value transfers, cryptography, and contracts including multi-party mechanisms. Each component is discussed, evaluated, and illustrated through the context of historical and contemporary economies, with featured applications in both developed economies

[7] Relatedly, the public sector is not limited to reacting to potential and actual private sector innovations, sometimes with the purpose of preventing them, as with fintechs in some countries, not to mention Bitcoin. The public sector can also take a lead role as stated in the BIS committee report. An example featuring the Canadian central bank is detailed later in this paper. Sweden's past and present modernization of its payments system provides other salient examples of public innovation, e.g., the first issue of paper currency, the highest contemporary user of e-transfers, its recently implemented fast payments system SWISH, and perhaps in the future the first with e-coins (Ingves 2016). Sweden naturally features its public and private partnerships (PPPs).

and emerging market countries. A recurrent focus is the general equilibrium impact of innovations and welfare gains from innovations featuring key components, which does not require that all components be introduced at the same time.

In section 2, ledgers are described in the context of various emerging market 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 household/SME level transactions. 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, with cash flow and balance sheets, allowed for greater accuracy at the individual level. This forges a connection between DLT and financial accounts more generally, and with this link comes a greater potential for DLT, though not yet realized – what we might term consensus categorization. Illustrative applications of the use of financial accounts are discussed in Appendix A.

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. The latter in the context of latency makes it impossible to simultaneously achieve all three reasonable objectives: consistency, availability and desired non-partitioning. There are also impossibility theorems for consensus in asynchronous systems. Yet systems which are periodically synchronized require costly communication and do not scale up, as at some point every node is connected to every other. Trusted third parties solve this problem, but this centralization not only requires trust but raises the issue of data integrity. Those who can write can also accidentally destroy or corrupt data. There are also data security issues and cyber risks. Hybrid systems with partial meshes, in which not all nodes connect to each other, may be best for many applications. These emerge in practice but not necessarily as a deliberate choice.

An illustrative example is Lightning Network, shifting small transactions to a cryptographically secure "off-chain" environment, so that only large netting transactions need to be directly settled on the blockchain. Roughly speaking, database management systems have not paid much attention to incentives among parties with conflict interests. Distributed ledgers for business applications err toward keeping everything secret, not solving a design problem. We thus point an arrow toward where we could go. An example from economics with transactions costs from linking is given as an illustrative optimal hybrid system.

Section 3 features the second component: e-messages and e-payments. Examples of e-money are compared and contrasted, looking at Thailand, with a dominant use of paper currency and little e-value transfer, and at Sweden, a dominant use of e-payments and almost no currency left. 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 this gain. The Kenyan systems use a trusted third party, and so what is to some a defining characteristic of DLT, no trusted third party, may not in fact be needed; the gains from the e-transfer component are large.

Some caveats bring us back to the featured issues of the paper. For one, in the Kenyan context with Safaricom, trust is less obvious than it might seem when taking into account the larger financial system, with banks indirectly holding accounts, the Kenyan shillings that were turned in for e-tokens. For another, a relatively undeveloped part of value exchange is the infrastructure for liquidity, a common shortcoming in Kenya with its dual currency system and in New York financial markets in securities and central bank reserves. Ironically, most exchanges for DLT cryptocurrency that provide liquidity in and out of fiat monies and other crypto currencies rely on traditional technology with trusted third parties and not on DLT. That is, the tokens which feature absence of third party trust rely on such trust for market exchange. The exchanges that are DLT-based are slow and illiquid, highlighting an area where further work is needed. However, we will not fall into the trap of insisting that all parts of a system be decentralized. Hybrids may well be desirable, though this again is a direction for where we should be going, to figure it out.

Cryptography is discussed in section 4. One immediate point is that this component of DLT is not new either. Examples are given of Mesopotamian tokens in sealed clay envelopes as manifests to accompany the shipping of goods, and split tally sticks in Medieval England as records of debt that circulated as money before being cashed in. More recently, but well before Bitcoin, we have secure multi-party communication, IBM crypto express cards, public vs. private keys, and zero-knowledge proofs and protocols, and even the blockchain.

A description of how contemporaneous e-systems work without a universal trusted third party begs the issue of consensus. However, there is no one single way to achieve consensus: Bitcoin cryptography with poof of work; Federated Byzantine protocols with layers of trust for fault tolerance; proof of stake with voting rights based on coin ownership; and proof of authority. The consensus part of DLT is often taken as a defining characteristic, but clearly hybrid systems involve aspects of trusted parties, if not a single universal central authority. One can find in the literature revealing language of semi-decentralized operations or semi-trusted third parties. Further, as noted, some transactions in e-value transfers systems are done offline, as in a partially connected mesh. Bitcoin and consensus algorithms have attracted much of the academic and industry interest, and there are interesting economic issues, but this is not the main focus of this paper.

Section 5 presents the fourth component, contracts and multi-agent arrangements. This section begins with 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. Or, to put this more bluntly, this is a plea to use alternative language for situations that are more nuanced. Mechanism design can distinguish full commitment vs. limited commitment, incentive compatibility for actions, truth telling for messages about initial and interim unobserved states, reneging and thus a restriction to time consistency, and collusion and non-uniqueness in implementation, which have remedies in better design. Reputation can be formalized as the utility consequence of not performing.

The key features of smart contracts that overcome obstacles are presented, namely commitment in entering into an agreement and in carrying it out, the immutability of terms. Smart contracts operate on distributed ledgers. The language is as in financial accounts, with states for the

balance sheet and flows as in cash flow, executed with commands. But this language covers virtually any contract that can be written. The contracts and mechanisms that economists envision in order 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, past messages as immutable history, and promised utilities as key state variables. There are hybrid smart contact systems with lots of possibilities and flexibility: unique and non-unique consensus; single or multiple trusting or non-trusted notaries; public vs. private nodes; oracles for public information; and broadcast vs. selectively private communication. Thus the link of DLT to contracts is forged. One can be excited about distributed ledgers if only because of this smart contract aspect.

This section concludes with a discussion of similarities and differences between economics and computer science in concepts of trust and some comments on integration. In computer science, nodes are trusted or not, and designs are centered around having a sufficient number of trusted nodes. Yet the decentralized Bitcoin validation system is not incentive compatible. There is, however, encouraging common ground. Contracts with costly state verification are literally contracts executed with messages, in which over some range of states, costly validation of messages is not used, that is, messages are not needed. Costly imperfect messaging can be incorporated into the mechanism design, or even no messages at all, yet versions of the revelation principle still apply. Though there are impossibility results regarding consensus and common knowledge in the economics literature, more 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 in going forward.

The remainder of the paper deals with specific applications. As an overview, contract theory is used to derive optimal arrangements, constrained only by obstacles to trade, featuring how the various aspects of ledgers can deepen infrastructure. Mechanism design and monetary theory are used to study public versus partitioned ledgers and allow improvements in payment systems. Micro- and macro-prudential regulation, rather than resisting innovation, can be improved with the use of distributed ledger technologies. The overall goal is to provide blueprints for the ex ante optimal design and regulation of financial systems, including not only choices at the end points of centralized versus decentralized system spectrum, as in the hype, but typically hybrid forms in between.

In more detail, section 6 focuses on implications of smart contracts and mechanism design for innovation in the design and regulation of financial infrastructure. Various key issues can then be discussed: Whether ledgers should be public or rather be private with deliberate randomization and concealment; delegation of decision-making to a third party as an optimal endogenous outcome to allow back and front loading of incentives in contacts; and, how to mitigate bank and market runs through time stamped and immutable records of the histories of transactions. Limitations are also presented, as in latency on networks, but with remedies. The implications of mechanism design depend on the underlying environment and the problem being analyzed. There are forces pushing towards decentralization and partitioned ledgers, in some cases, or for centralization and a role for hierarchical third parties, in other cases.

The Thai setting is featured in section 7, making the point that context matters. In Thailand, as in other countries in Southeast Asia, the reach of traditional financial infrastructure is extremely limited. The ironic virtue is that it is easier to start from scratch in implementing optimal designs, as gains are large and there are no coordination or consortia problems. From Asian Development Bank studies we know there are large gaps in services for credit, saving, payments, and insurance. From the Townsend Thai project we know that informal risk sharing is good for idiosyncratic risk. There are in effect village money markets replete with credit chains that resemble the sophistication of advance country systems. But there are shortfalls in reallocating risk across villages and regions, and management of cash is inefficient in rural areas. Interventions can make a difference. A government village fund program had positive impacts on consumption, borrowing, investment, profits, and intermediation, yet this was uneven. Those without kin did not benefit nearly so much, and villages remain largely disconnected from one another.

The point is that innovations using 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, as a common platform can provide structure for contract competition, 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 given: EvryNet, an intelligent financial automation operation system that provides open source banking services and financial contracts to unbanked and under-banked populations.

This does come with caveats. Competition among providers can fail to complete the financial system due to unexploited complementarities and lack of coordination. Traditional regulation by sector and product can exacerbate this problem. The timeline 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.

The paper then turns in section 8 to dealing with the implications of the mechanism design and monetary theory literatures for the design of payments systems, arguing for a blend of traditional payment systems with the need for credit and insurance, and for regulation. Described are various systems. One with synchronized ledgers is a relatively centralized implementation of solutions to optimal mechanism design problems. A more decentralized but often equivalent system features the role of tokens, both single tokens and multiple tokens, as in colored coins. Equivalently, ledgers could be partitioned with voluntary and costly disclosure. Distinguished as well is whether information in payments systems should be public or private. In the latter, one does not want tokens to convey full histories.

Another part of section 8 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 and 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 from others with whom contemporaneously-matched payment parties have not been matched previously. Likewise, there are potential crashes with crypto currencies 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 information. An encouraging aspect of the featured examples is that there are only key instances requiring public information on trade.

This part of the paper on payments then concludes in section 9 with two innovations, both using DLT. One was run as an experiment by a central bank for commercial banks in an advanced country, e.g., Project Jasper in Canada, and is replete with sophisticated algorithms implemented as smart multi-party contracts for queuing and clearing. The other application embraces the full set of possibilities for payments 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. Velo features an optimized liquidity management layer efficiently searching for offsetting cross-country fiat balances for expedient clearing. MTOs can be viewed as agents with varying income and balance sheets hit by shocks, the needs for trade of their customers, hence themselves needing credit and insurance. Velo also features a digital reserve bank. A supplementary discussion of the value tokens and cryptocurrency, and activist policy, is in Appendix B.

The paper concludes with a summary, section 10. 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.

2 Ledgers as Financial Accounts

Here we focus on the most obvious component of distributed ledger technology, namely, the ledgers themselves. The unique advantage of DLT as a ledger is that it can be held in common and shared. As a corollary, DLT provides an additional accounting check beyond double-entry book keeping on the reliability of recorded transactions. Further, links of DLT ledgers to financial accounts open up a vision for future innovations that could have great power.

2.1 Statement of cash flow as a ledger: From paper currency to distributed ledgers in a few steps

We make the link immediately to standard accounting concepts. Cash and paper currency transactions can be recorded on ledgers. For village economies, this is done in Samphantharak and Townsend (2009), where the statement of cash flow as a standard corporate account is created. A transaction log operates on the Townsend Thai monthly survey data [8] and records cash transactions that each household *i* has with any other household *j*. As with Bitcoin, there is an initial state, who holds what coins, that is modified by a transaction to deliver a new state. A

[8] http://townsend-thai.mit.edu/

difference with paper currency, though, is that currency is held by the household as part of its balance sheet and is not public. Currency is an actual portable physical token but not an electronic entry. But the accounting concepts underlying the use of currency and the use of coins are exactly the same.

A formal statement of cash flow goes a bit further. It distinguishes the purpose of the cash outflow (or inflow) and thus records cash used for consumption and investment, for production, and for financing as in borrowing and lending. The households in the Thai villages do not keep these cash flow accounts. However, we constructed one for each surveyed household from the Townsend Thai data. Such cash flow statements are essential for the study of liquidity, hinting at a data use for DLT that we revisit throughout this paper. Recording liquidity will be an essential feature of the mechanism design of distributed ledger systems.

We can now link the statement of cash flow to the new language of distributed ledgers with only a few conceptual steps, as it is not hard to imagine how the new technology could map onto the current paper currency systems. First, one could imagine in principle that accounts could be kept on a common account or centralized common ledger. To establish a proof of concept, this is being done with the Townsend Thai data. It's "just" a new integrated database that we are creating. A key step here is whether the transactions that would be recorded on the ledger are consistent with each other, which involves rechecking the database to quantify discrepancies. If *i* transacts with *j*, is *j* in the database and, if so, is *j* also reporting a transaction with *i*? Yet this uncovers discrepancies, and this is one of the main things DLT can remedy. Some of these discrepancies could be innocent measurement errors or honest mistakes in reporting. Unfortunately, when we were first gathering this data two decades years ago, we did not have the conceptualization of the common ledger as a check on the gathered monthly data. These and other types of discrepancies matter, for example as in the New York financial markets. One purpose of the common ledger component of DLT is to have consensus and avoid subsequent validation.

Distributed ledgers could also be done in practice – in real time – if transactions were out of e-wallet coins and hence recorded (electronic measurement more generally comes up in the following section). Next, the common ledger could be created and distributed among the households so that each has access to the common account or to their own identical copy (subject to privacy, which we come back to under mechanism design below). We could call this new common integrated database a *distributed ledger*.

In sum, the idea of ledgers as cash flow is not new. Yet when put on a common database, discrepancies can be readily spotted and corrected. Approved histories can be thought of as immutable. A limitation is that only subsets of transactions might be recorded, in which case coordinated statements of cash flow could be incomplete. On the other hand, the vision for further utility comes with the creation of complete and integrated accounts.

2.2 Financial accounts as ledgers more generally

A unified, more comprehensive measurement of the financial environment is represented by the entire set of financial accounts. Specifically, measured transactions can be used to create formal

financial statements, not only the statement of cash flow, but also the income statement and balance sheet. More specifically, an initial baseline survey can be used to enumerate financial and real assets, held at the beginning point in the timeline of the survey.[9] Items on the balance sheet would be amount of currency held, land, and other assets. Indeed, when on a common ledger, this links to the idea in cryptography of using ledgers as a registry of secure property titles. Likewise, cryptocurrency on the balance sheet would be an asset, hence termed a digital asset.

Of course, liabilities can be measured in the same way. Subtraction of liabilities from assets thus determines initial net worth. Then there are transactions over time. A household, for example, surrenders currency to buy another asset. Currency can be used to buy consumption, an expense on the income statement, and income is received as revenue on the income statement. The difference in revenues and expenses is saving, which, along with incoming gifts and remittances, must be equivalent with an increase in net assets. The statement of cash flow is similar to the income statement except that for the income statement one typically uses an accrued income concept. Expenses are booked only when there is revenue, as in finance, to measure profits as the return on assets used to operate projects. The distinction between accrual and cash flow methods allows for the distinction between productivity and liquidity, and is often essential.

A key point to note is that a given transaction in the data can and typically will enter multiple times across individual statements. Thus, the change in balance sheet and income statement must be consistent with each other. The books have to balance. This is the idea behind double-entry bookkeeping, done for accuracy at the individual entity level, which was a huge innovation at the time the concept, invented by Luca Pacioli, 1494, who is considered to be the "father of accounting," but could have been Benedetto Cotrugli, even earlier, in 1458. Distributed and common ledgers to reduce discrepancies is another layer on top, and maybe just as important an innovation as double entry book keeping was.[10] This is a more accurate system than accounts individually. All this comes from the log of transactions.

Of course, to create the complete financial accounts from a distributed ledger, certain meta data have to be recorded as part of measured transactions. Again, as an example, the code which creates the accounts for the Townsend Thai data operates on the underlying transactions data, pre-programed to recognize from the questions to which the transaction answers are given, where in the accounts particular transactions should be entered. Any entity, e.g., a large firm, is doing this with its own proprietary financial accounts, so the firm at least knows the nature of its own transactions through the lens of financial accounts. In contrast, a distributed ledger which records only transactions without categorization cannot be used to create complete financial accounts. The middle common ground is perhaps the most interesting: Transacting parties record the categorization, and reconciliation seeks to make the categorization common. This could be an additional advance made possible with common ledgers: *consensus categorization*. The common accounts component of DLT could allow this to be done while maintaining privacy, just as DLT can remove discrepancies in trade.

[9] We do not have a natural beginning point, unlike the genesis state in cryptography for the first e-coins, hence we need the measurement of a baseline.
[10] For a discussion, see Kestenbaum (2012).

To sum up, complete financial accounts can have value for the accuracy of measurement, for analysis of data, hence for the households and businesses themselves. The log of transactions can have value for policy.

Appendix A contains two example applications that convey the value of enhanced financial accounts for policy in emerging markets and the US, to track the impact of tariffs or liberalizations and to measure liquidity flows to build micro-founded macro models. Appendix B shows how a log of transactions can be used specifically as a basis for an activist cryptocurrency policy of a digital reserve bank.

2.3 DLT vs. traditional database: Limitations of distributed databases to be recognized and incorporated in designs

A ledger could be viewed as a traditional database, in which a user can Create, Read, Update, or Delete (CRUD) (Ray 2017). In contrast, with distributed ledgers a user can read and retrieve data, that is, audit records, and a user can write by adding more data, append only. Newly proposed transactions have to be validated in some way. Likewise, past validated histories are immutable. There is no updating of past transactions and no deletion.

But with the decentralized system of distributed ledgers comes known database problems. A theorem in computer science (the CAP theorem) states that it is impossible for a distributed data store to simultaneously provide more than two out of the following three guarantees: (i) Consistency – where every read receives the most recent write or an error; (ii) Availability – where every request receives a (non-error) response, without the guarantee that it contains the most recent write as with consistency; (iii) Partition Tolerance – in which the system continues to operate despite an arbitrary number of messages being dropped (or delayed) by the network between nodes, that is, a partition, for one reason or another.

To highlight further, in the presence of a partition tolerance, one has to choose between consistency and availability. Further, even when the system is running normally, there is a tension between consistency and latency. Latency is the amount of time a message takes to traverse a system, how much time it takes for a packet of data to get from one designated point to another (Wikipedia 2018b). These choices have to be part of the design. Bitcoin and other consensus algorithms make choices about how to address these for transactions in tokens.

More specifically, we come to the Fischer Consensus Problem of distributed computing (Fischer, Lynch and Paterson 1985), though we need some definitions first: “In computer science, synchronization refers to one of two distinct but related concepts: synchronization of processes, and synchronization of data. Process synchronization refers to the idea that multiple processes are to join up or handshake at a certain point, in order to reach an agreement or commit to a certain sequence of action. Data synchronization refers to the idea of keeping multiple copies of a dataset in coherence with one another, or to maintain data integrity. Process synchronization primitives are commonly used to implement data synchronization” (Wikipedia 2019). Otherwise the process is asynchronous.

Fischer proved that it is impossible to guarantee that any asynchronously connected set of communicating nodes can agree on even a single bit value – a devastating result. On the other hand, the Fischer consensus problem can be resolved simply by synchronizing from a single point, which is what the bitcoin mining protocol essentially does. However, doing so introduces a single point of centralization, ironic given the decentralized connotation of DLT as emphasized in the introduction. This centralization in turn can cause scaling problems, as every node must be connected to every other to achieve consensus, and the costs of messages rise exponentially with the number of nodes. This is one of the causes of congestion in Bitcoin: slow speed. Computer science/distributed systems bounce between these problems: CAP in asynchronous systems, and scaling and fault tolerance in synchronous systems. These features should drive choices as part of a constrained optimal design (Jacky Mallet, personal communication, January, 2019).

Among these choices, the existence of a trusted third party can greatly enhance speed and lower costs. Examples in practice include the use of a trusted third party to facilitate trade by eliminating reconciliation problems, as with the Digital Assets innovation for the Australian stock exchange. Adopting the language of Casey et al (2018), these systems are known as "permissioned" (or "private") blockchains, with a limited set of entities, or even a single organization, allowed to write to the blockchain. This can reduce scaling problems. Permissioned blockchain and partitioned ledgers with private information do make decentralized consensus harder.

Another example highlighted in Casey at al (2018), the Lightning Network, aims to greatly reduce cost and time constraints by shifting small transactions to a cryptographically secure "off-chain" environment, so that only large netting transactions need to be directly settled into a resource-constrained blockchain. With Hyperledger Fabric, a permissioned blockchain, a third-party auditor or regulator can obtain provably correct answers to queries about the system as a whole using zero-knowledge-proof concepts. Centralized DLT exchanges for crypto currency have "relayers," application interfaces that allow users to trade in a decentralized manner (Bronstein 2018).

Mallett (2009) compares strictly hierarchical/client servers with fully connected mesh networks and then speaks to the advantages of partial mesh networks. The point is to compare, and potentially select among network designs. The military uses a hierarchical system which suffers from lack of integration of local information but has the advantage of low latency, as one way commands from headquarters are obeyed. These hybrid designs need to be integrated further with economic systems, and, indeed, the industrial and management organization may be endogenous with the design selected. As a suggestive example, Townsend (1978) establishes with simple transaction costs arguments that optimal risk sharing arrangements partition agents into subgroups, despite ever decreasing per capita costs and ever increasing gains from having all agents in one mutual fund, due to portfolio diversification and the law of large numbers. More generally, the economic issue is whether to have OTC markets, centralized platforms, or a hybrid in between those.

The risk in having centralized control of a database is that anyone with sufficient access to it can destroy or corrupt data, so users are reliant on the security infrastructure of the database operator and have to trust those with write capabilities. In contrast, distributed ledgers use decentralized

data storage, in the sense that the ledgers are distributed among users. With cryptographic rules for change, security is inherent in this structure; there is no single copy. A key property of blockchains such as Bitcoin is that they do not rely on a single trusted third party as trustee or notary to intermediate transactions. The blockchain network enforces execution, giving this a social aspect. This is what Nakamoto (2008) meant by a system without a trusted third party.

One should conclude this section with the thought that there are tradeoffs in design: Which system might dominate is a function of the environment, including DLT technology, and goals, including the economics of outcomes.

3 e-Payments and e-messages

In this section we compare and contrast Thailand to Sweden and then put Kenya in between as a home of an impressive innovation in e-payments with measured, documented welfare gains. The point is that gains can be large even for mundane systems using components of DLT, gains that remain to be harvested in many low income and developing countries.

3.1 Thailand and the predominate use of paper currency

In emerging markets, such as those in Southeast Asia (SEA), 55% to 90% of all payment transactions are conducted through physical cash payments. The ratio of currency to GDP is 11. 37% for all of Thailand, the fourth highest among countries listed in a 2015 study (Rogoff 2016). Based on currency and coin outstanding, and measurement income and consumption in GDP, Thailand is estimated to have individual per capita currency holdings equal to 7 months of consumption, on average. Asian countries have a ratio of currency to GDP relative to the rest of the world, generally.

Alvarez, Pawasutipaisit, and Townsend (2018) use data that were gathered monthly, with consumption bi-weekly, in 16 villages, half in agrarian provinces in the Northeast and half in more industrialized or cash crop provinces near Bangkok. They find that typical households running small businesses use paper currency for small and large transactions, spending on consumption in normal times, with spikes in unusual times coming from durable goods, rotating savings and credit association (ROSCA) transactions. They receive paper currency from income in normal times, with spikes coming from land sales, loans, and gifts. The costs of cash mis-management are calculated to be of the order of magnitude of 2% to 9.5% of monthly consumption. The top end of that range corresponds with fitting a Miller-Orr (1966) model for business adding an ingredient, occasional free transactions, as for households in Alvarez and Lippi (2009). The businesses in the Thai setting are household-run small and medium enterprises (SMEs). The calculation uses an optimized value function from a dynamic program, the minimized discounted present value cost of holding cash. The lower end of the range of costs corresponds with the interest rate on bank accounts multiplied by average cash holding. Costs are non-trivial even at the low end of the range. This is far higher than estimated costs of business cycles, for example, and does not consider the costs of printing and distributing the currency. There are gains to be had from moving away from paper currency to electronic systems which could allow payment of interest.

