Pinna, Andrea; Ruttenberg, Wiebe

Research Report

Distributed ledger technologies in securities post-trading - Revolution or evolution?

ECB Occasional Paper, No. 172

Provided in Cooperation with:
European Central Bank (ECB)

Suggested Citation: Pinna, Andrea; Ruttenberg, Wiebe (2016) : Distributed ledger technologies in securities post-trading - Revolution or evolution?, ECB Occasional Paper, No. 172, ISBN 978-92-899-2335-4, European Central Bank (ECB), Frankfurt a. M., https://doi.org/10.2866/270533

This Version is available at:
https://hdl.handle.net/10419/154625

EUROPEAN CENTRAL BANK

EUROSYSTEM

Occasional Paper Series

Andrea Pinna
Wiebe Ruttenberg

Distributed ledger technologies in securities post-trading

Revolution or evolution?

No 172 / April 2016

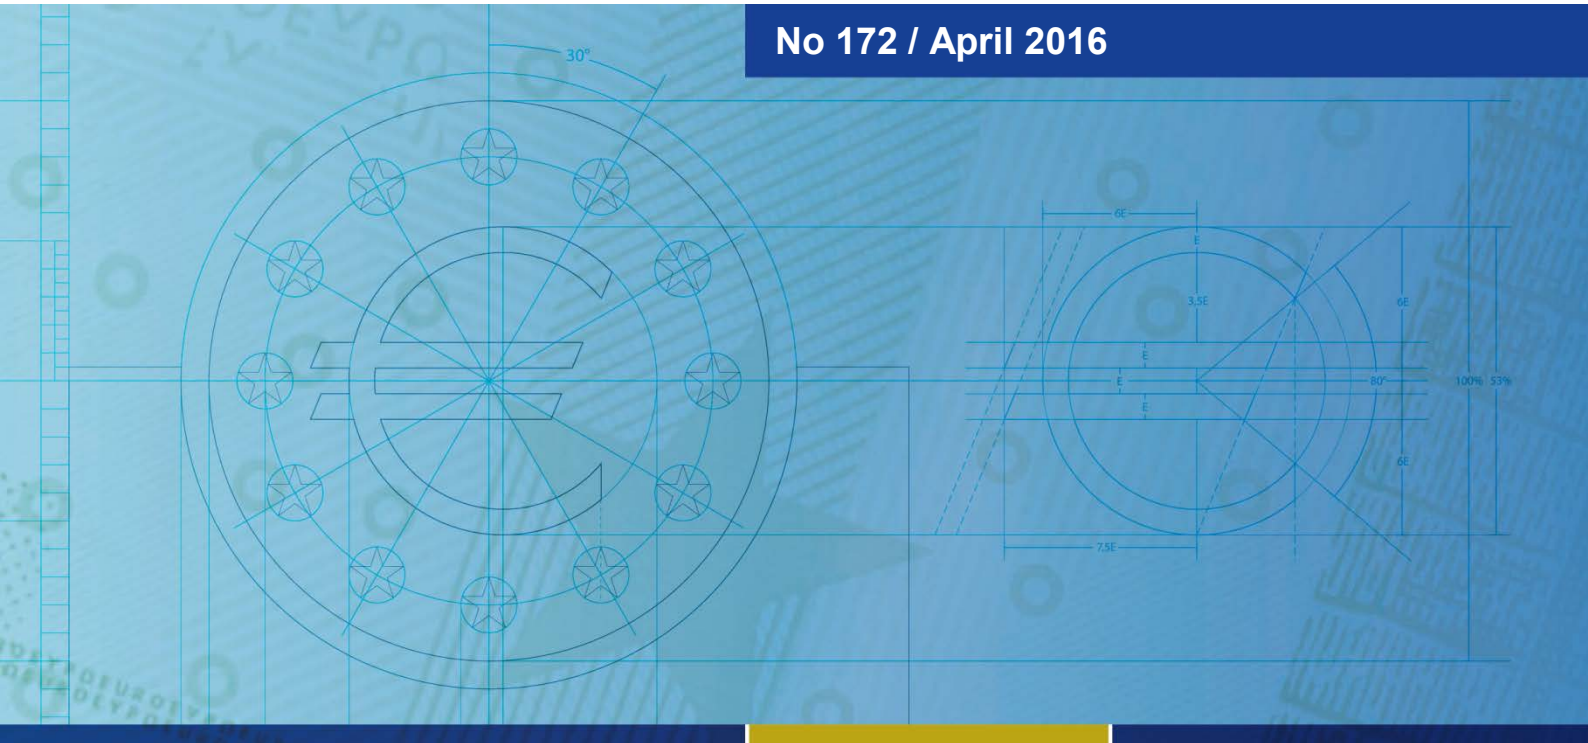

Note: This Occasional Paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB

Contents

Abstract

Over the last decade, information technology has contributed significantly to the evolution of financial markets, without, however, revolutionising the way in which financial institutions interact with one another. This may be about to change, as some market players are now predicting that new database technologies, such as blockchain and other distributed ledger technologies (DLTs), could be the source of an imminent revolution. This paper analyses the main features of DLTs that could influence their potential adoption by financial institutions and discusses how the use of these technologies could affect the European post-trade market for securities.

The original protocol underlying DLTs has its roots in the anarchic world of virtual currencies, which operate outside the conventional financial system. The public debate on DLTs has also been very much focused on the revolutionary potential of the technology. This paper concludes that, irrespective of the technology used and the market players involved, certain processes that feature in the post-trade market for securities will still need to be performed by institutions. DLTs could, however, stimulate a reorganisation of financial markets, which could in turn: (i) reduce reconciliation costs, (ii) streamline the post-trade value chain, and (iii) allow more efficient use to be made of collateral and regulatory capital. It should, nevertheless, be remembered that research into DLTs and their uses is at an early stage. The scope for financial institutions to adopt DLTs and their potential impact on mainstream financial markets are still unclear.

This paper discusses three potential models of how market players could adopt DLTs for performing core post-trade functions. The DLT could be adopted either: (i) in clusters, (ii) collectively, or (iii) peer to peer. The evaluation of the three adoption models assumes that they are all equally compatible with the regulatory framework. It shows that, assuming this to be the case, they would each have different advantages and costs.

JEL code: G21, G23, L15, O33

Keywords: Distributed ledger technologies, financial market infrastructures, fintech, settlement, clearing, blockchain, Bitcoin

Executive summary

Over the last decade, information technology has transformed the way people interact with one another. The same has not happened in financial markets, however, where intermediaries and the infrastructures they use to settle securities transactions often use proprietary databases, which cannot communicate with one another. Distributed ledger technologies (DLTs) are a type of technology that has emerged in the world of virtual currencies to allow users to share a common database. It is thought that DLTs might find their way into securities markets. The adoption of DLTs could, in theory, make post-trade processes more efficient, and could have a major impact on financial intermediaries that offer post-trading.

This paper analyses the main features of DLTs that could influence their adoption by financial institutions and discusses how use of these technologies could affect the European post-trade market for securities. The set-up of this market is currently still largely a legacy of the earlier domestic market infrastructures, which were responsible for developing business rules and technical standards before the creation of the Economic and Monetary Union, when integration across national markets was relatively unimportant. Little has changed in recent years in terms of how different levels of the market, e.g. custodians, cash correspondents, clearing members, collateral managers, settlement agents, central securities depositories (CSDs) and trade repositories, interact with one another. The lack of interoperability between proprietary databases is contributing to the ongoing use of siloed digital records of ownership, which restrict straight-through processing. Use of such records can also pose operational risks, and contributes to the continued existence of an uneven playing field. It prevents efficient use of collateral, and limits the potential for risk-sharing among European investors due to the higher cost of cross-border securities transactions.

By adopting a DLT, competing financial institutions would be able to share a common digital representation of asset holdings and to keep track of the execution, clearing and settlement of securities transactions outside their legacy proprietary databases, and without there needing to be any involvement of a central database management system. According to their proponents, DLTs have the potential to address many of the shortcomings identified in the post-trade market. These technologies are, however, still at an early stage of development and it is certainly too early to say what specific type of DLT will prevail, whether it will be widely adopted in the securities market, and whether its adoption will address the current market inefficiencies.

This paper considers some of the aspects of DLT models that might influence their potential adoption by financial market participants in the future and discusses the impact DLTs could have on post-trading. A distributed ledger can, for example, be open to any user, or access can be restricted to a set of authenticated participants. Its content can also be publicly available to anyone, or private and thus accessible only to a subset of network participants. Developers of DLTs are even investigating

the possibility that information could be made available only to users authorised by the account holder and to one or more supervisors/overseers. Restrictions on who can access and propose updates in the ledger, validate the updates, and read the information stored in the ledger, will be crucial in determining whether DLTs could be adopted in an institutional environment, such as stock and bond markets.

The original protocol underlying DLTs has its roots in the anarchic world of virtual currencies, which operate outside the conventional financial system. The public debate on DLTs has also been very much focused on the revolutionary potential of the technology. DLTs have relevance beyond the Bitcoin and its open blockchain model. Other types of DLT, such as restricted technologies and smart contracts, are better suited to the needs of financial institutions and could contribute to the development of safer, more reliable and more efficient post-trade processes.

