

The economics of distributed ledger technology for securities settlement

Evangelos Benos,⁽¹⁾ Rodney Garratt,⁽²⁾ Pedro Gurrola-Perez⁽³⁾

Abstract

Distributed ledger technology (DLT) is a database architecture which enables the keeping and sharing of records in a distributed and decentralized way, while ensuring its integrity through the use of consensus-based validation protocols and cryptographic signatures. In principle, DLT has the potential to reduce costs and increase the efficiency of securities settlement, the ultimate step of every security transaction. In this paper, we first examine to what extent DLT could add value and change securities settlement. We then characterize the innovation process in the post-trade industry and finally, we describe the economics of a hypothetical DLT-based security settlement industry. Our main conclusions are that: i) DLT has the potential to improve efficiency and reduce costs in securities settlement, but the technology is still evolving and it is uncertain at this point what form, if any, a DLT-based solution for securities settlement will ultimately take, ii) technological innovation in the post-trade industry is more likely to achieve its potential with some degree of coordination which could be facilitated by the relevant authorities, and iii) if DLT-based securities settlement becomes a reality, then it is likely to be concentrated among few providers which, in the absence of regulation, could result in inefficient monopoly pricing or efficient price discrimination with service providers capturing much of the market surplus.

Key words: Distributed ledger, securities settlement, CSD, post-trade industry.

JEL classification: L81, O32, G23.

(1) Bank of England. Email: Evangelos.benos@bankofengland.co.uk

(2) University of California Santa Barbara. Email: garratt@ucsb.edu

(3) Bank of England. Email: Pedro.gurrola-perez@bankofengland.co.uk

DISCLAIMER: This is a work in progress. The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees.

The authors would like to thank David Bailey, Paul Bedford, Richard Brown, Tatiana Cutts, Andrew Hauser, Todd Keister, Mark Manning, Irina Mnoghitei, Huy Nguyen, Chris Redmond, Simon Scorer, Karl Shell, Rhiannon Sowerbutts, Claire Sunderland Hay, Roger Willis, Graham Young, seminar participants at the Bank of England and participants of the 2017 joint Bank of England-LSE Workshop on Financial Market Infrastructures for many helpful comments and suggestions. During the time this paper was written Garratt consulted for Payments Canada and R3 on work related to Project Jasper. Any errors are our own.

1. Introduction

Over the last few years, a number of technological innovations with potentially transformative implications for the financial system, its intermediaries and users have been receiving close attention from researchers, industry participants, technology firms and regulators. In particular, in the case of the post-trade cycle, which includes clearing and settlement, the distributed ledger technology (or DLT) has often been seen as one of the more promising ones, with the potential to drastically simplify processes, reduce costs and increase efficiency and security. This has spurred a flurry of analysis and prototype development, exploring how exactly DLT could be implemented across the different layers of the post-trade chain.

Although “distributed ledger” (DL) is a term that is used to describe various technologies, it is usually understood as a decentralized dataset architecture which allows the keeping and sharing of records in a synchronized way, while ensuring their integrity through the use of consensus-based validation protocols and cryptographic signatures.¹ These technologies are at the base of Bitcoin and other digital currencies, but they may have a much wider range of applications.

In this paper, we focus on the industrial organisation and economic implications of the potential use of DLTs for securities settlement. Although DLT is only one of the possible technologies that could shape the next-generation of financial services infrastructures, we focus on it here because it is the technology which is currently attracting more attention in the specific case of securities settlement.² While the technological and implementation aspects of DL in relation to the post-trade process have been often discussed, the economics and the role of regulators have received less attention. Thus, our aim is to identify some of the key aspects of the economics of DLT, as it applies to securities settlement, to identify potential market failures as well as potential regulatory interventions that could help address these failures.

Overall, our paper makes three key points: First, while the current industrial structure of securities settlement appears to be inefficient and while DLT has the potential to eliminate some of these inefficiencies, there are still issues that need to be resolved in order to operationalize DLT and it is uncertain at this point how or if the technology will ultimately be used in securities settlement. Second, developing new technological know-how (whether it is DLT-based or not) with the aim of modernising securities settlement is akin to adding to a public good. For this reason, it is best if research and investment is done in a cooperative rather than competitive manner, which implies that public authorities and regulators can play an important role in facilitating such a cooperative outcome. Third, if a DL (or other) technological solution for settlement ultimately succeeds in replacing existing settlement methods, it is likely to be characterized by network externalities and decreasing average costs. This suggests that the industry may well retain its high degree of concentration. It is possible and likely that a small number of CSDs will be replaced by a small number of DL network providers and as a result future settlement services may be associated with

¹ Here, distributed data processing refers to a system where data is distributed across the participating nodes. This is in contrast to the existence of a single, centralized server that provides data processing services to all connected systems. Decentralisation means that the control of the database is shared across multiple nodes. In Section 3 we provide additional detail.

² To date, various DL developments have been disclosed, including a joint project of the Australian Securities Exchange (ASX) with Digital Asset Holdings (DAH), a project by the Japan Exchange Group (JPX), or a project by the Deutsche Börse in cooperation with Germany’s central bank, to mention just a few.

some form of monopoly pricing. This, in turn, means that regulators could play a role in either ensuring more competitive outcomes or otherwise in regulating the industry so as to minimize deadweight losses and/or the unequal distribution of economic surplus.

This paper follows a number of other publications, by public bodies, which have sought to highlight the issues and challenges of operationalizing DLT in the FMI space in general and in securities settlement in particular.³ A common conclusion reached by these studies is that it is unclear in what form DLT will be ultimately adopted as a solution for securities settlement, if at all. For example, Pinna and Ruttenberg (2016), from the European Central Bank (ECB), envisage three potential future scenarios: one in which DLT is adopted to improve cluster/internal efficiency while business practices largely remain as they currently are; another in which core players deploy DLT in specific markets, with some players becoming redundant; and a third one where DLT is fully implemented, allowing a peer-to-peer (P2P), largely dis-intermediated system for securities transactions. Similar ideas are also echoed in a number of discussion papers from a variety of private institutions with the views of market incumbents typically being more cautious and less optimistic.⁴

Many of the proposed DLT solutions for securities settlement are still evolving and uncertainty over the ability of new technological solutions to overcome inherent technical challenges, together with the industry inertia and monopolistic barriers, makes it difficult to predict whether, how and when DLT will be implemented in the post-trade space. The results will ultimately depend on the efficacy of new technological solutions, the economics of the post-trade industry, market forces and regulation. In this paper, we examine to what extent DLT could add value and change securities settlement and we describe the economics of a hypothetical DLT-based security settlement industry. However, it is important to note that our analysis on the economics of DLT innovation and industrial structure should apply equally well to any technological solution likely to be adopted to the extent that such a solution is characterized by network externalities and decreasing average costs.

The rest of the paper is organized as follows. In Section 2 we describe the current shape of the post-trade cycle along with its shortcomings. In Section 3 we describe the main attributes of DLT and in Section 4 we discuss the potential impact of DLT on the main processes and functions of the post-trade value chain. Section 5 discusses the economics of technological innovation in the post-trade industry, highlights the importance of cooperative arrangements, and explores the role of governments. In Section 6 we discuss the marginal costs of a DLT solution for settlement and argue that marginal costs will be low for DLT technologies that meet industry requirements. We also explore what the structure of a DL-based settlement industry could look like and what role regulators could play in maximising welfare and ensuring financial stability.

³ See for example CPMI (2017), ESMA (2017), FCA (2017), FINRA (2017), IOSCO (2017), Manning et al (2016), Mills et al (2016), Pinna and Ruttenberg (2016).

⁴ See for example Morgan Stanley (2016), Euroclear/Oliver Wyman (2016), SWIFT/Accenture (2016), responses to the survey of Mainelli and Milne (2016), Mills et al (2016), and World Economic Forum/Deloitte (2016).

2. Frictions in the current post-trade processes

In this section we look at the main players and processes in the post-trade chain and discuss some of its inefficiencies.

The post-trade cycle consists of all processes, intermediaries and infrastructures that are invoked from the time a trade in a financial security is agreed to the time when it is finally settled. It broadly consists of three core functions: order management (including trade validation), clearing (i.e. the calculation of counterparties' obligations), and settlement (i.e. the final transfer of assets).⁵ **Table 1** presents a summary of the main infrastructures and processes involved in a typical post-trade cycle.

Central Security Depositories, or CSDs, are solely responsible for facilitating settlement. In this role, they perform three main functions:

- a) A **notary function**, whereby they keep safe records of issued securities so as to ensure, for example, that no one fraudulently creates and trades non-existent securities (i.e. securities with no claims to real cash flows). This notary function is fundamental for maintaining confidence in the system and therefore it is critical for financial stability. It needs to be impartial and trusted by all participants, which is why a regulated third party is typically assigned to this role.
- b) **Settlement**, whereby they facilitate the transfer of legal ownership of securities from sellers to buyers, typically via delivery-versus-payment (DvP).⁶
- c) **Account maintenance**, whereby they hold and update ownership records following each transaction. CSDs hold and manage the (physical or dematerialized) immobilized securities and they provide accounts where securities are held so that ownership can be easily transferred through a book entry rather than by the transfer of physical certificates. They also provide vaults to hold physical certificates.

In addition to these functions, CSDs may also provide services related to the securities they hold such as custody, asset servicing, financing, reporting (both internally and to regulators and trade repositories), or securities lending.⁷ It is common for domestic CSDs to also perform clearing functions for domestic markets, while international CSDs (ICSDs), EuroClear and Clearstream, also clear and settle cross-border securities.⁸

Because of their importance for the correct functioning of financial markets, CSDs operate in a highly regulated environment. They are subject to national laws on securities issuance, settlement and safekeeping, while being supervised by the relevant authorities, typically the securities or banking regulator, and subject to the oversight of the relevant central bank(s). Since 2014, CSDs in the European Union have to comply with the CSD Regulation (CSDR), which stipulates prudential standards and various other requirements as a basis for a common authorisation process. CSDs are

⁵ See, for example, CPMI-IOSCO *Principles for Financial Market Infrastructures* (CPMI-IOSCO, 2012).