3.2 Sweden as an almost cashless society

The Riksbank began in the 1980s to make systematic efforts to shift a large part of the cost of managing paper currency to the private sector so that the private sector would bear and internalize cash management. The number of central bank branches was successively reduced, from one in each province to 20 nationally, and now they're down to one cash distribution center staffed by eight people. Price distortions were corrected as banks were asked to pay transport costs and also received interest on cash held overnight. The Riksbank's role is limited to printing, transportation to the single cash center, and the destruction of defective and obsolete notes and coins. The private sector has coordinated in allowing inter-operability: one credit card network for clearing, one single bank ID, and one mobile application (Swish) for low-value payments, with the single central bank cash center operating as a decentralized wholesaler between banks and the Riksbank (Ingves, 2016; Skingsley 2016).

Sweden is currently down to less than 2.5 as a ratio of paper currency per GDP, one of the lowest in the world. It is a highly digitized country with most transactions occurring in electronic form, under debit cards, credit cards, and e-transfers, as reported in a Riksbank survey (2018). Card payments per person are among the highest in the world. There are various electronic clearing systems with financial institutions as key nodes intermediating payments: 160 million transactions yearly in the data clearing system (owned by Swedish Bankers' Association); 180 million transactions in the Swish system in 2018; 800 million transactions yearly in the Bankgirot system (batch); and 2.2 billion card payments.[11] Not all of these data are public but they exist in electronic form, obviously.

3.3 Kenya: M-Pesa as an e-money innovation with large social gains

Kenya lies midway between currency-intensive Thailand and virtually-cashless Sweden. In this context, e-money can have social value, especially for certain segments of the population. More generally, the potential of new technologies to transform traditional systems is significant, and e-money systems have been endorsed by the G-20 as an opportunity to build financial markets, constructing new financial systems that increase financial access for large unbanked populations in developing countries (G20 Research Group 2013).

M-Pesa is an e-money implemented by Safaricom. Households can go to a company agent and exchange Kenyan shillings for cell phone credits, which can then be used for purchases or money transfers. For example, a migrant worker in Nairobi can send cell credits back to relatives in the village, which an agent there cashes out back into shillings on request. This is a functional and comprehensive value transfer system in the context of the actual rugged environment of the economy that gave birth to it.

[11] Ironically, The Swedish Post and Telecom Authority (2017) writes in a report that currency infrastructure is being depleted. Cash is becoming less accessible. Many banks have stopped managing cash with ATMs, which are disappearing. In a growing number of stores around the country, it is not possible to pay with cash today. There is a common worry that people who use only currency may become isolated outside the payment system unless they find the support and help they need. See also Erlandsson and Guibourg (2018).

M-Pesa functions as a "stable coin," that is, with a fixed local exchange rate to fiat currency. Notably, the exchange between cell accounts and Kenyan shillings are 1-1 apart from a schedule of pre-specified transaction costs. These costs are quite low: 6% for tiny values and falling to less than 1% for larger values. This is one-sixth of Western Union's rates and one-twelfth of Postal Pay's. In Kenya, Safaricom is a single trusted third party, keeping all of the accounts, though customers can see their own accounts. The technology uses relatively inexpensive cell phones. Adoption in Kenya among those without bank accounts rose from 20% in 2008 to 90% by 2014.[12]

There is social value to M-Pesa. Studies have shown that M-Pesa aids in economy-wide risk-sharing (Suri 2017). The staggered nature of the roll-out allowed a quasi-natural experimental evaluation. Consumption is smoother, more immune from households' specific income shocks. Value can be transferred from households running budget surpluses to those running deficits, for example, and transferred across regions. Mobile money has also allowed a more efficient allocation of labor and resulted in a meaningful reduction of poverty (Jack and Suri, 2014).

Non-bank fintechs such as M-Shwari use M-Pesa to lend to this low-income population, accessing both the record of transactions in a scoring system and using M-Pesa as the payment/repayment medium. There are now over twenty digital credit providers in Kenya.[13] The point is that the scoring systems are using the transaction data that are recorded in M-Pesa.

Yet to be emphasized here, and key to the discussion earlier, Safaricom does not refer to its system as a distributed ledger system. The ledgers are not distributed. They are owned and operated by Safaricom with customers permitted to see individual pieces and make associated approved transactions. Put another way, customers see Safaricom accounts for balances of their ownership of M-Pesa cell-credits and can verify transfers. Customers could but likely do not keep accounts of their currency. Safaricom has complete accounts for cell-credits for all customers and is the trusted third party running that database. Likewise M-Pesa is not categorized as a cryptocurrency.

This can be fine. In Kenya, it has worked well so far. One does not need to incorporate all the components of DLT in order for implementation of a subset of components to have value.

One cannot help but note, though, that Safaricom could in principle lend its funds, and thus would look ever more like a bank, with fractional reserve banking. Actually, its funds are put on deposit in commercial banks, and the interest is contributed to charity. This is what allowed the advent of M-Pesa and approval by the Ministry of Finance, that Safaricom not be classified as a commercial bank. Yet in countries such as Kenya, bank runs and failures are commonplace. For this reason, as they and others began to think about these risks, Safaricom switched to putting its funds in multiple banks. The point: There are limits to trusted third parties, if not direct then indirect. In some contexts third party trust is a real issue.

[12] See Jack, Suri, and Townsend (2010) for this discussion.
[13] See Carlson (2017) for a study of one of them, and more recently Bharadwaj, Jack and Suri (2018).

3.4 The role of broker-dealers, shortages, thin markets, and common concerns about liquidity in various disparate contexts: Concerns to be incorporated in subsequent design

Dealers in private e-money and paper currency face shortages of liquidity, of one object or another, and this can show up in various ways. Returning to M-Pesa and the example of Kenya, Jack and Suri survey households that use M-Pesa and the agents that are contractual spatial outposts for Safaricom (Jack and Suri 2011). As reported in Jack, Suri, and Townsend (2010), agents run out of one object or the other on a regular basis. Over 60% of agents run out of e-money anywhere from approximately once a month to multiple times a day. Likewise, close to 50% of dealers run out of Kenyan shillings. Recall that the exchange is guaranteed to be 1-1 with no variation in prices or transactions fees. Shortages typically occur with fixed prices, of course. In other situations, one might imagine varying prices but with the potentially lingering problems of thin markets (that is, not many participants).

New systems emerge to cope with these challenges. In Kenya, there are transfers and borrowing/lending among Safaricom agents in a kind of informal market, including gifts. Inter-agent markets could be formalized and potentially improved upon, though of course subject to the obstacles of the environment. Here, with costly transport of fiat paper currency, spatial ingredients play an inherent role. While the e-part is virtually instantaneous, paper currency has to get to the agent. No formal system has as yet been designed.

In a very different context, in value and location, but quite close conceptually, consider the New York financial market system. There is inter-bank borrowing and lending of excess reserves, and broker-dealers provide liquidity to this market. As documented in Cocco, Gomes, and Martins (2009), the relationships of traders with dealers who have low correlation in liquidity shocks allow insurance against shortage of funds. Lagos and Zhang (2018) note the role of liquidity in monetary policy.[14] In this sense, the shortages are a driving force. In this New York context though, unlike the Kenyan example, the market is entirely in e-objects. There are continuing innovations; see Li and Schürhoff (2012) and Hendershott and Madhavan (2015). Still, problems remain and there is considerable scope for improvements in the e-infrastructure systems that are used today.

Likewise, the functioning cryptocurrency exchange platforms should be integral rather than peripheral to the debate about tokens and distributed ledgers. Brokers can provide liquidity through implicit insurance but are potentially charging abusive markups and committing fraud. The most popular cryptocurrency exchanges, such as Coinbase, Binance and Kraken, are implemented as centralized exchanges, thus again the same irony. These crypto exchanges also rely on traditional technology, where customers can access and trade using e-mail and simple passwords. This is what has led to hacking episodes. However, contemporary decentralized exchanges using DLT technology, which include 0x, Protocol, AirSwap, and OmiseGO, are difficult to use, have limited capability, and are low volume (Glazer 2018).

In conclusion, innovation in financial infrastructure may be possible. On the other hand, there may not be inherent contradictions as tokens and exchanges fulfill different economic functions, one to provide record of ownership and the other to facilitate exchange. It is an advantage of

[14] See also Lagos and Wright (2005).

economic analysis that we can draw these distinctions, getting beneath the hype. A core issue is whether or not new distributed ledger-based trading platforms lower clearly identified costs.

4 Encryption

Cryptocurrencies are e-message systems similar in their basics to M-Pesa but differing in the amount of cryptography used, diverse methods for validation of transactions, and the public nature of the ledgers. One might have thought this cryptography component is what makes DLT new, but that is not the case.

The use of cryptography goes back to at least to the Mesopotamians, where it was key a key part of their economic and messaging systems. TCP/IP protocol is reminiscent of that, where pieces are put into "envelopes" where they are encased by beginning and ending bits, and then disseminated.

About 7500-3500 BC, in Mesopotamia, the code for communication consisted of tokens of some 6 types, distinct shapes representing particular commodities. Then around 3500-3100 BC, new complex tokens were covered with lines or dots conferring qualitative and quantitative information (Schmandt-Besserat, Denise. 2014). Eventually tokens were put in clay envelopes as a manifest for shipping goods, sealed so that tampering with the manifest would be evident on arrival, as would theft of cargo, as a check of the actual cargo inventory against the manifest would reveal. Writing on the clay manifest envelopes to convey contents of the message (and cargo) is what gave birth to Cuneiform writing (Trubek 2015).

Another historical example: Tally Sticks as messages and proof of contract emerged as monies in Medieval England and were used for centuries, including as a money means of payment (Harford 2017). Tallies were a way of recording debts with a system: Willow sticks recorded the original debt transaction and then were split in half; with a distinctive grain, the two halves would match only each other, the requisite proof. The lender's half of the tally stock was used as a safe and convenient form of payment (hence the word "stock"). When cashed in, the two halves were checked (hence the word "check").[15]

This historical discussion leads into key issues in computer science and algorithms. A central component is public-key cryptography. Keys come in pairs, public keys which may be disseminated widely, and private keys which are known only to the owner. The generation of such keys depends on cryptographic algorithms based on mathematical problems to produce one-way functions. Effective security only requires keeping the private key private. In such a system, any person can encrypt a message using the public key, but that encrypted message can only be deciphered with the private key. Roughly, it is known who sent the message but not what the

[15] "Those tally sticks met an unfortunate end. The system was finally abolished and replaced by paper ledgers in 1834. To celebrate, it was decided to burn the sticks - six centuries of irreplaceable monetary records - in a coal-fired stove in the House of Lords, rather than letting parliamentary staff take them home for firewood. Burning a cartload or two of tally sticks in a coal-fired stove is a wonderful way to start a raging chimney fire. So it was that the House of Lords, then the House of Commons, and almost the entire Palace of Westminster - a building as old as the tally stick system itself - was burned to the ground." (Harford 2017)

message is. Public key algorithms are fundamental security ingredients assuring the confidentiality, authenticity and non- repudiable electronic communications and data storage. They underpin various Internet standards.

IBM was putting crypto express cards into its mainframes by 2009. A key distinction now is the perfectly opaque computational systems in which participants cannot look in at all versus zero knowledge proof systems which do allow participants to look in and find proof that each subset of code ran as intended.

4.1 Bitcoin

Transactions in Bitcoin are encoded messages. The keys ensure no one can transact on someone else's ID, impersonating a node. Because the message or transaction can only be created with the key combination, it is known the spender wishes to unlock and spend the coin, plus there is commitment to the transaction, so it cannot be undone or reneged upon later.

Double spending would be possible if two messages from a given node were able to spend the same coin. In fact, with internet latency, the problem mentioned earlier, it would be hard to know which transaction came first and should be valid in principle, as time stamps are not necessarily chronological. For Bitcoin, blocks of individual transactions are broadcast to the entire network, or at least to those listening among the community of users. Blocks economize on messages and costs of validation. Anonymous nodes verify blocks of new transactions of which they are aware, transactions that have been accumulating as candidates on individual copies of ledgers over the previous 10 minutes or so. The messages broadcast to the community of users consists of these new series of transactions on the block, concatenated with a randomly chosen number, and then hashed into a 256-bit encrypted message. (Bitcoin is credited with the use of public randomization devices.) Hashing is essentially a one-way function. To dis-encrypt the hash back into the message, one needs to proceed by trial and error, though solutions are evident once found.

One potential problem is that nodes as bad actors could propose the latest version of the ledger that they would like to become immutable, e.g., containing the second of the double-spend transactions, while the first was already used to acquire something else. To thwart this, under the Bitcoin system, it is as if one node were selected at random to certify a current new candidate ledger; and key to Bitcoin, this certification requires time and energy. This limits entry into validation. A proof-of-work algorithm requires a selectable amount of work to find the random number that, when added to the set of transactions, created the hash. Difficulty is controlled by having to match the first string of bits of the hash only, e.g., 30 initial digits as zeros. The discovered random number is then added to the bottom of the block as the proof, a certification of work done that all can confirm easily.

In fact, a group of miners are all simultaneously running code to decipher, and the miner that finds the solution first finds it essentially randomly, as everyone is attempting by simple trial and error. The premise of Bitcoin, and computer science more generally, is that most nodes are honest. Temporary multiplicity or fraud is possible if another branch containing new blocks is created. But the convention is that the longest interim chain is considered to be the valid one, and

to reinforce this, miners have incentives to mine the longest chain as they are rewarded in Bitcoin only if the block of transactions they validate becomes, eventually, part of the immutable history. There could be threats of a double-spending run or extortion if a group of miners acquired 51% of computing power. All of the blocks of ledgers in a chain are linked together, given that the top of each ledger contains the hash of the previous ledger. Cryptography with proof-of-work by these miners, implemented up to six times, ideally sends the probability of malfunction or fraud asymptotically to zero.

There is no single trusted third party in this public distributed ledger. Yet that is a bit misleading. The fundamental premise is that most nodes are trustworthy, and incentives of the mechanism which in fact is not collusion-proof are relied on heavily in the algorithm. The synchronization also requires a kind of centralization, as noted earlier. Some recent economics literature has focused on critiques of Bitcoin.[16] Others argue the platform is more robust and malleable than it might seem, (Casey et al 2018).

Many new coins have emerged with alternative verification systems which, like Bitcoin, also do not rely on single trusted third parties.

4.2 Ripple

Closely related to the Kenyan environment where we featured transfers of value from city to town are cryptographic e-money systems for the purchase and sale of national currencies, including transfers of value across international borders. It is at present relatively costly and slow to do this through commercial bank systems. Ripple is a for-profit entity that has revolutionized this environment by working with mainstream large financial institutions.

Users of Ripple buy the coin "XRP" and make payments among each other using cryptographically-signed transactions denominated in XRP. But transactions in fiat currencies and other objects are frequent and tokens represent these fiat monies on the ledgers. Ripple is essentially a payments protocol for fiat money through these tokens. XRP has value, but part of the reason for the coin is a design to prevent hacking. Flooding servers with innumerable transactions causes the cost of transactions to rise exponentially and so is not a cost-effective strategy. Put differently, simply being malicious can be exorbitantly costly.

For XRP-denominated transactions, Ripple can make use of its internal ledger. Its primary function is inter-bank transfer, using a decentralized trust system on a permissioned ledger. Banks are key nodes trusted by users that hold funds and issue balances on behalf of customers. Users have to specify the other users they trust and to what amount. When a payment is made between two users that trust each other, the balance of the mutual credit line is adjusted, subject to limits set by each user. The user paying out to a customer in the country of destination is effectively lending value to the originator of the transaction, trusting to get it back. This is similar to the medieval *hawala* system among merchants and correspondent banks, a popular and informal value transfer system based not on the movement of cash, or on telegraph or

[16] See for example, Athey et al 2016, Budish 2018, Foley, Karlsen and Putnins 2018, Griffin and Shams 2018, Prat and Walter 2018, Biais et al (2018) and Cong, He and Li (2018).

computer network wire transfers between banks, but instead on the performance and honor of a huge network of money brokers (known as *hawaladars*) (Wikipedia 2018a).

In order to send assets between users that have not directly established a trust relationship, the protocol tries to find a path between the two users such that each link of the path is between two users that do have a trust relationship, again subject to caps.[17] In principle, if there is not a trust path, then XRP can be used to balance the transaction in the other direction. Likewise, credit lines readjust to earlier levels if there are flows among trusted users going the other way. Outstanding IOUs are on a public ledger of accounts.

4.3 Stellar

The Stellar Development Foundation is a not-for-profit organization that provides greater access and inclusion by connecting people to low-cost financial services. Stellar is open source and a public ledger: it sees expansion into underserved populations as its primary mission. Stellar does not rely on mainstream financial institutions. Individual users are not necessarily large financial institutions such as bank, are allowed to be non-bank and small, such as money transfer operators.

Stellar Consensus Protocol (SCP) (Mazieres 2016) utilizes a Federated Byzantine Agreement that allows more universal access, akin to the internet, as a way to interact among strangers. The aptly-named Byzantine Generals Problem is illustrative and comes from distributed computing. A subset of a group of generals are potential traitors (Lamport, Shostak and Pease 1982; Robinson 2009), bad nodes either malicious or sending error ridden messages. The decision to be made by the generals is whether to attack or withdraw, that is, approve transactions or not, and this requires consensus. The generals are exchanging messages with each other. If the number of potential traitors (faults) is known, and all other nodes tell the truth, then, intuitively, cross-checking a sufficient number of messages is sufficient. However, the degree of difficulty is a function of primitive assumptions. If a coordinator is assumed to always send honest messages, then things are easier, as one only needs to check a small sample of other nodes – a bit of centralization. If the coordinator could be faulty, more cross-checks are needed. If nodes somehow cannot lie about what they have heard, that is helpful, though this requires more rounds. If nodes start repeating what they have heard from others, the group may have to abandon the primary coordinator node.

In the Stellar protocol for validation, each participant names others it considers important and requires that the majority of these others agree to any batch of transactions. Yet those other important participants do not agree until the participants they consider important also agree. Each transaction requires a majority of nodes designated as "important" by both traders (Ray 2018). The system is thus designed to be open to new entrants naming their own trust network. Unlike Bitcoin, Stellar forms a decentralized consensus among a group of nodes that are transitively connected to each other by trust. This is a way to limit the congestion and time cost of cross checking messages.

[17] See also Dandekar, et al (2012).

Stellar, like Ripple, can transfer value across virtually any object, e.g., creating fiat money tokens that are then cashed out. Stellar can also transfer the tokens of others who wish to design their own platform with their own coin, linking to Stellar for value transfers.

Stellar uses anchors and market makers to intermediate the exchange of individual parties. An anchor is typically a highly regarded financial institution, namely a commercial bank. To make an international exchange transaction, for example, a customer makes a deposit with the anchor in fiat currency, of the country of origination for example, and thus issues an IOU, a debt, to the depositor. This is termed a base account. The fiat money is then converted 1-1 to a token equivalent amount, apart from fees, and these tokens are termed assets, in a base account. The anchor contacts a market maker, a user, who, like broker dealers in other contexts, posts bid-ask spreads and carries some inventory of a variety of assets. If not holding the fiat token of the country of destination, a second broker dealer is found as a go-between. An algorithm searches for minimal paths. Potentially finding the optimal path is an NP hard problem, but typically only two steps are used, at most. The user broker-dealer needs to establish a trust line with the anchor to ensure itself the deposit with the originating fiat money is there and that the asset is backed in that sense.

4.4 Proof of stake based cryptocurrencies, and others

Proof of stake (PoS) is an alternative for validating the transactions executed by miners. In contrast to proof of work, miners put up a number of their coins on a block. A miner is then chosen deterministically by an algorithm based on such posted coins and how long owned, among other criteria. Other variations include Delegated Proof of Stake and Weighted Proof of Stake.

Related is Practical Byzantine Fault Tolerance which allows a distributed computer network to function as desired and correctly reach a sufficient consensus despite malicious components (nodes) of the system failing or propagating incorrect information to other peers. This has been extensively researched and optimized with a diverse set of solutions in practice (Curran 2018).[18] A final variation is Proof of Authority, relatively centralized, as the name connotes.[19]

In sum, the mostly decentralized part of cryptocurrency transfers is the validation procedure, which does not make use of a single, common trusted third party. To some, this is the hallmark of DLT. However, it too has deep historical roots.

4.5 Distributed Consensus

The roots of distributed ledgers come from distributed consensus, which has been studied for decades in computer science. The traditional application is reliability in distributed computing systems. Narayanan (2016) defines a distributed consensus protocol: There are *n* nodes that have an input value. Some of these nodes are faulty or malicious. A distributed consensus protocol has

[18] For recent implementation and a clear statement, see Tendermint https://tendermint.com.
[19] For a listing of each category with associated coins and ledgers, see https://hackernoon.com/a-hitchhikers-guide-to-consensus-algorithms-d81aae3eb0e3.

two properties: It must terminate with all honest nodes in agreement on value; and the value must have been generated by honest nodes. This, again, is the basis of secure multiparty computation. Ripple and Stellar are, in this language, distributed consensus protocols but with identified nodes, which Narayanan distinguishes from Bitcoin where nodes are anonymous, and so he does not want to refer to their tokens as crypto currencies.

Another clear definition of distributed ledgers: Multiple, distrusting organizations who run a protocol to create an append-only log in which all participants can verify the integrity of entries appended to the log (Neha Narula, personal correspondence, January, 2019). This definition has the advantage that it does not refer to trusted third parties, as otherwise it is ambiguous whether there is one common trusted party, or layers of trust. It is sufficient that there be some distrust, and note that it is among organizations, not nodes, recognizing conflict of interest. For example, Digital Asset designed a ledger for the Australian stock exchange which falls under this definition. A buyer and seller of an equity agree to a trade on the Exchange, but before it is recorded, the Exchange as an entity posts an encrypted message, sending an alert to each party for them to verify and confirm, each party then putting the trade in their proprietary Contract Store, and finally the entry is confirmed as an immutably (encrypted) item on the common ledger. Note that the Exchange is playing a role like a notary as a named party. The buyer and seller do not trust one another and do not have to consent to the posted trade listed on the Exchange, either. Once all three parties agree in the protocol, we are done. Notably, the trade does not require a larger consensus protocol the way Bitcoin, Stellar, Ripple and others do.

Finally, zkLedger is a protocol which allows outside auditors, for example for regulation, to verify accurate information while protecting privacy through zero-knowledge proofs of cryptography (Narula, Vasquez and Virza 2018).

5 Contract theory and smart contracts: Mechanism design

This section on contract theory provides the opportunity to review and be specific about the various meanings of the word *trust*, as with trust in third parties, needed or not. We can distinguish trust in contracting with a time lag in performance; parties trust each other; parties trust 3rd party intermediary; and parties trust the incentives under which their counterparty is acting, to name some examples. The relevance of contracts comes to life and is super-charged with the possibilities created by the distributed ledger technology which allows “smart contracts”. Likewise, we can be more precise in this section about underlying frictions and how smart contracts can deal with these. Finally, we return to the diverse perspectives of mechanism design and computer science, yet find some unexploited common ground which could be used in subsequent designs.

5.1 Contract theory and the meaning of trust

The simplest of contracts between two parties is the purchase and sale of a commodity or asset. Indeed, if done as a spot market exchange, then we could avoid the language of contracts altogether. But here it is clarifying to think of the contract as an agreement of the buyer to

surrender cash and the seller to surrender the commodity. More generally, the buyer is debited and the seller is credited. When there are lags between the time of the agreement to trade and the eventual payment, then there can be issues. If there is no collateral, then the parties trust each other to carry out their part. This is a full commitment contract, under which promises are made and honored, as is commonly assumed in the contract literature, hence, one notion of trust. However, an alternative is for a trusted third party to stand between the traders making trade possible among strangers. PayPal and Alibaba/Ant Financial are examples.