The current debate on DLTs is very much focused on the technology itself, and on how disruptive the roll-out of this technology could be for the current post-trade market and its market players. It should, however, be kept in mind that, irrespective of the technology deployed, certain functions present in the post-trade market for securities will always need to be performed by institutions. This limits the potential disruption caused by introducing DLTs, in particular in view of the fact that the performance of some post-trade functions is heavily regulated. A clear example is the notary function. The clearing function will also continue to be necessary for trades in which derivatives are involved. If smart contracts were widely adopted, some financial intermediaries such as custodians could, meanwhile, see their role disappearing.

The adoption of DLTs could take place in any one of a number of different ways, depending on the strategic decisions taken by market players and on the role played by public institutions as legislator, regulator and catalyst. The financial industry is a network industry. At present, various legal and operational barriers are hindering the straight-through processing of transactions. The adoption of DLTs will only help to address the shortcomings identified in the post-trade market if there is suitable technical standardisation of the technologies, and if all market participants are subject to common business rules and sound governance arrangements. Nonetheless, even proprietary technologies might lead to more efficient post-trading, at least within clusters of financial institutions.

Neither market players nor DLT initiatives have given significant attention to the question of how the cash leg of a trade could be linked to the securities leg. Whereas it is impossible, at the current stage, to speculate on whether central bank money will ever be available on a distributed ledger, cash accounts need to be updated when securities transfers take place, in order to allow straight-through processes in the delivery-versus-payment (DVP) model currently in use. In terms of financial stability and oversight, settlement in commercial bank money has its limitations, due to the default risk of commercial banks. Clarification is also needed with regard to the legal status of the ledger, the legal enforceability of smart contracts, and the legal validity of the method used for authenticating DLT users.

In conclusion, it is true that distributed ledger technology has potential, and at the same time, innovation is welcome in the European post-trade market, wherever it can bring safety and efficiency. A number of factors could, however, pose potential barriers to the widespread uptake and use of DLTs. First, the technology is not yet mature; second, the clarification of critical legal, operational and governance issues will take time; and third, even were DLTs to be adopted widely, certain post-trade functions will continue to be necessary. It is not yet, therefore, clear whether DLTs will cause a major change in mainstream financial markets or whether their use will remain limited to particular niches. It is possible that a DLT may find its way into the mainstream market, but should this happen, it is more likely to cause a gradual change in processes, rather than a revolution in the market.

1 Introduction

Distributed ledger technologies (DLTs) allow users to modify records in a shared database, i.e. the ledger, without necessarily needing to use a central validation system that imposes its own standards and processes. A database of this kind could potentially be used in financial markets to create a distributed ledger that settles trades (e.g. securities transactions) for any given set of assets and their holders. This paper describes the main features of DLTs and discusses how use of these technologies could affect the landscape of the European post-trade market for securities.

The set-up of the European post-trade market is currently still largely a legacy of the earlier domestic market infrastructures, which were responsible for developing business rules and technical standards before the creation of the Economic and Monetary Union, when integration across national financial markets was relatively unimportant. Since its establishment, the Eurosystem has pursued a pan-European approach, with the aim of creating a single market for financial services. It has introduced a number of measures in cooperation with European public institutions and market players, including, e.g. TARGET2 (the second generation of the Trans-European Automated Real-time Gross settlement Express Transfer system), TARGET2-Securities, SEPA (the Single Euro Payments Area), and other joint efforts with European public institutions and market players.

Focussing on the post-trade of securities transactions in particular, the harmonisation measures related to TARGET2-Securities have achieved unprecedented results at the level of central securities depositories (CSDs), and are thus making a significant contribution to the creation of a single market.[1] Little has changed, however, in terms of how different levels of the market, e.g. custodians, cash correspondents, clearing members, collateral managers, settlement agents, CSDs and trade repositories, interact with one another. The lack of interoperability between proprietary databases is contributing to the ongoing use of siloed digital records of ownership, which require manual updating to reconcile them with any change that occurred in the records of counterparties at different levels of the post-trade value chain. These extra back-office procedures restrict straight-through processing, pose operational risks and contribute to the continued existence of an uneven playing field. They also prevent the efficient use of collateral, and limit the potential for risk-sharing among European investors, due to the higher cost of cross-border securities transactions.

This paper analyses how distributed ledger technologies (DLTs), first developed for virtual currencies, could be employed in securities transactions, and thus potentially transform the aforementioned post-trade landscape. The adoption of DLTs would

1 TARGET2-Securities (T2S) introduced a common platform where connected CSDs can hold their accounts. The balances reflected in their ledgers are automatically updated and realigned but the final ledger remains legally within the remit of each CSD's regulatory and legal framework.

mean that competing financial institutions would be able to share a common digital representation of asset holdings and keep track of the execution, clearing and settlement of trades outside their legacy proprietary databases, and without there needing to be any involvement of a central database management system – providing that safety, efficiency, and regulatory compliance are ensured. Whether such a potentially revolutionary change would be viable is not yet clear, but further insight may be provided by a number of ongoing DLT projects that are being developed by fintech start-ups and incumbent financial institutions.

The distributed databases currently being used by the majority of organisations allow authorised users to read records from copies of a master database that is duplicated in multiple locations. Any transaction, i.e. any update to the database, is, however, submitted for validation to a central database management system. By adopting a DLT instead, users would become peers in a shared database, which they could rely on to record transfers of assets and to perform additional related activities involving multiple parties, e.g. trading, clearing and asset servicing. DLT users can propose new transactions and, depending on the approach to operating the DLT chosen, they can either contribute to validation collectively or have a subset of users responsible for this task. In either case, a transaction is validated when a specified proportion of the network's validators have reached a consensus as to its legitimacy. Changes to the shared database are then reflected instantaneously in its digitally signed versions, which users can store locally (either in their entirety or with only a subset of transactions/accounts visible). Users can then extract the updated information they need for conducting their respective businesses from these locally stored databases.

Beyond the recent hype around some of the potential uses of DLTs[2] in a range of areas, some post-trade institutions are assessing the possibility of adopting the new technology for their purposes. That could have far-reaching consequences for the efficiency and safety of European markets for securities, collateral, derivatives and cash. Depending on their governance and on the approach to adopting the technology chosen by market players, DLTs could either solve various issues relating to interoperability, or could induce a new layer of fragmentation between different clusters of institutions, each working with different distributed ledgers.

The purpose of this paper is: (a) to give an overview of some of the DLTs in existence at the time of writing, and (b) to assess their potential impact on the post-trade industry in a set of abstract future scenarios, which range from the most restricted to the widest potential future adoption of DLTs in securities trading.[3]

The paper is structured as follows: Section 2 presents the specifications of some DLTs that are being developed; Section 3 presents the current post-trade industry; Section 4 discusses the potential impact of distributed ledgers on the market; and Section 5 draws conclusions on the potential impact of DLTs on the European post-trade industry.

---

2 Almost $500m of venture capital was invested in DLT-related businesses during the first 11 months of 2015, the majority of which was Bitcoin-related (Source: coindesk.com http://www.coindesk.com/bitcoin-venture-capital).

3 DLT applications outside the securities market (e.g. payments, syndicated loans and foreign exchange transactions) are beyond the scope of this paper.

2 Distributed ledger technologies (DLTs) and their specifications

2.1 What are DLTs and why their specifications matter

DLTs allow their users to store and access information relating to a given set of assets and their holders in a shared database of either transactions or account balances. This information is distributed among users, who could then use it to settle their transfers of, e.g. securities and cash, without needing to rely on a trusted central validation system.

In financial markets, the substantial dematerialisation of securities and cash has progressively shifted the settlement of a trade from the physical delivery and paper-based recording, to a system of book transfers in digital databases.[4] What remains unchanged is the need for an authoritative "golden record" of holdings to be kept by specific financial market infrastructures, and for intermediaries involved in the settlement process to update their individual databases by communicating with the other institutions involved, at the different levels of post-trading, in order to be able to reflect the changes in each other's records. The high cost of this type of reconciliation process has led many market players to consider distributed ledgers as an alternative to central validation systems – currently one per institution (internal records of outstanding positions) or per cluster of institutions (e.g. interoperable market infrastructure) – to keep their reciprocal records updated.

Sharing a database with no central validation system can create difficulties when different users have conflicting incentives. Due to the latency of communication via a network, a malicious user may lead a counterparty to believe that some cash/securities have been credited to its account, whilst other users believe these same assets to be recorded in another account – which might be that of the malicious user or of a third user.[5] A similarly undesirable outcome might also arise when a dysfunction in the entry of information by a bona fide user in the network results in inconsistent information being recorded in the copies of the ledger held by the various users. Before all distributed copies of the ledger have been reconciled, the beneficiary of a transaction may perform what erroneously appears to be a delivery-versus-payment. There needs to be a way of avoiding or addressing such errors, and thus achieving consistency across different copies of the ledger.