⁶ Delivery-versus-payment means that securities and cash are exchanged simultaneously and only after they have both been deposited with the CSD. This minimizes settlement risk, i.e. instances where a security buyer makes a cash payment but does not receive the security (or vice versa) because the seller has in the meantime defaulted.

⁷ CSDs may lend securities to a security seller to enable him to settle his transaction.

⁸ The term "international CSD" is used to characterize CSDs which are also depositories of Eurobonds.

also subject to [Principles for Financial Market Infrastructures](#) (PFMIs) issued by the CPMI, IOSCO and the FSB (CPMI-IOSCO, 2012).

			Sell-side	Exchanges, trading platforms	Clearing agents	CCPs	CSD	Custodians
		Functions/Processes						
Account management, reporting, error management, identity management, security, messaging	Order management	Trade data input (booking)						
		Trade validation: scrutiny, cross-reference, data enrichment						
		Calculation and collection of fees and duties (broker fees, taxes, exchange charges, commissions, etc.)						
		Margin financing						
		Matching, confirmation, SSI enrichment and error resolution.						
	Clearing	Netting						
		Margins calculation						
		Novation						
		Risk management						
	Settlement	Exchange of payment and securities						
		Reconciliation of positions, cash and collateral						
		Risk management						
	Notary function							
	Depository							
	Asset servicing	Management of corporate actions, close-out of positions, exercise/delivery notices, etc.						
	Custody	Maintaining indirect ownership records						
Management of changes during the lifecycle								
Collateral management	Access and movement of collateral on behalf of clients							
Cross-border intermediation								
Securities lending	Lending securities to facilitate settlement							

Table 1: Typical processes, intermediaries and infrastructures in the post-trade cycle

The processes occurring after the trade has been cleared (e.g. settlement, custody, asset servicing, etc.) generate a significant amount of revenue in the form of commissions and fees. On the flipside, they represent a substantial cost for market end-users. According to Oliver Wyman/SWIFT (2014), revenue from settlement, custody and collateral management amounted to \$40-45 billion in 2013, which represented approximately 13% of the total trade value chain (from execution to settlement). Of these, approximately \$3bn corresponded to (I)CSDs and \$39bn to custodians. On the other hand, Broadridge (2015) estimates that the industry spends \$17bn to \$24bn per year in core post-trade processing, reference data, reconciliations, trade expense management, client life-cycle management, corporate actions, tax and regulatory reporting. For the most standardized asset classes - equities and fixed income, excluding OTC derivatives, costs amount to \$6bn to \$9bn annually.

These costs are partly the result of the participation of several intermediaries in handling a single transaction, each of whom uses its own systems for processing, sending and receiving instructions, reconciling data, managing errors, etc.; with each intermediary in the chain needing to keep an updated record of the transaction. This, together with the lack of standardisation, gives rise to costly

back-office reconciliation processes. Some analysts estimate that by standardising post-trade processing systems and rendering them interoperable, market participants could benefit from economies of scale and save \$2 - \$4 billion annually (Broadridge, 2015).

The reasons for this complexity are largely historical (e.g. use of paper certificates, industry consolidations, etc.), but also a consequence of inertia and aging technology. There may also be legal reasons; for example, direct security ownership has been often replaced by indirect ownership, where investors hold securities through custodians or sometimes through a chain of custodians (a typical example being cross-border investments). Other sources of inefficiencies include lack of competition due to entry barriers and lack of interoperability across providers. As a consequence, these processes translate into higher costs for end-users.

3. Distributed ledger technologies (DLT)

In this section we provide an overview of DL and blockchain technologies and of their main design features.

In traditional centralized database architecture, information is located, stored, and maintained in a single location and is controlled by a central administrator, who ensures its integrity. The data is usually stored in raw form and there is some sort of security perimeter which protects the database from external attacks.

Distributed databases provide an alternative configuration where the database (or copies of it) is spread across the network and is stored at different physical locations. Although the database is shared across nodes, the control of it is usually centralized and the integrity of the database relies on a central administrator. A centralized application is used to manage the database and synchronize the data. The nodes are inside the security perimeter, are trusted, and access is controlled by the network administrator.

A *distributed ledger* (DL)⁹ is a distributed database, in the sense that each node has a synchronized copy of the data, but departs from the traditional distributed database architectures in three important ways:

1. *Decentralisation*: The control of the database (read/write access) is *decentralized*; that is, it is performed by multiple (or all) network participants. There is no need for a central administrator to ensure the integrity of the data or its consistency across nodes. Instead, this is achieved through some consensus mechanism or validation protocol.
2. *It is reliable in trust-less environments*: The consensus mechanism ensures the consistency and integrity of the database even if the parties involved do not fully trust each other.
3. *Cryptographic encryption*: The ledger uses cryptographic encryption tools to deliver (1) and (2) above.

When compared with traditional distributed databases, these three characteristics have made DLT an attractive candidate for applications in the post-trade industry. While distributed databases can

⁹ Sometimes also called "mutual distributed ledger" (Mainelli and Milne, 2015), to emphasize the importance of mutualized ownership.

also provide shared, synchronized copies of the ledger, and this could reduce the number of reconciliation processes, decentralization would allow the nodes a greater control over their own information and could further reduce costs by facilitating the disintermediation of the post-trade process. Finally, while traditional distributed databases rely on trusted nodes and keep the copies of the ledger within a security perimeter controlled by the central administrator, the ability of operating safely in a competing environment with untrusted third parties is a critical component of any potential application in the financial industry.

The *blockchain* is a particular architecture of a DL whereby data are batched into a sequence of blocks linked to each other using cryptographic tools. This dependence of the various blocks of data on the previous blocks makes it extremely difficult to retroactively alter the database.¹⁰ Thus, the blockchain forms a perpetual chain of immutable records.

Access and validation rights

There are different types of DLs according to their accessibility restrictions and the type of consensus processes employed to ensure the integrity of the ledger. In terms of who has access to the data, DLs can be:

- **Public**, whereby any user is allowed to read/view the ledger, or
- **Private**, whereby only approved participants have access to the data.

Similarly, depending on who is allowed to validate the ledger, DLs can be:

- **Permissioned**, whereby only a specific group of trusted users can validate or modify entries to the ledger or,
- **Permission-less**, whereby anyone is allowed to build and verify the ledger. The fact that no central authority is needed to validate entries to the ledger makes it resistant to censorship;¹¹ that is, no actor is able to prevent a transaction from being added to the ledger.¹²

Accessibility is closely related to privacy and confidentiality. While in a public ledger, like Bitcoin, anyone can see, for example, that A has paid an amount X to B (although the true identities of A and B may be unknown), security transactions usually require confidentiality. With the exception of regulators and trade repositories, in most cases only those parties with a legitimate need to know should be able to see the details of a transaction. Ensuring confidentiality in a DLT implementation is still a challenge and constitutes an important research question. Moreover, financial markets are subject to regulatory oversight and so permission-less structures could pose a challenge, given their potential opacity. For these reasons, the settlement industry is mainly focussing on private,

¹⁰ This is because if one wanted to alter data in a block, one would have to also alter all the subsequent blocks.

¹¹ Swanson (2015) makes an interesting distinction between centralisation in governance and centralisation in validation. In principle, validation can be de-centralized while at the same time governance and rule-making authority may rest with a single entity. The opposite may also be true.

¹² It should be noted, however, that some validation processes may be susceptible to censorship by a small minority of participants. See the Liberty Street Blog post:

<http://libertystreeteconomics.newyorkfed.org/2015/09/a-distributed-version-of-repugnance-as-a-constraint-on-markets.html>

permissioned ledgers, where validation rights are limited to a trusted set of participants, and where nodes can be assigned different roles, with different access rights.

Consensus mechanisms

As with any database, to ensure the integrity of the ledger, modifications to the ledger have to be validated. For each new transaction, for example, somebody has to check whether the counterparties have enough assets (or cash) in their accounts and to make sure these same assets (or cash) are not being used in more than one transaction (i.e., that there is no double-spending).¹³ To solve this problem without depending on a central authority for trade validation, DL protocols typically rely on a consensus mechanism by which a selected subset (or even all) of the nodes agree on which trades are valid and which is the most recent version of the truth. The design of the consensus mechanism will largely depend on whether the DL is permissioned or not.

In permission-less DLs the identity of the participants is unknown and therefore a consensus mechanism is required which does not rely on trusting the participants or knowing their identity. Such a validation protocol cannot be based on simple majority (each node has a vote), as it would allow an attacker to alter the ledger by creating nodes and controlling more than half of the votes. So the question is how participants can share validation rights without significantly compromising the security of the system because of the risk of Sybil and denial-of-service attacks.¹⁴ Bitcoin's solution was a consensus mechanism based on **"Proof-of-Work" (PoW)**; such "proof" is achieved by solving a computationally intensive mathematical problem which makes it economically expensive for a single attacker to have under his control enough computing power to alter the ledger.¹⁵ In theory, under the rule of one-CPU-one-vote, an attacker would need to control more than 50% of computing power to alter the ledger; however, under certain circumstances, controlling even 30% of computing power might be sufficient to alter the ledger (Eyal and Sirer, 2014).

Another type of validation system is the **"Proof-of-Stake" (PoS)** consensus process. This assigns a participant a share of validation rights according to its stakes in the system. The criteria to measure the validators' stake in the system can vary. They can be measured in terms of internal tokens, of off-ledger assets pledged as collateral, or can be based on the reputation of the validator in the network (if his identity is known). Some PoS processes require voters to place "bets" on the true state of the ledger, so would-be attackers trying to claim something false would be in a minority with their prediction and lose their bets.¹⁶

Permissioned systems do not have to solve the problem of an untrusted network and are therefore capable of processing transactions faster and are cheaper to maintain than capital-intensive

¹³ Validation functions can be performed separately. Ensuring a participant has funds available to initiate a transaction can be done locally, but prevention of double spending requires a global perspective (see section 6.1).

¹⁴ A Sybil attack is a traditional way of undermining a P2P network by creating large amounts of fake identities in order to gain a disproportional amount of influence or votes in the network.

¹⁵ Of course how much this is an inhibitor, will also depend on the values at stake.

¹⁶ For example, in Ethereum's proposed "Casper" protocol, anyone can participate in block production by placing a security deposit ("posting a bond"). After posting a bond one has an opportunity to bet on which block will be included next. The incentives are such that one makes money by betting with the eventual consensus and lose money by betting against the consensus.

permission-less systems. On the other hand, permissioned systems are arguably more open to censorship and reversibility.