With borrowing and lending, trust issues are even more apparent, as the lender either trusts the borrower to repay, or otherwise there should be collateral backing the loan put in escrow. In the latter case, a transaction among strangers is possible, implemented with a conditional if-then statement to define what happens with and without default (see below for smart contracts). (see Geanakoplos 2003; Kilenthong and Townsend 2018 for a literature on securities as contracts which embed collateral as a contract characteristic).

When there is private information known by only one of the parties, then others may trust the informed party to send a message and reveal it. Essentially, this is assumed in much though not all of the computer science literature, though there remain issues of capacity and latency of message systems, creating challenging problems in their own right. In contrast, in the contract mechanism design literature, the contract has to be designed in such a way as to induce the informed party to tell the truth.[20] The message is endogenous. Mechanism design enables agents to believe in the truth of others because they "trust" or believe that others will respond rationally to incentives built into the system. The mechanism or game need not be imposed per se; it can be derived as constrained-optimal, that is, given information constraints. Likewise, where there is unobserved effort or other unobserved actions, actions are recommended. Either a party is trusted to take appropriate actions or otherwise must be given an incentive to do so. This is the classic moral hazard problem (Harris and Raviv 1979). Likewise, though information is private, constrained-optimal arrangements cannot only allow reporting but also some revelation, hence information becomes public. These latter mechanisms are not black and white, that is, not all private nor all public.

A third example: Output from a project is privately observed. Verification of project output by a lender or insurer is possible but costly. This is exactly costly verification of messages and hence in principle could be part of computer science designs. Over a range of outputs, repayment of a loan is constant, resembling debt. Actual outcomes need not be known, and in effect there is no message. For low outputs, however, claims are not trusted but verified at a cost, as if validating financial statements, for example. This is costly state verification (Townsend 1979, Gale and Hellwig 1985). The notion of verification, however costly or infrequent, is central to execution of actual contracts, a point made in Gans (2018).

The general point here is that the word "trust" could be confounded with information and incentive issues. But, on the other hand, from contract theory we know there are a variety of uses of the word trust that are separate from one another. In particular, incentives are separated from full commitment as distinct concepts.

[20] See the revelation principle in Harris and Townsend (1981), Myerson (1982; 1986), Dasgupta, Hammond, and Maskin, 1979.

Limited commitment has a specific meaning in contract theory. One party may wish to withdraw and go their own way, in which case contracts with limited commitment make sure that future rewards do not fall below certain thresholds, to retain participation. This is loaded into the contract. In the autarky version of this, there is implicit trust that a banishment penalty can be imposed were withdrawal to happen. Ironically, is takes this off-equilibrium commitment to deal with the original limited commitment problem.[21] Likewise, in the competitive market version of this, we trust a party not to break the relationship even though it might be mutually advantageous to one of the original parties and an entrant. For example, this can occur when one party of a contract is drawn off by a competitor or a third party that has not entered in the multi-agent arrangement (Jacklin 1987). Relatedly, both parties to a contract may wish to renegotiate and start over, but that would be bad for incentives ex ante. So, either the parties trust that will not happen, as in full commitment models, or a time consistency constraint is appended to the underlying contract problem, so that time consistency is loaded in and there is no temptation.[22] Enhanced commitment to deal with single party or multiple party deviations is possible under the smart contract technology.

Otherwise, if contracts are incomplete and penalties cannot be imposed internally as part of the contract, then there is a way to formalize a role for reputation when there is no trust, and, on the other hand, a way to design and commit to optimal social penalties through a scoring function (Lehnert, Ligon, and Townsend 1999). This kind of indirect scoring does improve commitment in the broader sense. The notion that we trust third parties because we know they are worried about their reputation is formalized in this literature. It is not taken for granted. It is modeled and part of the design.

We have been discussing how to find constrained efficient ex ante agreements, via maximization subject to truth telling constraints. Though with assumptions a solution exists and constitutes a valid (Bayesian) Nash equilibrium, typically the solution is not immune to collusion. That is, players need not act in the Nash sense of taking the strategies of others as given, but rather could all act in concert to enhance their welfare at the expense of the principal. An example would be auctions where agents can collude in their bids. So either we trust parties not to collude or additional constraints on the multi-party contracts need to be imposed. Again, these kinds of multi-party conditioning statements can be programmed in under the smart contract technologies. The outcome is an explicit, contracted function of what each player says or does, in potentially complicated ways, so that deviations from a proposed collusion solution dominates, i.e., eliminates collusion.[23, 24]

Contracts interface with ledgers and standard financial accounts. An asset as collateral in escrow is not really the same as an individually-owned asset, as another party has a contingent claim on it. A special accounting is required. Likewise, it's natural to think of the flow of future wage payments as human capital on the balance sheet, a standard view in some sense (Aiyagari 1994;

[21] See Ligon, Thomas and Worrall (2002) for penalties.
[22] See Phelan and Townsend (1991) for an example of how to impose this and what difference it makes for the optimal contract design.
[23] See section 10 in Harris and Townsend (1981).
[24] See Moore and Repullo (1990), Palfrey and Srivastava (1989) for examples of an implementation literature.

Huggett and Kaplan 2015). But when such flows can be collateralized, because the flow is in escrow, then the modified accounting needs to be implemented. Smart contracts on the distributed ledger allow an immutable commitment that future income flows be sequestered as collateral. To this we now turn explicitly. The point here is that standard financial accounting needs to be altered. These two components of distributed ledgers—accounts and contracts—are interacting.

5.2 Smart contracts

At their most basic level, as noted, cryptocurrencies such as Bitcoin use language familiar to accountants, economists and computer scientists. The state of the system is the current ownership of a digital asset, the *stock*, and a transaction or transition is the change in its ownership, the *flow*. These stocks and flows are on ledgers. Bitcoin is all about verifying and validating flows back to the genesis state where assets were originally created. This ties flows to stocks and requires public validation or verification of information.

More generally though, as in the smart contract composer Corda for example, the concept of ledgers is generalized to mean simply lists of "facts". Everyone has a ledger, but it is synchronized and held in common only for shared facts. Consensus is broken into two pieces. Validity consensus means that a transition is accepted, that is, has the required signatures, both for the current proposed transaction and for every transaction that led up to the proposed transaction. This is like the crypto asset example. Unique consensus is different and is the key to the generalization: A given party may not have a record of every single transaction, and that is not always required. There is not a consensus. On the other hand, a party could potentially request missing transaction information from notaries. The latter is necessary to thwart the double spending problem, for example. The point here is that what is needed depends on the underlying environment and what the ledger is trying to accomplish, ideally as part of a constrained-optimal arrangement. More on notaries momentarily.

A contract is entered into by multiple parties. Parties are nodes. Identities of nodes could be anonymous, as in Bitcoin, but this is not required. Identities could be named and public, for example the legal identity of an organization or the service identity of a network service. Note that trusted parties such as banks are allowed to be there, and be public, but so also are others, as well as strangers. A node writing contracts is an app providing a service, also allowed to be a universal service, if desired. The permissioned set of nodes for a given contract still have their access controlled by a doorman, so in that sense all contracts are only semi-private and there is a sense of centralization.

The contract is a code which is validated initially. It either works or not (and one can imagine several independent validators of code, dealing with potentially malfunctioning nodes, which also links back to thwarting potential collusion). A contract agreement is made via public and private keys. After this initial validation, it becomes immutable. In this sense there is no reneging on whether agreements have been entered into, nor claims that they are written in a different way. There is no need for trust on these particular dimensions.

A contract specifies states at points in time, such as current ownership for example or other facts. Communication under a contract is node to node, not necessarily broadcast to the entire public, as in Bitcoin, but on a need to know basis, as pre-specified in the contract. An oracle, a computer science term used to denote a function or node which knows the answer, is used to verify known facts that are states of a contract. Commands initiate transfers and result in output states.

Corda allows a notary service for validation of communication and proposed transactions. Upon being sent a proposed transaction, the notary will either accept it, if the notary has not already signed other conflicting transactions, or will reject it, as would happen in an attempt to double spend. Every state has an appointed notary and a notary will only notarize a transaction if it is the appointed notary of all the transaction's input states. A notary may be a single network node, in which case this part is quite centralized, and has a trusted-third-party aspect. Alternatively, there can be a cluster of mutually trusting nodes to deal with faults, or mutually distrusting nodes to deal with incentives. Notaries can choose a consensus algorithm based on privacy, scalability, legal system compatibility, and algorithmic agility. A notary could decide not to provide validity consensus, though in some contexts this runs the risk of denial of state attacks.

The key capabilities of smart contracts is that they overcome underlying frictions. Smart contracts allow full commitment (as opposed to limited commitment), immutability (no reneging), conditionality (as in the collateral example), observability (as in information in public view), verifiability (one can check past histories or verify states), and enforceability (as penalties can be incorporated or evidence taken to an outside legal system).

Szabo, the inventor of smart contracts, states the advantage succinctly: Smart contracts would enable both parties to observe the other's performance of the contract, guarantee that only the details necessary for completion of the contract are revealed to both parties, and be self-enforcing to eliminate the time spent policing the contract (Gord 2016).

They use digital signatures: Private cryptographic keys held by each party to verify participation and assent to agreed-upon terms. A smart contract will take actions, e.g., disperse payments, without further action by the counterparties, and they can access or refer to outside information or data to trigger actions.

Smart contracts can be stored and executed on a distributed ledger, an electronic record that is updated in real-time and intended to be maintained on geographically dispersed servers or nodes. Through decentralization, evidence of the smart contract can be deployed to all nodes on a network, which effectively prevents modifications not authorized or agreed to by the parties.

Smart contracts fit naturally with elements of mechanism design, that is, they allow the execution of the kinds of contracts and mechanisms that economists have largely taken for granted: the revelation principle, enduring relationships, promised utilities, and resolutions of the hold-up problem. We now elaborate on these in turn.

5.2.1 Messages: Revelation Principle

A game or mechanism among players is a specification of messages that can be sent, message spaces, and an allocation rule mapping realized messages onto outcomes. Consider a Bayesian Nash equilibrium of such a mechanism. The revelation principle asserts there is an alternative game in messages reduced to statements about underlying unobserved facts, private information, in which agents have incentives to tell the truth. This is without loss of generality. The point is that messages would be transmitted by actors or nodes and put on distributed ledgers, e.g., Harris and Townsend (1981). A related point is that the messages become endogenous objects. We have moved from how to design database and communication systems to the incentives underlying the data to be transmitted.

5.2.2 Impact of enduring relationships (duration): Past history of messages becomes committed and creates a new state as part of the determination of contemporary outcomes

Distributed ledgers housing smart contracts can help to implement formal and enduring relationships (Townsend 1982). The idea as indicated earlier is to give a household, firm, or trader an incentive to reveal private information. Sometimes we take this for granted, as a household with low but privately observed income will have an incentive to implicitly announce or reveal low income by borrowing, assuming here that loans must be paid back next period. Conditionality with collateral can assure repayment, but there are other means. The opposite is true for households with high income today and investing, with more funds coming back, available for the next period.

We can imagine an economy with two agents, one borrowing/investor with variable income and a second willing to enter into a contract. The first agent has variable endowments over time and is risk averse, caring about ex ante expected utility. The second is risk neutral with essentially unlimited deterministic resources. Income realizations of the first agent are not public. The pair agree and enter into a contract at some initial date, and it is then carried out over time. Though borrowing and lending is an incentive compatible contract with private information, it is not the best that can be done. The optimal information-constrained arrangement is a blend of borrowing/lending and insurance, with more risk contingencies in the contracts, and at lower inter-temporal interest rates. Longer contracts are better.[25]

Again, a long term relationship contract is implemented using messages of unobserved states that are put on a consensus ledger.

5.2.3 Promised utility as the state

The contract payoff can be divided into two pieces. One piece is a contemporary reward or penalty today, with goods, money, and/or tokens changing hands as a function of the message today. The other piece is the discounted expected utility of continued participation from the next period onward, varying with the outcome and action.[26] This expected utility is the key state variable to be entered on ledgers for tomorrow triggered by states as messages today. The

[25] See also Green and Oh (1991) and Karaivanov and Townsend (2014).
[26] See Abreu, Pearce, and Stacchetti (1990), Green (1987), Phelan and Townsend (1991), and Spear and Srivastava (1987) for examples.

messages, or actions, satisfy truth telling or incentive constraints, so there is no need for trust other than presumed maximization. These utility numbers are under full control of the long-term contract, and here in this example, they are public, with the conditional possibilities for future periods a function of facts/states. The idea is to create incentives to announce income states truthfully. A higher utility number tomorrow comes when the income is valued more in the future than it is today. Fernandes and Phelan (2000) generalize with unobserved past states to include yet another dimension: utility threats. There are upper bounds for what one can get if one deviated in the past and contemplated lying or disobedience now, threat bounds high enough to keep agents on the equilibrium path so that there is no deviation in the first place.

If the underlying economic environment is captured accurately, these expected utility numbers or threats will be indeed the actual realized subjective utility an agent would experience in that environment under the incentive scheme. This is also a qualification. If the approximation of the model to the actual economic environment is poor, then promised utilities in the contract are poor approximations of actual subjective utility. It is also possible that the contracts of the theory may seem relatively complicated for users.

5.2.4 Implementation through sequential play: Loading in commitment that was missing previously

Holden and Malani (2018) examine how to use the blockchain mechanism to resolve the 'hold-up' problem in economics: ex post bargaining when the contract is incomplete, without commitment. In their scheme, choices that cannot be verified by a third party can be resolved by a blockchain mechanism that can commit the parties to not engage in renegotiations. The keys here are posting of agreements, having them notarized by multiple parties as witnesses, and the security and inalterability of the blockchain. Moore and Repullo (1988) and Maskin and Tirole (1999) resolve an implementation problem using simple sequential mechanism, as in Moore (1992): when one party tells the truth, they earn more than if they lie. Joshua Gans (2018) shows how this implementation mechanism can be achieved with smart contracts to solve a trade problem.

5.3 Smart contracts in computer science and incentives in economics: Contrasts, similarities, and blends to be incorporated in subsequent design

Though the language of smart contracts and mechanism design fit remarkably well with each other, as noted in the previous section, there are some key differences. In economics we try to avoid a priori classifications of trust. Individual agents are self-interested and optimize given their information and incentives, whereas nodes in computer sciences are trustworthy or not. In economics it is part of the design to get agents to be honest and obedient, as a function of the social context, that is, the purpose for which one is using the ledgers. In economics, writing out contracts as fully as possible is largely regarded as desirable, if it were possible.[27] Whereas in computer science, codes are acknowledged to be buggy and hence it can be argued that simplicity has a special virtue. In economics, commitment is desirable if attainable, and

[27] But see Hart and Moore (2007).

renegotiation and bargaining is likely to change outcomes. In computer science, the need to seek consensus and validate transactions as with decentralized distributed ledgers opens the door to limited commitment, hence more limited arrangements.

As in Narayanan and Weinberg (2018), there is a sparsity of hybrid models between the middle ground of Nakamoto's dichotomy, honest as opposed to malicious nodes, as is standard in computer science, versus all players being strategic, as in economics. Narayanan and Weinberg think of cryptocurrencies as mechanisms and propose a protocol which must incentivize compliance for miners. They note that, for the most part, existing protocols are not incentive compatible. This is a remarkable gap that needs to be filled.

On the other side, smart contract coding can contain errors and has been difficult to use. Remedies are under development, for example: Agrello is an initiative for human readable smart contracts; iOlite is how to write smart contracts in natural language.

Fortunately, some reconciliation is possible. There is more common ground than one might realize a priori, though these distinct literatures are not cross-referencing one another even though the underlying environment and problems can be strikingly similar.

A version of the revelation principle works even when communication is limited, as in Prescott (2003), i.e., no message is received. New here, the revelation principle can be established even though message transmission and receipt is noisy; despite the noise, one may as well let the computer transmit the message to the network the way it would have been sent originally, under the original game and message space. The implementation problem of Harris and Townsend (1981) is a way of letting common private information be public by having multiple agents announce underlying states. Without noise in messages, at least, having two agents with common information announce is sufficient. Intuition would suggest that if messages are noisy, more agents announcing is better, but of increasingly limited value. Related again is the notion of default tolerance which allows, for example, up to 1/3 of nodes to communicate inaccurately, deliberately or through machine error.

Rubinstein (1989) formalized some starkly damaging results in the context of (another) generals problem. Despite multiple repeat messages, coordination of attacks can fail to happen. Two generals have to coordinate an attack in order for it to be successful. One sends a message to the other but cannot be sure the message arrives. So the second general, knowing this, sends a confirmation, but in turn cannot be sure it arrived. There is no natural truncation that gives common knowledge.

Yet Coles and Shorrer (2012) subsequently showed that the multiple channels may permit collective action: Parties may be able to coordinate their actions when messages' arrivals at their destinations are sufficiently correlated. This also permits cutoff equilibria, where players take action after receiving a minimum number of confirmations.

6 Application: Multiparty smart contracts and the design of financial infrastructure in high-valued systems

This section considers how implementing smart contracts can help with design under three topics: private information and partitioned ledgers; delegation of authority to a third party platform; and mitigation of bank runs.

6.1 Permissioned private ledgers and gains from concealment

In many environments, unique consensus is not desirable even if this were technologically possible. This is a generic implication of private information for optimal contract design. Though private information is effectively reported via messages or indirectly by choice and display options, this does not mean that such internal message data should be made public on the ledger. Often, the opposite is true: Messages should be kept private. Of course, cryptography with partitioned ledgers makes this possible. Messages are states, but they do not have to be seen by everyone. DLT allows partitioning but it does not have to be done in a mechanical and more extreme way of keeping all proprietary information entirely private.

Suppose shocks are to preferences and not to endowments. That is, agents are urgent or patient to consume. This is exactly or approximately the standard specification in modeling financial institutions or trade in financial markets. The Diamond and Dybvig (1983) model of banks and runs has versions of this, exactly. The Duffie, Gârleanu, and Pederson (2005) model of OTC trade in securities and consumption has stochastically varying security holding costs, which motivates the trade. More generally, banks and traders face shocks to their portfolios that come from the needs of their customers.

As an example, consider an economic environment with two parties to a contract with varying and random preferences to consume (Townsend, 1988). There is correlation in these shocks over time and over the agents, and here all agents are risk averse (not just two of them as in the earlier example). However, shocks are private to only one agent at each date and the identity of that agent alternates over time. That is, only one of the two parties in an initial period is announcing its urgency today in the first period, hopefully compensated by more goods or value today if urgent, and the other, second party, announces tomorrow in the second period. The agents announce their states to the ledger, not the other party, then the ledger algorithm determines what is made public. If the announcement of the first agent in the first date were public, it would undercut insurance possibilities for tomorrow. Typically, with allocations and histories of announcements known, there can be no insurance for the second agent in the second period. More is preferred to less, so it is hard to engineer a tradeoff. However, if agent two in the second period were unsure of what was announced by agent one in the first period, something which the ledger can keep secret, then agent two has to weigh the consequences of lying, announcing he is urgent and ending up with very little consumption. Insurance and truth telling is achieved by having the mechanism randomize over the consumption allocations in the first period as a function of agent one's message. Allocations are seen by all but the message is not. With risk aversion, this randomization is a welfare loss, but it is outweighed by the overall insurance benefit.

We shall return to the discussion of randomization and concealment in the discussion of payments platforms, below.

6.2 Delegation of portfolios to a third party: Platforms as custodians

Distributed ledger technology allows commitment to a multi-party smart contract in which awards are allowed to vary over time as a function of shocks. When there are both private unobserved idiosyncratic shocks and publicly seen aggregate shocks, it can make sense for households to delegate portfolio decisions and commit as if to a third party custodian (Townsend 1988), a kind of endogenous centralization. This can happen in principle with village funds, cooperatives, wealth managers, or exchange traded funds that can be undone only by a restricted set of designated participants. This allows front loading as when incentive constraints in future periods bind, limiting the value of having resources then, so more value is paid out contemporaneously; or back loading, strengthening inter-temporal incentives by having more at stake in future periods contemporaneously as a function of what is said today. China has implemented regulation to prevent this kind of delegation, that platforms be pass-though only, though motivated by fraud. DLT can facilitate implementation of a multi-party contract with commitment to the longer term, including sequestered funds to prevent withdrawal from the arrangement. The third party custodian could appear as an implementing node, a financial institution or reserve bank for example, in a smart multi-party contract.

6.3 Mitigating runs on banks and markets

DLT can improve upon current technology used by banks and markets to mitigate runs. Diamond and Dybvig (1983) is a seminal paper on bank runs. There is a more recent literature on market runs that draws an analogy, Martin, Skeie and Von Thadden (2013, 2014). Though this remains an important rationale for intervention, the problem of runs has a partial if not complete solution that can be found within the mechanism design literature itself, though as Martin, Skeie and Von Thadden (2013, 2014) anticipate, micro structure is known to matter. A simple version of this is suspended convertibility, thus reassuring investors that some of their money would remain, regardless. More sophisticated sequential service models treat customers differently depending on when they arrive to the trading window, that is, on the history of traders and trades before them, determining their incentives to announce privately observed shocks (Green and Lin 2003).

This is where the DLT is featured. Messages in distributed ledgers are essentially time stamped and immutable, so it is quite natural to think about history based on previous reported transactions in the blockchain being used to determine consequences for actions and messages today. With independent private values, there is a unique Bayesian Nash equilibrium which eliminates runs entirely and is a dominant strategy equilibrium. With correlated private values, some run-like phenomena remain since the history of traders at any moment in time is, on the one hand, self-reported and yet, on the other, desired as a key statistic, as it is revealing of aggregates. But nevertheless, the risk of runs can be substantially mitigated.

These mechanisms could be used in practice as a private sector alternative to public liquidity. This could dramatically mitigate the central bank moral hazard problem, that bailouts ex post create perverse private sector incentives ex ante. This could also alleviate the concern that private platforms might somehow pose a systemic risk. They would be designed precisely to guard against that. Optimal regulation would then amount to verifying that platforms implement these insights and are designed properly.

Here of course there is a political economy issue, getting diverse parties to agree to these changes. Another qualification stems from the latency in networks, which confounds the chronological ordering. A modification to treat identically orders received over small intervals of time will need to be designed. Budish, Cramton and Shim (2015) propose such a mechanism to deal with high frequency traders.

7 Application: Building financial infrastructure in developing country contexts

In the first part of this section two sources are used to document that existing financial infrastructure is quite limited in developing countries, featuring an application in Southeast Asia, Thailand in particular. Examples are then given of smart contracts that would help. A general equilibrium perspective on decentralized contract competition then follows. An example of a proposed DLT innovation is featured. Caveats are noted.

7.1 Limited financial infrastructure

Banks in low income and emerging market countries use legacy infrastructure. A large number of them do not deliver reliable and flexible banking solutions for the low-income households and SMEs. According to a study commissioned by the Asian Development Bank on financial access and digital solutions, the percent of needs met by formal financial service providers in Indonesia, Philippines, Cambodia, and Myanmar ranges from 16-74% for savings, 48-72% for credit, 1-4% for insurance, and 11-35% for payments (Asian Development Bank 2017). The report also attaches US dollar estimates of these supply gaps, typically in the billions, and the percent that could be met with digital finance, which ranges from 19-44%. It is not uncommon that households in SEA cannot access basic financial products or financial services critical to improving their standard of living. More than 70% of people living in Asia have no access to basic banking products, such as bank loans. In some contexts, credit cards are not prevalent and asset management services rare indeed. While peer-to-peer lending services are emerging, the underlying custodian banking services are still lacking.