DLTs allow their users to reach consensus on a particular version of the distributed ledger, in particular on the sequential order of transactions. This means that there

4 Some jurisdictions still require there to be a physical certificates These certificates are stored in depositaries that are linked to a settlement system where the transfer of ownership takes place electronically.

5 This issue of "double spending" cannot happen where physical cash is used or where a central authority is responsible for authorising transactions. In the first case, once the bearer banknote has been handed to a beneficiary it cannot be replicated and used in a second transaction. In the second case, a central authority would reject the second of the two transactions.

cannot be any doubt as to the users' respective holdings. Central validation is replaced in a DLT by a set of cryptographic solutions and economic incentives that combine to prevent illicit updates and reconcile discrepancies. The ledger produced can thus be considered authoritative, although its management is shared among users with conflicting incentives.

The specifications of a DLT determine its ability to carry out the tasks typically performed by the current financial market infrastructures and intermediaries, and which are necessary to ensure public confidence in financial markets. This paper focuses on the following post-trading functions: keeping accounts at the top-tier level (depositary/registrar service) and at lower level (custody service); checking the entitlement of an investor to the control of an asset (know-your-customer and anti-money-laundering obligations); transmitting and reconciling transfer orders prior to settlement (clearing function); facilitating settlement and hedging settlement risk until the transfer of funds and securities are final and enforceable (netting and risk management); discharging participants' obligations through the transfer of funds and securities (settlement function); ensuring the integrity of an issuance and avoiding the unwarranted creation of securities (notary function); avoiding theft of private information, malicious updates and the denial of service (cybersecurity); and managing events initiated by an issuer of securities and their impact on end-investors (asset servicing).

The ability of a DLT to carry out the various tasks listed above depends on its technical and operational specifications, irrespective of the entity performing the task and provided that it acts in compliance with the regulatory framework.[6] Only certain DLTs would be able to carry out these tasks, and thus to make the institutions currently responsible for them redundant. This depends on their technical design and the economic incentives they provide to validators. This chapter discusses the different DLT specifications, with reference to three criteria:

- participation in the ledger (restricted versus unrestricted) – see Section 2.2;
- validation method (allocation of votes to users) – see Section 2.3; and
- data structure of the shared database (unspent transaction outputs versus consensus ledgers) – see Section 2.4.

Although there is currently no official taxonomy for DLTs, this paper uses the term blockchain to refer to a database structure that can only be updated by appending a new set (or block) of valid transactions to the log of previous transactions. The DLT protocol is designed such that consensus is reached on transactions involving "unspent transaction outputs", i.e. the set of assets available to the initiator of a transaction. Other DLTs are referred to as consensus ledgers, as they do not keep track of the history of transactions but instead operate on the basis of consensus reached on a ledger of accounts, which are updated with new transactions at each validation round. Smart contracts, i.e. self-executable updates that are stored in the distributed ledger and triggered by particular events happening either within or

6 Discussion of the compliance of DLTs with current regulation is outside the scope of this paper.

outside the ledger, are discussed as a standalone technology, as they can be coded into different types of distributed ledgers.

Some of the characteristic features of DLTs can be found in earlier database technologies that have been developed since the 1990s, e.g. in the field of master-master replication. These technologies allow a number of parties to update records in a common database, with conflicts being resolved by some form of consensus algorithm. It is possible that there will be a renewal of interest in technologies of this type, as a result of the current focus on shared databases (and distributed ledgers), and that they will, as a result, be updated and come to represent a competing alternative to DLTs. They could allow otherwise traditional databases designed to be shared among financial institutions to be updated non-centrally. Although the focus of this paper is on DLTs, the analysis of their impact on the post-trade market also applies to any technology which would allow financial intermediaries to share a common database that records their respective positions.

2.2 Unrestricted versus restricted DLTs

In terms of users' participation, DLTs can be divided into those which are restricted and those which are unrestricted. Restricted DLTs are closed systems whose members are identified and accountable entities. Ledger updates can only be proposed and validated by authorised participants. In unrestricted DLTs, by contrast, any entity can access the database and, depending on the specific validation method used, may be able to contribute to updating the ledger or to submit spam transactions to cause a denial of service.[7]

Consensus may seem to be a redundant feature in a network where only authorised entities are able to make updates, as any changes made to the ledger illegitimately could be punished as a breach of contractual obligations. It could, nonetheless, be useful as a way of increasing the resilience of a clearing or settlement system, in order to protect it against the possibility of malevolent users or external attackers trying to disrupt its functioning, i.e. as a way of improving its cyber resilience. A DLT could do this by ensuring that the distributed ledger is correctly updated even if a (limited) number of validation nodes are off-line, experience a bug, or fall under the control of an unauthorised entity.

The idea of a blockchain was introduced in 2008 as a basis for the virtual currency Bitcoin, which is an example of unrestricted DLT. The distributed ledger is "authoritative" because every user agrees on it, even though in Bitcoin and some other DLT initiatives there are no central, regulated/authorised institutions playing any role in the process. The technology behind Bitcoin represents a remarkable achievement, in that it combines the idea of sequentially linking records in an

7 Distributed ledgers can be public or private, depending on who can read the data. An unrestricted DLT has, by definition, a public ledger, whereas a restricted DLT may have either a public or a private ledger.

immutable "hash-chain"[8] with anti-spam algorithms and novel economic incentives, in a way that allows users (who are authenticated by means of pseudonyms) to transfer tokens between themselves while avoiding central authorities and censorship tools.[9]

The financial industry has developed over time as a network of mutually trusting institutions, with legal agreements and regulated procedures in place in order to avoid risks, such as operational and counterparty risk, that are not directly related to the business of a securities issuer. Each institution trades with accountable and authorised counterparties, under the supervision and oversight of regulators.[10] This creates scope for the implementation of restricted DLTs among market players. The characteristics of blockchain technology that were so important for the purposes of the Bitcoin network – pseudonymity of market participants, immunity from supervisors, copies of the ledger being accessible to anybody all over the world, and irreversibility of unlawful transactions – are not relevant to the financial industry. Market players instead need a system that is compatible with the standards they are required to meet – implementation of know-your-customer (KYC) rules, transparency and accountability vis-à-vis regulators, respect of the rule of law and confidentiality of trading strategies – and that is relatively cheap to maintain.

The rules on participation in a DLT have an effect on its maintenance costs, as the range of tools that can be used to engender truthful behaviour by validators is more limited if the technology is unrestricted. As discussed in the following subsection, DLTs may make use of game-theoretic tools, in addition to public-key cryptography,[11] in order to address cyber resilience and ensure that malicious updates are either bound to be unsuccessful or would be economically disadvantageous from the point of view of any user or external attacker.

In a restricted distributed ledger the identity of participants is known, at least by its governance body. This implies that any wrongdoer can be identified and his misbehaviour can be punished in the case of future activity in the ledger. Restricted distributed ledgers also expose the conduct of any participants in the DLT network to the set of rules and law-enforcement measures that typically apply to off-ledger activity. By definition, users of unrestricted DLTs cannot be held accountable outside the distributed ledger for their activity in the network. Off-ledger incentives such as fines or loss of reputation cannot therefore be used to induce participants to comply with the rules of the ledger if their identity is unknown. Due to the lack of off-ledger incentives, unrestricted DLTs need participants to deploy their own resources within the ledger to deter the validation of illicit transactions. This can either be done by

8 The hash value is a string of data that is univocally linked to a specific version of the distributed ledger or to parts of it. See Haber and Stornetta (1991). For an explanation of the motivation for storing updates in "blocks", see Benaloh and de Mare (1991) and Bayer, Haber and Stornetta (1993).

9 See Nakamoto (2008).

10 Authentication processes are required by regulation and would therefore need to be a feature of all DLTs. How they are designed will be important, particularly if a DLT is to be applied at global level.

11 Public-key cryptography allows the holder of a specific pair of private and public keys to: a) sign a message with a private key to allow any network participant to check – by using the corresponding public key – that he was the author, and b) receive a message signed by network participants with his public key and be the only person able to read that message.

means of an ex-ante investment in computational power, or ex-post in the form of assets that are native to the ledger and are posted as collateral. This type of security suffers due to its economic model. It is limited by the fact that a one-off investment is sufficient to allow any attacker to tamper with the distributed ledger, and that any form of coordination among unidentified participants is difficult to achieve on occasions where it is necessary to recover the truthful version of a ledger. The governance of a DLT is therefore another characteristic which may make its adoption more or less likely.

2.3 Validation methods

Once a new set of transactions submitted by DLT users has been validated by a specific number of validators in the network, other users receive the update and can change their local copy of the ledger accordingly. Validators need to check that the assets that would be moved by a transaction are available to the transaction originator according to their most recent information. The originator must therefore provide the hash value of the latest version of the distributed ledger, in addition to the standard information on how many assets are being moved and which users and accounts are involved. Validation rules attempt to find a satisfactory compromise in the trade-off between cyber resilience and efficiency. They rely on both economic incentives and cryptographic security features, which are specific to the individual DLT.