Figure 2 illustrates the relation between different accessibility options and the types of validation protocols.




Figure 2: Accessibility and validation DL configurations.

Smart contracts

An important new feature of DLT is that it is not only a database, but can also include executable code as part of the ledger to create what has been fancifully labelled as a *smart contract*. Perhaps the best way to understand smart contracts is to contrast them to what occurs in a Bitcoin transaction.¹⁷ Bitcoin transactions are push notifications: bitcoins represent *unspent transaction outputs* (UTXOs) and validation requires checking that a person sending a bitcoin does in fact have that UTXO. Once an UTXO is sent it is destroyed and a new UTXO is generated, and the process of spending and validation is repeated. In order to implement smart contracts more information about the current “state” of the network is required. For example, a smart contract that is used to automatically calculate and pay periodic coupon payments would require information of the current and past states to execute and keep track of the payments.

In the context of post-trade processes, the concept of smart contract goes beyond record keeping and data sharing and may have much deeper implications. For example, the embedding of pre-specified code in the ledger will presumably allow legal agreements to automatically execute their legal clauses. This could reduce the ability of counterparties to ex-post strategically renege on their contractual obligations which, in turn, could affect other contract parameters such as the price and

¹⁷ See https://www.ethtrade.org/Understanding_ethereum.pdf.

could also lead to smart contracts being used as a means of signalling high credit quality in case this is not otherwise easily verifiable.

However, there are still significant questions around how this could be implemented in practice, including how to translate the legal contract terms into code and what the legal implications might be, or how a smart contract will interact with an external source of data (e.g. with LIBOR rates data). On the other hand, the development of smart contracts could help the industry reach common standards in the same way that ISDA master legal agreements simplified derivatives trading.

4. Is DLT the solution to post-trade frictions?

In this section we discuss how DLT could potentially add-value to the post-trade cycle and what some of the main risks and challenges are.

There has been a significant amount of effort dedicated to understanding what potential benefits DLT could bring to the post-trade cycle, what the main risks are, and what challenges it faces (e.g. Pinna and Ruttenberg (2016) and Mills et al. (2016)). What emerges from these analyses is that, while DLT has great potential to add value to different layers of the post-trade process, it should not be seen as a panacea and still has to overcome challenges before it can deliver some of these benefits. Implementation and operational risks (e.g. cyber risks and the risk of fraudulent activities) and the risks derived from increased interconnectedness are often cited as major priorities to be considered (e.g. ESMA, 2017). On the other hand, as we point out below, some of the benefits usually associated with DLT, like shorter settlement times, are not necessarily hindered by technology, but by current market practices.

In general, we think DLT might impact the post-trade cycle in a number of ways. We first discuss the potential benefits:

- **Reducing reconciliation and data management costs:** Having a distributed, synchronized record of security ownership which is shared by the relevant participants can simplify and automate many of the processes currently involved in the post-trade cycle, reducing the need of reconciliation, providing a common standard and eliminating the need of centrally maintaining back-up systems. It is hard to estimate by how much these costs would be reduced by the DL technologies but, using estimates on cost reduction resulting from supply chain automation in non-financial industries, Mainelli and Milne (2016) estimate a potential reduction of 50% or more of costs in financial security transactions. Santander estimates that more efficient digital ledgers could cut costs in the banking industry by up to \$20 billion a year (Santander/Oliver Wyman/Anthemis, 2015).
- **Flexible settlement times:** DL technology could reduce the duration of the settlement cycle for securities transactions. However, most current settlement cycles are not long because of technological limitations but because of banks' back-office processes, legal arrangements and liquidity management practices.¹⁸ In the case of Gilts (T+1) or equities (T+2), for

¹⁸ A somehow different consideration may apply to other asset classes, where the processes are still substandard. One typical example is that of syndicated loans, where the post-trade cycle is inefficient, slow, costly and unpredictable, frequently stretching to T+21, far from the T+7 promoted by the industry. However,

example, participants have at least a full working day to prepare for settlement and borrow securities or cash if needed. A cycle of T+0 would nearly eliminate settlement risk, but would also require prepositioning cash or securities in advance of the trade which would have implications for liquidity management.¹⁹ Because of this, current DLT implementations are offering flexibility in settlement times,²⁰ allowing for faster settlement times for those participants who want to reduce their settlement risk exposure.

- **Automated clearing:** According to the PFMI, clearing is the “computation of counterparties’ obligations to make deliveries or payments on the settlement date”. In a DL environment for cash securities settlement, once the trade is agreed between the two counterparties the calculation of obligations could happen simultaneously, since all the relevant information would already be shared between participants. However, the absence of the netting benefits that clearing provides would increase the liquidity demands for settlement. Therefore, any DL solution would also need to include netting algorithms. If that were the case, DL could potentially eliminate the need for a clearing agent, thereby reducing costs and operational risk. Nevertheless, the case for clearing would likely remain where there is still counterparty credit risk, like in derivatives transactions.
- **Direct ownership:** In the current market design, it is common to find that investors are not the direct owners of their securities, but only hold them indirectly through chains of custodians that operate between issuers and investors. This is partly a legacy from the time where securities were issued as paper certificates and had to be immobilized to facilitate their trading through book-entry transfers. However, custody chains reduce investor rights to the least favorable custody term operating in the chain (Micheler, 2015). In principle, a DL solution could facilitate direct ownership, reducing legal and operational risks and intermediation costs for the investors. It could also increase transparency, as investors could have direct access to their holdings and issuers could keep track of beneficial owners.
- **Traceability and transparency:** To the extent that DL provides an immutable record one can trace back the history of any flow of funds or securities. Also, any amendments to the ledger are traceable.
- **Enhanced security and resilience:** A decentralized system with no single point of failure is more resistant to a single- (or even multiple-) node attack. And having a distributed database enables faster recovery in the event of such attacks. The use of cryptographic signatures to access data and the encryption of the elements of the ledger will also enhance security.

On the other hand, there remain significant challenges to the implementation of DLT and there are also important risks associated with these innovations. Some of these issues are particularly relevant in the securities settlement space and would largely determine which DLT configurations will be adopted, if any at all.

- **The notary function:** In its role as a notary, the CSD ensures and guarantees that the issued securities and any transfer of ownership are correctly recorded. In a DL environment, while

even in these cases delays in settlement are largely driven by legal complications associated with transfer of loan participations (Mainelli and Milne, 2015).

¹⁹ See Khapko and Zoican (2017) for a formal exposition of this trade-off.

²⁰ Morgan Stanley expects windows to shorten, but not to T+0. Some platforms are building a “T+ whatever you want” which would also allow users to price for the liquidity they want (Morgan Stanley 2016, p. 7).

the integrity of the records is ensured by the integrity of the ledger itself, a trusted body may still be needed to guarantee the validity of securities issued. In that case, the question is whether this function will be the responsibility of a single entity (potentially a CSD) or of a set of (trusted) validators. Some important legal issues would still need to be resolved, for example, with regards to territorial requirements concerning restrictions as to where data can be physically maintained, or to existing regulations which require asset issuance and registration to be performed by a CSD.

- **The depository function:** Migration to a DL would require integrating the ledger with legacy assets currently held by custodians or depositories. One solution could be to use digital tokens to represent off-ledger assets, but such a “tokenisation” mechanism would perhaps require a trusted entity guaranteeing a faithful correspondence between tokens and assets and would also need to consider what the legal status of the tokens could be. New securities, on the other hand, could potentially be designed as purely digital assets. In both cases, changes to the regulatory and legal frameworks may be needed.
- **Delivery versus payment (DvP):** To provide DvP, the DL should be able to interact simultaneously with cash accounts. And depending on how the DL interacts with the cash leg, we can think of two possible options:
 - The DL interacts internally with the cash leg. In other words, the DL moves digital cash. This implies that there are digital currency accounts with the Central Bank and that payments can be made in digital central bank money through the DL.²¹
 - The DL interacts with external cash accounts (or a digital cash ledger). This would require an interface between the DL and the cash accounts of the participants. Such interface could be one or more entities linking the DL with the participants’ external accounts. This would not be very different from the functions currently performed by settlement banks.
- **Settlement finality.** This is a critical element for risk management. There should be a clear legal basis regarding when settlement finality occurs in order to define when key financial risks are transferred in the system, including the point at which transactions are irrevocable (CPMI-IOSCO, 2012). A known problem with DLs is that consensus mechanisms like the one used by Bitcoin can only guarantee probabilistic finality, due to forking.²² Although this probability converges to 1 as time passes, settlement finality requires certainty about the point at which the trade becomes final. Non-PoW protocols exist which avoid this problem.
- **Legal ownership:** Settlement requires a formal (i.e. legally defined) indication of transfer of ownership once securities and cash have changed hands. Thus, an important issue that needs to be addressed is whether a DL entry can legally constitute proof of ownership.

²¹ The Bank of Canada recently released a study (Chapman et al. 2017) that assessed their work to build a proof of concept wholesale interbank payment system using a distributed ledger. The study concluded that at this time a “DLT wholesale payment system is unlikely to match the net benefits of a centralized wholesale payment system”, but that “[Over-the-counter markets (for stocks, bonds and derivatives), syndicated loans and trade finance] could be significantly improved using a DLT-based platform if they could be integrated with a core wholesale payment system, resulting in the transfer of cash payments using central bank money.” (p. 10)

²² In blockchain-type of consensus, appending a block at a point which is not considered the last block of the valid chain could lead to a bifurcation (a “fork”). This is resolved by the agreement that the longest chain is always the valid one.