SMEs can only access a limited number of financial products, often at highly prohibitive costs, even though SMEs constitute as much as 89-100% of all enterprises and account for 52-97% of total employment in ASEAN countries (ASEAN 2017). SMEs are regularly credit-constrained as financial institutions favor seemingly reputable and larger enterprises, such as those that are state-owned and owned by business magnates.

7.2 Analysis on the ground: Townsend Thai Project

The Townsend Thai Project with its data allows for an in-depth study of financial access and informal markets.[28] Overall, the sharing of risk so as mitigate the impact of shocks onto consumption and production is good but not perfect. Idiosyncratic household and business

[28] See Townsend (2016).

specific shocks are large and occur frequently; they constitute the bulk of the risk in the village economy. But notably, these idiosyncratic shocks are mostly, though not entirely, pooled away (Chiappori et al 2014).

The mechanisms are local village or *tambon* (county) level "money markets". Households borrow to repay other pre-existing loans, and borrow to lend, making new loans; these behaviors are shown to be one source of consumption smoothing (Sripakdeevong and Townsend 2018). Delayed repayment on loans goes back through the credit chain, causing delays for those who have borrowed to lend. But early repayment also goes back through the credit chain, even more so quantitatively. Network maps trace out kinship and transaction-based links (Kinnan et al 2018; Kinnan and Townsend 2012).

Within this, one detects shortfalls and things not working well. Risk sharing depends more on family/kin or how closely related in an ancestry tree. Trust may be an issue. Aggregate village-level or *tambon*-level shocks by definition are not diversifiable at the village level and, as confirmed in the data, end up in consumption. As in sophisticated financial markets, households with businesses running projects with self-financed or borrowed assets have rates of return which reflect that risk premia: Idiosyncratic shocks are largely pooled and have low risk premia, and village-level aggregate shocks have higher risk premia. But by the same standard, aggregate shocks across villages could be but are not often pooled (Samphantharak and Townsend 2018). Thus there are geographic limitations.

7.2.1 Impact of interventions: The Million Baht Fund with scope for improvement

In 2002, the government of Thailand put a Million Baht Fund in every village, creating a quasi-formal village-level institution to be run by a local committee. Pre-existing baseline and post-intervention panel data allow insights as to how this influx of capital impacted the villages and households. Levels of consumption increased, among other variables (Kaboski and Townsend 2011, 2012). Investment and profits of those already in non-agricultural businesses increase if they had access to credit and productivity (Banerjee et al 2018).

The informal sector acted as a catalyst, playing an augmenting, complementary role. For the relatively poor, the lowest wealth quartile initially, the best-fitting financial regime shifted from a buffer-stock-saving regime to a costly-state-verification regime. In the latter, output can be claimed but only low outputs which result in indemnities are verified. It is as if there were simple debt contracts with constant transfer back to investors, except that default with lower repayment is possible if underlying conditions are verified. Verification costs are lower when there are kinship households in the village (Ru and Townsend 2018).

But there is still room for improvement. Kinship ties can be a mixed blessing – what if no kin live nearby? Village fund committees did not allocate funds efficiently, as those who got loans were not found to have higher rates of return. The allocation was based more on family and political connections (Vera-Cossio 2018). Some of the village fund loan money did make its way indirectly to other non-direct-borrowing households, again through kin connections, as in the network credit chains. So this was helpful in mitigating efficiency losses, but it was not enough.

Village funds did not reach, directly or indirectly, all of those who should have been ex ante optimal targets. This is a limitation that could be overcome with DLT, as will be explained below. Improvement on existing systems is, of course, an implementation problem in its own right (Roth and Shorrer 2017). More work is needed on designs that will work in practice, in this context. Finally, there is relatively little financial infrastructure that works at the inter-village level.

7.3 Examples of smart contracts needed in this context

7.3.1 Escrow

As a lead example of conditionality and what is feasible with smart contracts on distributed ledgers, collateral can be put in escrow and released on certain verifiable conditions. This conditionality allows the use of tokens as collateral for loans, for example, but other assets can be used as well. Collateral in turn allows more distant parties who know little about one another to lend. There need be no trust per se.

7.3.2 Letters of Credit

De facto bank notes and letters of credit require not a bank per se, as with traditional and contemporaneous infrastructure, but only some third-party validation and/or novation. In a letter of guarantee, some third party underwrites the risk of the loan by providing escrow accounts of coin or other collateral that back the credit.

7.3.3 Waterfall payment

This refers to agents who are linked to each other via economic transactions. For an example of what is possible for SMEs, receivable revenue can be contracted and secured in a trusted account. This can be used to buy inputs, for example. It could also be used to fund employee payroll. Employee accounts, when secured and sequestered, can be used in turn as collateral for borrowing. As another example, for Village Funds, the chains previously mentioned linking borrowing and repayment can be automated and, more to the point, extended beyond village boundaries. This can improve welfare relative to the currently used technology which relies at best on promises, without a contracting technology that securitizes payments streams.

7.3.4 Savings products

A simple savings product can also take advantage of conditionality. Thai households need to convert their cash to investments, as noted. This becomes a less acute problem when currency income is converted somehow to coin/token in an e-wallet, because at that point, tokens can pay interest and the opportunity cost of holding cash could be virtually zero. Still, one could then take the next step and create a portfolio management tool as a smart contract that either transfers

funds automatically across accounts conditioned on states, according to some liquidity versus return trade off, or an app that at least sends encrypted messages back to the saver as reminders. Brokers with algorithms could provide this service; if written up as smart contracts then effectively the service is implemented on DLT.

Commitment savings could earn higher returns in tokens, paid into the wallet. Commitment accounts would be investments into escrow that limit withdrawals, if the savers wish to enter into these accounts. A predetermined quantity of tokens could be locked in a smart contract, which is programmed to release the tokens to the appropriate parties via airdrops when a set of conditions are met, e.g., at the maturity date, after a lock-up period, or if certain programmed system conditions mandate a liquidity infusion

7.3.5 Insurance

For community level insurance, the initial contributions are prepaid premiums into the mutual fund. Participants can choose among various options such as duration of the agreement, frequency and amounts of payments in, and if/how to cash out. These are social rules which, when agreed upon, become the smart contract.

For the payouts, the indemnities part of the community insurance, suppose a participant has low income and this could potentially be evident if, for example, wages are in coin. Then a simple conditional code will work well. To make the indemnity known to everyone in the fund, as the level of the remaining fund, with benefits, would be lower, there could be automated conditional messages to others, so as to achieve a public record of payouts. Messages are part of the distributed ledger and would be in sync, here uniqueness. Another contingency: Low or ill-timed rainfall could be the source of low income, and rainfall is verified by sending a message to a data source, an oracle, and getting a trusted message in return. The message, rainfall as a fact, is on the ledger and a key state variable. The oracle might need to be compensated.

With private information on income, an agent can send a message about its underlying situation, which could trigger an indemnity. The message could be published as in a database if these messages are to be public, though we return below to the optimal design of published data, conditionality, and where consensus is not unique. The optimal incentive compatible messages for a contract can be found as the solution to a dynamic mechanism design problem. In equilibrium there is no problem with respect to what is reported by the agent because the agent by design has the right incentives and at this point the code cannot be rewritten. The incentives come from a careful weighing of benefits versus costs as a function of outcome and reports, working out the optimal contracting part, putting the design into the smart contract and ledger. The presumption is that these improvements in insurance would outweigh costs of messages and memory.

7.3.6 Loans

With loans, an issue arises about how to grant them. One way to allow improved contracts is to allow other households to send messages about the underlying situation of that initial applicant agent. Studies have shown there is good information in the community on who should get loans,

as some borrowers have higher expected returns (Hussam, Rigol and Roth 2017). On the other hand, when households know that what they report has consequences for the initial would-be borrowing agent, who could be friend or kin, then these other households are dishonest or even could collude in these reports. A possible solution is to pay for truthfulness using a peer prediction index of subsequent events such as profits or of repayment to the lender. This history of prediction performance in the past could be used for future rewards and penalties. While public reporting of messages alone appears to give households incentives to predict accurately, suffering loss of face for discrepancies between predicted and actual events, tokens might be able to do better, to monetize loss of reputation. While some of these features are implemented in micro credit initiatives, they are far from universal and systematic. DLT provides a relatively inexpensive technology for implementation.

Multiple agents could also be sending reports simultaneously, as many contracts are multi-lateral. The design of the contract is such that each has an incentive to report truthfully if others are doing the same. A problem of collusion thus arises as noted: Agents coordinate in reports to get off-equilibrium beneficial outcomes. Yet, as noted earlier in Harris and Townsend (1981), the implementation literature has insights on how to thwart this threat, through multilateral conditionality. So this too can be implemented as a smart contract.

Distinguishing commodities and securities, points in time, and multiple agents, virtually any multi-party contract can be written as a smart contract.

7.4 The general equilibrium perspective on the provision of financial services

Here we take a standard general equilibrium perspective but with the inclusion of DLT. An economy consists of standard items, but with an innovation on the provision of financial services. The commodity space of a well-specified economy consists of factors of production such as land, labor, capital, and intermediate and final produced inputs and outputs including capital and consumption goods. The commodity space also distinguishes all of these by: Time, as if there were time stamps; location, allowing for distinct geographies and shipping; and uncertainty, conditioning on states of the world, both exogenous and endogenous. Key actors are the users, households and traditional firms, including micro and SMEs. Some households run micro and small-medium enterprises and are thus both households and firms combined. More generally, there is an ownership structure that specifies the shares that a household has in particular enterprises. Finally, an economy also specifies who knows what, at least initially, though signals can be generated, and some observe what others do or can make inferences. There are information partitions that capture these ideas.

The generalized commodity space becomes the space of incentive compatible contracts written on top of the underlying traditional commodity space, but subject to contracting costs, as is enumerated below.

Efficiency of an economy is judged by the Pareto criterion. An allocation of feasible contracts is said to be efficient, or constrained efficient if there are underlying obstacles, if there does not exist an alternative feasible set of contracts which makes some households and firms better off

and others no worse off. The current system with its symptoms of limited access to financial products is inefficient in this sense.

We could go on to specify, as would be traditional, banks and markets as primitives of the economy, but it is here that we depart from the traditional model and break new ground. All possible forms of financial intermediation are on the table. Prescott and Townsend (1984a; 1984b) show how to decentralize such environments with a competitive broker-dealer sector for a multiplicity of information specifications. Prescott and Townsend (2006a) embed individual and multilateral contracts, as enumerated above into an entire economy with a nexus of activities under contracts connected to general equilibrium flows. Competition among broker-dealers and intermediaries drives profits to zero, but these intermediaries remain essential for pooling, hence the language of disintermediation should not be used.

Relatedly, private information per se is not a rationale for regulation. Many economies decentralize in the sense of the welfare theorems. Joaquim, Townsend and Zhorin (2018) model imperfect competition among financial service providers including adverse selection. In this context frictions interact with profits. Improved contracting technologies which do not also bring more entrants can decrease household and SME welfare.

7.5 General equilibrium caveats: Coordination and regulation

Two caveats are important to the question of competition and decentralization. First, following Pesendorfer (1995) and Makowski (1980), one can begin in an economy with incomplete markets and contracts. Then, one can imagine that financial service providers are allowed to innovate. However, due to complementarities in uncoordinated innovation, the outcome of this process does not necessarily make models complete. More coordination would be needed. A related point, ill-informed or outdated regulation can segment industry and make the needed coordination impossible.

Second, markets and access should not be entirely open. When there is private information and minimal scales of operation, then some forms of competition can undercut incentives and optimal diversification. Equity markets should not free-ride on existing infrastructure, for example. Certain kinds of exclusivity are needed. Competition should be ex ante for the right to provide services, not ex post to draw off customers from contracts. As in Acemoglu and Zilibotti (1997), minimum scales of operation across sectors means that risk sharing is incomplete. Stretching the extensive margin to high minimum scale projects is good, but this means that such high end projects receive relatively more funding than zero or lower end projects, that is, portfolios are not balanced in funding. Unrestricted trade into equities issued by firms would undercut the optimum. One intermediary should do all the packaging. Retrading, as in Jacklin (1987), is another negative factor. Townsend and Xandri (2018) provide blueprints for market design and regulation in this context.

Blueprints are needed as otherwise the system may develop piecemeal and not achieve a constrained efficient outcome. Regulators need to see and understand the big picture.

7.6 Featured example of innovation: EvryNet

EvryNet is an intelligent financial automation operating system that aims to provide open-source banking services and financial contracts to unbanked and underbanked populations. It is being built on the need-smart-contract premise and incorporates the general equilibrium perspective. The system is under still under development as of this writing.

EvryNet is creating an interoperable smart contract platform that enables not only traditional banks but also micro-finance institutions and others to initiate and execute banking products and financial contracts. The envisioned provision of contracts should be at competitive prices due to competition across providers in the provision of computer memory (storage) and computation power. A rating system tracks performance of providers in validation.

The financial services have a multi-tiered architecture. The core component of the platform is a financial service portal which users and institutions can utilize to build financial contracts, either standard or customized. The portal is underpinned by a smart contract composer, which enables smart contract creation using distributed ledger technology. EvryNet also offers encryption to store smart contracts privately. Once the user selects a smart contract template, specifies necessary inputs, and selects nodes based on trust or reputation score, the smart contract will be processed through EvryNet's virtual machine.

The smart contract can optionally check for compliance or necessary regulations. For instance, the EvryNet platform can allow the relevant organizations to certify by signing the contract digitally or even executing the compliance-check code to ensure regulation compliance. It allows event hooking in smart contracts to seamlessly receive relevant events from external entities. Many real-world smart contracts may need external inputs to complete the conditional transactions, for example, the confirmation of a shipment.

In summary, rather than take as given the current set of institutions and markets, the vision of EvryNet is to recreate them through the new distributed ledger e-transfer and contract technologies, including executing nodes as a new production section. However, the implications of costs of coding, validating and memory for optimal contract design have yet to be determined.

8 Application: Distributed ledgers and payments systems, private vs. public and the role of tokens

In this section we consider payments systems which are designed to be constrained optimal in support of trade, credit, and insurance.

The essential idea is that distributed ledgers could keep track of messages as a part of the execution of a multi-period, multi-commodity, multi-agent smart contract and thus optimally allocate underlying risk while facilitating trade and exchange. Featured in this context is the use of tokens, both single and multi-colored coins. The role of tokens is twofold. First, tokens are one way to interpret ledger entries on a centralized system, as colored entries. Second, tokens as real objects can allow a decentralized implementation of the same allocation, a hybrid system

that can mitigate the scaling problems of centralized systems. If tokens are held in private, then incentives for voluntary disclosure need to be included, though this is not always a binding constraint. How well these various accounting and token systems function depends on the size of the message space relative to the needs for credit, insurance, and trade coming from the underlying economic environments. In some instances, information should be kept private, so more limited message systems are actually preferred, as constrained-optimal.

8.1 Tokens on ledgers as a way to achieve unique consensus, an insurance example with voluntary disclosure

As in Townsend (1987) suppose there are four agents in spatially separated locations and some subset of the agents travel. More specifically, a risk averse agent a is paired with a risk neutral agent b initially, in the first period at one of the two locations, such that the risk neutral agent can insure the risk averse one; likewise for agents a' and b' at a second location in the first period. The pairings switch in the second period. Agents a and a' make announcements of their urgency to consume, and to induce truth telling, if urgent today, they receive the good but at the expense of getting less of it in the second period. Likewise, if patient today, they receive less of the good today at the benefit of getting more of it in the second period. A centralized public ledger recording all messages is one way to implement this.

However, in reality, if there were no record of announced preference shocks and no record of allocations in the first period, then there could be no link of the first period to the second period, and so, essentially, no insurance can be obtained in either period for agents a and a'.

The introduction of tokens as a hybrid system can solve this problem. Announced patient agents in the first period receive more tokens than urgent agents. Tokens could be carried literally as coins or, alternatively, be private but immutable ledger entries. In the second period, agents with more tokens can display them in order to be on the receiving end of goods. Tokens or DLT entries are equivalent and convey the necessary history.

8.1.1 Multiple colored tokens and distinguished histories: Trade with insurance

We need not rely on pure insurance examples. In environments with two or more goods there are exchanges of one good for another in each date, driven by the usual motives for spot trade, but those desires to trade are driven by preference shocks impacting inter-temporal tradeoffs, and those are private.

In a hybrid decentralized system, a portable concealable token could be used to keep track of trades in the first period. An agent may give up one good and purchase the other and be expected to reverse the situation in the future. That would be fine, as with the earlier examples, portable tokens can handle this. But now we introduce additional shocks to inter-temporal discount rates, different for different goods. After these shocks, agents may have ex post regrets and wish to reverse the trade in the first period in order to get the good they most prefer in the second. One possible solution is to have multiple colored tokens, or equivalent multiple digital assets on a private ledger, so as to have more dimensions in which to keep track of more detailed histories.

These ideas in economics have tight links to cryptography and the idea of colored coins.[29] For example, ordinary Federal Reserve bank notes can be given bar codes so they can be used as tickets to Yankee baseball games. The team signs a message that includes a specific game date, seat number and serial number of the bill, with signature of issuer, all stamped on the bill. The advantage of using pre-existing bank notes is that they cannot be easily counterfeited, so there would be no need to print new tickets. And such a system is decentralized. Alternatively, the stadium could check a central database for information when a "ticketholder" enters the gate with a note having a certain serial number. The serial number links back to the primitive transaction, the purchase of the ticket. Either way, metadata is being attached to the note.

The point is that coins have public verified histories, to trace ownership. This history can be made meaningful and put to other uses. Coins that "originate" in certain transactions can have associated extra metadata, which act like a color. The colors are a metaphor of course, as they would simply be bit strings. It is important of course that all participants understand the rules of this payments system, so they know how to interpret the colors.[30]

More generally, and simply, and to link back to the economics, distributed ledgers could keep track of messages as a part of the execution of a multi-period, multi-commodity, multi-agent smart contract, and thus optimally allocate underlying risk while facilitating trade and exchange. Participants do need to know that specific keys are used to sign valid transactions, that is, they must know and understand the mechanism and its rules.

8.2 Permissioned private ledgers, when consensus is not unique due to optimally kept secrets

While multi-colored tokens convey more history, one should not jump to the conclusion that more information is preferred to less. We gave an example of private information and randomization earlier, with two agents, two dates and one good. Shocks are private to only one agent at each date and the identity of that agent alternates over time. If the announcement of the first agent in the first date were public, it would undercut insurance possibilities for tomorrow. More is preferred to less, so it is hard to engineer a tradeoff. Randomization can help. Allocations generated by lotteries are seen by all, but the original message is not. With risk aversion this randomization in allocations is a welfare loss, but it is outweighed by the overall insurance benefit.

If we generalize this example to four agents and imagine that agents cannot know what happened at the first date in a different location, then tokens can be allocated and tied to the random consumption allocation. Tokens can be carried into the second date, so the public part remains public. Yet tokens need not reveal entire histories of message. That is, if tokens are colored data entries on private systems, there should not be unique consensus in such environments.

The insurance example may seem a bit counterintuitive but the same idea shows up in applied work in finance and with the same elements: risk aversion and information asymmetry. Lyon

[29] See Narayanan et al (2016).
[30] One clarification, the idea of coloring coins came about as a way to overcome initial limitations in Bitcoin. With the current version of Ethereum these are no longer needed, but the point about distinguishing histories remains.

(1996) analyzes the optimal transparency of order flow information as in foreign exchange markets, arguing that slower revelation of information, which could reveal market-wide order flow, improves risk sharing among dealers facing unavoidable position disturbances. Garrett, Martin, Lee and Townsend (2018) show in a similar but distinct context that some post-trade information disclosure can improve liquidity, but revelation by a self-interested platform is a worse outcome than no information at all. There are other examples of optimally limited shared information. In Prescott and Townsend (2006b), auditors make incentive compatible announcements of underlying states and then depart, making way for an incoming and relatively uninformed agent, assigned there as a solution to the mechanism design problem. The auditor would be akin to verifying underlying states or objects on a ledger. For the incoming uninformed agents, they do not know what path they are on and face tradeoffs in making announcements or taking actions.

All these optimally designed systems require commitment to the design, including the control of information. Leakage is a potential problem in practice and privacy is a concern.

8.3 Lessons from monetary theory for payments systems: The need for coordination

The issue of public vs. private information comes up as well in thinking about coordination. The issue is how to achieve an ideal Pareto optimal allocation and what information is required. Likewise, lack of key common information can lead to market crashes. Either way, there are clear implications for micro/macro prudential regulation. Ironically, DLT, rather than being a regulatory concern, offers in the context of this problem an obvious solution. Of note, these systems are centralized.

8.3.1 The impossibility of decentralized exchange

A natural objective is to try to achieve the Pareto optimal allocation associated with a Walrasian competitive equilibrium. Ostroy and Starr (1974) ask whether this can be done under a decentralized exchange when information is limited, or whether centralization of information is required.

This happens despite the fact that in the model many important items are simply taken as given so that there are as few obstacles as possible: The target Walrasian allocation is given; the prices as common marginal rates of substitution in the Walrasian allocation are given and fixed; and paths of matched agents are known in advance. Agents simply act as computers implementing code. The question: How to write the code to implement a social optimum? The answer to the question of decentralization: Impossible to always do this. The key insight: Information on the distributed ledger should be public in some instances.

In the Ostroy-Starr model, money plays the role of unit of account. There are potentially many agents and many underlying commodities. Agents start in the model with underlying endowments of commodities, as if assets on a ledger. But the model would apply equally well to endowments of securities, various possible fiat monies, or combinations of all these. Actual payments systems handle retail and wholesale trade, securities settlement, and cross border currency flows.

Agents in the model meet pairwise and then trade. The point is that not everyone is together all the time in one spot -- they are matched in, for example, an OTC market. Trade in the model is monetized as a payment order in a quid pro quo condition: When goods are supplied, the supplier is given unit of account credits, in monetary terms. Likewise, when goods are purchased, the purchaser is assigned debits in unit of account, in monetary terms. Under a natural quid pro quo spot trade condition, the value of the purchases must match the value of sales in each and every contemporary transaction. Here, then, there is no credit. The prices used to value commodities and securities come from the target Walrasian allocation. These transaction values are placed on the ledgers as flows and result in new commodity/asset positions as stocks. Of course, as emphasized from the outset of this paper and from the discussion of accounts and ledgers, the corresponding trade ledgers (flow) and asset ledgers (change in stocks) must be consistent.

In a key example provided by Ostroy and Starr (1974) and corrected by Kim (2015), it is shown that the appropriate trades across agents when they meet in subsets can require centralized knowledge of the underlying environment and trade histories. In particular, knowledge of identities of agents, histories of trade and initial excess demands are needed, not only pairwise, of those contemporaneously matched, but also of others with whom the contemporaneous set of matched traders have not been matched previously. The idea is straightforward: Implementation is both forward and backward looking. One has to know where the system should be headed, and the remaining options in the future to get there, hence what trades need to have been accomplished in the past in order to make it feasible. Sometimes there are multiple choices among a given contemporaneous pairing and guidance is needed, from the forward and backward perspective. In the key example, there is only one instance in which information would need to be shared.

In the Ostroy-Starr (1974)/Kim (2015) example, private information about initial excess demand is the source of the potential problem. However, one can think of some initial trading period appended onto the beginning of the Ostroy-Starr model, differing from the Ostroy-Starr initial period – perhaps earlier trading rounds which determine excess demand at the time we tune into the Ostroy-Starr initial period. These initial excess demands are the deviations remaining from the ultimate target. If there were a common consensus verified ledger to which all traders have access, then this required information would be known.