In an unrestricted network, it is impossible to share validation rights on a “one-head one-vote” basis. Any user can create multiple network addresses (a strategy known as a “Sybil attack”), either in order to cast a number of votes large enough to allow unilateral validation of updates, or to flood the network with an unmanageable number of requests, and thus cause there to be a denial of service for licit transactions. DLTs use a number of different consensus processes to address these issues. We discuss two of these – proof of work (PoW) and proof of stake (PoS) – in detail below.

A proof of work (PoW) system is based on mathematical problems involving non-invertible hash functions, i.e. mathematical functions that are difficult to solve but whose solution, once provided, can be easily checked. Finding the solution to these problems is a matter of luck and “work”, i.e. computational effort, rather than mathematical ability, as they can only be solved by iteration. In a proof of work system, each validator selects a set of pending transactions that can be shown to be licit as they abide by standard bookkeeping rules, e.g. only accounts that are available to the initiator are updated, and none of the transactions will result in a deficit. The validator adds the set of transactions to the existing trustworthy ledger (in Bitcoin, “the consensus blockchain”) and uses the newly proposed version of the ledger as an input for the hash function. Once a solution to the mathematical problem has been found, it is sent to the rest of the network together with the new proposed ledger. All other network participants may then check that: a) the update contains only licit transactions, and b) the validator worked to find a correct solution to the mathematical problem. If both conditions are satisfied, the validators agree on

the new version of the ledger and use it as a starting point to validate the other pending transactions.[12]

Validators bear the cost of the hardware and the electricity required to solve hash functions. They take on this task on a competitive basis, as a reward is credited to the validator who finds the solution first. This takes the form of tokens native to the ledger (e.g. newly minted bitcoins in the case of Bitcoin DLT) and possibly fees paid by the initiators of the transactions they helped to validate. The mathematical problems set in PoW systems have the sole purpose of making validation expensive and assigning votes on the basis of computational power. The authoritative distributed ledger is set to be the one, from amongst those published by network users, whose history of validated updates is longest, hence the one whose hash function required the greatest computational capacity to solve. The cost of validation is adjusted regularly to deter any participant from attempting to validate new transactions at a pace such that its malicious version of the ledger remains the longest in terms of the number of validations. In expected terms, a malicious user can only falsify the distributed ledger if it can validate new transactions faster than the rest of the network, i.e. only by holding more than half of the total computational power deployed by all validators.[13]

Bitcoin transactions, which are validated through PoW, are relatively safe, under the non-negligible assumption that attacks to the ledger would only be motivated by economic considerations. When the value of the cryptocurrency increases and the expected revenue from misappropriation thus also grows, the cost of reaching any given percentage of the total computational power in the network increases by the same factor. It is true that potential attackers are induced to invest a higher amount of resources to control the majority of computational capacity in the network when the value of the bitcoins they can obtain grows. Also the optimal investment undertaken by validators to increase their computational capacity and receive bitcoins as a reward for correct validation, is, however, also proportional to the value of each bitcoin.[14] Thus, when bitcoins appreciate in value and the expected revenue from falsifying the blockchain increases, the cost of acquiring 50% of the total computational capacity held by the network also grows, assuming that bona fide investors continue to increase their investment in validation capacity proportionally. Providing the distributed ledger only records transactions in the native asset or token, the cost of controlling the network will increase in proportion to the increase in the expected revenue that would be earned from gaining control. Any attempt to falsify the blockchain is therefore unprofitable from a purely economic point of view.

It may happen that some validators solve the PoW for a set of transactions they wish to add to the ledger after another validator has successfully completed a separate

---

12 When multiple validators validate two updates at the same moment, two versions of the distributed ledger are created, whereby different network participants, momentarily, trust two different sets of information. Any disagreement is then reconciled by the fact that a new update is validated and added to either of the two versions by a validator, depending on which particular version that validator happened to have worked on.

13 It is actually sufficient for a malicious user to hold 25% of computational power to be able to attack the network, by changing the incentives of other validators. See Eyal and Sirer (2014).

14 Validators of blocks that are repeatedly considered invalid by the rest of the network lose their rewards.

modification to the previous state of the ledger (which the first validator was also taking as the basis for their additions). The transactions of this first validator cannot then be added to the ledger, as the ledger has now changed. They are therefore returned to the set of pending transactions, and will be validated by creating a new mathematical problem based on the new updated ledger. Their original validator, however, loses out on its reward. Modifications to the standard PoW are being developed in order to allow computational power used in the unsuccessful validation to be used to increase the amount of computational power necessary for an attack, rather than being wasted.[15]

A second type of validation system is a proof of stake (PoS) consensus process. This assigns shares of validation rights to users according to their stake in the system. How a validator's stake is to be measured is thus a critical aspect of a system of this type, and different DLTs take different approaches. Some possible criteria used to measure a validator's stake are the amount of native tokens owned, the amount of particular native tokens or off-ledger assets escrowed in the ledger as collateral, or the reputation of the validator in a restricted DLT (known as "proof of identity").

The varying specifications that DLTs have in terms of participation in the ledger can indirectly affect the efficiency of the network. Instead of requiring every participant to invest a vast amount of resources in the maintenance of the ledger, responsibility for maintenance can be left to participants who will act in good faith, knowing that illicit behaviour would be punished and the truthful state of the ledger re-established by agreement between accountable participants. Developers of restricted blockchain technologies can thus choose to use less expensive consensus algorithms than those that are necessary in unrestricted DLTs. In this way, in restricted DLTs, validation is not made artificially difficult or costly for all users, on an ongoing basis, but is instead made costly for attackers, and only when there is an attack.

There are also sidechains that can provide a layer of permission on top of that provided by the Bitcoin blockchain, as a way of shifting at least part of the cost to validators from the expensive computational investment needed for PoW, to reputational cost or collateral. Other DLTs provide a standalone blockchain with its own protocol, and allow sub-ledgers to be created, for which users can decide the consensus mechanism to be applied for validation. Sub-ledgers with different validation methods cannot interact directly, but inter-ledger functionalities are under development that could potentially make interoperability possible in the future, at least from a technical point of view.

A PoS system with collateral allows one or more parties to keep assets of the participants in escrow (possibly outside the ledger, with a trusted party). If a participant is proven to have attempted (whether successfully or otherwise) to validate an illicit transaction, they will then forfeit their collateral. In "punitive" PoS systems of this type, when a lack of consensus emerges in the network, the digital

[15] See Sompolinsky and Zohar (2013). Zamfir also describes simplified implementation on the Ethereum blog.

signatures of the guilty users are blacklisted and their collateral passes into the ownership of the system.

In a closed network such as that formed by financial market infrastructures and intermediaries, the two-thirds threshold of bona fide validators required in many consensus systems to ensure correct validation of new transactions may be considered relatively safe. A party planning to launch an attack on the network would need to control the credentials of over a third of the financial institutions acting as validators in order to be able to interfere with the correct recording of new transactions in the distributed ledger. In terms of cyber resilience, it is not clear whether this type of security is stronger or weaker than that provided by centralised databases. It depends on the relative probability of a cyberattack succeeding in: i) taking control of one central dedicated system, and ii) controlling over a third of validators, who access the network of DLT participants from computers (validation nodes) that may have different security standards than a central server.[16] Some consensus protocols automatically downgrade nodes who disagree with the majority of the network, in order to lower the risk of nodes being progressively taken over by an attacker aiming to control the system.

Once new transactions have been validated and added to the most recent version of the ledger, they may or may not be considered as settled (i.e. finally and irrevocably agreed upon by all participants) or as pending, depending on the specification of the DLT. In some cases, pending ledger updates may be reversed during a specified time period after validation has taken place. Some developers of unrestricted DLTs have announced that they will switch to restricted technologies in order to meet the needs of financial institutions, which are particularly interested in the finality of settlements and in the cost of maintaining a ledger, rather than in its openness to unknown parties.

2.4 Blockchain as a database structure

The impact of DLTs on latency, throughput and the use of resources varies. The blockchain technology allows a number of participants in a restricted or unrestricted peer-to-peer network to validate new transactions or blocks of new transactions and append them to the chain of previously validated transactions or blocks of transactions. The chain of validated blocks then constitutes a blockchain, which is updated and distributed to all participants, in order to ensure consistency of information. Any party that joins the network receives either the entire latest version of a blockchain file or its hash.

In blockchain DLTs, such as the well-known Bitcoin blockchain, information on the ownership of an asset is structured as a chain of transactions. In particular, each user owns a series of unspent transaction outputs (UTXO) which represent a number of assets (e.g. bitcoins). Validators agree to assign different addresses in the

[16] The answer to this question requires further analysis and is beyond the scope of this paper.

network (i.e. owners) to each of these outputs.[17] An UTXO represents assets (e.g. bitcoins) that can be sent by their owner to any set of network participants. An UTXO may be seen as the digital representation of a banknote. It cannot be split into parts but is spent as a whole, and any change the payer expects to receive is actually another banknote (i.e. UTXO) of lower value. Once the payee receives the UTXO, that UTXO is spent and cannot be used anymore. It generates two new UTXOs, each of which represent a lower amount of assets (e.g. bitcoins) and are available to the payee and the payer, respectively. UTXOs are used in many ongoing DLT projects whose data structure is based on blocks of individual transactions forming the blockchain, rather than on users' accounts balances. Each blockchain user has portions of a large number of UTXOs linked to their address. Whenever a user wants to send some of their tokens, they have to specify which UTXO they are using as an input for the transaction, which part of the UTXO is being sent to the beneficiary, and which part is being sent back to their own address.[18]

In DLTs that use tokens such as bitcoins, every token is characterised by a unique address. This means that the token could potentially be associated with a particular asset, and be transferred as a bearer of that asset by updating the blockchain. If a trusted party declares that a particular token address bears an asset that can be freely transferred between counterparties, the transfer of that particular token from one investor to another would then also transfer the property rights attached to it.[19] Tokens that have property rights attached are said to be "coloured" tokens, to distinguish them from other, apparently identical, tokens that are used solely as money to make electronic payments, i.e. as a cryptocurrency. The holder of a coloured token – in practical terms, the owner of a private cryptographic key that gives access to a public key that has the address of the coloured token – would then be entitled to use the corresponding asset and may claim any payments due in relation to its ownership until it is transferred to another network participant.