- **Trade matching and error management.** It has been pointed out that there are still significant obstacles to using DL for trade matching²³ (see, for example, DTCC 2016). This is because DL does not necessarily have the functionality to compare different data domains, to address contract mismatches or to process exceptions. The case of Bitcoin is a very simple (one-dimensional) case compared with security transactions where matching may be required over a big number of attributes with complex rules and cross-dependencies. Some commentators believe that central matching may continue to be required as pre-ledger processing. On the other hand, the already intricate aspects of exception management in financial transactions are further complicated by the immutability of the ledger and the lack of a central administrator.
- **Confidentiality:** Regardless of the rules implemented to reach consensus, the validation typically involves multiple participants and therefore appears to be in conflict with the confidentiality required in security transactions, where the contents of a transaction need to be invisible to everyone except to the parties of the transaction. To solve this problem, one option could be to adopt a DL design where the consensus mechanism only involves a trusted authority and the transacting parties. In some platforms, consensus between firms is achieved at the level of individual deals, not the level of the system, and so transactions are validated by parties to the transaction rather than a broader pool of unrelated validators.²⁴ Another solution proposed is to isolate in segregated ledgers the transaction data from the data needed for validation, keeping the ledgers linked through a “synchronization log”.²⁵ Alternatively, the consensus protocol could include some mechanism, like zero-knowledge proofs (ZKPs) protocols, for example, by which other parties can participate in the validation process without knowing the actual contents of the transaction. To date however, there is no implementation of ZKPs that is scalable enough to be used in a production environment.
- **Identity management:** Private permissioned structures seem the most favoured by the industry to be adopted for securities settlement. Identity and access management is currently centralized in a few trusted entities (CSDs, CCPs). These same entities could perform a similar function using a centralized distributed ledger. They would validate identities and verify credentials. But if the ledger is to be decentralized, then it is critical to find processes by which identity and access management can be effectively controlled and protected from attacks. Additionally, if identity verification is done through cryptographic keys, there should be a trusted mechanism for issuing and replacing such keys.
- **Scalability:** Securities settlement requires the capacity of processing a large number of transactions in a secure and reliable way. Moreover, any settlement platform should have sufficient capacity to handle potential increases in volume in times of market stress. This places further restrictions on the choice of validation protocol, to the extent it can limit the rate at which transactions can be processed.

In summary, the extent to which the benefits, risks and costs associated with DLT could materialize will depend on how the above issues are addressed and which DL configuration is adopted. At one extreme, DLs could simply improve on existing technical aspects of database management, with the

²³ Trade matching is the comparison of the trade attributes (e.g. date, price, quantity, etc.) as submitted separately by the security buyer and seller. This comparison is done electronically and a trade cannot be settled unless the buyer and seller-submitted attributes match.

²⁴ See, for example, <http://www.r3cev.com/blog/2016/4/4/introducing-r3-corda-a-distributed-ledger-designed-for-financial-services> .

²⁵ See, for example, http://digitalasset.com/static/documents/The_Global_Synchronization_Log.pdf

current actors playing essentially their same roles. At the other extreme, DL could potentially allow for a complete disintermediation of the whole post-trade process, from security issuance to settlement, enabling a pure P2P transaction structure. Many of the challenges we have mentioned point to intermediate solutions, where different nodes may have different roles and where a subset of nodes continue performing the same core functions.

How the above questions are addressed will also determine which processes in the post-trade value chain could become redundant (if any), what services would still be required or what new roles may emerge. Custodians, for example, manage their clients' securities or collect investment income generated by client assets, functions that could potentially be automated with smart contracts. But they also offer solutions beyond DLT, like cross-border intermediation, where an institutional investor may benefit from the link between its local custodian and a foreign custodian to overcome differences in trading systems and regulations when trading a foreign security. Similarly, some of the current processes performed by CSDs may not be required under some DL configurations, but CSDs may still be needed for securities issuance or to manage access and permissions in the network, for example.

5. DLT innovation and the role of governments

In this section we discuss how DL technologies might evolve in the current post-trade industry landscape. We argue that innovation in securities settlement may be best achieved through cooperation and that governments can play a role in encouraging cooperative research.

We start by characterising the current structure of the securities settlement industry. We then discuss how various features of DLT may determine whether innovation will take place by incumbents, start-ups or consortia. We also discuss features of the current landscape that indicate why the current environment may be ripe for innovation. Lastly, we discuss the role government might play in achieving an efficient market outcome. It should be noted at this point that our arguments on how best to facilitate innovation are not exclusive to DL technology.

5.1 The industrial organisation of securities settlement

The settlement industry tends to be monopolistic at a domestic level, with little or no competition among providers. This is because the core functions of a CSD strongly benefit from network externalities and low marginal costs and therefore naturally gravitate into monopolistic structures. In the EU, for example, most countries have one CSD to service their domestic securities.²⁶ The situation is similar at a global level, with most CPMI countries having a single domestic CSD.²⁷ There are two international CSDs (ICSDs), Euroclear and Clearstream.

In terms of their business model, the revenue for CSDs and custodians arises from fees and commissions for the services they provide.²⁸ These fees are usually charged per trade (in the case of

²⁶ As of October 2014, there were 41 CSDs in 37 countries covering the whole geographical area of Europe. In the 28 countries of the European Union, there are 31 CSDs plus 6 securities settlement systems operated by national central banks (in Belgium, Bulgaria, the Czech Republic, Greece, Poland and Romania), which primarily settle government securities, according to the "CSD Factbook" (ECSDA, 2015)

²⁷ See the Bank of International Settlements Red Book on Market infrastructures 2015 (CPMI, 2015)

equities clearing and settlement), per record per year (in the case of registrar services) or per message (in the case of SWIFT). Although CSD services are regulated, they are in general not subject to price caps.

With regards to their ownership structure, there are different types of arrangements. Some CSDs belong to publicly listed companies while others are user-owned. Some CSDs are vertically integrated into a corporate group that includes an exchange and/or a CCP, while others operate separately from the trading and clearing infrastructures. CSD membership also varies across jurisdictions. While in most cases CSD members are large banks and dealers, in some European markets (Nordic countries and Greece, among others), there is a “direct holding” structure, which means that end-users (e.g. the beneficial owners of a security) have individual accounts at the CSD level.

5.2 The timing of DLT innovation

DLT comes at a time when new regulations and changing economic conditions are reshaping the financial services industry. Financial institutions are facing greater scrutiny, a challenging macro-economic environment, increased capital constraints and wide-ranging regulatory changes, all of which have exerted pressure on profit margins.

As such, there are stronger incentives for technological innovations that either generate new revenue or reduce costs. An area with great potential for cost reductions is post-trade processing. As Morgan Stanley puts it, “for many banks, especially investment banks in 2016, a radical reduction and simplification in processing costs will be a blessing” (Morgan Stanley, 2016; p.15). In this context, it is not surprising that banks are seriously considering the prospect of developing DLT to reduce their back-office costs.

One should be aware, however, that different participants have different incentives for the adoption of DLT in settlement and face different risks depending on the final configuration adopted. While the sell-side may be primarily interested in reducing costs, an important incentive for clearing agents, CSDs and messaging services is to influence the innovation process so as to retain or conveniently redefine their market position. For the sell-side, the risks are linked to the high initial costs of R&D and of replacing the legacy systems and modifying business models. For other incumbents, however, the adoption of fully-fledged DL-based settlement could be a threat to their monopolistic position and business models.

5.3 The economics of technological innovation in securities settlement

Given that the settlement industry is highly concentrated, it satisfies the assumptions of the Schumpeterian Hypothesis, which predicts that technological innovation is more likely to occur in industries characterized by large firms with market power. The simple logic behind firm size is that large firms have more money to invest (or can obtain financing more easily) and can obtain larger

²⁸ For example, according to Euroclear’s Annual Report 2016, 85% of total consolidated operating income was from fees and commissions.

gains to innovation because they typically have a larger market share. An alternative way of expressing this is to say that larger firms can spread the costs of innovation over more units and therefore make a more profitable use of this technology. For example, if firm A produces 100 cars per year and firm B 10,000 cars, then only firm B will invest in a \$10,000 technology that will cut production costs by \$1 per car, per year. That is because firm B will recoup the R&D cost in just one year, whereas firm A will recoup it in 100 years.

Market power is important because it impacts the ability of the firm to appropriate the gains to innovation. This theme is addressed by Arrow (1962) who emphasizes that the “peculiar attributes” of knowledge as a commodity may make monopoly power difficult to exert. Technological innovations are, by their nature easily duplicated, and hence will have low appropriability. The know-how associated with these innovations is therefore likely to become a public good since it is both *non-rivalrous* - i.e., use by one party does not preclude simultaneous use by another – and, importantly, *non-excludable* i.e., once the technology is developed, it becomes widely known and available for anyone to use it. This is important because, in such cases, investment in acquiring know-how is likely to be sub-optimally low as market participants will fail to internalize the benefits (and profits) that accrue to others by using the new technology. This underinvestment tends to be greater for what Arrow refers to as “basic research” – research that is applicable to multiple innovative processes.

Arrow discusses the role of patents as a solution to underinvestment in basic research and identifies a trade-off. On one hand, as Arrow emphasizes, the cost of transmitting information is typically near zero and hence the distribution of know-how should be unlimited so that everyone can benefit from it. On the other hand, if access to know-how is unlimited with no (intellectual) property rights attached to it, no single party will have sufficient incentives to invest in acquiring it in the first place. As Arrow states, “precisely to the extent that [the attainment of property rights] is successful, there is an underutilization of the information.” (p. 617).

Despite being over a half-century old, Arrow’s work is particularly relevant to the discussion here because DLT know-how is essentially information in the form of computer code.²⁹ In particular, one of the features of DLT is that new developments tend to be open-source. Hence, it is arguably best to model innovation in DLT as adding to a public good.

Arrow suggested numerous remedies to the problem of under-investment in public goods such as technological know-how.³⁰ Research performed at public universities is one solution. He also advocated the use of different organizational structures for research, such as joint ventures by private firms. This is a recommendation that seems to be right on the mark for innovation in today’s securities settlement industry. Consortia and communities such as R3, the Post-Trade Distributed Ledger Group, the Hyperledger Project and the MIT Digital Currency Initiative have recently sprung up that seek to facilitate collaboration, including that on basic research. Aspects of their structure and likely impact on the supply of innovation are captured by the literature on R&D for the production of public goods.

²⁹ It is important to note that we are not referring here to an up-and-running DLT settlement network, but rather to the technology that underlies such a network.