Ostroy and Starr do also discuss alternatives that mitigate the need for centralized common information. They describe what could be termed a monetary solution, sufficient ex ante liquidity in one of the commodities, termed the money good, ample enough so that expenditures of commodities or asset purchases of any agent can be financed out of this liquidity regardless of who meets whom when and regardless of underlying economy-wide efficient targets for trade. But that amount of liquidity, whether a commodity, or fiat money, or tokens, is large and potentially costly, as liquidity held for this purpose is not invested in economic activities. This is made more explicit in other models.[31] Real time gross settlement systems are in fact frequently run as hybrids with liquidity saving mechanisms made possible by computer algorithms and queues (Martin and McAndrews 2008).

[31] See Townsend (1980) for example, and a more general discussion of why money holding is a distortion and money should bear interest.

Another Ostroy-Starr alternative is a central warehouse, such as a very large broker-dealer with whom all can trade, though such a large entity could bring in other distortions, namely, market power. A third alternative is credit, as if from a Walrasian banker, of the kind witnessed in early trade fairs (Townsend 1990). In the Ostroy-Starr model, however, this requires two more rounds of trade, which in their model violates the desired criterion, to achieve all trade in one round of pairings only.

It is tempting to think of the Walrasian banker as a central bank or digital reserve bank. In practice, for good reasons, central banks worry about intraday exposure and thus require good collateral. Central banks also worry about the ultimate motive for trades, such as interbank borrowing transactions that are passing through their payments system but are not designated as such. Central banks would like to know more, given macro prudential concerns. Finally, for coordination and macro prudential concerns, certain transactions need to be public and hence entered on the public part of distributed ledgers. The earlier section dealt with how to make a decision on this from the standpoint of mechanism design.

8.3.2 Information problem with private monies: Circulating private debt and multiple media of exchange equilibria

A somewhat related setting is found in Townsend and Wallace (1987), which focuses on payments made via high velocity circulating, privately-issued debt. Securities can serve as payments devices and circulate (remember the tally sticks). The point here is that this is an issue with e-securities, e-assets, and likewise. E-tokens are designed to facilitate liquidity and trade but there can be problems.

To motivate the set-up, the idea that securities can serve as a payments device should be familiar. In New York markets for example, brokers experience shortages of various securities. Under rehypothecation of collateral, a lender who gave up cash for securities as collateral becomes a borrower in turn, passing the collateral on along a chain. Singh (2011) finds the velocity of circulation of treasuries is now higher than M2. This can compete with fiat currency, as in Muley (2016).

Globally, Carlson et al (2016), Greenwood, Hanson, and Stein (2016), Krishnamurthy and Vissing-Jorgensen (2012) and others are all consistent in finding a liquidity premium for those treasuries. The treasuries have become the money-like assets, more like money than money itself. This has policy implications, though what to do about it depends on one's point of view.

Townsend and Wallace (1987) provide an instructive example environment. There are four agents, four periods, and two locations. Agents have endowments of a consumption good that varies over time, but again, as earlier, one can generalize and imagine these are other objects. One can trace out chains of named debt from the issuer at the issue date passing through third parties to the issuer at the redemption date (here no reneging is allowed). This circulating private debt is the medium of exchange in contemporaneous transactions, supporting trade in other short-term non-circulating securities and the consumption commodity. One can take the consumption good as the numeraire, but again, as earlier, this could be generalized.

The point of the model is that there is a coordination problem. There are a large number of potential equilibria, each of which achieves the same target Pareto optimal, complete markets equilibrium real allocation. But these equilibria vary in who is issuing the debt initially, hence what objects are circulating, that is, what objects are providing the payments device. In the key example, to be specific, agents 1 and 2 are matched in location one and agents 3 and 4 are matched at location two. Agents 1 and 4 stay put at their respective locations, but agents 2 and 3 keep switching back and forth from one period to the next. There are many equilibria that achieve the Pareto optimal target: Either all the debts that are allowed to circulate could be issued by initial parties in one of the two locations; or by the parties in the other, second location; or, they could be issued in various particular convex combinations. But by assumption, in the informationally decentralized model environment, there is no way for traders in one location to know what is going on in the other. Too much or too little debt as liquidity could be issued.

This can cause problems later in subsequent markets. Some circulating debts would be "over-issued" resulting in a precipitous drop in their prices later on in the trading cycle. Agent-traders suffer from excessive fluctuations in their intertemporal consumption profile. It seems failure to achieve coordination links up with observed chaotic conditions. Bills of exchange were traded in London money markets, and these crashed, leading to arguments for the creation of a central bank.

DLT keeping track and verifying initial issues of long-term debt in exchange for consumption or other objects, if public, would achieve in the example environment the necessary coordination. A related point: Not all information need be shared all the time. Here it is only information on initial security issues. Interestingly, and a warning to policymakers, there are no liquidity premia associated with the circulating private debt, to pick up from the data, yet the coordination problem remains. An understanding of the environment and data tracking will be needed.

8.4 The public versus private decision

Mechanism design and monetary theory have provided examples of when information sets should remain private, to achieve an information-constrained credit and insurance system, and when it is socially valuable to have them be public, for coordination reasons. In practice, these two motives may intersect with each other, and more work would be needed in this context for hybrid design.

9 Payments systems on distributed ledgers

We now present two designs for payments systems that utilize distributed ledgers. One was implemented by the Bank of Canada for commercial banks for use in the interbank market. This is largely about how to potentially improve on existing systems or at least validate that DLT payments can do as well as traditional systems. A second is under development in Thailand for money transfer operators that engage in international fiat money exchange. This is about creating and implementing something new, including constrained optimal designs for trade, credit, and insurance in coordination with a digital reserve bank.

9.1 Interbank payments: Project Jasper

To summarize from the Payments Canada (2017) white paper, Project Jasper is an experiment in private permissioned distributed ledgers which allow for the exchange of central bank issued digital assets. The goal is to transform the payments structure in Canada. If an improved ecosystem could be built, there could be significant benefits for the whole financial sector, and the economy overall. It is important to keep pace with the shift to digital commerce and remove impediments, and to try to get to a near frictionless end-to-end customer experience. In experimental simulations, Project Jasper could indeed handle the high volume of Canada's large-value interbank transfer system.

Project Jasper first explored Ethereum, and then moved to R3's Corda platform to allow for improvements in settlement finality, scalability, and privacy. As noted in the white paper, DLT allows improved back-office payment processing and reconciliation with and across participating financial institutions, reducing the likelihood of costly errors and improved automation through the use of smart contracts.[32] Another goal for Project Jasper was cybersecurity, creating backup ledgers which eliminated single points of failure.

As noted in the report as well, it is possible to limit information on a database when privacy and confidentiality concerns among parties are paramount. For example, in Project Jasper, parties see only their own activity. The role of the Corda notary node is played by the Bank of Canada, though Corda can eliminate the need for such a single trusted database operator.

Netting promotes funding efficiency and smoother intraday payments flow. Phase 2 of Project Jasper appears to be one of the first instances of a central queue within a DLT platform for payments. For example, a participant's account gives permission for a bounded, specified amount of value to be placed into a queuing option. This can be changed, but not when the queuing algorithm is running; during that time the participant is blocked as codes search over best transfers.[33] Phase 3 extended the Phase 2's proof-of-concept to the settlement of exchange traded equities.

The central bank maintains a commitment to settle accounts but has risk exposure in doing so, and so collateral is posted by participants. Technically there are two versions that differ in collateral and the loss allocation procedure: The defaulter pays, and the survivors pay through pooled risk.[34, 35]

---

[32] Reconciliation and reliable registries were mentioned in the introduction to this paper.

[33] See Güntzer, Jungnickel and Leclerc (1998) and Kuussaari (1996) for further discussion of algorithms for clearing and risks.

[34] See Monnet and Nellen (2014) for theoretical modeling of CCP clearing.

[35] Also related, from Monetary Authority of Singapore (2017, page 6), running its own innovation center: "The prototypes successfully demonstrate several points. Firstly, that key functions of a RTGS system such as fund transfer, queueing mechanism and gridlock resolution can be achieved through different techniques and solution designs. Secondly, decentralizing the key functions of a RTGS system may not only mitigate the inherent risks of a centralized system, such as single point of failure, but may also affirm the promised benefits of DLT, for example cryptographic security and immutability."

In any event, Jasper is illustrative of sophisticated multiparty smart contracts, showing what it is feasible to do with distributed ledger technology. Other potential examples of DLT-based payments systems come from the Monetary Authority of Singapore, a think tank. Nevertheless, there may be growing skepticism among central bankers concerning such interbank systems for large value payments. Are they really needed? In the Project Jasper application, it seems it succeeded in recreating the functionality of existing traditional payments infrastructure, that is, on one the one hand, it did as well, but on the other it did not dominate.

The Project Jasper white paper does note that DLT is particularly relevant for cross-border payments, as current infrastructure still relies on financial institutions keeping their own databases, and so counterparty risk is high. Subsequent collaboration among Bank of England, Bank of Canada and the Monetary Authority of Singapore explored in detail the limits of current inter-bank transfer systems and the possible use of distributed ledger technology to issue central bank digital currency as the key central bank liability to be used in wholesale payments systems.

9.2 Optimized design of cross-border payments: Velo

Consider the case of Cambodian or Myanmar migrants in Thailand, migrants who both need to fund their trip (now regulated), and to send money back home to family. Remittances in Southeast Asia amounted to $63.9 billion in 2016 while transfer fees have grown 3% over the six years leading up to 2016 to 7.1% (Leong 2017). The high transfer fees are partly due to legacy technology in the formal sector and limited access to formal currency exchange markets. Traditionally, a Money Transfer Operator (MTO) conducts cross-border transactions using a single server to record transactions among its international subsidiaries, going through traditional financial institutions, or entering into a bilateral agreement with an MTO in another country. This limits the number of partners they can have and results in high transaction costs.

Velo proposes to create for MTOs in Southeast Asia a highly liquid decentralized settlement layer on a permissioned blockchain. It is a hybrid in the sense that it utilizes, and effectively transfers, fiat money, but it does this with its own indigenous token and with fiat tokens.

MTOs can be viewed as agents with varying underlying balance sheets hit by the needs of their customers for trade. To trace out an example transaction, a client, to be called Alan, asks an MTO in Thailand to transfer money to Alice in Cambodia. Alan is essentially lending money as a deposit to the Thai MTO, which will be extinguished when Alice in Cambodia receives her money. At some point in the timeline, Alice gets a text message that her money from Alan is in. Alice then goes to a Payment Gateway, e.g., a 7-Eleven or a registered point-of-sale (POS) store, and shows a QR code from the text message. The 7-Eleven staff scans it and gives Alice the money. At this point, the Cambodian MTO is lending money to the Thai MTO. This is to be repaid, ideally, by a transaction in the reverse direction. More generally, netted transactions among a larger community of MTO operators happen in a multi-lateral clearing operation.

When the Thai MTO first enters the Velo Network, it enters into a smart contract and converts Thai fiat money into Velo tokens. The Velo token is then locked up with a digital reserve bank, a key institution of which Velo is a part. The digital reserve bank has a charter and rules for MTO participants. In return for its deposit of Velo, the MTO gets an intraday Thai Fiat credit balance

in Thai fiat tokens. This serves as an upper bound for outgoing remittance transactions on any given day. The Thai MTO then initiates a transaction for Alan by submitting the order to Stellar through the Velo Network, converting Thai baht tokens to Cambodian riel tokens. Once executed, that transaction amount is deducted from the available within-day Thai MTO's fiat credit balance.

Since the Velo Network is built on Stellar, it can provide interoperability. Velo hopes to achieve transaction throughput of 1,000+ per second, with each transaction taking 3 to 5 seconds at a very low cost, a tiny fraction of a penny. Velo maintains an off-chain order book, and groups and offsets transactions before sending batched orders to Stellar at regular intervals. Order book matching maintains an accurate source of transactional references between any two transacting MTOs, providing reconciliation. This reduces the probability of disputes and delays, a feature of DLT commented on earlier in this paper

At the end of the day, as anticipated above, the many transactions that occur among all MTOs will be netted using the digital reserve system's smart contract distributed ledger algorithm. Each MTO will be notified how much it needs in order to settle its credit line. If the MTO successfully settles its end-of-day balance, its fiat credit balances get replenished, also resulting in an increase in their credit rating. Alternatively, if any MTO fails to settle the balance, the Digital Reserve Bank will liquidate the MTO's Velo tokens that had been locked as collateral, executing the conditionality of the contract, and slashing the MTO's credit rating.[36]

Under Velo, counterparties can conduct transactions with multiple pairings in different countries without having to trust each other, while leaving clearing and settlement roles to the digital reserve bank and the Black Apple Foundation, respectively (fiat money is held by the Black Apple Foundation in each country). The premise of the system is that flows will ultimately balance, e.g., in Thailand, Thai baht paid in will equal Thai baht paid out, or that the Black Apple Foundation is flexible enough in its cross-country holdings that it can cover fiat money deficits.

Some MTOs are already engaged in a contractual bilateral relationship, with both of their fiat credit lines supporting their financial obligations bound by a bilateral agreement. In principle, depending on these relationships they can handle credit and insurance with each other, at least bilaterally. Going one step further, Velo could allow the implementation of constrained-optimal contracts implemented on DLT as envisioned above in the examples on information constrained credit and insurance.

10 Summary and Conclusion

Distributed ledgers have the potential to transform economic organization and financial structure, yet the subject is embroiled in controversy, hype, and lack of consistent terminology. Rather than get embroiled in proper definitions, what is or is not a distributed ledger, and how broadly or narrowly the term distributed ledgers should be interpreted, we focus instead on the economics of what distributed ledgers can do by breaking concepts down and analyzing key individual

[36] One vision for Velo has the token as a utility at a fixed price. If there are markets for Velo exchanges against fiat, new issues are raised, and addressed, in Appendix B.

components. We also compare and contrast the economics framework with the frameworks of computer science and data management disciplines, clarifying the terminology, and the technology, where possible.

The first four sections of this paper described distributed ledgers in terms of those key, though familiar, component parts: ledgers as accounts, e-messages and e-transfers of value, cryptography, and contracts including multi-party mechanism. We put each in context and provided examples, emphasizing what these parts of DLT brings to the table, one at a time.

For the first component, ledgers, we link the ledgers of cryptocurrencies to a statement of currency flow as a standard financial account and to currency as a balance sheet item. Thai villages' cash flow accounts are presented as an example, as in this context currency flows are reasonably well measured and have been used in analysis. The ledgers of DLT are then put into this context, taking an additional step, which in essence is simply creating a common integrated database of cash flows across households and the associated balance sheet entries. With this, one uncovers discrepancies, not unlike how the use of DLT can help in reconciliation of trades. A related analogy: Standard financial accounts were an invention associated with double-entry book-keeping, for more accurate measurement, a huge innovation at the time. Likewise, DLT enhances measurement across diverse parties.

Relatedly, the DLT ledgers could be used with transaction data to create standard financial accounts, namely the income statement, balance sheet, and statement of cash flow. These accounts are useful in analysis, and for policy, as illustrated by two examples in Appendix A.

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 comes impossibility theorems with regards to: consistency, accuracy, partitioning. In the presence of a partition tolerance, one has to choose between consistency and availability. Further, 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 who can write can also (accidentally) destroy or corrupt data and there is data security, cyber risk. Hybrid systems between the end points of strictly hierarchical database systems ignoring incentives and fully connected network meshes which exacerbate privacy concerns emerge in practice, though not necessarily as deliberate choices among the universe of possibilities. The economics literature provides clear examples of how costly connections among agents can lead to constrained optimal partitioning.

For the second component of DLT, e-transfers and e-messages, we compare and contrast two countries, Thailand and Sweden, from high to low ratios of currency to GDP, respectively. For Thailand, there are large welfare losses to the current reliance on paper currency. For Sweden one worry, ironically, is the other way, at least in the near term—groups vulnerable to the disappearance of currency. We then feature an innovation, e-money for a country in between,

Kenya, where M-Pesa has generated large welfare gains. This is done with a trusted third party, Safaricom, and without crypto currency, though the ledgers of Safaricom resemble tokens in some ways. The key is allowing easy ins and outs from Kenyan shillings to cell credits, facilitating transfers from urban migrants back to villages, for better risk sharing and poverty reduction.

This section on e-value transfers concludes with an important discussion of potential limitations of trust and the centralization/decentralization issue. In the Kenyan context with Safaricom, trust is less obvious than it might seem when taking into account the larger financial system, with the banks holding Safaricom, and hence customer shilling accounts, they are subject to runs. More generally, cryptocurrency systems rely on distributed ledgers with some kind of consensus validation. Cryptocurrency exchanges are largely built as traditional systems with brokers or centralized exchanges. There are increasing exceptions, but decentralized exchanges using distributed ledgers seem to suffer from limited capacity and liquidity, with low volume. This is identified as an area where more work is needed. It also makes the point that hybrid systems like this may not only persist but may potentially be part of an optimal design.

E-transfers naturally raise issues related to cryptography and verification of messages if there is not universal trust in a single third party, as in the subsequent section of the paper. These ideas are ancient, dating back to Mesopotamia with sealed clay envelopes containing tokens as a manifest of shipments and Medieval tally sticks split into pieces, the stock circulating as money that uniquely matched the foil held by the borrower, checked on redemption.

More recently, but before Bitcoin, we have secure multi-party communication, IBM crypto express cards, public vs. private keys, and zero-knowledge proofs and protocols. This section on cryptography contains an important discussion of the various validation systems used, from Bitcoin's proof of work, to proof of stake, to federated decentralized systems of trust, which require a listing of trusted nodes but a list that is different across different nodes, to proof-of-authority systems, with algorithms to achieve consensus. To some, the term distributed ledger refers uniquely to some if not all of these decentralized validation systems, though concentration on that topic takes away from the necessary focus on distributed ledgers.

For contracts, we highlight from standard contract literature the needed distinctions to get at the meaning of the word *trust* as in "trusted third parties," as featured in the introduction. These distinctions include full commitment versus limited commitment; contract theory allows the latter but this does not preclude some trades. Likewise, from contract theory and mechanism design we can speak of incentive compatibility for actions and truth-telling for messages for initial and interim unobserved states. Lack of trust in this specific sense is crucial in the design but is not an insuperable barrier. Related, it is key to prevent parties from pulling out or not performing. Default has remedies in collateral and/or reputation made explicit under long term contracting. Reneging and a restriction to time consistency can be remedied with commitment. Problems of collusion and non-uniqueness in implementation have remedies in better design. We need to get specific across the various trust aspects, and contract theory helps us do that, especially in hybrid systems. As for illustrative contexts, contracts and mechanisms are ubiquitous, and we come back to this with applications, below.

For the contract part of DLT innovation, we feature the key technical capabilities of smart contracts, highlighting the commitment in entering into the agreement and carrying it out, the immutability of its terms, and the conditionality aspects that help resolve the various trust problems, as agreed to options are executed automatically as a function of the state. Smart contracts help to mitigate underlying frictions in the environment. Key concepts here also include states of the system and transitions, commands, validity consensus, unique consensus, notaries and non-unique consensus, multiple trusting or non-trusted notaries, public and private nodes, oracles, and broadcast communication versus selectively private communication.

A key point is that the various aspects of trust and incentives that come from mechanism design can be implemented on smart contracts, from multiple sequential messages to costly state verification, incentive compatible actions and reports for the platforms, and beyond. Some specific examples are given, promised utility as a key state variable and how to build in commitment to undercut ex post bargaining and hold ups. Economists have had these mechanisms in mind, and so in some sense have taken smart contracts for granted without realizing it.

This topic of contracts and mechanism design contains within it and concludes with the centralization/decentralization issue. The "planner" problem of mechanism is a metaphor for economists, not taken literally per se, but in the context here is it clear that one has to specify how contracts and multi-party agreements are entered into and validated, and how they are executed over time. The doorkeepers and notaries involved with smart contract execution, if single entities, are hierarchical features, though competing non-trusting notaries conjure up the image of more decentralized systems.

Mostly, computer science and economics take different points of view regarding the meaning of the words trust. Symptomatic, the decentralized Bitcoin validation system is not incentive compatible. But there is encouraging common ground. Contracts with costly state verification are literally contracts executed with messages where over some range of states costly messages are not used. A version of the revelation principle works with no communication at all and with noisy messages, so one need not abandon mechanism design when facing the reality of imperfect messaging. Though there are impossibility results regarding consensus in the economics literature, there are also contributions on the effectiveness of multiple repeat messages and how iterations of decentralized validation can be truncated and achieve coordination.

A second part of the paper features specific implications of smart contracts and mechanism design for the design and the regulation of financial infrastructure.

The discussion begins with high value systems. The implications of mechanism design are a function of the underlying environment and the problem being analyzed. There is a push toward decentralization and partitioned ledgers, on the one hand, and for centralization and a role for hierarchical third parties, on the other. Specifically, mechanism design theory can be used with distributed ledgers to consider: permissioned and private ledgers and the gain from concealment, keeping secrets; delegation to third party platforms to deal with private and public shocks, that allow back- and front-loading in contracts; and design of market ledgers which mitigate if not eliminate bank and market runs through time-stamped and immutable records of the histories of

transactions. For that last result on runs, caveats from computer science come into play, namely, latency on networks, but potential remedies are noted.

This section then goes on to feature an example where innovation would allow large gains. In Thailand and other countries in Southeast Asia, traditional formal financial infrastructure is currently extremely limited. From Asian Development Bank studies, there are reportedly large gaps in services 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, as the risk premium is low for idiosyncratic risk, due to good pooling, but high for aggregate risk, even though aggregates differ across villages and could be pooled and hence better insured. Interventions have helped, but more is 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 increase in assets for high TFP SMEs which got funds. But loans were more readily available for village committee members and connections to them; some kinship pathways mitigated distortions, but evidently these rely on pre-existing trust. High TFP SMEs without kin could have benefited more. Smart contracts on distributed ledgers can help overcome these trust issues, allowing trade among strangers.

Innovations in this Southeast Asia context can come in two related forms. First, individual smart contracts are given as examples of how to take advantage of distributed ledger capability: escrow 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. A second form is contract competition with open access to providers and free entry, as in general equilibrium models with an intermediary broker sector.

However, the limits of contract competition are noted: Competition among providers can fail to complete the financial system due to unexploited complementarities and lack of coordination in a Nash equilibrium, even in environments in which all players are tiny, that is no player has any mass. A key issue for regulation is the timeline of competition and when to impose exclusivity, for example, free competition ex ante could be fine but with exclusivity and restrictions on trade ex post.

Featured as an example of innovation is EvryNet, an intelligent financial automation operating system that aims to provide open-source banking services and financial contracts to unbanked and underbanked populations. An interoperable smart contract platform enables not only traditional banks but also micro-finance institutions and others to initiate and execute banking products and financial contracts. Contracts can be provided at competitive prices for computer memory (storage) and computation power. A rating system tracks performance of providers.

A third part of the paper deals with the implications of the mechanism design and monetary theory literatures for the design of payments systems. A system which records histories of trade – which DLT can provide – 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 avoiding the problem 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 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, 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 public verified histories, to trace ownership.

However, the ingredients of some environments argue for partitioned ledgers with sustained private information, necessary to attain a constrained optimum, as when randomization is needed and messages should be kept secret. This last system is a hybrid, centralized/decentralized system, centralized in so far as all messages are put on a common ledger but decentralized in that that the ledger is not distributed. Notably, economic outcomes under contracts, data acquired, and data disseminated are all endogenous and co-determined.