Some blockchain technologies build upon the computational capacity of Bitcoin in order to provide additional functions, either through new platforms, known as "sidechains", or by adding an additional asset known as a "metacoin" to the network.[20]

Token-based DLTs, such as Bitcoin, that rely on PoW are only effective in a cryptocurrency setup, i.e. where the amounts earned by a potential attacker and by validators are both multiples of the value of a bitcoin. The same validation method does not protect against wrongdoing by users in systems where the settlement of transactions involves any assets other than bitcoin tokens, e.g. securities. The value of specific bitcoins that bear ownership rights to a security (control rights and the

17 See the Bitcoin developer guide for additional details.

18 The process described here, of how a blockchain transaction works, is invisible to the user, as all steps in the handling of UTXOs are performed automatically by client applications.

19 The transfer is cost-free from the perspective of the end-user, notwithstanding the fact that the payment of a fee to unknown validators may be necessary to ensure prompt validation of a transaction. The transaction is, however, by no means free in terms of the cost it creates for society. In the case of Bitcoin, transactions of this type generate high costs in computational power (investment in hardware and high energy consumption).

20 See Swanson (2015).

right to future cash flows) would be greater than that of standard bitcoins paid to validators who invest resources in ledger maintenance. Validators acting in good faith would only receive the bearer asset, whose value is that of a single bitcoin. In such a situation, there may thus be an economic incentive for potential attackers to invest in computational capacity and electricity, in order to counterfeit the transfer of bitcoins giving control rights to securities and rights to the future dividends associated with them.

2.5 Consensus ledgers

Users of blockchains can retrieve account balances by aggregating the complete history of past transactions and verifying the UTXOs associated with their address in the network.[21] The approach taken in consensus ledgers is more similar to that used in standard bookkeeping, and is based on individual accounts. Instead of monitoring successive transfers of tokens, consensus ledgers are updated each time there is a round of validations, and the ledgers then update the balance of users' accounts.[22] The updating of consensus ledgers is crucial to being able to restore the system after network outages. It means that nodes that had been off-line or that have experienced a loss of data do not need to download the whole history from other users to recover their own records. Their accounts will be in line with the rest of the network again as soon as the next updated state of the ledger is agreed upon by the rest of the network and sent to them. Although this may not be of great importance in a DLT with a large number of validating nodes, it is an advantage when the network is restricted and only a small number of nodes are responsible for validating updates.

A new version of the ledger is agreed upon every few seconds by a list of validators, who can be elected by users to validate a set of pending transactions. Different nodes may initially select different sets of transactions to be validated. The first step is then to reach a consensus on the list of transactions to be processed in a given round. This avoids the situation where some nodes could be wasting resources trying to validate a set of transactions which then ultimately go back to the pool of pending transactions because the number of nodes working on them is not sufficient to reach the minimum threshold for validation. After reaching consensus on the list of transactions to be processed, the nodes then exchange views on them in order to: a) collectively discard any transactions that do not have sufficient support among validators; and b) validate transactions that have the required pre-determined majority of votes in favour.

21 Computers known as "light clients" are connected to a blockchain network and store part of the history of the blockchain. To check past activity and be able to validate transactions with confidence, users, however, need to take into account the full history of the blockchain.

22 Previous versions of the consensus ledger are kept by only a small number of nodes, for the purpose of ensuring auditability. This is similar to current practice for many traditional databases.

2.6 Smart contracts

Smart contracts are a way of transposing the contractual obligations imposed on users into the digital distributed ledger. The purpose is to ensure that all provisions are complied with by means of automatic updates to the users' accounts. A smart contract can have access to a number of accounts and can transfer assets according to the terms of the contract as soon as an event (either within or outside the chain) triggers the application of these terms.

Smart contracts can thus provide for automatic transactions – such as crediting a dividend or coupon payment, issuing and reacting to margin calls, or optimising the use of collateral – to take place in the ledger in response to a specific corporate action or market event.[23] Since smart contracts are written in the ledger, validation of their execution follows the same procedure as any other transaction. If a DLT is able to ensure that ledger updates are cyber resilient, the execution of its smart contracts is therefore also similarly protected from attack. This is the main difference between smart contracts and similar procedures stored in non-DLT databases.

Once one party has created a smart contract by adding some lines of code to the ledger, other participants can accept the contract and make it executable. The parties' agreement to the terms of contract is then validated, and as of this point, the contract can no longer be retracted. The consequences of a smart contract cannot be ignored, as the code has a direct and immediate effect on securities and cash accounts in the ledger when an event triggers its execution.[24]

Were DLTs to be adopted in financial markets, smart contracts may prove to be the element that causes real change. As explained in the following chapter, the limitations of unrestricted ledgers and the need to ensure compliance with regulation mean that intermediaries are likely to continue to play a role, at least at some levels of the post-trade value chain. Smart contracts could, however, perform a number of the duties that are currently typically carried out by incumbent post-trade institutions.

Anecdotal evidence suggests that some technical issues relating to the consistency of updates to the ledger have not yet been fully resolved. When a smart contract is executed, it may modify any account it has access to. Moreover, the execution of a smart contract may require complex computations. If smart contracts are to be used, it may therefore be necessary to pay an opportunity cost, determined by a market pricing mechanism. This approach is already being used in some DLT solutions to avoid the performance of other validations, e.g. transaction settlements, being worsened, and to prevent "denial of service" attacks where the network is made unavailable by an unmanageable flow of requests.

---

23 Some developers envisage the possibility of creating smart contracts by means of application programme interfaces (APIs). This would allow any participant to code a smart contract in the ledger by filling out a form with pre-determined fields.

24 This notwithstanding the fact that it is not clear how contracts written in programming language might be read and enforced by the courts.

3 The current post-trade landscape

For a new transaction to be validated in a DLT, users are required to exchange messages in order to reach a consensus, and the distributed ledger is then updated. Having a process such as this, where consensus needs to be reached among a number of parties, limits the speed of settlement compared to that achieved by centrally managed database technology.[25] The above-mentioned lack of interoperability makes post-trading processes slow, however, even in systems using current centralised database technologies. The possibility of adopting DLTs for use in securities transactions should not be assessed only on the throughput of the system, but also on the basis of the potential impact on current business practices.

The lack of interoperability between centralised database systems restricts straight-through processing for a range of non-vertically integrated financial institutions. In addition to widening the settlement cycle and increasing the cost of back-office procedures, the need to reconcile accounts kept by different intermediaries creates certain risks, such as chains of settlement failures (as delayed settlement of one transaction may affect the settlement of trades with third parties), human errors (the system sometimes being reconciled manually), and limited collateral fluidity.

At present, financial intermediaries keep multiple separated records of the same information. There is a misperception that DLTs could save back office costs by avoiding the duplication of data. In fact, redundancy is maximised in DLTs and may even be useful for reacting to cyberattacks, as different nodes keep a copy of the ledger or part of it. It should be emphasised once again that the driver of high back-office costs in securities markets is not the cost of storing redundant data but rather the redundancy of business processes.[26]

Current market conditions – with low interest rates and the expectation of an increase in collateral demand due to regulatory changes – have squeezed financial intermediaries' margins, making the fixed costs of non-profit making back-office procedures more of a burden. Financial intermediaries are therefore hoping that DLTs may allow them to avoid certain reconciliation processes and reduce the amount of collateral and capital bound to the settlement cycle.

Financial intermediaries have to update their own accounts every time a new transaction takes place. They are then required to send any relevant results of this exercise to the interested parties, at different levels of the post-trade industry, in order that they can reconcile their own accounts to reflect the new situation and inform their interested parties of any changes. The securities flow between trading

25 The same does not apply to reading records in the distributed ledger. Read speed in DLTs is comparable to that in centrally managed databases, where a master copy is distributed to users after each update.

26 A study conducted by the consultancy firm Oliver Wyman for SWIFT estimated there to be global costs of $5-10bn for clearing activities, $40-45bn for settlement, custody, and collateral management (of which $39bn is paid to market players in the custody chain), and $20-25bn for post-trade data and analytics.

and settlement takes time, although execution of the matched settlement instructions at the settlement level can be instantaneous.