³⁰ Arrow focused on more than just appropriability in his paper. He also emphasized uncertainty and increasing returns as reasons why a free enterprise economy will underinvest in invention and research (p. 619)

Bozeman et al. (1986) provide a public goods model of cooperative R&D which formalizes the efficiency gains to collaboration. The model shows that individual expenditures on R&D will be too low in the absence of cooperative agreements, as a result of positive externalities. Because research output is a public good and rents cannot be appropriated, individual developers do not internalize the full social benefit of their R&D. However, cooperation between developers achieves an optimal level of R&D investment.

These ideas can be formalized in a simple mathematical model that captures the ideas expressed verbally and graphically in the Bozeman et al. paper. Suppose there are n firms and these firms have fixed budgets for R & D, b_i , $i = 1, \dots, n$. Firms decide how to allocate their fixed budgets between an applied component x_i that reaps internal private benefits and contributions to a public component y that benefits all market participants. Assume that each firm's (gross) profits are increasing in the quantities of public and applied research and that firms have convex iso-profit curves that are described by functions $\Pi_i(x_i, y) = k$ for any positive constant k .³¹ Moreover, assume, for simplicity, that the marginal cost of producing a unit of the public good in terms of the applied research component (i.e., the marginal rate of transformation) is equal to 1 for all firms. Then each firm i , acting independently, chooses x_i to maximize $\Pi_i(x_i, y)$ subject to $0 \leq x_i \leq b_i$ and $y = \sum_{i=1}^n b_i - x_i$.

A Nash equilibrium is given by a profile (x_1^*, \dots, x_n^*) such that, for each firm,

$$x_i^* = \operatorname{argmax}_{x_i} \Pi_i(x_i, y(x_i, x_{-i}^*)).^{32}$$

In an interior Nash equilibrium³³ each firm chooses a quantity of applied research such that the marginal rate of transformation (which equals 1) equals the marginal rate of substitution between the private and public goods. Formally, the first-order necessary condition for an interior solution is

$$\frac{d\Pi}{dx_i} = \frac{\partial \Pi_i(x_i, y(x_i, x_{-i}^*))}{\partial x_i} + \frac{\partial \Pi_i(x_i, y(x_i, x_{-i}^*))}{\partial y} \frac{dy(x_i, x_{-i}^*)}{dx_i} = 0,$$

which, at an optimal solution, yields the equilibrium conditions described above:

$$MRS_i(x^*) = \frac{\frac{\partial \Pi_i(x_i^*, y(x^*))}{\partial x_i}}{\frac{\partial \Pi_i(x_i^*, y(x^*))}{\partial y}} = 1, \quad \text{Eq. (1)}$$

$i = 1, \dots, n$.

Next we turn to the socially optimal level of investment. Since it is reasonable to assume profits can be redistributed between firms, the efficient level of investment in the public good can be found by maximizing the joint profit of the firms. Formally, the joint optimization problem is

³¹ Convexity is convenient, but not necessary, as will be clear from the arguments that follow.

³² We abuse notation slightly and let $y(x) = \sum_{i=1}^n b_i - x_i$.

³³ Existence and uniqueness of a Nash Equilibrium for this game is established, under additional assumptions, in Bergstrom et al. (1986).

$$\max_{(x_1, \dots, x_n)} \sum_{i=1}^n \Pi_i(x_i, y)$$

subject to:

$$y = \sum_{i=1}^n b_i - x_i.$$

Suppose an interior solution $x^{**} = (x_1^{**}, \dots, x_n^{**})$ exists in which all firms conduct some applied research and no firm conducts only applied research. The first-order necessary conditions describing an interior solution now require that the marginal increase in firm i 's profit from increasing their amount of applied research equal the marginal increase in joint profit from adding to the public good:

$$\frac{\partial \Pi_i(x_i^{**}, y(x^{**}))}{\partial x_i} = \sum_{i=1}^n \frac{\partial \Pi_i(x_i^{**}, y(x^{**}))}{\partial y},$$

$i = 1, \dots, n.$

Simple manipulation of these equations delivers the standard Samuelson condition:

$$\sum_{i=1}^n MRS_i(x^{**}) = 1. \quad \text{Eq. (2)}$$

Clearly, at the interior voluntary provision solution described by Eq. (1), the Samuelson condition is violated: specifically, $\sum_{i=1}^n MRS_i(x^*) = n > 1$. Firms substitute too little into the public good because they do not internalize the benefits this substitution has for others.

More generally, given differentiable profit functions that are non-decreasing in public and private research and strictly increasing for more than one firm, then, whether there is convexity or not, the supply of public research in an interior, voluntary-provision Nash equilibrium will be too small. In the interior Nash equilibrium, the marginal cost of the public good in terms of the private good equals the marginal rate of substitution between the public and private goods of any single contributor. But, if at least two firms strictly benefit from both private and public research, then the sum of the marginal rates of substitution over the population in the voluntary solution must exceed the marginal cost.

If the sum of marginal rates of substitution between the public and private goods is greater than the marginal cost of the public good, then a Pareto improvement is possible. This is so because if the total amount that people are willing to pay for a marginal unit of the public good exceeds the marginal cost, then there is some way of increasing the amount of the public good by a little bit while dividing its cost so as to make everyone better off.

In Box 1 we illustrate the Bozeman et al. model for the case of two firms. This box is intended only to illustrate how strategic interaction leads to inefficiency. We are not suggesting that the simple Cobb-Douglas technology used in the illustration accurately captures the interaction of applied and basic research in generating firm profits.

Box 1: An illustration of the Bozeman et al. (1986) model

Consider two firms, A and B, with fixed budgets for R&D denoted by a and b , respectively, where $a > b > a/2$.³⁴ As before, firms decide how to allocate this budget between an applied component x_i and contributions to the public component y . Assume each firm's iso-profit curve is described by the Cobb-Douglas function $x_i y = k$ for any positive constant k , $i = A, B$.³⁵ Moreover, as before assume that the marginal cost of producing a unit of the public good in terms of the applied research component is equal to 1. Then firm A, acting independently, chooses x_A to maximize $x_A y$ subject to $x_A \leq a$ and $y = a - x_A + b - x_B$, and likewise for firm B.

The mathematical structure generates a simple reaction functions for each firm. The reaction function of firm A is:

$$x_A(x_B) = \frac{a + b - x_B}{2}$$

and the reaction function of firm B is:

$$x_B(x_A) = \begin{cases} \frac{a + b - x_A}{2} & \text{if } x_A \geq a - b \\ b & \text{otherwise} \end{cases}$$

The resulting Nash equilibrium choices of x_A and x_B are characterized by the intersection of these reaction curves, as shown in Figure 1.




Figure 1: A public goods model of cooperative R & D: x_i denotes the quantity of applied research for firm i , $i = A, B$. $x_A(x_B)$ and $x_B(x_A)$ are the reaction functions of firm A and B, respectively. NE denotes Nash equilibrium quantities of applied research. E identifies a socially optimal level of applied research.

At the Nash equilibrium solution $x_A = x_B = (a + b)/3$ and $y = (a + b)/3$. However, this solution is

³⁴ We change notation slightly to avoid using subscripts. The restrictions on the relative magnitudes of a and b ensure a Nash equilibrium exists in which both firms contribute to the public good.

³⁵ We can think of this as a production function multiplied by an output price which is normalized to 1.

not efficient. The cooperative solution is achieved is by maximizing joint profits $x_A y + x_B y$ subject to the constraints $x_A \leq a$, $x_B \leq b$ and $y = a - x_A + b - x_B$. The profit maximizing joint venture has $x = y = (a + b)/2$, where x is the combined private production which can be allocated in any way; a natural division would be $x_A = x_B = (a + b)/4$, as shown by point E. This exactly splits the gains to cooperation without requiring transfers. Under cooperation the positive externality of each firm's public good investment is internalized, there are lower levels of applied research and a higher level of production of the public good, making both firms better off. This conclusion does not depend on there being only two firms.

5.4 A case for cooperation

The above analysis suggests that innovation in securities settlement may be best achieved through cooperation, but it is not entirely clear who in reality might actually contribute to innovation. In the mathematical models presented earlier, innovators are generic firms. In reality, however, there are many different service providers at different points of the post-trade value chain as well as customers with whom these providers interact. For this reason, it is worth asking whether collaboration in innovation should involve all or a subset of these providers.

We think that collaboration in DLT innovation would achieve the most if it involved participants from the entire spectrum of post-trade service provision. This is because DLT can be considered both a *radical* innovation (i.e., one that has a high potential impact) as well as a *systemic* one (i.e., one that affects the entire post-trade value chain and not just one or a few of its components).

A radical innovation (as opposed to an incremental one) can fundamentally alter the character of the processes involved and potentially undermine the position of current market incumbents. For instance, in a DL-based system where transactions are cleared and settled on a bilateral basis between ultimate buyers and sellers there will likely no longer be central settlement service providers per se but instead providers of technological solutions that facilitate bilateral settlement. In such instances of radical innovation, cooperation between younger and innovative firms on one hand and larger incumbent ones on the other may be useful. This is because while older incumbent firms may have a better understanding of the legal and economic dimensions of post-trade processes they also tend to be constrained in their ability to innovate due to structural inertia resulting from sunk costs, limited information on behalf of their CEOs, etc. or external constraints such as high exit costs. Structural inertia and sunk costs have been for many years one of the main obstacles for streamlining post-trade processes. In addition, banks have been reluctant to undertake technological changes where the benefit is more long term. For example, Morgan Stanley reports that banks only invest in new technologies if the initial investment is recouped by the resulting cost savings within 3 years (Morgan Stanley, 2016).

A systemic innovation, on the other hand, is one that affects a wide range of pre-existing processes regardless of whether it affects them in a substantial or less substantial manner. In addition to being radical, a full DL-based solution for settlement will also be systemic to the extent that it has the potential to change (and possibly make redundant) many or all of the existing post-trade processes. And when it comes to systemic innovations, there is usually a need for tight coordination between various industry participants in order to fully exploit these innovations (Teece, 1996). This is another reason for incumbents (banks, custodians, CSDs) to seek to participate in joint projects which include FinTech start-ups.

Overall, collaboration between incumbent firms and FinTech start-ups may provide a strong route for DL-based innovation to proceed. Without the support of incumbent firms, FinTech firms may not succeed on their own as they lack the incumbents' clientele. Alternatively, incumbent firms may be unable and/or unwilling to innovate on their own due to both internal and external constraints.