Other model environments make clear that consensus and public ledgers can be needed for coordination and prudential regulation. A theorem in economics on the impossibility of decentralized monetary exchange is quite relevant. Knowledge of identities of agents, histories of trade and payments, and initial excess demands are needed for implementation of optimal allocations, not simply pairwise knowledge, of those contemporaneously matched, but also from others with whom the contemporaneous set of matched traders, the payment parties, have not been matched previously. Put simply, the history of trades needs to be on the common immutable ledger. Implementation of a Walrasian optimum in a decentralized way is both forward and backward looking. One has to know where the system should be headed, and the remaining options in the future to get there, hence what trades have been accomplished in the past and what trades are needed now in order to make this feasible.

For another example, circulating private debt is the medium of exchange for contemporaneous transactions in both short-term non-circulating securities and consumption commodities. Yet there can be many equilibria that achieve the Pareto optimal target, for example, either all the debts that are allowed to circulate could be issued by initial parties in one of the two locations, or by the parties in the other second location. But by assumption, in the informationally decentralized model environment, there is no way for traders in one location to know what is going on in the other. Mismatch is likely, too much or too little debt issued, with resulting crashes later. These conflicts and the need for coordination are likely to arise with multiple cryptocurrencies.

A problem with using distributed ledgers for coordination and regulation reemerges: the problem of scale. The two example environments above suggest that only key pieces of information need to be shared, which is encouraging. This of course will need to be worked out in practice.

Featured in this third part of the paper are two payments systems on distributed ledgers. One, an experiment, is Project Jasper of the Bank of Canada, which is designed for domestic interbank

payment settlement and is focused on payments (see Payments Canada 2017). The role of notary node is played by the Bank of Canada. Phase 2 of Jasper appears to be one of the first instances of a central queue within a DLT platform for payments, for matching and netting. A matching algorithm on a distributed ledger platform employs the language of states and conditionality. Netting promotes funding efficiency and smoother intraday payments flow. In Project Jasper, parties see only their own activity, a traditional point of view, while the role of the unique notary is played by the Bank of Canada. It is a sophisticated multi-party platform taking advantage of smart contract possibilities.

The second featured payments system is Velo, creating for money transfer operators (MTOs) in Southeast Asia a highly liquid decentralized settlement layer on a permissioned blockchain. Remittances in fiat money in SEA have transfer fees currently at 7.1%. The high transfer fees are partly due to legacy technology in the formal sector and limited access to formal currency exchange markets. The Velo Network is a settlement layer that avoids direct transfers of fiat money yet enables participants to efficiently conduct cross-border transactions. An optimized liquidity management layer will efficiently search for offsetting cross-country balances to allow for expedient clearing while minimizing risk. MTOs can be viewed as agents with varying underlying balance sheets hit by the needs for trade of their customers, hence needing credit and insurance. Some MTOs are already engaged in a contractual bilateral relationship, with both of their fiat credit lines supporting their financial obligations bound by a bilateral agreement. Velo can allow the implementation of constrained-optimal contract arrangements implemented on distributed ledgers, as envisioned in the examples on credit and insurance described earlier. Optimal information partitions are envisioned as part of the design.

11 References

Abel, Andrew B., N. Gregory Mankiw, Lawrence H. Summers, and Richard J. Zeckhauser. 1989. "Assessing Dynamic Efficiency: Theory and Evidence." *Review of Economic Studies* 56 (1): 1–19.

Abreu, Dilip, David Pearce, and Ennio Stacchetti. 1990.) "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica* 58 (5): 1041-1063.

Acemoglu, Daron, and Fabrizio Zilibotti. 1997. "Was Prometheus Unbound by Chance? Risk, Diversification, and Growth." *Journal of Political Economy* 105 (4):709-751.

Adrian, Tobias and Hyun Song Shin. 2009. "Money, Liquidity, and Monetary Policy." Federal Reserve Bank of New York, Staff Report No. 360.

Aiyagari, S. Rao. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics* 109(3):659-684.

Alvarez, Fernando and Francesco Lippi. 2009. "Financial Innovation and the Transactions Demand for Cash." *Econometrica* 77 (2): 363-402.

Alvarez, Fernando, Anan Pawasutipaisit, and Robert M. Townsend. 2018. "Cash Management in Village Thailand: Positive and Normative Implications." MIT Working Paper.

Amberg, Niklas, Tor Jacobson, Erik von Schedvin, and Robert M. Townsend. 2016. "Curbing Shocks to Corporate Liquidity: The Role of Trade Credit." National Bureau of Economic Research Working Paper, No. 22286.

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas. 2016. "Public Debt as Private Liquidity: Optimal Policy." National Bureau of Economic Research Working Paper, No. 22794.

ASEAN. 2017. "ASEAN Statistical Leaflet – Selected Key Indicators 2017." Pamphlet published by ASEAN Secretariat, Jakarta, October 2017. Accessed October 19, 2018 at: https://www.aseanstats.org/wp-content/uploads/2017/11/ASEAN-Statistical-Leaflet-2017_Final.pdf

Asian Development Bank (2017) "Accelerating Financial Inclusion in South-East Asia with Digital Finance." ADB.org, accessed November 30, 2018 at: http://dx.doi.org/10.22617/RPT178622-2.

Athey, Susan, Ivo Parashkevov, Vishnu Sarukkai, and Jing Xia. 2016. "Bitcoin Pricing, Adoption, and Usage: Theory and Evidence." Stanford University Graduate School of Business Research Paper 16-42. Accessed online: https://ssrn.com/abstract=2826674.

Auclert, Adrien. 2017. "Monetary Policy and the Redistribution Channel." Working paper.

Banerjee, Abhijit, Emily Breza, Robert M. Townsend and Diego Vera-Cossio. 2018. "Access to Credit and Productivity: Evidence from Thai Villages." Working paper, University of California, San Diego.

Bank for International Settlements (BIS). 2017a. "Distributed ledger technology in payment, clearing and settlement: An analytical framework." Committee on Payments and Market Infrastructures Papers No. 157 (February, 2017).

Bank for International Settlements (BIS). 2017b. "Distributed Ledgers In Payment, Clearing And Settlement Carry Promise As Well As Risks." Press release, 27 February 2017. Accessed November 13, 2018. https://www.bis.org/press/p170227.htm.

Bank for International Settlements (BIS). 2018. "Cryptocurrencies: Looking beyond the hype." *BIS Annual Economic Report*, pp. 91-109. Accessed November 30, 2018 at: www.bis.org/publ/arpdf/ar2018e.htm.

Bech, Morten and Rodney Garratt. 2017. "Central bank cryptocurrencies." BIS Quarterly Review (Sept. 2017): 55-70. Accessed October 22, 2018: https://www.bis.org/publ/qtrpdf/r_qt1709f.pdf

Bersem, Mario, Enrico Perotti and Ernst-Ludwig von Thadden. 2012. "Sand in the Wheels of Capitalism, On the Political Economy of Capital Market Frictions." Working Paper, Copenhagen Business School.

Bewley, Truman. 1983. "A Difficulty with the Optimum Quantity of Money." *Econometrica* 51(5):1485-1504.

Bharadwaj, Prashant, William Jack, and Tavneet Suri. 2018. "Can Digital Loans Deliver? Take Up and Impacts of Digital Loans in Kenya." Working paper.

Biais, Bruno, Christophe Bisiere, Matthieu Bouvard, Catherine Casamatta and Albert Menkveld. 2018. "Equilibrium Bitcoin Pricing." Working paper, Toulouse School of Economics.

Bronstein, Max. 2018. "Mapping the Decentralized Financial System: More transparent, open, and programmable financial services." TokenEconomy.co, posted Aug 7, 2018. https://tokeneconomy.co/mapping-the-decentralized-financial-system-7c5af65e0335

Browning, Martin, Thomas F. Crossley, and Joachim Winter. 2014. "The Measurement of Household Consumption Expenditures." *Annual Review of Economics* 6 (6): 475–501.

Bryant, John, and Neil Wallace. 1984. "A price discrimination analysis of monetary policy." *Review of Economic Studies* 51 (2): 279-288.

Budish, Eric. 2018. "The Economic Limits of Bitcoin and the Blockchain." Working paper, University of Chicago.

Budish, Eric, Peter Cramton and John Shim. 2015. "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response." *Quarterly Journal of Economics* 130 (4): 1547–1621.

Bureau of Economic Analysis. 2017. *NIPA Handbook: Concepts and Methods of the U.S. National Income and Product Accounts*. U.S. Department of Commerce, https://www.bea.gov/sites/default/files/methodologies/nipa-handbook-all-chapters.pdf

Campbell-Kelly, Martin. 2010. "Victorian Data Processing." *Communications of the ACM* 53 (10): 19-21.

Carlson, Mark, Burcu Duygan-Bump, Fabio Natalucci, Bill Nelson, Marcelo Ochoa, and Jeremy Stein. 2016. "The Demand for Short-Term, Safe Assets and Financial Stability: Some Evidence and Implications for Central Bank Policies." *International Journal of Central Banking* 12 (4): 307-333.

Carlson, Stacy. 2017. "Essays in Financial Innovation and Development." Ph.D. Diss., MIT.

Casey, Michael, Jonah Crane, Gary Gensler, Simon Johnson and Neha Narula. 2018. *The Impact of Blockchain Technology on Finance: A Catalyst for Change*. Geneva: International Center for Monetary and Banking Studies, Volume 21 of Geneva Reports on the World Economy. Last accessed April 11, 2019: https://www.cimb.ch/uploads/1/1/5/4/115414161/geneva21_1.pdf.

Cass, David, Masahiro Okuno, and Itzhak Zilcha. 1978. "The Role of Money in Supporting the Pareto Optimality of Competitive Equilibrium in Consumption Loan Type Models." In *Models of Monetary Economies: Proceedings and Contributions from Participants of a December 1978 Conference.* Federal Reserve Bank of Minneapolis Conference, pp. 13-48.

Cao, Sean S., Lin Cong, and Baozhong Yang. 2018. "Auditing and Blockchains: Pricing, Misstatements, and Regulation." Working paper, Georgia State University.

Chandrasekhar, Arun G., Robert M. Townsend, and Juan Pablo Xandri. 2018. "Financial Centrality and Liquidity Provision." Working paper, MIT.

Chiappori, Pierre-André, Krislert Samphantharak, Sam Schulhofer-Wohl, and Robert M. Townsend. 2014. "Heterogeneity and Risk Sharing in Village Economies." *Quantitative Economics* 5(1):1-27.

Cipolla, Carlo M. 1956. *Money, Prices, and Civilization in the Mediterranean World, Fifth to Seventeenth Century*. Princeton, NJ: Princeton University Press.

Clower, Robert. 1967. "A Reconsideration of the Microfoundations of Monetary Theory." *Economic Inquiry* 6 (1): 1-8.

Cocco, João F., Francisco J. Gomes, and Nuno C. Martins. 2009. "Lending Relationships in the Interbank Market." *Journal of Financial Intermediation* 18 (2009) 24-48.

Coles, Peter A., and Ran Shorrer. 2012. "Correlation in the Multiplayer Electronic Mail Game." *BE Journal of Theoretical Economics* 12 (1). Published online 2012-05-04, doi:10.1515/1935-1704.1576.

Columbus, Louis. 2019. "Top 10 Ways Internet Of Things And Blockchain Strengthen Supply Chains." Forbes.com, posted Jan 13, 2019 11:17 am. https://www.forbes.com/sites/louiscolumbus/2019/01/13/top-10-ways-internet-of-things-and-blockchain-strengthen-supply-chains/#4fea18ac5e4e

Cong, Lin and Zhiguo He. Forthcoming. "Blockchain Disruption and Smart Contracts." Forthcoming in *Review of Financial Studies*.

Cong, Lin, Zhiguo He and Jiasun Li. 2018. "Decentralized Mining in Centralized Pools." George Mason University School of Business Research Paper No. 18-9.

Crouzet, Nicolas, Apoorv Gupta and Filippo Mezzanotti. 2019. "Shocks and Technology Adoption: Evidence from Electronic Payment Systems." Working paper, Northwestern University.

Curran, Brian. 2018. "What is Practical Byzantine Fault Tolerance? Complete Beginner's Guide." Blockonomi.com. Last updated May 11, 2018. https://blockonomi.com/practical-byzantine-fault-tolerance/

Dandekar, Pranav, Ashish Goel, Ramesh Govindan and Ian Post. 2012. "Liquidity in Credit Networks: A Little Trust Goes a Long Way." Working Paper, Stanford University.

Dasgupta, Partha, Peter Hammond, and Eric Maskin. 1979. "The Implementation of Social Choice Rules: Some General Results on Incentive Compatibility." *Review of Economic Studies* 46 (2): 185-216.

Denison, Erin, Michael Lee, and Antoine Martin. 2016. "What do cryptocurrencies do?" Federal Reserve Bank of New York Conference Paper.

Diamond, Douglas W. and Philip H. Dybvig. 1983. "Bank Runs, Deposit Insurance, and Liquidity." *Journal of Political Economy* 91 (3): 401-419.

Diamond, Peter A. 1965. "National Debt in a Neoclassical Growth Model." *American Economic Review* 55 (5) Part 1: 1126-1150.

Doepke, Matthias and Martin Schneider. 2006. "Inflation and the Redistribution of Nominal Wealth." *Journal of Political Economy* 114 (6): 1069-1097.

Doepke, Matthias and Robert M. Townsend. 2006. "Dynamic Mechanism Design with Hidden Income and Hidden Actions." *Journal of Economic Theory,* 126 (1): 235-285.

Domeij, David and Tore Ellingsen. 2018. "Monetary Policy in Incomplete Markets." Working paper, Stockholm School of Economics.

Duffie, Darrell, Nicolae Gârleanu, and Lasse Heje Pedersen. 2005. "Over-the-Counter Markets." *Econometrica* 73 (6):1815-1847.

Falkon, Samuel. 2017. "The Story of the DAO—Its History and Consequences." Medium.com, posted Dec 24, 2017, last accessed February 7, 2019. https://medium.com/swlh/the-story-of-the-dao-its-history-and-consequences-71e6a8a551ee.

Fernandes, Ana and Christopher Phelan. 2000. "A Recursive Formulation for Repeated Agency with History Dependence." *Journal of Economic Theory* 91 (2): 223–47.

Fernández-Villaverde, Jesús and Daniel Sanches. 2016. "Can Currency Competition Work?" NBER Working Paper 22157.

Fischer, Michael J., Nancy A. Lynch and Michael S. Paterson. 1985. "Impossibility of Distributed Consensus with One Faulty Process." *Journal of the Association for Computing Machinery* 32 (2): 374-382.

Foley, Sean, Jonathan R. Karlsen, and Talis J. Putnins (2018) "Sex, Drugs, and Bitcoin: How Much Illegal Activity Is Financed Through Cryptocurrencies?" *Review of Financial Studies*, Forthcoming. SSRN, last revised 15 Dec 2018, https://ssrn.com/abstract=3102645.

Freeman, Scott. 1996. "The Payments System, Liquidity, and Rediscounting." *American Economic Review* 86 (5): 1126-1138.

Fulford, Scott, and Scott Schuh. 2017. "Credit Card Utilization and Consumption over the Life Cycle and Business Cycle." Federal Reserve Bank of Boston Working Paper 17–14.

G20 Research Group. 2013. "The 2013 G20 St. Petersburg Summit Commitments." Accessed on October 1, 2018. http://www.g20.utoronto.ca/analysis/commitments-13-stpetersburg.html

Gale, Douglas and Martin Hellwig. 1985. "Incentive-Compatible Debt Contracts: The One-Period Problem." *Review of Economic Studies* 52 (4): 647-663.

Gans, Joshua S. 2018. "The Fine Print in Smart Contracts." Working paper, University of Toronto.

Garratt, Rodney, Antoine Martin, Michael Junho Lee and Robert M. Townsend. 2018. "Endogenous Liquidity and Interdealer Trading in Over-The-Counter Markets." Working paper, Federal Reserve Bank of New York.

Garratt, Rodney and Neil Wallace. 2018. “Bitcoin 1, Bitcoin 2, ... : An Experiment in Privately Issued Outside Monies.” *Economic Inquiry* 56 (3): 1887–1897.

Geanakoplos, John. 2003. “Liquidity, Default, and Crashes, Endogenous Contracts in General Equilibrium.” In M. Dewatripont, L. P. Hansen, and S. J. Turnovsky, eds., *Advances in Economics and Econometrics: Theory and Applications, Eighth World Conference Volume II*. Cambridge, UK: Cambridge University Press, pp. 170-205

Geerolf, François. 2017. “Reassessing Dynamic Efficiency.” Working Paper, UCLA.

Glazer, Phil. 2018. “Decentralized Cryptocurrency Exchanges.” HackerNoon.com, updated Mar 5, 2018. https://hackernoon.com/decentralized-cryptocurrency-exchanges-93039613eeb7

Gord, Michael. 2016. “Smart Contracts Described by Nick Szabo 20 Years Ago Now Becoming Reality.” Bitcoin Magazine, last updated Apr 26, 2016 2:02 PM EST. https://bitcoinmagazine.com/articles/smart-contracts-described-by-nick-szabo-years-ago-now-becoming-reality-1461693751/

Green, Edward J. 1987. “Lending and the Smoothing of Uninsurable Income.” In *Contractual Arrangements for Intertemporal Trade,* eds., Edward C. Prescott and Neil Wallace. Minneapolis, MN: University of Minnesota Press, pp. 3-25.

Green, Edward J. 1999. “Money and Debt in the Structure of Payments.” *Federal Reserve Bank of Minneapolis Quarterly Review* 23 (Spring): 13–29.

Green, Edward J. and Ping Lin. 2003. “Implementing Efficient Allocations in a Model of Financial Intermediation,” *Journal of Economic Theory* 109 (1): 1–23.

Green, Edward J. and Soo-Nam Oh. 1991. “Contracts, Constraints and Consumption.” *Review of Economic Studies* 58 (5): 883-99.

Green, Edward J. and Ruilin Zhou. 2005. “Money as a Mechanism in a Bewley Economy.” *International Economic Review* 46 (2):351-371.

Greenwood, Robin, Samuel G. Hanson, and Jeremy C. Stein. 2016. “The Federal Reserve's Balance Sheet as a Financial-Stability Tool.” 2016 Economic Policy Symposium Proceedings. Jackson Hole: Federal Reserve Bank of Kansas City.

Griffin, John and Amin Shams. 2018. “Is Bitcoin Really Un-Tethered.” Working paper, University of Texas Austin.

Güntzer, Michael M., Dieter Jungnickel and Matthias Leclerc. 1998. “Efficient algorithms for the clearing of interbank payments.” *European Journal of Operational Research* 106 (1): 212-219.

Harford, Tim. 2017. “What tally sticks tell us about how money works.” BBC.com, last updated 10 July 2017. https://www.bbc.com/news/business-40189959

Harris, Milton and Artur Raviv. 1979. “Optimal Incentive Contracts with Imperfect Information.” *Journal of Economic Theory* 20 (2): 231-259.

Harris, Milton and Robert M. Townsend. 1981. “Resource Allocation under Asymmetric Information.” *Econometrica* 49 (1): 33-64.

Hart, Oliver and John Moore. 2007. "Incomplete Contracts and Ownership: Some New Thoughts." *American Economic Review* 97 (2):182-186

Hendershott, Terrence and Ananth Madhavan. 2015. "Click or Call: Auction vs. Search in the OTC Market." *Journal of Finance* 70 (1): 419-447.

Hinzen, Franz J., Kose John and Fahad Saleh. 2019. "Proof-of-Work's Limited Adoption Problem." Working paper, NYU Stern School of Business.

Holden, Richard and Anup Malani. 2018. "Can Blockchains Solve the Holdup Problem in Contracts." Working Paper, No.2018-12, Becker-Friedman Institute, Chicago.

Huggett, Mark and Greg Kaplan. 2015. "How Large is the Stock Component of Human Capital?" NBER Working Paper 21238.

Hussam, Reshmaan, Natalia Rigol, and Benjamin Roth. 2017. "Targeting High Ability Entrepreneurs Using Community Information: Mechanism Design in The Field." Working Paper, Harvard University.

Iansiti, Marco and Karim R. Lakhani. 2017. "The Truth about Blockchain." *Harvard Business Review* 95 (1): 118-127.

Ingves, Stefan. 2016. "Efficient Payment Systems and the Riksbank's Approach to Cash Distribution." Money and Banking Conference, Central Bank of Argentina.

Jack, William and Tavneet Suri. 2011. "Mobile Money: The Economics of M-PESA." NBER Working Paper No. 16721.

Jack, William and Tavneet Suri. 2014. "Risk Sharing and Transactions Costs: Evidence from Kenya's Mobile Money Revolution." *American Economic Review,* 104 (1):183-223.

Jack, William, Tavneet Suri, and Robert M. Townsend. 2010. "Monetary Theory and Electronic Money: Reflections on the Kenyan Experience." *Economic Quarterly* 96 (1): 83–122.

Jacklin, Charles J. 1987. "Demand Deposits, Trading Restrictions, and Risk Sharing." In Prescott, Edward C. and Neil Wallace (eds.) *Contractual Arrangements for Intertemporal Trade*. Minneapolis, MN: University of Minnesota Press, pp. 26–47.

Janin, Simon, Arthur Gervais and Akaki Mamageishvili. 2019. "FileBounty: Secure and Efficient File Exchange in Rational Adversarial Environment." Working paper, ETH Zurich.

Joaquim, Gustavo, Robert M. Townsend and Victor Zhorin. 2018. "Optimal Contracting and Imperfect Competition among Financial Service Providers." Working paper, MIT.

Kaboski, Joseph P. and Robert M. Townsend. 2011. "A Structural Evaluation of Large-Scale Quasi-experimental Microfinance Initiative." *Econometrica* 79 (5):1357-1406

Kaboski, Joseph P. and Townsend Robert M. 2012. "The Impact of Credit on Village Economies." *American Economic Journal: Applied Economics* 4 (2): 98-133.

Kaplan, Greg and Giovanni L. Violante. 2014. "A Model of the Consumption Response to Fiscal Stimulus Payments." *Econometrica* 82 (4):1199-1239.

Karaivanov, Alexander, and Robert M Townsend. 2014. “Dynamic Financial Constraints: Distinguishing Mechanism Design from Exogenously Incomplete Regimes.” *Econometrica* 82 (3): 887-959.

Kareken, John, and Neil Wallace. 1981. “On the indeterminacy of equilibrium exchange rates.” *Quarterly Journal of Economics* 96 (2): 207-222.

Kehoe, Timothy J., David K. Levine, and Michael Woodford. 1990. “The Optimum Quantity of Money Revisited.” FRB-Minneapolis Working Paper 404.

Kestenbaum, David. 2012. “The Accountant Who Changed the World.” All Things Considered. National Public Radio (NPR). Broadcast October 4, 2012.

Kilenthong, Weerachart T. and Robert M. Townsend. 2018. “A Market Based Solution for Fire Sales and Other Pecuniary Externalities.” Working paper, MIT.

Kim, Kyungmin. 2015. “Summary of ‘Money and the Decentralization of Exchange’ and Some Comments.” Working paper, Board of Governors of the Federal Reserve System.

Kinnan, Cynthia, Krislert Samphantharak, Robert M. Townsend, and Diego Vera-Cossio. 2018. “Networks and Risk Sharing in Village Economies.” Working paper, MIT.

Kinnan, Cynthia, and Robert Townsend. 2012. "Kinship and Financial Networks, Formal Financial Access, and Risk Reduction." *American Economic Review* 102 (3): 289-93.

Klein, Benjamin. 1976. “Competing Monies: Comment.” *Journal of Money* 8 (4): 512-519.

Kocherlakota, Narayana R. 1998. “Money Is Memory.” *Journal of Economic Theory* 81 (2): 232-251.

Krishnamurthy, Arvind and Annette Vissing-Jorgensen. 2012. “The Aggregate Demand for Treasury Debt.” *Journal of Political Economy* 120.2: 233-267.