Diagram 1

Post-trade processes in the securities leg of current transactions

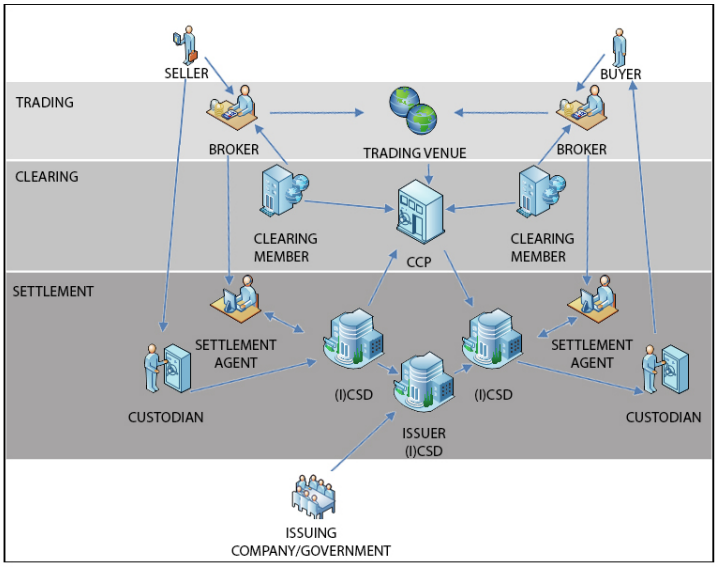

Note: (I)CSD = (international) central securities depository, CCP = central counterparty.

Diagram 1 presents a stylised picture of the security leg of a securities transaction between end-investors.[27] The buyer and seller have to instruct their respective brokers as to their willingness to trade. The orders are routed to a trading venue where they can "cross" in the order book or on an alternative trading system and make a trade. The details of the trade are often sent to a clearing house that reconciles orders, possibly netting them with other pending instructions in order to lower the outstanding positions of its members. The clearing house may thus, in some cases, become the central counterparty (CCP) to both final investors (in which case the process is known as "netting by novation"). Possibly at the same time, the members of the clearing house inform their respective brokers of their obligations and the brokers instruct their settlement agents. The settlement agent of the seller's broker receives the securities from the seller's custodian into its account, and credits them to the clearing house – which, for simplicity, we assume to have accounts in both the investors' central securities depositories (CSDs). The clearing house then issues an instruction for the securities to be credited to the account of the buyer's settlement agent, who credits them to the buyer's custodian. It may be necessary to carry out a reconciliation between the investors' CSDs and the issuer's CSD, e.g. to allow the execution of the notary function and of asset servicing. Each of these steps

27 The process may be simplified by using either internalised settlement (when both end-investors have their accounts with the same custodian) or via consolidation among intermediaries, e.g. the clearing member and the settlement agent could be one and the same entity, which sometimes also acts as a custodian for end-investors. Consolidation may, however, cause a lack of interoperability to lead to anticompetitive behaviour, particularly in the case of market infrastructures which constitute natural or regulatory monopolists.

may require a party's records to be reconciled with those of other parties at different levels of the value chain.

4 DLTs and the future of post-trading in securities markets

The impact of DLTs on post-trade financial institutions depends on at least three factors:

- the level of the post-trade value chain being considered – see Section 4.1;
- the type of governance they would be subject to – see Section 4.2; and
- the way in which the core incumbent institutions are willing and allowed to implement the innovation – see Section 4.3 where three stylised scenarios are discussed.

When a new transaction has been validated and the corresponding update to the distributed ledger can no longer be reversed, the transaction is settled in the ledger. It is important to mention that the consequences off-ledger are, however, unclear.

Blockchains and consensus ledgers can either record tokens (assets that are native to the ledger and can be used as part of users' incentives) or claims (a set of IOUs referring to off-ledger assets). It is only where tokens are used that settlement in the distributed ledger is sufficient to discharge the two parties from their obligations. If the blockchain represents off-ledger assets, the consequences of an update to the distributed ledger are unclear, in terms of their finality and the effect on ownership rights. In what follows, it is assumed that these important details will be addressed by legislators and that a DLT could be used to transfer property rights on financial securities.

4.1 Potential impact of DLTs in different layers of post-trading

When DLTs first started to be widely talked off (mid-2015), many developers were only offering products for one or a small number of the layers of the securities value chain. The situation has now changed and an increasing number of DLT solutions are being developed whose functionalities cover different layers of the securities transactions value chain.[28]

A number of innovators have also already been able to supply their products to financial institutions. Ripple, for example, is present in the foreign exchange market where it links banks acting as gateways in its consensus ledger; NASDAQ and Symbiont have issued private stock on private and public blockchains, respectively; and Overstock (a publicly listed company) has requested approval from the Securities and Exchange Commission to issue part of its shares on a proprietary blockchain, which would be fully publicly distributed.

[28] A small number of DLT solutions also cover the price-formation layer which relates to trading activity.

The incumbent players have joined forces, via various consortia, with companies involved in the application of new technologies to financial intermediation (known as fintech companies) in order to develop shared technologies. The range of specifications that DLTs offer could in fact be to the detriment of standardisation and interoperability, and even lead to increasing fragmentation, if market participants or clusters each adopt their own approaches. Nonetheless, irrespective of whether DLTs are adopted in clusters or across all financial institutions, the impact on post-trade processes is heavily dependent on the technological and functional implementation market that users choose.

This section discusses each layer of the whole post-trade value chain (clearing, settlement and asset servicing) separately, assessing, in each case, the possible consequences of DLTs.

4.1.1 Potential impact on the settlement layer

The importance of the notary function for financial stability is such that many national markets entrust a regulated monopolist with the task of ensuring the integrity of security issuances and the settlement of its trades. This role will also remain crucial if DLTs are adopted, as a reliable institution will be needed to ensure that the number of securities recorded in the distributed ledger corresponds to the description of the issuance given in, e.g. a global certificate, possibly stored as an immutable blockchain entry.

The notary role cannot be delegated to issuing companies or governments due to the misalignment of economic incentives this would create. It is necessary that a third party perform this role in order to ensure that only the amount of securities publicly issued is traded, with no unwarranted dilution of investors' claims. Confidence in the claims that securities bear over the underlying real assets is fundamental to allowing both corporate financing in the primary market, and maturity transformation, storage of value and hedging in secondary markets. It would therefore, most likely, be unrealistic to propose peer-to-peer issuance or settlement between issuing entities and end-investors participating in a distributed ledger for mainstream securities markets. The involvement of regulated entities is generally required, irrespective of the technology adopted.

Validation, and therefore settlement, of transactions could be delegated to nodes operated by a range of market participants. This, however, raises issues related to confidentiality, given the current DLT technology at least, and transactions may therefore need to be treated confidentially and validated by third-party institutions, such as the current financial market infrastructures.[29] The difficulty is caused by the fact that efficient validation of transactions in a DLT currently requires validators to access the details of the trade (possibly indirectly) in order to check its validity. Market players, however, have an inherent interest in keeping their trading strategies

[29] Regulatory issues are beyond the scope of this paper. It should be noted that the CSD regulation sets a series of rules on entities allowed to settle transactions and on their reporting to regulators.

confidential, and a range of new DLT features is therefore being developed in order to allow confidentiality to be maintained during the validation process, e.g. zero-knowledge validation.

The cash leg of a settlement represents a distinct area for DLT technology, separate from the securities leg. To date, market players have given little attention to the bridge between technologies used in the securities leg and in the cash leg of a trade. Whereas it is impossible at the current stage to predict whether and when central bank money will ever be available on a distributed ledger, cash accounts need to keep up with securities transfers to allow straight-through processes in the delivery-versus-payment (DVP) model currently in use, since sellers will likely continue to expect to be paid in real-world money rather than in virtual currencies. Any distributed ledger for securities therefore requires at least one reliable entity to link the securities and cash accounts of different participants. This could theoretically be achieved using commercial bank money (CoBM), although this approach would pose risks, as the cash received would actually be a claim on the private institution whose cash account has been debited, and could therefore be affected by bankruptcy procedures involving this institution. Pre-funded accounts for settlement in CoBM have been envisaged as one possible solution, but their use would freeze liquid assets and thus limits the benefits of distributed ledgers in terms of collateral fluidity.

In the DLTs currently available, individual securities accounts would need to be opened for the issuer and the investors. They would need to be managed either by a designated ledger administrator or directly by regulated participants that are able to perform identity checks.

The concept of settlement finality is crucial to ensuring the smooth functioning of financial markets. How settlement finality is achieved when recording securities in a distributed ledger may vary between DLTs. Different technologies include different functionalities in terms of when a transaction can be recalled, whether the ledger is interfaced or integrated in off-ledger settlement systems, and whether ledger updates represent claims on off-ledger assets or the asset is itself encoded in the ledger as a token or smart security. The securities currently being traded could be put in escrow to allow them to be settled in their digital form on the distributed ledger. Some countries already require materialised securities to be held, with only the bookkeeping being carried out electronically. The possibility that, at some point in time, outstanding security issuances could be migrated en masse to a distributed ledger is unrealistic, but an interim dual system with transactions being settled both in traditional and DLT systems by their current operators is technically feasible.