5.5 The change from a permanent to a contestable monopoly

In a permanently monopolistic market, firms are typically profitable and although they have the capacity to innovate (on account of their profitability) they have little incentive to do so as they already have the entire market. This is evident in the case of CSDs where the shortcomings of core post-trade processes largely derive from a lack of incentives to pursue more efficient solutions.³⁶

However, when FinTech³⁷ start-ups entered into the post-trade space with DL proposals that promised lower costs and increased efficiency for the end users, the monopoly became contestable (with the possibility of end-users migrating to lower-cost solutions) and the incumbents became incentivized to innovate even if it was just to keep away potential competitors. As Morgan Stanley points out, "it took the threat of new entrants funded by VCs at growth-tech multiples (not bank multiples) to get the incumbents together to discuss how to deliver more speed and efficiency at lower cost" (Morgan Stanley, 2016; p.5).

Since not all incumbents have the resources to conduct their own R&D and because some are already technologically behind their challengers, they may opt to buy or form alliances with those competitors. This way they can bring the new know-how in-house while at the same time eliminating potential competition. The need for coordination and standardisation which is required for DL technologies also favours forming such strategic alliances.

There is also the view that organisational determinants of the direction of innovation cannot be reduced to a discussion of the virtues of competition and monopoly (Teece, 1996). In fact, perfect competition can sometimes be seen as incompatible with innovation.³⁸ A good example where the link between market power and innovation is undone is precisely the case of FinTech firms, which are not cost intensive and may get support from venture capital.

A large part of economic theory takes the view that technology changes tend to imply persistent dominance and concentration. The creation of de-facto standards, for example, will usually benefit the incumbent. However, there is also a line of thought which supports the idea that drastic innovations can give the entrant an advantage over the incumbents (e.g. Reinganum, 1985). In this setting, major innovations would occur at early stages and would be produced by new entrants, while the incumbents will incorporate at a later stage and would only produce minor incremental

³⁶ Although CSDs have faced some competition, this has been mostly around the provision of ancillary services.

³⁷ See Schindler (2016) for a compelling account of why the FinTech revolution is occurring at this time. Many of the buy and sell side arguments he provides apply here.

³⁸ This is the conclusion drawn in Shell (1966 and 1967), and further explored in Shell (1973) and Antinolfi, Keister and Shell (2001), in which technical knowledge is modelled as a non-conventional factor of production that complements capital and labor in ways that are invariant to scale. Increases in all factors, including technical knowledge, imply increasing returns to scale which implies that the pure competitive provision of technical knowledge is not possible (this is an immediate consequence of Euler's theorem).

improvements. DLT's classification as a radical and systemic innovation makes these latter arguments seem more compelling.

5.6 The role of governments in post-trade process innovation

Governments have long recognized the benefits of encouraging cooperative research and the role government can play in addressing positive externalities, discouraging free riding and reducing potential antitrust liabilities of joint research. The National Co-operative Research and Production Act of 1993 in the United States is one example where the government encourages joint research by clarifying antitrust rules, reporting procedures and by possibly providing legal aid.³⁹

More direct involvement could come in the form of tax adjustments that punish free riders and reward cooperative research participants. Bozeman et al. (1986) elaborates on this idea and argues that incentives to free-ride are higher, the smaller the relative size of the firm. Their analysis suggests that a tax incentive to "relatively small firms should be disproportionately greater than that for the relatively larger firms in order to induce the optimal amount of R&D." (p. 265)

Governments can also directly fund collaborate research efforts. This is common in the US electrical utility industry, for example, where the Electric Power Research Institute receives government funding in the form of grants and engages in collaborative research. Alternatively, government agencies can fund grants that are sought by members of research consortia. The Seibel Energy Institute in the United States is a consortium of public and private agencies that provide seed money to researchers who hope to obtain larger scale government funding for energy research. But in encouraging and incentivising innovation, it is important that governments do not to pick winners, which could ultimately lead to a misallocation of resources.

Innovations in technologies for securities settlement have the potential to impact not only private banks, messaging services and clearing entities, but also financial market infrastructures (FMIs). FMIs are typically run or at least supervised by central banks. Hence there is a role for central banks and other government agencies to participate directly in collaborative research efforts with private institutions. Such initiatives are already under way. The central banks in Canada, Singapore and Hong Kong are collaborating directly with groups of their domestic private banks and private FinTech firm R3 to explore potential uses of DLT.⁴⁰ The Bank of England and the Federal Reserve Bank of Boston joined the Hyperledger Project. Deutsche Bundesbank and Deutsche Börse joined forces to build a functional prototype of a blockchain-based securities settlement system. The FICC in the United States is working with Digital Asset. The central bank of Russia coordinated the development of Masterchain, which is being used by private Russian banks to explore DLT applications. Finally, the European Central Bank and the Bank of Japan are collaborating on research into the use of DLT for financial market infrastructure. Given the high degree of interconnectedness of private and public infrastructures inherent in many DLT applications such cooperation seems essential.⁴¹

³⁹ <https://www.justice.gov/atr/national-cooperative-research-and-production-act-1993>

⁴⁰ As mentioned in footnote 21, the Canadian efforts are summarized in Chapman et al. (2017).

⁴¹ The idea of government-sponsored collaboration is of course not uniquely applicable to DLT. The same rationale would apply to any kind of know-how with applications to security settlement.

Government agencies can also play a role in facilitating the success of private R&D. Clarification of industry rules and the regulatory framework for new technologies is very important. For instance, the UK Financial Conduct Authority (FCA) has created the Regulatory Sandbox which, according to the FCA website, provides “a ‘safe space’ in which businesses can test innovative products, services, business models and delivery mechanisms in a live environment without immediately incurring all the normal regulatory consequences of engaging in the activity in question.”⁴² In fact, the industry needs more than regulatory clarification. One common problem in financial market infrastructures is the lack of industry standards for reporting financial market transactions. This problem is currently being addressed by the wide-scale adoption of ISO 20022, but similar coordination may be required in the development of new technologies for the post-trade process. Governments, alongside the industry and international standards authorities can play a role in this by providing guidance or even proposing standards.

6. The economics of a fully decentralized DL-based settlement industry

DLT’s ultimate promise is to allow end-users to engage in bilateral, decentralized settlement. In this section we analyse the likely economic dynamics of a theoretical, decentralized, DLT-based, post-trade industrial landscape. In doing so, we examine whether these dynamics might lead to instances of market power and what the implications for financial stability could be. We also identify potential areas of regulatory intervention.

It is not a-priori clear how DLT might affect the structure of the settlement industry. In the near term, it is more likely that existing CSDs and other intermediaries might use DLs in an attempt to lower their operating costs without significantly changing their business models. In this case, settlement costs for market end-users may not change substantially as any savings resulting from DLs could be internalized and simply translate into higher profits for incumbents. However, given that DLT can in principle facilitate peer-to-peer settlement, in theory it has the potential to fundamentally alter the structure of the industry over the longer term by decentralising it. For this reason, we explore here the economic properties of a hypothetical decentralized DLT-based settlement environment.

The key difference between the current structure of the settlement industry and a decentralized, DL-based one, is that the latter does away with intermediaries: securities transactions are cleared and settled on a bilateral basis between ultimate buyers and sellers who themselves are part of the network of distributed ledgers. The information on executed transactions is maintained in each and every node of the network and the validation of new transactions is done via a consensus mechanism.⁴³

In this new environment, there would likely be no settlement service providers per se but instead providers of technological solutions that facilitate bilateral settlement. In other words, end-users would purchase software, technical support and participation rights in a given DL network, all of

⁴² See <https://www.fca.org.uk/firms/project-innovate-innovation-hub/regulatory-sandbox>

⁴³ This setting corresponds to one of the potential scenarios described in Pinna and Ruttenberg (2016). That is, a scenario where the current post-trade processes are superseded by automated clearing and settlement, taking place among a network of security-issuing entities and final investors.

which would enable them to engage in bilateral, decentralized settlement. As such, participation in a DL settlement network would likely be excludable and therefore, contrary to the DL know-how itself, constitute a *private* good.

From an economic standpoint, the key advantage of this decentralized system is that it would economize on intermediation costs and the associated fees charged by traditional settlement service providers, as explained earlier. The elimination of these costs would create an economic surplus which, in a competitive environment, benefits end-users by reducing their settlement costs. However, whether a decentralized industry for DL-based settlement services ends up being a competitive one, largely depends on its characteristics. We next review these characteristics in an attempt to gauge how likely it is for the industry to be competitive.

6.1 The cost of processing transactions in DL systems

We start by considering the costs of DL-based settlement. We argue that, while in traditional blockchain systems average costs are likely to be increasing in the number of transactions, in the DL systems which are most likely to be adopted for securities settlement, average costs do not increase significantly with the size of the network. To show why this is the case, we distinguish between the two functions of the DL confirmation process, namely transaction validation and uniqueness.

In all settlement systems, transaction processing is done in two steps. The first step is to check whether a transaction is, in principle, valid, i.e. it satisfies a set of necessary criteria and conditions which however are not sufficient on their own to ensure transaction finality. Transaction validation involves ensuring that the format and required fields of the transaction information are correct, that the signatures check out, that the transaction is compliant with any constraints/rules that pertain to it and that the security buyer has a legitimate claim to the funds she is spending.

The second step is to confirm the transaction by ensuring that cash and securities included in the transaction are not associated with any other pending transaction. This step is critical because, in principle, there could be two (or more) equally *valid* but *conflicting* transactions which cannot be simultaneously settled.⁴⁴

The traditional banking analogue of these two steps might be the processing of a cheque a customer writes and signs against her bank account.⁴⁵ First, there is the validation stage. Is the signature correct? Has she filled out the cheque correctly? Did she use ink and not a pencil? A teller can complete this step simply by looking at the check and he can use his local computer terminal to check if the customer *appears to have* sufficient funds in her account *at the time the cheque is presented*. But he doesn't know if other cheques have been presented at other branches or if the payer has just hit "OK" on an ATM transaction to withdraw all the cash. It is only when everything is brought back and processed centrally that the determination can be made that the cheque has indeed cleared (i.e. is "confirmed" in the blockchain sense).

⁴⁴ Because of the double-spending problem, for instance.

⁴⁵ The authors thank Roger Willis for suggesting this analogy.