Kuussaari, Harri. 1996. “Systemic Risk in the Finnish Payment System: An Empirical Investigation.” Bank of Finland Discussion Paper No. 3/96.

Lagos, Ricardo and Randall Wright. 2005. “A Unified Framework for Theory and Policy Analysis.” *Journal of Political Economy* 113 (3): 463-484.

Lagos, Ricardo and Shengxing Zhang. 2018. “Turnover Liquidity and the Transmission of Monetary Policy.” Federal Reserve Bank of Minneapolis Working Paper No. 734.

Lamport, Leslie, Robert Shostak, and Marshall Pease. 1982. “The Byzantine Generals Problem.” *ACM Transactions on Programming Languages and Systems* 4 (3): 382-401.

Lehnert, Andreas, Ethan Ligon and Robert M. Townsend. 1999. “Liquidity Constraints and Incentive Contracts.” *Macroeconomic Dynamics* 3 (1):1-47.

Levine, David K. 1989. “Efficiency and the Value of Money.” *Review of Economic Studies* 56 (1): 77-88.

Leong, Nicholas. 2017. “Remittances are ripping off migrant workers in ASEAN.” EastAsiaForum.org, last updated 8 November 2017.

https://www.eastasiaforum.org/2017/11/08/remittances-are-ripping-off-migrant-workers-in-asean/.

Li, Dan and Norman Schürhoff. 2012. "Dealer Networks: Market Quality in Over-The-Counter Markets." Working paper, Federal Reserve Bank of New York.

Ligon, Ethan, Jonathan P. Thomas, and Tim Worrall. 2002. "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies." *Review of Economic Studies* 69 (1): 209–44.

Lim, Youngjae and Robert M. Townsend. 1998. "General Equilibrium Models of Financial Systems: Theory and Measurement in Village Economies." *Review of Economic Dynamics* 1 (1): 59-118.

Lucas, Robert E. and Nancy L. Stokey. 1983. "Optimal Fiscal and Monetary Policy in an Economy without Capital." *Journal of Monetary Economics* 12 (1): 55-93.

Lyon, Richard K. 1996. "Optimal Transparency in a Dealer Market with An Application to Foreign Exchange." *Journal of Financial Intermediation*, 5 (3): 225-254.

Makowski, Louis. 1980. "A Characterization of Perfectly Competitive Economies with Production," *Journal of Economic Theory* 22 (2), 208–221.

Mallett, Jacky. 2009. "Limits on the Communication of Knowledge in Human Organisations." *Studies in Emergent Order* 2: 1-18.

Manuelli, Rodolfo and Thomas J. Sargent. 1988. "Longer Trading Periods in a Townsend Turnpike Model." Working Paper, Stanford University.

Manuelli, Rodolfo and Thomas J. Sargent. 2010. "Alternative Monetary Policies in a Turnpike Economy." *Macroeconomic Dynamics* 14 (5): 727-762.

Martin, Antoine and James McAndrews. 2008. "An Economic Analysis of Liquidity-Saving Mechanism." *Economic Policy Review* 14 (2): 25-39.

Martin, Antoine, David Skeie, and Ernst-Ludwig Von Thadden. 2013. "The Fragility of Short-term Secured Funding Markets." Federal Reserve Bank of New York Staff Reports No. 630.

Martin, Antoine, David Skeie, and Ernst-Ludwig von Thadden. 2014. "Repo Runs." *Review of Financial Studies* 27 (4): 957-989

Maskin, Eric and Jean Tirole. 1999. "Unforeseen contingencies and incomplete contracts." *Review of Economic Studies* 66 (1): 83-114.

Mazieres, David. 2016. "The Stellar Consensus Protocol: A Federated Model for Internet-level Consciousness." Stellar.org, accessed April 12, 2019: https://www.stellar.org/papers/stellar-consensus-protocol.pdf.

McAndrews, James and William Roberds. 1999. "Payment Intermediation and the Origins of Banking." Federal Reserve Bank of Atlanta Working Paper 99-11.

Miller, Merton, H. and Daniel Orr. 1966. "A Model of the Demand for Money by Firms." *Quarterly Journal of Economics* 80: 413-435.

Mills, David, Kathy Wang, Brendan Malone, Anjana Ravi, Jeff Marquardt, Clinton Chen, Anton Badev, Timothy Brezinski, Linda Fahy, Kimberley Liao, Vanessa Kargenian, Max Ellithorpe, Wendy Ng, and Maria Baird (2016) "Distributed ledger technology in payments, clearing, and settlement", Federal Reserve Board Finance and Economics Discussion Series, No. 2016-095.

Monetary Authority of Singapore and the Association of Banks in Singapore. 2017. "Project Ubin Phase 2: Re-imagining Interbank Real-Time Gross Settlement System Using Distributed Ledger Technologies." Project report. Accessed November 15, 2018. https://www.accenture.com/t20171116T081243Z__w__/sg-en/_acnmedia/PDF-66/Accenture-Project-Ubin-Phase-2.pdf

Monetary Authority of Singapore. 2018. "Cross-Border Interbank Payments and Settlements: Emerging opportunities for digital transformation." Project report, Accessed March 22, 2019.http://www.mas.gov.sg/~/media/ProjectUbin/Cross%20Border%20Interbank%20Payments%20and%20Settlements.pdf

Monnet, Cyril and Thomas Nellen. 2014. "The Collateral Costs of Clearing." SNB Working Papers, 4/2014.

Moore, John. 1992. "Implementation, contracts, and renegotiation in environments with complete information." In Jean-Jacques Laffont (Ed.) Advances in Economic Theory: Sixth World Congress. Volume 1. Cambridge: Cambridge University Press, pp. 182–282.

Moore, John and Rafael Repullo. 1988. "Subgame perfect implementation," *Econometrica* 56(5):1191-1220.

Moore, John and Rafael Repullo. 1990. "Nash Implementation: A Full Characterization." *Econometrica* 58 (5): 1083-1099.

Moskov, Phillip. 2018. "What Is Bit Gold? The Brainchild of Blockchain Pioneer Nick Szabo," Coincentral.com, posted 22 May 2018, accessed 2/14/18, https://coincentral.com/what-is-bit-gold-the-brainchild-of-blockchain-pioneer-nick-szabo/

Muley, Ameya. 2016. "Collateral Reuse in Shadow Banking and Monetary Policy." Job Market Paper, MIT.

Myerson, Roger B. 1982. "Optimal Coordination Mechanisms in Generalized Principal-Agent Problems." *Journal of Mathematical Economics* 10 (1): 67-81.

Myerson, Roger B. 1986. "Multistage Games with Communication." *Econometrica* 54 (2): 323-58.

Nakamoto, Satoshi. 2008. "Bitcoin: A Peer-to-Peer Electronic Cash System." White paper, accessed October 28, 2018, https://bitcoin.org/bitcoin.pdf.

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

Narayanan, Arvind and Matt Weinberg. 2018. "Survey of Algorithmic Game Theory Research Questions for Crypto/Blockchain." Lecture presented at "Cryptocurrencies and Blockchains," Becker-Friedman Institute, University of Chicago, November, 2018.

Narula, Neha, Willy Vasquez and Madars Virza. 2018. "zkLedger: Privacy-Preserving Auditing for Distributed Ledgers." In *Proceedings of the 15th USENIX Symposium on Networked Systems Design and Implementation*, Renton, WA, April 9–11, 2018. Renton, WA: USENIX, pp. 65-80.

Ostroy, Joseph M. and Ross M. Starr. 1974. "Money and the Decentralization of Exchange." *Econometrica* 42 (6): 1093-1113.

Palfrey, Thomas R. and Sanjay Srivastava. 1989. "Mechanism Design with Incomplete Information: A Solution to the Implementation Problem." *Journal of Political Economy* 97 (3):668-691.

Paweenawat, Archawa, and Robert M. Townsend. 2012. "Village Economic Accounts: Real and Financial Intertwined." *American Economics Review* 102 (3): 441-46.

Paweenawat, Archawa, and Robert M. Townsend. 2018. "The Impact of Regional Isolationism: Disentangling Real and Financial Factors," Working Paper, MIT.

Payments Canada. 2017. "Project Jasper: A Canadian Experiment with Distributed Ledger Technology for Domestic Interbank Payments Settlement." White paper prepared by Payments Canada, Bank of Canada and R3, September 29, 2017. Accessed on October 19, 2018. https://www.payments.ca/sites/default/files/29-Sep-17/jasper_report_eng.pdf

Pesendorfer, W. 1995. Financial innovation in a general equilibrium model." *Journal of Economic Theory* 65 (1), 79−116.

Phelan, Christopher and Robert M. Townsend. 1991. "Computing Multiperiod Information Constrained Optima." *Review of Economic Studies*. 58 (5): 853-883.

Piazzesi, Monika and Martin Schneider. 2018. "Payments, Credit and Asset Prices." Working Paper, Stanford University.

Prat, Julien and Benjamin Walter. 2018. "An Equilibrium Model of the Market for Bitcoin Mining." Working paper, École Polytechnique.

Prescott, Edward S. 2003. "Communication in Private-Information Models: Theory and Computation." *Geneva Papers on Risk and Insurance Theory* 28 (2): 105-30.

Prescott, Edward C. and Robert M. Townsend. 1984a. "Pareto Optima and Competitive Equilibria with Adverse Selection and Moral Hazard." Econometrica 52 (1), 21–45.

Prescott, Edward C. and Robert M. Townsend. 1984b. "General Competitive Analysis in an Economy with Private Information." International Economic Review 25 (1):1-20.

Prescott, Edward S. and Robert M. Townsend. 2006a. "Firms as Clubs in Walrasian Markets with Private Information." *Journal of Political Economy* 114 (4), 644–671.

Prescott, Edward S. and Robert M. Townsend. 2006b. "Private Information and Intertemporal Job Assignments." *Review of Economic Studies* 73 (2): 531–48.

Ray, Shaan. 2017. “Blockchains versus Traditional Databases.” Emerging Technology Blog, HackerNoon.com, updated November 5, 2017. https://hackernoon.com/blockchains-versus-traditional-databases-c1a728159f79.

Ray, Shaan. 2018. “Federated Byzantine Agreement.” Emerging Technology Blog, TowardsDataScience.com, updated April 8, 2018, https://towardsdatascience.com/federated-byzantine-agreement-24ec57bf36e0.

Reese, Frederick. 2017. “Land registry: A Big Blockchain Explored.” CoinDesk.com. Accessed on October 1, 2018. www.coindesk.com/blockchain-land-registry-solution-seeking-problem/

Robinson, Henry. 2009. “Barbara Liskov's Turing Award, and Byzantine Fault Tolerance.” The Paper Trail Blog, posted March 30, 2009. https://www.the-paper-trail.org/post/2009-03-30-barbara-liskovs-turing-award-and-byzantine-fault-tolerance/

Rogoff, Kenneth S. 2016. *The Curse of Cash*. Princeton: Princeton University Press

Roth, Benjamin N. and Ran I. Shorrer. 2017. “Making it Safe to Use Centralized Marketplaces: Dominant Individual Rationality and Applications to Market Design.” Working paper, Harvard Business School.

Rubinstein, Ariel. 1989. “The Electronic Mail Game: Strategic Behavior Under ‘Almost Common Knowledge’.” *American Economic Review* 79 (3): 385-91.

Ru, Hong and Robert M. Townsend. 2018. “Narrowing the GAP: The Costly State Verification Regime in Rural Thailand.” Working paper, MIT.

Samphantharak, Krislert, Scott Schuh, and Robert M. Townsend. 2016. “Integrated Household Surveys: An Assessment of U.S. Methods and an Innovation.” Working paper, Federal Reserve Bank of Boston.

Samphantharak, Krislert and Robert M. Townsend. 2009. *Households as Corporate Firms: An Analysis of Household Finance Using Integrated Constructing Financial Statements from Integrated Household Surveys and Corporate Financial Accounting.* Cambridge; New York: Cambridge University Press.

Samphantharak, Krislert, and Robert M. Townsend. 2018. “Risk and Return in Village Economies.” *American Economic Journal-Microeconomics* 10 (1): 1-40**.**

Samuelson, Paul A. 1958. “An Exact Consumption-Loan Model of Interest With or Without the Social Contrivance of Money.” *Journal of Political Economy* 66 (6): 467-482.

Schilling, L. and Harold Uhlig. 2018. “Some Simple Bitcoin Economics.” Becker Friedman Institute for Research in Economics Working Paper 2018-21.

Schmandt-Besserat, Denise. 2014. “Tokens: their Significance for the Origin of Counting and Writing.” UTexas.edu, last updated 3/2/14, accessed 1/29/19: https://sites.utexas.edu/dsb/tokens/tokens/

Schuh, Scott and Joanna Stavins. 2014. “The 2011 and 2012 Surveys of Consumer Payment Choice.” Federal Reserve Bank of Boston Research Data Reports No. 14-1.

Shin, Hyun Song. 2018. “Cryptocurrencies and the economics of money.” Speech given at Bank for International Settlements’ Annual General Meeting, Basel, 24 June 2018.

Singh, Manmohan. 2011. “Velocity of Pledged Collateral: Analysis and Implications.” Working paper, IMF.

Skingsley, Cecilia. 2016. “Should the Riksbank Issue e-Krona?” Lecture presented at FinTech Stockholm, Berns, Sweden 16 November 2016.

Spear, Stephen E. and Sanjay Srivastava. 1987. “On Repeated Moral Hazard with Discounting.” *Review of Economic Studies.* 54 (4): 599-617.

Sprague, Oliver M. W. 1910. “History of Crises under the National Banking System.” National Monetary Commission (1910 [1968]), 51st Congress, Second Session, Senate Document No. 538.

Sripakdeevong, Parit and Robert M. Townsend (2017) “The Village Money Market Revealed: Credit chains and shadow banking,” Working paper, MIT

Starr, Ross M. 1974. “The Price of Money in a Pure Exchange Monetary Economy with Taxation,” *Econometrica* 42 (1): 45-54

Suri, Tavneet. 2017. “Mobile Money.” *Annual Review of Economics* 9: 497-520.

Sveriges Riksbank. 2018. “Payment patterns in Sweden 2018.” Riksbank.se. Accessed on October 1, 2018. https://www.riksbank.se/globalassets/media/statistik/betalningsstatistik/2018/payments-patterns-in-sweden-2018.pdf.

Swedish Post and Telecom Authority. 2017. “Grundläggande betaltjänster i en digitaliserad framtid (Essential Payment Services in a Digitalised Future).” Stockholm. *Swedish Post and Telecom Authority Report No. 20.* PTS-ER-2017:20. Accessed on October 1, 2018. https://www.pts.se/globalassets/startpage/dokument/icke-legala-dokument/rapporter/2017/post/grundlaggande-betaltjanster-i-en-digitaliserad-framtid---pts-er-2017_20.pdf.

Szabo, Nick. 1998. “Secure Property Titles with Owner Authority.” Satoshi Nakamoto Institute. Accessed 2/19/2019: https://nakamotoinstitute.org/secure-property-titles/.

Tirole, Jean. 1985. “Asset Bubbles and Overlapping Generations.” *Econometrica* 53 (6): 1499-1528.

Tobin, James. 1978. “Discussion.” In *Models of Monetary Economies: Proceedings and Contributions from Participants of a December 1978 Conference.* Federal Reserve Bank of Minneapolis Conference, pp. 83-90.

Townsend, Robert M. 1977. “The Eventual Failure of Price Fixing Schemes.” *Journal of Economic Theory* 14 (3): 190-99.

Townsend, Robert M. 1978. “Intermediation with Costly Bilateral Exchange.” *Review of Economic Studies* 45 (3): 417-25.

Townsend, Robert M. 1979. “Optimal contracts and competitive markets with costly state verification.” Journal of Economic Theory 21 (2): 265-293

Townsend, Robert M. 1980. “Models of Money with Spatially Separated Agents.” In *Models of Monetary Economies,* edited by John Kareken and Neil Wallace. Federal Reserve Bank of Minneapolis: 265-303.

Townsend, Robert M. 1982. “Optimal Multiperiod Contracts and the Gain from Enduring Relationships under Private Information.” *Journal of Political Economy* 90 (6): 1166-1186.

Townsend, Robert M. 1987. “Economic Organization with Limited Communication.” *American Economic Review* 77 (5): 954-971.

Townsend, Robert M. 1988. “Information Constrained Insurance: The Revelation Principle Extended.” *Journal of Monetary Economics* 21 (2/3): 411-450.

Townsend, Robert M. 1989. “Currency and Credit in a Private Information Economy.” *Journal of Political Economy* 97 (6):1323-1344.

Townsend, Robert M. 1990. *Financial Structure and Economic Organization*. Cambridge, MA: Blackwell.

Townsend, Robert M. 2016. “Village and Larger Economies: The Theory and Measurement of the Townsend Thai Project.” *Journal of Economic Perspectives* 30 (4):199-220.

Townsend, Robert M. and Neil Wallace. 1987. “Circulating Private Debt: An Example with a Coordination Problem.” In *Contractual Arrangements for Intertemporal Trade* edited by Edward C. Prescott and Neil Wallace. University of Minnesota Press.

Townsend, Robert, and Juan Pablo Xandri. 2018. “Regulation and the Optimal Design of Financial Markets” Working paper, MIT.

Trubek, Anne. 2015. “What the Heck is Cuneiform, Anyway?” Smithsonian.com. last updated October 20, 2015: https://www.smithsonianmag.com/history/what-heck-cuneiform-anyway-180956999/#dcLC5HSIETSzgobb.03.

Tucker, Paul. 2014. “Are Clearing Houses the New Central Banks?” Over-the-Counter Derivatives Symposium, Chicago.

Vera-Cossio, Diego A. 2018. “Targeting Credit through Community Members.” Working paper, University of California, San Diego.

Wallace, Neil. 2014. “Optimal Money Creation in 'Pure Currency' Economies: A Conjecture.” *Quarterly Journal of Economics* 129 (1): 259-274.

Wikipedia contributors. 2018a. “Informal value transfer system,” Wikipedia, The Free Encyclopedia, accessed November 6, 2018 https://en.wikipedia.org/w/index.php?title=Informal_value_transfer_system&oldid=840911091.

Wikipedia contributors. 2018b. “Distributed data store,” Wikipedia, The Free Encyclopedia, accessed November 19, 2018. https://en.wikipedia.org/wiki/Distributed_data_store.

Wikipedia contributors. 2019. “Synchronization (computer science).” Wikipedia, The Free Encyclopedia, accessed January 28, 2019. https://en.wikipedia.org/wiki/Synchronization_(computer_science).

Woodford, Michael. 1990. “Public Debt as Private Liquidity.” *American Economic Review Papers and Proceedings* 80 (2): 382-388.

Xandri, Juan Pablo. 2016. “Credible Reforms: A Robust Implementation Approach.” Working paper, Princeton University.

Appendix A Enhanced Financial Accounts

Here we take up supplementary material on the use of enhanced financial accounts.

A.1 An example of the use of village and community-level financial accounts: Tariffs and flow of funds

There is huge interest in the impact of tariffs in the US, under the Trump presidency, and further, it would be useful for trade and financial flows to be recorded in real time. Likewise, in reverse, one could examine the impact of trade and financial liberalizations in emerging markets.

Paweenawat and Townsend (2012) follow the Bureau of Economic Analysis (2017) guidelines and show first how to reconfigure household and business financial statements, and second, how to aggregate up to create the set of national income and product accounts (NIPA), with the economy as the village. The income statement is transformed, being careful about value added to the production account, and the balance sheet is transformed by taking time differences to create the savings/investment account. Flow of funds accounts measure net acquisition of financial assets, assets minus net incurrence of liabilities, and this is equivalent to gross savings less expenditures on real capital. The balance of payments account of the village economy follows, thus how villages, and regions, interact with each other.

A.1.1 A counterfactual policy analysis

Paweenawat and Townsend (2018) calibrate a model that integrates real and financial sectors, allowing for occupation choice, trade in goods across manufacturing and agricultural sectors, and external borrowing and lending. The model has judiciously chosen obstacles to trade, namely transactions costs for commodity trade and collateral requirements for credit. After fitting the model-generated village paths to the data, one can examine simultaneously and consistently the activities of the featured case study of sampled households and businesses along with selected aggregates. That is, one can determine what is happening over time at the household level with their own financial accounts and what is happening at the economy-wide, village level with the NIPA accounts. One can also distinguish movements of real capital from movements in paper currency.

It then becomes possible to conduct counterfactual policy analysis: What if trade and capital flows had not been allowed to liberalize or, alternatively, what would happen if there were further innovations due to enhanced financial infrastructure? If, for example, there had been a push for protection in the past, and somehow trade across regions had been more restricted than in reality, then a wedge would move relative prices. Likewise, one can examine counterfactual restrictions on inter-regional flows of capital if more savings had been targeted to be invested at home. The model predicts what would have happened to interest rates, wages, and prices; to occupation choice, production, profits, and earnings; and, finally, to the trade balance, the current account, and the balancing flows of borrowing/lending.

The impact of altered policy is not homogeneous. Removing an obstacle at the village level, through DLT or other means, is not the same thing as increasing social value. Likewise, imposing obstacles can be welfare improving for some households as a function of balance sheets and income flows.[37] The application is of course not specific to Thailand. One can imagine examining the impacts of tariffs and flows of funds in the US.

A.2 Generalized statements of liquidity accounts in the US

Nothing in these examples is particular to Thailand and the predominant use of paper currency as the medium of exchange there. Samphantharak, Schuh, and Townsend (2016) show how a conventional statement of cash flows can be made for advanced countries such as the US. Using actual data, the statement of cash flow for households is disaggregated into item-by-item liquidity accounts: the inflows and outflows to and from demand deposits; credit, debit and prepaid cards; and paper currency. Likewise, though not yet well measured in the Boston Fed surveys, the statement of liquidity accounts links conceptually to the other financial accounts and thus to the variation in income and long-term financial assets.

Of course, distributed ledger technology is not yet the source of data, though one can envision how we might get there. The Boston Federal Reserve bank survey uses data from survey questionnaires and data from diaries. Likewise, data from the Survey of Consumer Finances (SCF) and Panel Study of Income Dynamics (PSID) are interview-based. Still, there is increasing use of administrative data as a cross-check on consumer responses, data which is electronic. For example: Are households reporting bank transactions consistent with data from corresponding banks? Browning, Crossley, and Winter (2014) seek to integrate the collection of wealth, income and spending data in the British Household Panel Survey so that for each household the intertemporal budget constraint holds. An Office for National Statistics (ONS) Economic Expert Working Group (EEWG) envisions using web surveys, mobile surveys, and phone apps to scan barcodes and till receipts. There is now also exclusive electronic data surrendered voluntarily by customers, as with Mint, and the use of commercial bank information by information aggregators. Use of DLT to create complete financial accounts is not as far as away as it might seem, a priori.

How could this be used? As emphasized in this section, we could see and understand better the role of liquidity in an economy. Work emphasizing liquidity and payments and links to monetary policy is a bit sparse but increasing. Significant recent contributions include Kaplan and Violante (2014), Piazzesi and Schneider (2018), Adrian and Shin (2009), Doepke and Schneider (2006), Auclert (2017) and Fulford and Schuh (2017). DLT creates the capability of providing enhanced measurement over and above current US surveys. Currently there are discrepancies between the cash flows associated with aggregated income statements and the cash flows associated with changes in aggregated balance sheets.[38]

[37] See also Bersem, Perotti, and von Thadden (2012) for a related discussion of welfare comparisons.
[38] See Samphantharak, Schuh, and Townsend (2016).