The impact of DLT on cyber resilience does not depend only on the validation methods used in the ledger. The ability of individual participants to resist an attack from outside the network will also play a major role. If each node had the security features that centralised databases currently used for settlement have, an attacker would have to control a number of these IT systems (rather than just one) in order to affect the system. This would thus represent an improvement to security. If, however, the validation role is distributed among less cyber-aware institutions, the overall improvement or otherwise depends on the relative negative and positive effects of, respectively, the lower security of the individual nodes, and the fact of having a

number of nodes, rather than just one. In either case, the redundancy of the copies of the distributed ledger kept by different participating institutions is likely to make recovery easier, should there be a failure in the system.

4.1.2 Potential impact on the custody layer

A distributed ledger could allow securities to be held directly by final investors or via their point of contact with the financial market, e.g. a broker or commercial bank that can enforce know-your-customer (KYC) and anti-money laundering (AML) provisions. Moreover, smart contracts can ensure that corporate actions and actions necessary to ensure compliance with requirements on collateral eligibility and optimisation are executed automatically.

Smart contracts could be seen as a possible way of facilitating asset servicing. Any event that can be recorded or communicated in digital form can form part of a self-executing algorithm that updates accounts automatically. Examples of the types of asset servicing that could potentially be performed by smart contracts include: collecting income from the issuer's cash account to credit automatically calculated dividends/coupons on shareholders'/bondholders' accounts at their due date; withholding or reclaiming tax; splitting or redeeming stocks; and, in general, any crediting or debiting of accounts in the ledger that is carried out under instructions triggered by an event that can be verified either in the ledger or through a trusted off-ledger institution.

Once the ledger has been programmed to perform activities such as those mentioned above, there are few tasks left to the chain of custodians. The identification of final investors and issuers and controlling their access to the ledger are the only current tasks where human interaction might continue to be required. This type of "gatekeeping" role could be performed by the entity managing the ledger, but the need for identity checks may make physical interaction with final investors necessary at some point. A commercial bank or public authority could be responsible for this aspect. The development of smart contracts that can handle corporate events and other ancillary services, such as collateral management and securities lending, may then see the current custodian role become redundant.

4.1.3 Potential impact on the clearing layer

The impact of DLTs on trade enrichment, confirmation, and matching depends on whether and how trading platforms could be integrated with the distributed ledger. Were these two separated systems, the impact of DLT would be limited to a simplification of the netting and risk management procedures. If transaction messages on the ledger were themselves settlement instructions for cash transactions, however, a series of typical back-office procedures that take place between trade capture and the instruction of settlement might disappear.

Distributed ledger technologies have the potential to allow trading and settlement of securities to take place at almost the same time – not only on the same day (within the T+0 cycle), but even potentially with instantaneous settlement. Where trading platforms are connected to a distributed ledger, it is possible to make the posting of orders to the trading venue or over the counter by buyers and sellers dependent on the integrated DLT system having ascertained the availability of securities in the account of the seller and cash in that of the buyer.[30] The technology needed to deliver instant settlement is available in current traditional systems, but it would be unfeasible to implement it globally for the vast amount of transactions taking place, while the business processes and siloed databases remain as they are.

The possibility of instant settlement, i.e. the ledger being updated immediately after a trade has been executed and validated, would affect the role of clearing for cash transactions.[31] Being able to require the availability of securities and cash would eliminate liquidity and credit risk from any trade executed for immediate delivery, although the impact of eliminating netting on market liquidity and price informativeness has yet to be assessed.

Where the distributed ledger has some latency, or a trade involves derivative contracts which would be executed at a later stage, clearing would still be necessary in order to hedge risk until the exchange of securities and/or cash is final and irrevocable. Smart contracts allow automatic netting and, if collateral management systems are linked to the distributed ledger, also margin calls. Even in cases such as these, the impact of distributed ledgers on incumbent clearing houses nonetheless depends on the level of technology being implemented, and the extent to which regulators are willing to delegate the clearing process to smart contracts.

In DLTs where settlement is not instantaneous, and clearing would thus still be necessary before settlement of the underlying securities transaction can take place, smart contracts could potentially change the way in which netting and collateral are managed. By encoding smart contracts in the ledger, a CCP could make margin calls automatically executable in the accounts of its clearing members. Minimum settlement cycles vary between DLTs, ranging from a few seconds to an hour.

It would also be conceivable for traders to have to borrow securities or cash from a third party as a prerequisite to executing the trade. Under this type of system, an order could only be posted once the securities or cash had been able to be sourced from the accounts of a lending party, and any required collateral blocked in the account of the borrower. The execution of the trade would then transfer the securities and cash simultaneously, providing that the securities and cash belonging to the market participants are accessible to the set of smart contracts executing the trade in the ledger. This type of innovation could make clearing unnecessary for both matching and risk management purposes. This would free up the collateral that is currently immobilised either with a central counterparty or for hedging the settlement risk bilaterally.

---

30 See, e.g. Patent application US 2015/0332395 filed by Goldman Sachs & Co.

31 For derivative transactions, the risk exposure would need to be hedged until execution of the contract.

Some developers claim that DLTs for settlement could be constantly active, i.e. in operation 24 hours a day, seven days a week. Although there are no technical limitations on the operating time of DLTs, it is not yet clear if this claim will hold, particularly given that a distributed ledger may need to interact with various off-ledger entities.

CCPs are recognised as low-risk counterparties, and netting allows the counterparties of qualified CCPs to face lower capital charges. This might dissuade some financial intermediaries from avoiding clearing. Instant settlement on the distributed ledger could, nonetheless, still be a favourable option, if banking regulation will treat it as virtually risk-free. A major factor in this will undoubtedly be the decision as to whether securities and cash registered in the ledger are considered as only a representation of the off-ledger world, or their settlement is also considered to be final off-ledger.

4.2 Governance of distributed ledgers

The financial industry is a network industry. The adoption of DLTs, in whatever form, for use in the post-trade of securities transactions will face the same issues as are confronting the current post-trade set-up, i.e. the need for technical standardisation, common business rules and sound governance arrangements. One deciding factor will be whether the control of any particular, widely used DLT specification is in the hands of one single private entity (irrespective of whether this entity is a fintech start-up or a well-established financial institution), or of the community that is using that specific DLT. Some market players have joined forces to develop a common stance on distributed ledger technologies. Many institutions are, however, also investing resources in developing their own solutions. These may reduce interoperability if a common standard does not prevail.

One major operational obstacle to the adoption and use of distributed ledgers could be a problem that has also prevented greater interoperability between post-trading institutions: a lack of harmonisation and standardisation. Common technical standards and business rules are a prerequisite if financial markets are to reap the full benefits of the new technology, although even proprietary technologies would allow the efficiency of post-trading to be increased, at least for clusters of financial institutions. Harmonisation is therefore still needed. To have market-wide gains in terms of safety and efficiency, every system in the market might need to be able to communicate with all the different DLTs adopted.[32]

[32] Some developers of DLTs envisage the use of SWIFT/FIX/ISO20022-compliant messages to allow communication between different systems.

4.3 Overall impact of different levels of implementation on the post-trade landscape for securities

Given the impact of distributed ledgers at different layers of the value chain, any decisions taken by regulated core institutions, such as registrars and CSDs, on the adoption of DLT are likely to determine, to a large extent, the consequences for the post-trading landscape. Although fintech companies may develop innovative solutions, they are unlikely to be able to compete in such a highly regulated market unless regulation changes. Large financial institutions might decide to use DLTs to internalise settlement, but it is unclear whether they would be able to achieve the volumes necessary to benefit from the IT investment. There are at least three stylised scenarios that can be considered:

Scenario 1) the incumbent institutions embrace the new technology to improve cluster/internal efficiency, leaving business practice "as it is".

If different institutions succeed, separately, in developing their own preferred technologies, this will bring benefits for their internal efficiency, but with little possibility of linking the different market infrastructures and participants to a fast and interoperable securities settlement system for Europe and beyond. In this scenario, the existing business arrangements can be imagined to be replaced by an equivalent system made up of different distributed ledgers, with little effect on the institutions involved (see Diagram 2 for an illustration of the case of a distributed ledger shared between intermediaries, e.g. on the buy side of a transaction).