Costs in traditional blockchain systems

The above-described two steps in transaction processing also determine processing costs. In traditional blockchain systems, such as the Bitcoin protocol, transaction validation is performed by the *nodes* in the system and, also usually (but not always), by the miners. Critically, because these systems work on the basis of maintaining a single global record, all transactions must be validated by all nodes. Transaction uniqueness, on the other hand, is guaranteed by *miners* in public systems like Bitcoin and Ethereum, who perform a proof of work (PoW) problem. PoW involves solving a hash problem,⁴⁶ the difficulty of which does not depend on the size of the transaction block. This process is massively computationally intensive, but the computational effort in mining does not vary significantly with transaction throughput. Therefore, in traditional blockchain systems the cost of processing transactions includes a variable component associated with transaction validation and an, at least approximately, fixed component associated with transaction uniqueness.

Thus, to see how costs vary with scale, we take a closer look at the costs associated with transaction validation alone. For a fixed number of nodes, as the number of transactions being created in the system grows, the load on each node increases. This implies that the total costs of transaction validation are proportional to the size of the network.

The number of transactions can also increase because new participants join the network. In particular, the addition of a participant to a network with N existing participants leads to an additional $2N$ potential types of transactions.⁴⁷ The need to validate additional transactions adds costs to existing nodes in the network and makes the costs for entering nodes higher as well. Adding new participants, however, may result in an exponential growth of transactions (greater than the potential $2N$ transaction opportunities that are created by adding a participant to a system with N nodes) because of network effects. This means that as the network grows (adds more nodes) validation costs could grow exponentially.

Costs in DL-based settlement systems

It is highly unlikely, however, that future DLT systems for securities clearing and settlement will utilize traditional blockchain technology. The primary reasons have to do with privacy requirements of financial market applications and the need for settlement finality.⁴⁸ Moreover, there is general recognition that the benefits that proof-of-work provides, such as trust-less validation and censorship resistance, are neither required nor desired when it comes to securities settlement.

As such, alternatives to proof-of-work such as Corda's "Notary" architecture have now emerged that are driven by requirements around scalability, privacy and usage assumptions. The notary design differs from other consensus models in that it utilizes a trusted authority, which is acceptable in

⁴⁶ A hash function is a function that takes input of any length and produces output of fixed length that is deterministic, but not predictable. As such, it is impossible to predict what the output will be if there is any change to the input. The hash problem involves hashing information about new blockchain transactions, past transactions and the miner's public address and combining it with a nonce (i.e. a meaningless additional information that changes the hash output), in order to obtain a result with certain characteristics. The only way to achieve a desired hash result is through trial and error, trying literally billions of different nonces until a solution is achieved. It is very computationally expensive to find a solution, but once one is found it is easy to check that the solution is valid.

⁴⁷ N buys and N sells between the new participant and each of the existing participants.

⁴⁸ Finality in PoW based blockchain systems is probabilistic due to the possibility of forking. See the relevant discussion in Section 4.

many financial market applications. This allows consensus to be reached on an individual transaction basis, rather than in blocks, with limited information sharing. Other non-PoW protocols such as Hyperledger Fabric or Intel's Sawtooth Lake "Proof of Elapsed Time" algorithm preserve the blockchain structure, but modify the consensus process in order to allow enhanced confidentiality and scalability.

DL systems that do not rely on proof-of-work have significantly lower costs of ensuring transaction uniqueness than traditional blockchain systems. Moreover, transaction validation can be done using a partial history of DL transactions and hence the costs to perform this function do not increase proportionally with the size of the network or the number of transactions.

6.2 Cost function and network externalities

We next consider the shape of the cost function associated with the provision of DL-based settlement solutions. Those firms or entities engaged in the development of DL-based solutions for settlement purposes are likely to face high costs associated with the relevant R&D expenses. However, once a DL-based solution is developed and firms move to the implementation phase, R&D costs are sunk and do not factor into production decisions. As discussed above, in private, permissioned systems there is an administrative role that the system provider would either offer or at least facilitate which would be independent of transaction volume. Hence, there would likely be some fixed costs associated with a DL-based securities settlement system.

Access to a commercialized DL-based solution could likely be offered to interested parties at virtually no cost as the settlement service provider would likely only have to sell software and network access rights. And as we have argued in Section 6.1, verification of an ever-increasing number of transactions is unlikely to be a factor in the provider's cost function both because the computing power required is not large per se, but also because each additional user would contribute their own computing power and therefore bear part of the additional verification costs. Therefore, once the DL settlement solution is rolled out, it should be able to handle large settlement volumes at low and decreasing average costs. Decreasing average costs is an important conclusion because it implies that there is no economic limit to the size of the distributed ledger.

A DL-based securities settlement solution would also likely be characterized by network externalities. This means that the more market participants adopt a given solution, the more valuable this solution becomes to existing and potential new users. This is because a larger DL-based settlement network allows users to settle trades with more counterparties the same way a larger telecommunications network allows users to contact more people thus rendering their own participation to the network more valuable. For this reason, whatever DL-based solution is adopted first, would likely be increasingly difficult to challenge later on. This, in turn, means that early entrants in this industry may have a significant first-mover advantage.

6.3 Industrial structure, pricing and welfare

The above characteristics suggest that it is likely that a DL-based settlement industry, if it arises, will be concentrated, possibly becoming a monopoly or an oligopoly. This is because, on one hand, the decreasing average costs could induce several firms to pool resources, thus creating one or few providers with substantial market share. In addition, network externalities will tend to push the industry toward a single (or few) solution(s). Once a provider has established themselves, they may be increasingly difficult to challenge due to their first-mover advantage. Thus, one conclusion so far is that despite the promise of dis-intermediation that the DL technology brings about, the industry may well simply transition from a CSD-centred monopoly to a DL-provider-centred one.

The main problem with a potentially concentrated DL settlement industry is the same as in any concentrated industry, namely that welfare is potentially reduced. This can happen when pricing policies designed to reap monopoly profits prevent full extraction of market surplus, a problem commonly referred to as deadweight loss.

The potential for deadweight loss in a future DL-based settlement industry in part depends upon the ownership structure of the distributed ledger platform. To be clear, we have argued that R&D in DL technology is a public good, but that a working platform will ultimately be owned and operated by someone. We expect this “someone” to be a large consortium, but it could be a single entity. In the former case, pricing would likely relate to membership in the consortium. That is, access to the ledger could be allowed for non-members at a price or it could require purchasing membership. New membership pricing could take the form of a two-part tariff where new members are charged (an estimate of) their surplus to participating on the ledger and then required to pay marginal cost for future transactions. This would be efficient, but it would involve transfers of surplus from new members to the members that originally invested in the DL technology. In fact, it seems plausible that anticipation of this type of pricing policy is a factor that motivates new members into the consortium at the R&D stage.

A monopoly ownership structure would arise if a single entity provided the successful platform or if a successful group decided to operate as a monopoly. Predicting the form of monopoly pricing is difficult without a deeper assessment of market demand characteristics. However, as is the case with any monopolist facing a downward-sloping market demand curve, profit maximization without perfect price discrimination results in a reduction of output and (yet again) deadweight loss. We next elaborate on the possible pricing schemes for the case of a pure monopoly.

a) Simple pricing

If the settlement service provider charges a fixed price per unit of settlement services sold, then they would optimally charge a mark-up over their marginal costs which would be suboptimal from a welfare perspective, since consumers willing to purchase a certain amount of services at above marginal cost prices would not be able to do so.⁴⁹ We cannot predict what share of the market might be priced out of this service or what clearing and settlement arrangement they would use. Presumably they would continue to use existing structures although these might become more expensive on a per unit basis if there is a large migration of business to the new technology. It is also possible that regulators would impose pricing restrictions on a monopolist

⁴⁹The monopolist mark-up would be higher the more inelastic demand is.

DL-based settlement provider to ensure broader market participation, as is common in the case of natural monopoly cost structures.

b) Non-simple pricing

Alternatively, a monopolistic provider of DL-based settlement services might engage in non-simple pricing in an attempt to extract more economic surplus. Examples of non-simple pricing schemes include:

- i) *Declining block pricing*: The monopolist charges each client a declining price for each quantity “block” of settlement services provided. Here there is no limit in the final supply of services provided⁵⁰ but the first “blocks” of purchased services are substantially more expensive than in simple pricing.⁵¹ Although the last “blocks” of purchased services (assuming large enough demand) are cheaper than in simple pricing, the average price across all blocks may be higher than in simple pricing. Thus, the monopolist may appropriate more economic surplus than in a simple pricing scheme.
- ii) *Two-part tariff*: In this scheme, the monopolist charges an initial upfront fixed fee for access to the DL network and then charges a variable fee for each unit of settlement services sold. The optimal way for a monopolist to do this is to set the variable fee equal to its marginal cost and the fixed fee equal to the entire consumer surplus. As discussed above, we think it is likely that a consortium would apply this pricing policy to new members.
- iii) *Third Degree Price discrimination*: Given that in reality not all consumers value equally each unit of services purchased (i.e. have different demand curves), the monopolist can maximize its profit by charging a different price to different types of market participants. In the DL-based settlement landscape this could mean for example that retail investors pay a different price than institutional ones. It is worth noting that price discrimination could be implemented in combination with any of the above pricing schemes.

The pricing schemes described above (i.e. block pricing, two-part tariffs and third degree price discrimination) would all result in reduced (or zero) deadweight loss so in this respect they could be efficient. However, in those instances, most of the economic surplus would accrue to the monopolist, giving rise to an unequal distribution of economic surplus.

6.4 The role of regulators

If the above description, of what a DL-based settlement industry landscape might look like, is accurate, then the question that arises as to what the role of regulators should be. We think that competition and prudential regulators might have a role to play, both because of the likely structure of the industry and because of financial stability concerns. Assuming that the DL-based settlement landscape is concentrated (i.e. a monopoly or oligopoly), competition regulators might need to consider whether to regulate prices in a manner that minimizes deadweight losses (but still allows settlement service providers to recoup their costs).⁵² Additionally, the regulator may require that DL

⁵⁰ The monopolist continues to sell as long as his (declining) price remains above marginal costs.