Appendix B Cryptocurrency: The role and value of tokens in economies with distributed ledger systems

This appendix focuses on tokens and cryptocurrency, reviewing the velocity and frequent-use-in-payments definitions of money and the prevalence of multiple media of exchange in many economies, including monies indirectly backed. One implication of mechanism design theory pinpoints the special role of tokens relative to fiat money fungibility, as fiat money hurts implementation. Implications from monetary theory for valued fiat money are reviewed, distinguishing two branches: One in which for various models the value of money is endogenous and another in which money has value only due to taxes or legal stipulations. But regardless of what gives fiat monies value, tokens in hybrid systems can play a role and have value, built on top of the fiat structure. Indeterminacy in token values in this context has remedies in the roots of monetary theory: interest, use requirements, and central bank commitment. A digital reserve bank can implement these features along with activist token policies armed with transactions data from the distributed ledger. The ideas for doing so come from the various monetary models, with their explicit micro underpinnings. But to reiterate, these digital token policies are recommended even when the rest of the economy is monetized with fiat money, as tokens can deal with the gaps that remain.

Two points are worth making at the outset, one to reiterate and another which has been only implicit. First, tokens and cryptocurrency are not a necessary part of distributed ledgers. Digital Asset does not use coins in their implementations in stock markets; the same for ledgers in land titles, Walmart's tracking system, and the logistics of Maersk shipping. Second, even if there were tokens in any one of these or other systems, this per se should not be surprising given the multiplicity of monies we see in practice in many economies.

This section also reviews the implications of mechanism design for the relationship of tokens to fiat money; uncontrolled use of the latter can undercut the incentives of former. There follows a review of the implications of monetary theory for the role of tokens in hybrid systems, regardless of what is giving fiat monies their underlying value. Implications for the indeterminacy of token values and the divergence of private and social values are reviewed, as evidenced in the loss of the fundamental welfare theorems, with empirical tests. Remedies for indeterminacy are given, as implemented in the envisioned innovation with MTOs. Finally, optimal activist token policies come from monetary models combined with data from transactions on distributed ledgers, gathered as part of intrinsic operations of the digital reserve bank.

B.1 Media of exchange, definitions of money

Bech and Garratt (2017) display as a taxonomy flower the various kinds of money in existence. For us here, rather than stress a distinction between fiat money and cryptocurrency, as if there were room for only one or the other, likewise for outside money versus private credit, we can adopt standard definitions and go to data for measurement. This provides an impartial and unitary treatment of what we mean by money – simply something used frequently in transactions. The conclusion is that, even within a given economy, there are multiple media of

exchange, and so we would not be surprised to find that tokens of various kinds, depending on the purposes, can play complementary roles.

The standard definitions have to do with velocity and frequency of use in payment. Velocity of an object is the average amount traded per unit time divided by the stock outstanding. The velocity of fiat money is a key object, though alternative measures of the stock run from base money, M0, to larger aggregates including commercial banks deposits, M1, M2, and so on.

A payment matrix enumerates in its rows and columns what is exchanged for what. An object in a given row may have a high proportion of value in use to acquire several other objects specified in each of the columns. If the numbers are high, then such an object can be called a money. Within a given economy, multiple devices can be used, in addition to the associated country's fiat money. The extent of this varies across economies but multiplicity is not uncommon.

B.2 Multiple media are typical

In contrast to thinking of fiat currency or Bitcoin as competing entities, a payments matrix for ICRISAT village economies in India shows prevalent use of both grain and fiat currency (rupees) (Lim and Townsend 1998). Grain is used to pay labor, for example. In country-level economies with advanced payments systems, such as the US, household surveys from the Federal Reserve Bank of Boston show the use of fiat money currency but also the use of checks, debit, and credit cards (Bitcoin is negligible so far) (Schuh and Stavins 2014). A related point: A given agent typically uses multiple media over relatively short periods of time.

Of course, an object, good, or security can be useful and have value, even if it has low velocity. Relatedly, we often see layering, with some payment devices backed by others. Some systems operate offline. The indirect claims can have high velocity, while the underlying backing need not. Fiat money, for example, is relatively recent. Previously, government monies were backed by gold (or silver), and, in principle, gold certificates as government bank liabilities could be redeemed. It was simply easier to trade paper claims than cash in and transport the gold which otherwise was held in depositories. Currently, commercial banks are required to hold some central bank reserves, hence to hold fiat money. In turn, banks issue claims such as demand deposits, so that again payment devices are indirectly backed.

But, private debt can also be money in a given economy, an asset with high velocity. Suppose that agents trust the entity issuing the liability to always honor redemptions even when presented by third parties. In economic terminology, there is full commitment on the part of the issuer to repay at full face value. Paper checks are written instructions by a customer to its trusted banker to transfer value to a third party. Countersigned post-dated checks circulated as money in Paraguay when the penalty for overdrawing was prison. Historically, bills of exchange drawn on the issuer have circulated as media of exchange. Trade credit can also act as liquidity (Amberg et al. 2016).

Nowadays, fiat tokens are a leading example of backing and relevant to the discussion here. They represent a combination of public monies and private monies backed by trust. A named

trusted entity in Stellar acts as an anchor. A customer deposits a fiat money with the anchor, either paper currency or a claim on another commercial bank. The anchor issues the customer an IOU for the deposit. The anchor then issues a token, which is a claim on that fiat deposit. These tokens go on to be traded in Stellar Marketplace through the order books of broker dealers. These fiat tokens could have high velocity, as when they are cashed in. For example, flows come back the other way from the country of destination to the country of origin, balancing the original transaction, giving another customer ownership of the fiat deposit.

If there is something special about crypto currencies, it is the irony that they can function too well, as a means of payment with virtually instantaneous and almost costless execution. The costs are low and interconnectedness high, so one can get into and out of coins almost instantaneously. In principle, we get close to a world without frictions in which at most tokens become a unit of account. With no one holding their coin, the value is not pinned down. As a unit of account, transactions can be denominated in tokens, but this is an arbitrary convention and does not pin down values per se.

That is, we know from basic Walrasian theory that the unit of account is arbitrary. Further, pure accounting systems can clear markets with sequential or delayed payment, allowing purchases of commodities and securities before sales. This unit of account role of money has been displayed historically in trade fairs and the prevalence of ghost currencies (Cipolla 1956; Townsend 1990). Despite sequential trade in goods and the presence of banks, little actual coin was used in trade or deposited into accounts. Revealingly, accounts could be kept in defunct, devalued coinage. Money transfer ledger systems among sitting local bankers were prevalent in early Medieval Europe (McAndrews and Roberds 1999).

B.3 Lessons from mechanism design for tokens

In some economic environments, tokens can increase trade and welfare, as in the earlier sections where we noted in the discussion of ledgers the role that tokens can play in implementation. Further multiple colored tokens can be used in some economic environments, so that there would be a welfare loss from collapsing to one. Kocherlakota (1998) makes the same point as he emphasizes that (a single) money serves only as partial memory.

Relatedly, in mechanism design, unobserved actions such as saving can partially undercut trade and insurance systems. When savings is observed and controlled, the contract allows high powered incentives, altering consumption and payoff streams so that incentive and truth telling constraints cause as little damage as possible. Doepke and Townsend (2006) show how to incorporate unobserved savings, in essence another set of incentive constraints, but these are further restrictions on the design problem and cause a loss in welfare relative to full observability. These losses can be substantial. Observability can be recovered with tokens as immutable histories and rules for their use. In particular, one would not want to allow full convertibility of tokens to fiat, as the information regained with those high-powered incentives

would then be lost.[39] On the other hand, if parts of the system are not incorporated into the design, then loss of welfare will result, but one can still solve the design problem.

B.4 Lessons from monetary theory in Walrasian, competitive markets

A lesson from monetary theory is that fiat money can have value even if intrinsically worthless. The same arguments can apply to some cryptocurrencies though we need a derivation, as below, to get to that conclusion.

A key idea in monetary theory is that when strangers meet one another there can be an absence of double coincidence of wants. In a pairwise meeting, one party has something the other wants, half of the basis for a trade, but not the other way around. Money can serve as a medium of exchange in this instance; it serves as the other half of the trade, taken on by the party giving up something of intrinsic value, only in order to be able to use it in the future, when the situation is reversed.

Economic models try to simplify as much as possible to make this point and to be able to go on to consider regulation and policy. The first series of these models is outlined here and describes how intrinsically useless objects as money can have value, whether paper currency and fiat money generally or tokens and cryptocurrency. The theory here does not yet make a distinction between the two.

B.4.1 Models of money with endogenous valuation

In Townsend (1980) an agent has fluctuating periodic endowments of a single good, for example 1,0,1,0 and that agent is paired at each date with someone on the opposite side of the sequence, 0,1,0,1. But a given agent is never paired with the same person twice, hence capturing the idea that agents are meeting strangers. Related, no IOU as a promissory note to pay in the future, given the model construction of the separation in space and timing, can come back to the issuer.

Endowments of the model should not be taken literally but represent wage earnings for laborers and migrants, crop harvests for farmers, or profits for business, including trade as a business. Of note also in this model is that "endowments" average out to a constant, both other time for a given agent, though agents discount the future, and also across agents at a point in time. Money is a balancing item, offsetting ups and downs for a given agent, or across agents in the cross section from those who are well endowed to those who are not. In a monetary equilibrium, worthless pieces of paper or tokens have value in buying goods; in the US, this would be the USD price of consumption. Money is in this sense a unit of account. Prices are in monetary terms. Finally, even in this simple model, money is being stored at any given moment by some agents over time from one period to the next, thus money is also a store of value. However, in a monetary equilibrium without any other intervention, with stable prices, there remains unexploited gains from trade, and consumption is moving with income (see below).

[39] Agents may be tempted to make deals on the side, a threat that could reduce multiple currencies to one. However, for there to be an effect, agents would have to hand over their private keys to a third party, and unless this is all in the same family, this type of out of off-path deviation could be anticipated to be minimal.

Borrowing and lending among agents in the above model does not necessarily complete the space of trades. As in Manuelli and Sargent (1988, 2010), there is a bond market for borrowers and lenders who meet sometimes, but not always and not forever. In addition, there are aggregate shocks. Given the timing of feasible trades, things do not balance out with debt alone. Money still plays a crucial role in smoothing the remaining residual fluctuations. In Woodford (1990), as in the model of Townsend (1980), public debt plays the role of providing liquidity. Such models are also used in Domeij and Ellingsen (2018).[40]

In a related model, the net incomes of households, producers or traders are not deterministic, but rather are drawn as random variables. One might suppose realizations of income are drawn independently over time for a given agent and independently over agents at a point in time, as in Bewley (1983). Money is acquired when income is high via sales and given up to buy goods when income is low. Insurance contracts could have served the purpose of smoothing and, indeed, could do even better, especially if there is a large continuum number of agents, so that the average state of the economy over agents at a point in time is constant. But in Bewley (1983), such contracts are ruled out a priori as simply not available, which can happen in reality when there is limited access to financial infrastructure (or lack of coordination). In this economic model, money is, if anything, an even more obvious store of value, a buffer stock which can be used against future shocks for self-insurance needs. An agent would always like to have some of it on hand, as income can always turn out to be low. Running out of money can be disastrous if incomes are near or at zero. So here again, money has value even if intrinsically useless. Again, as it stands, without further intervention no finite amount of money is ever enough. We cannot replicate the full insurance solution. In Kehoe, Levine and Woodford (1990), the stochastic variables are such that there should be inflation.

Other models limit trade through the demographics of finite lives. Agents earn wage income when young, but at some point in the life cycle as they age they have zero labor earnings and need to rely on previous savings. Money can serve as social security, a bridge across the various generations, which was an insight of Samuelson (1958). Indeed, in this setting, stable money can achieve a full optimum. Adding more goods and variety in incomes, even within generations, allows more realism, as in Freeman (1996) and Green (1999). Money allows old debtors to repay old creditors from whom they previously borrowed and allows old creditors to buy a good y from the young. Young debtors borrow to get good x. Here money is, even more obviously, a unit of account as it is essential for trade. Given the environment, it must be on the other side of each commodity trade, in goods x and y. If there are further obstacles to trade such as difficulty getting to market, then again there is scope for intervention (see below). That said, Tobin (1978) is skeptical when models such as overlapping generations are presented in the abstract as models of money.

It is enough in any of these models to keep track of who meets whom when and what is exchanged. This is all summarized in each model by a sequence of budget constraints, actions taken, and market clearing conditions. Tokens, fiat or crypto, have a social value in each model.

[40] See also Angeletos, Collard, and Dellas (2016) and Diamond (1965).

B.4.2 Testing for inefficiency using national income data

Arguments for social value beyond private value are not simply abstract. Monetary theory has an empirical context and has been tested in actual economies. In particular, the overlapping generations and public debt literatures mentioned earlier have been taken to macro data from national income accounts to judge whether economies are on efficient paths. If there is an over-accumulation of real capital, then the interest rate is lower than the natural rate of growth, which is clearly an inefficient outcome. Alternatively, if money has value as a valid bubble in equilibrium, less savings is put into productive assets and more is put into the bubble, raising the interest rate. Abel, Mankiw, Summers, and Zeckhauser (1989) propose a criterion for evaluating dynamic inefficiency that only involves comparisons of cash profits from capital, as a rate of return, with output coming from the history of the level of investment. They find that one cannot reject efficiency. Geerolf (2017) adjusts for land rent, taking it out, and adjusts profits of entrepreneurs, a fraction of which is arguably simply opportunity-cost wages and not a real return. He finds some economies may have been on inefficient paths. One notes in the reported results that some Asian miracle economies are among the excessive-investment economies.

Related and substantively important, in every monetary model there is also an equilibrium in which money does not have value due to self-fulfilling bad expectations (Cass, Okuno, and Zilcha 1978; Green and Zhou 2005; and Levine 1989).

B.5 The value of money comes from cash-in-advance or payment of taxes

There are other ways to give money an economic value. Suppose money must be used in trade or must be used to make certain required payments owed to the government. These arguments were given for fiat money but can also be applied to tokens.

In an early paper on monetary economics, Starr (1974) specified, for tractability, a finite horizon model. This raised the problem that in the last period, money could not have value because nobody would want it – there is no future with it. However, money had a value, nevertheless, because agents were required to pay taxes to the government with it at the end. Thus, money was always demanded in equilibrium, and the price could not fall to zero. Some see this as a realistic setting and one reason why fiat monies, as compared with tokens, have value in actual economies.

Likewise, the so-called cash-in-advance model of Clower (1967) specifies that fiat money is legal tender, and thus has to be acquired at least one period in advance in order to allow purchases with it in a given period. One could say that money has value in these models due to a legal stipulation requiring its use.[41] In the model of Lucas and Stokey (1983), some designated goods do not require cash in advance (the credit goods), but other goods do require cash in advance (the cash goods). Such models with credit goods allow real consumption loans and exchange in goods without money, but if there is at least a subset of crucial cash goods, however

[41] See also Bryant and Wallace (1984).

small, then the price of money cannot go to zero. Money must have a positive value in equilibrium.

B.6 A hybrid model of positive token values

In sum, there are two ways of making money have value: Endogenously in an equilibrium given the underlying environment; and exogenously by legal or other restrictions. Layering the two ways of modeling money delivers endogenous valuation of tokens in realistic economic settings. In this way, we can talk about innovations in payment systems using tokens which, nevertheless, take as given and utilize valued fiat money. It is not an "either/or" proposition. In contrast, much of the literature thinks of Bitcoin and other cryptocurrencies as competing with fiat money.

Specifically, one could start with an environment in which fiat money is endogenously valued in one of the equilibria or has to be used for some purpose. Either way, fiat money then has value. In equilibrium, we can describe net earnings and profits of households, firms and traders in fiat money terms, not as real commodities per se. Agents would be modeled as having indirect utility functions over money, or wealth more generally, conditioned on prevailing prices, rather than direct utility over commodities. Generously, this is one interpretation of Klein's (1976) demand for money.[42] With this outcome as a new starting point, a new underlying environment as it were, there may be remaining gaps to be filled, as mentioned in the various models in the section above on endogenous money. Thus, tokens could potentially help fill these gaps and can have endogenously determined positive value.

B.6.1 Implementation in practice

With all of this, we now have set a more complicated but realistic stage for understanding the use of tokens in some distributed ledger systems.

We illustrate as an example the role of Velo tokens in the Velo system described earlier. To review, trades of MTOs are placed on the Stellar Marketplace for exchange with token fiat monies. MTOs are traders in fiat monies to be thought of as firms. They have their own fluctuating and time varying profits, as do their potential customers, and hence would be looking for a mechanism of exchange that allows smoothing these fluctuations among them, within and across borders so to speak, as if in one integrated economy. Indeed, some customers have a direct stake in what is going on in another country through migrants or business. The objects of trade among MTOs are not goods per se but these multiple fiat monies and associated tokens.

B.7 The value of money and cryptocurrency: Social and private values can diverge

As is clear from the discussion above, valued money is a bubble in the sense that it can have social value and be priced even though intrinsically useless. The fundamental welfare theorems

[42] An exception, Klein (1976), talks of a money to which other monies are tied. Indeed, we can introduce the idea of having multiple-country fiat monies, with value by fiat for example, and think of the dollar standard across currencies.

of economics fail. Competitive equilibria without fiat money exist, but they are not typically optimal. Money is a social contrivance that allows bridges across generations, links among spatially separated agents, or buffers idiosyncratic shocks. However, decentralized social consensus is required. Moreover, valued bubbles bring, in addition, the issues of indeterminate or unstable values, a multiplicity of equivalent monies, and hence unstable exchange rates and a multiplicity of equilibrium paths. Both of these points are lessons from monetary theory.

For example, Kareken and Wallace (1981) noted that without obstacles or fiat restrictions, the exchange rate between country-specific fiat monies is indeterminate. In the Samuelson (1958) model, as posited by Tirole (1985), for example, there are knife-edge good paths leading to optimal equilibria but many other paths leading to autarky with increasing inflation. Autarky itself is a valid equilibrium. If no one expects money to have value, it will not. More recently, Garrett and Wallace (2018) repeat these themes in a context featuring co-existence of Bitcoin and fiat money. They also show how there can be a bubble in Bitcoin that can break at any moment with positive probability. Schilling and Uhlig (2018) also feature a version of the exchange rate indeterminacy of the Kareken-Wallace (1981) model, but in a different model, rationalized with spatial separation and obtained as an additional "speculative" condition, implying that Bitcoin prices could form convergent supermartingales or submartingales.

B.7.1 How can one remove indeterminacy of token values and eliminate bad equilibria? The Velo system as an example of three ways to do it

The interest on Velo tokens can be generated by investing collateral fiat money in bonds that pay dividends in fiat currency. Likewise, one could view the Velo system as creating real capital for infrastructure with a positive rate of return, which in Fernández-Villaverde and Sanches (2016) creates scope for policy to achieve an optimal allocation. As Garratt and Wallace (2018) put it, indeterminacy in the value of a money is removed in all models if real dividends on money are paid. Velo does pay interest on lock up accounts. These pin down the coin to fiat money exchange just as in the monetary model such policies on fiat money pin down its real value.

A second device is to impose some required demand. The Velo token is used as collateral backing trades in fiat monies, so it is required for trade similarly to the way money is required for purchases in cash-in-advance. The value of transfer in fiat is required to be commensurate with the value of Velo tokens. Likewise, some cryptocurrencies, including EvryNet, require tokens to pay for the "gas," that is, to pay costs, to do contract validation. Of course, this imposes some distortions which in the ideal equilibrium would not be needed. But still indeterminacy can be a big problem, so there seems to be a tradeoff, as yet unmodeled.

A third way to mitigate indeterminacy is an argument in its own right, as follows.

B.7.2 The need for commitment in cryptocurrency design

Programming the operations of the digital reserve bank is a key part of managing expectations and avoiding bad equilibria. Monetary policy can be time inconsistent when money is modeled as cash-in-advance, as in Lucas and Stokey (1983). Essentially, the government is always

tempted to tax money balances once they are there, via inflation. Subsequent literature has tried to alleviate this by management of real and nominal assets and the maturity structure of debt. The conditionality and pre-programmed nature of distributed ledger protocols of Velo is a direct advantage in this context. Xandri (2016) makes this point, how a central bank can acquire a good reputation, and Fernández-Villaverde and Sanches (2016) hint at this in their frequent reference to automata, immutability, and quasi commitment. Credible commitments can be loaded in up front as part of the algorithmic design of the digital reserve bank, as smart contracts entered on ledgers. Central banks can change rules, though in practice in OECD countries central banks handle time inconsistency issues well. That said, political authorities try to push them off of it.

Stable coins which alter the level of tokens in the system is an example of a pre-programmed central bank algorithm. However, the economics literature has argued that price fixing schemes are doomed to failure (Townsend 1977).

B.8 Interest rate policy for the digital reserve bank, insights from the monetary models

Social systems designed to achieve optimal allocation can require interventionist, activist liquidity policy. This becomes clear through the lens of the models of money enumerated at the outset. These lessons then apply to the token management of the digital reserve bank.

Spatial separation is a key friction realistic in many settings, giving rise to limited trade, as noted in Townsend (1980). Due to timing and frictions, agents in equilibrium with valued money will periodically run out of previously accumulated cash. That is, there are binding constraints. Agents are optimizing, but marginal rates of inter-temporal substitution are not equated over agents. To achieve a full Pareto optimum, a situation in which no one can be made better off without making someone else worse off, marginal rates of substitution must be equated to the natural discount rate. For this candidate for a new equilibrium to be valid, those holding tokens must find that tokens effectively bear interest at a rate that is the same natural rate of preference discount. Tokens are still valued, but the obstacle to trade, the cost of carrying tokens, is being removed. If the effective interest rate, the real return, is achieved via deflation of tokens, then the real balance value of tokens is increasing. So, some tokens must be taxed out of the system. This tax in the model is imposed lump sum and applies only when agents hold positive token balances, known in advance. In practice, it will be crucial to find taxes that do not influence behavior on the margin. Interest should be paid and taxes implemented, whether in fiat or in tokens.

Fees on transactions that reflect true costs of transactions, in technology or risk, should be implemented. This works toward equating marginal costs with marginal benefits. Likewise, interest on tokens encourages users to hold more of them, rather than assets that otherwise dominate.

A related example of an active policy comes from historical experience. The Federal Reserve Bank in the US was created initially not for full employment nor price stability, but to provide an elastic currency (Sprague 1910). There were heavy seasonal movements in the demand for currency. For example, currency was needed to buy grain from farmers at harvest, to move the

crop. But bank liabilities were backed by gold, so withdrawals of deposits from banks put stress on the system. Banks in agricultural regions could not create their own money-like objects. Instead they had to find value by calling in loans from their correspondent banks in NY. In turn, city banks which had funded margin accounts on the NY stock market called in these loans, and this led to financial crises.

In the corresponding analogue models of Freeman (1996) and Green (1999), creditors and debtors do not necessarily meet at the "right" time. A fraction of debtors arrives late, and a fraction of creditors leaves early. This naturally complicates the debt settlement process and leads to inefficiencies and fluctuations in interest rates. It is resolved by having a central bank, or the digital reserve bank, engage in open market operations, or in security transactions with the securitized notes that they issue (based on segregated collateral of correspondents), to control the requisite means of payment and smooth the interest rate.[43]

In another version of these models (Chandrasekhar, Townsend, and Xandri 2018), individual-level market participation is determined as if by an exogenous random shock. Entities with high value are those that are in the borrowing/lending market in Velo when the market is thin, when risk aversion is high, when risk from variable returns is high, and when the remaining players are judged to be important (high Pareto weight). Reserve bank liquidity should be directed ex-ante to these key players (Wallace 2014, and Townsend 1982).

For each of these examples of policy, micro data is needed on transactions. Wallace (2014) argues, for example, that one can know the policy ought to be active, not laissez-faire, but not know without more information which way it goes, inflation or deflation. With the digital reserve bank as a notary as in Velo for example, data would be available. In principle, this could be organized by financial accounts, tying into an earlier discussion.

[43] See also Townsend (1989).