⁵¹ Decreasing block pricing is currently used in Fedwire Funds, the US large value payment system. See Copeland and Garratt (2016) for a detailed description of Fedwire pricing and for a discussion of how this pricing scheme meets the mandated cost recovery objectives.

⁵² This type of intervention is also known as “Ramsey pricing”.

settlement technology be compatible with alternative DL technologies so that other DL providers are incentivized to innovate and develop competitive solutions. This could partially alleviate the role of the network externality in concentrating activity.

From a financial stability perspective, we note that if there is a first mover advantage in the provision of DL technology, and time is of essence, then firms may rush to roll out a product that is not yet fully developed. As such, financial stability risks could ensue if some market participants start using a solution with design flaws or without having performed extensive testing including against cyber risks. This could be particularly problematic if the industry is concentrated and therefore largely reliant on a single DL network.

7. Concluding Remarks

The distributed ledger technology (DLT) is a mechanism by which data can be stored, shared and synchronized across multiple locations without the need of a single administrator. Although DLT became more widely known through its application on virtual currencies, such as Bitcoin, it has potential applications in other areas of finance. Post-trade processing, in general, and securities settlement in particular, is one area where DLT is often seen as having the potential to reduce back-office costs, increase security and simplify processes, bringing net benefits to the end-user. Although it is unclear how or whether DLT will finally be adopted, one potential scenario would be a DLT-based solution that facilitates peer-to-peer clearing and settlement and dis-intermediates the industry. If this scenario materializes, the obvious question that then arises is what the new post-trade industrial landscape might look like and what the role of regulators might be.

To answer this, we first discuss how technological innovation in post-trade processing could benefit from collaboration among industry participants. In the particular case of DLT (or similar technological innovations), we argue that technological know-how resembles a public good in that it is both expensive to develop and difficult to appropriate. This, in turn, means that there are unlikely to be sufficient private incentives for innovation in a competitive environment and that by facilitating some degree of coordination among market participants, relevant authorities could help foster the optimal amount of innovation.

We then argue that in DLT-based settlement, participation in a settlement network constitutes a private good offered by whoever develops the technology faster and better. Since DLT-based settlement is a technology characterized by network externalities, by a first-mover advantage and by decreasing average costs, in equilibrium there will likely be one or only a few providers of DLT-based settlement. This means that the economic surplus obtained by technological improvements and by dis-intermediating traditional post-trade processing risks either being offset by deadweight loss or being entirely consumed by the new network provider(s), depending on the pricing scheme that these provider(s) choose to employ. In either case, absent any regulatory safeguards, end-users may not benefit from reduced costs.

Finally, from a financial stability perspective, there could be significant implementation risks if market participants, under the pressure of gaining a first mover advantage, deploy solutions which have not been sufficiently tested, including against cyber risk, or which suffer from other design flaws.

8. References

- Antinolfi, G., Keister, T., and Shell, K. (2001), "Growth dynamics and returns to scale: Bifurcation analysis". *Journal of Economic Theory* Vol. 96, No. 1-2, 2001, pp. 70-96. doi:10.1006/jeth.1999.2632.
- Arrow, K.J. (1962), Economic welfare and the allocation of resources from inventions, in R.R. Nelson (ed.), *The rate and direction of inventive activity*, Princeton University Press, Princeton, NJ.
- Bergstrom, T., Blume, L., and Varian, H. (1986), "On the private provision of public goods", *Journal of Public Economics* (29), 25-49.
- Bozeman, B., Link, A., and Zardkoohi, A. (1986), "An economic analysis of R&D joint ventures", *Manage. Decis. Econ.*, 7, 263–266. doi:10.1002/mde.4090070408
- Broadridge (2015). "Charting a path to a post-trade utility. How mutualized trade processing can reduce costs and help rebuild global bank ROE". Broadridge White Paper. <http://www.broadridge.com/broadridge-insights/Charting-a-Path-to-a-Post-Trade-Utility.html>
- Chapman, J., Garratt, R., Hendry, S., McCormack, A. and McMahan, W. (2017), "Project Jasper: Are distributed wholesale payment systems feasible yet?". *Bank of Canada, Financial System Review*, June 2017.
- Copeland, A. and Garratt, R. (2015), "Nonlinear pricing with competition: The market for settling payments", Federal Reserve Bank of New York Staff Reports, no. 737. August 2015.
- CPMI (2015). "Red Book Statistical Update". Committee on Payments and Market Infrastructures (CPMI), December 2015 (<http://www.bis.org/cpmi/publ/d155.pdf>)
- CPMI (2017). "Distributed ledger technology in payment, clearing and settlement. An analytical framework". Committee on Payments and Market Infrastructures (CPMI), February 2017. <http://www.bis.org/cpmi/publ/d157.pdf>
- CPMI-IOSCO (2012), "Principles for Financial Market Infrastructures". <http://www.bis.org/cpmi/publ/d101a.pdf>
- DTCC (2016), "Embracing disruption. A white paper to the industry", Depository Trust & Clearing Corporation (DTCC) White Paper, January 2016. <http://www.dtcc.com/news/2016/january/25/blockchain-white-paper>
- ECSDA (2015), "CSD Factbook", European Central Securities Depositories Association (ECSDA). http://ecsd.eu/wp-content/uploads/2015_CSD_Factbook.pdf
- ESMA (2017). "The distributed ledger technology applied to securities markets", European Securities Markets Authority (ESMA) Report, 7 January 2017. https://www.esma.europa.eu/sites/default/files/library/dlt_report_-_esma50-1121423017-285.pdf
- Euroclear (2016). Euroclear Annual Report 2016. <https://www.euroclear.com/dam/annualreport/2016/EuroclearPlcAnnualReport2016.pdf>
- Euroclear/Oliver Wyman (2016). "Blockchain in capital markets. The prize and the journey", February 2016. <http://www.oliverwyman.com/content/dam/oliver-wyman/global/en/2016/feb/BlockChain-In-Capital-Markets.pdf>

- Eyal, I. and Sirer, E.G., "Majority is not enough: Bitcoin mining is vulnerable".
http://fc14.ifca.ai/papers/fc14_submission_82.pdf
- FCA (2017). Discussion paper on distributed ledger technology. UK Financial Conduct Authority Discussion Paper 17/3. April 2017.
<https://www.fca.org.uk/publication/discussion/dp17-03.pdf>
- FINRA (2017). "Distributed Ledger technology: implications of blockchain for the securities industry." A report from the Financial Industry Regulatory Authority. January 2017.
- IOSCO (2017), Research report on financial technologies.
<https://www.iosco.org/library/pubdocs/pdf/IOSCOPD554.pdf>
- Khapko, M. and Zoican, M. (2017), " 'Smart' Settlement". Working paper (June 13, 2017). Available at SSRN: <https://ssrn.com/abstract=2881331> or <http://dx.doi.org/10.2139/ssrn.2881331>
- Mainelli, M. and Milne, A. (2016). "The Impact and potential of blockchain on the securities transaction lifecycle", SWIFT Institute Working Paper No. 2015-007, May 2016.
- Manning, M., Sutton M. and Zhou, J. (2016), "Distributed Ledger Technology in securities clearing and settlement: Some issues", *JASSA The Finsia Journal of Applied Finance*, Issue 3.
- Micheler, E. (2015). "Custody chains and asset values. Why crypto-securities are worth contemplating". *Cambridge Law Journal*, 74 (3), November, pp. 505-533.
- Mills, D., Wang, K., Malone, B., Ravi, A., Marquardt, J., Chen, C., Badev, A., Brezinski, T., Fahy, L., Liao, K., Kargenian, V., Ellithorpe, M., Ng, W., and Baird, M. (2016). "Distributed ledger technology in payments, clearing, and settlement," *Finance and Economics Discussion Series 2016-095*. Board of Governors of the Federal Reserve System, <https://doi.org/10.17016/FEDS.2016.095> .
- Morgan Stanley (2016). "Blockchain in banking. Disruptive threat or tool". *Morgan Stanley Global Insight*, April 20, 2016. <http://www.the-blockchain.com/docs/Morgan-Stanley-blockchain-report.pdf>
- Oliver Wyman/SWIFT (2014). "The Capital markets Industry. The time they are a-changin'".
- Pinna, A., and Ruttenberg W. (2016), "Distributed Ledger Technologies in securities post-trading. Revolution or evolution", European Central Bank (ECB) Occasional Paper Series no. 172, April.
- Reinganum, J., (1985), "Innovation and industry evolution", *The Quarterly Journal of Economics*, 100, issue 1, p. 81-99, <http://EconPapers.repec.org/RePEc:oup:qjecon:v:100:y:1985:i:1:p:81-99> .
- Santander/Oliver Wyman/Anthemis (2015), "The FinTech 2.0 Paper: rebooting financial services".
http://www.oliverwyman.com/content/dam/oliver-wyman/global/en/2015/jun/The_Fintech_2_0_Paper_Final_PV.pdf
- Schindler, J. (2016), "FinTech and financial innovation: Drivers and depth," *Finance and Economics Discussion Series (FEDS) working paper*, forthcoming.
- Shell, K. (1966), "Toward a theory of inventive activity and capital accumulation", *The American Economic Review*, Vol. 56, No. 1/2 (March 1966), pp. 62-68.
- Shell, K. (1967), "A model of inventive activity and capital accumulation". in *Essays on the Theory of Optimal Economic Growth* (K. Shell, ed.), Cambridge, Massachusetts: MIT Press, 1967, Chapter IV, 67-85.

- Shell, K. (1973). "Inventive activity, industrial organisation and economic growth". in *Models of Economic Growth* (J.A. Mirrlees and N. Stern, eds.), London: Macmillan, and New York: Halsted (John Wiley & Sons), 77-100.
- Swanson, T. (2015), "Consensus-as-a-service: a brief report on the emergence of permissioned, distributed ledger systems", April 6, 2015.
- SWIFT/Accenture (2016), "SWIFT on Distributed ledger technologies", Position paper. April 2016.
- Teece, D. J. (1996). "Firm organization, industrial structure, and technological innovation". *Journal of Economic behaviour & Organization*, vol 31 (1996), 193-224.
- World Economic Forum/Deloitte (2016), "The future of financial infrastructure. An ambitious look at how blockchain can reshape financial services". August 2016.
http://www3.weforum.org/docs/WEF_The_future_of_financial_infrastructure.pdf