Diagram 2

How a distributed ledger may affect the efficiency of post-trade in the securities market, assuming current business practice continues

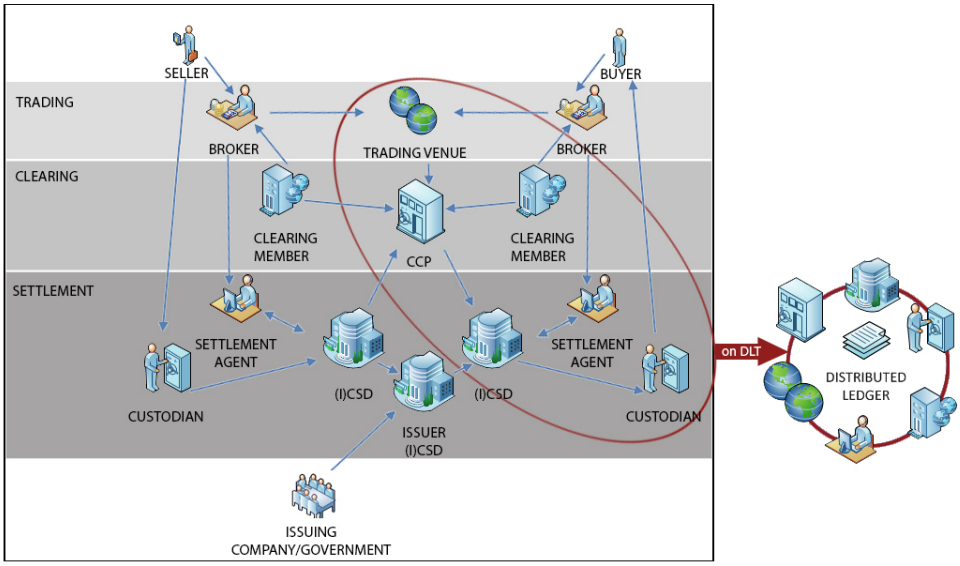

A set of post-trade institutions (in this diagram, on the buy side) may develop their own DLT for internal use. All business relations inside the red circle would then take place as straight-through processes on the distributed ledger.

In this scenario, institutions' motivation for adopting DLTs would be to save reconciliation costs that the AITE group estimates at $1.20 billion for 2016.[33] The only breakthrough would be that any change in the ownership of a security would be processed as part of the near-real time updating carried out using the information available to the cluster of intermediaries that have direct access to the same ledger.

The impact of this change would be small in terms of the number of players involved. With the various clusters of institutions continuing to use siloed distributed ledgers, the need to bridge these ledgers in order to reconcile accounts along the chain of intermediaries would remain unchanged. It would be possible, under this scenario, to achieve, on a small scale, the long-sought interoperability between (international) central securities depository ((I)CSDs), central counterparties (CCPs) and collateral management service providers, at least among institutions that adopt a common DLT. This would be an improvement in terms of efficiency, but it would not represent a significant development in terms of market integration.

Scenario 2) Core players, such as CSDs, adopt market-wide distributed ledgers. In this "adoption model" scenario, at least some peripheral players might become redundant.

The distributed ledger would make the current post-trade landscape faster and more automated, with securities and trades being serviced and settled in a common distributed ledger rather than between siloed databases.

Diagram 3

How a market-wide distributed ledger may affect the post-trade landscape of securities markets

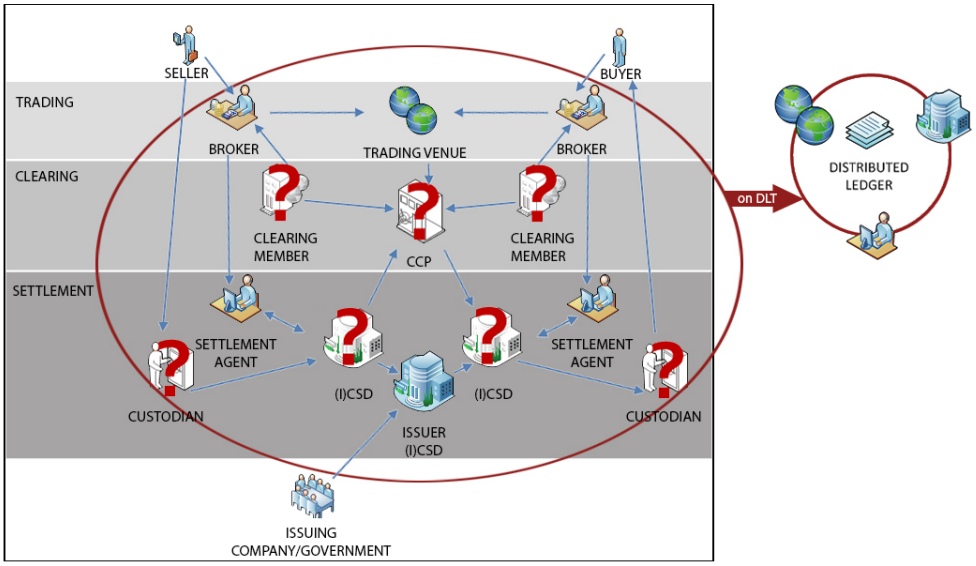

If the whole post-trade industry migrated to a distributed ledger settlement process, securities accounts would be updated automatically. Depending on the extent of the implementation of smart contracts, some layers of the industry could become redundant.

33 See http://aitegroup.com/report/reconciliation-technology-solutions-2014-recs-get-ready-rumble-%E2%80%A6.

Diagram 3 shows the case of a distributed ledger where either a trading venue or the settlement agents of a number of brokers allow external trading parties to access the distributed ledger. Under this “adoption model” scenario, the crediting and debiting of investors’ securities accounts could be performed with the same cost- and time-efficiency as that with which internalised settlement is currently carried out in the accounts of custodian banks. It would, however, take place in the distributed ledger, meaning that the segregation of securities in individual investors’ accounts would be carried out at no additional cost.

Reporting would take place automatically and in near-real time. Regulators could keep track of each transaction and of outstanding positions, thus complying with the provisions of the draft securities financing transactions (SFTs) regulation. The long-sought interoperability between financial intermediaries and market infrastructures, which might still act as a link between securities and national cash accounts, would be achieved.

The use of smart contracts, potentially to be implemented at a subsequent stage, could make intermediaries such as custodians and CCPs redundant, as the tasks they perform could become automated. The efficiency gain offered by this approach would be particularly important in markets where interoperability is currently low, and where the lack of automation means that days of post-trade processing are required to settle a transaction.

Securities transactions executed over the counter (OTC) would also become more transparent, in so far as regulators have access to the ledger, since the latter would be populated in real time with the same type of information as is currently routed to trade repositories for some segments of the capital market.

Scenario 3) Issuing companies, governments or fintech companies take the lead in implementing peer-to-peer systems for securities transactions, thus taking the post-trade industry into a “new world”.

In this extreme scenario, the current post-trade processes would be superseded by automated clearing and settlement, taking place among a network of issuing entities and final investors (see Diagram 4). This type of change has the potential to promote financing from small investors and SMEs, in a niche market that could operate in parallel to the more structured market for larger-scale issuances. Companies and governments could issue their financial instruments directly on the ledger, which would be of particular interest to start-up companies issuing shares by private placement, whilst smart contracts would execute any corporate action automatically.

KYC and AML rules constitute a major obstacle to the realisation of this "new world" scenario, where price discovery takes place on an open trading platform and settlement takes place automatically on the distributed ledger. Innovations in the area of e-identity may, however, facilitate the major changes that would be needed. In general, trusted parties need to certify the identity of potential investors, in order to allow them to open an account and trade on a distributed ledger in compliance with current regulation. The e-identity does not need to be granted by a financial

intermediary and can be part of the e-government tools that are currently being developed, including by countries such as Estonia and Finland.

Diagram 4

How a peer-to-peer market for securities based on DLTs could affect the post-trade landscape

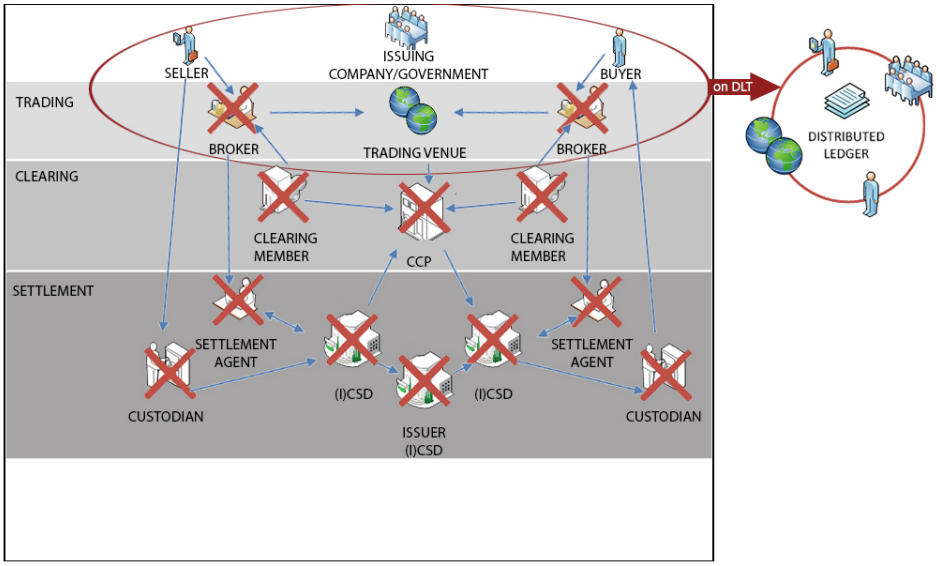

If capital markets were to migrate to a peer-to-peer model, the whole chain of intermediaries would become redundant, and companies or governments could issue their own securities on the distributed ledger.

The validation technologies currently in place do not protect privacy, and offering direct access to the distributed ledger could allow sophisticated investors to pursue unfair trading. By tracking the activity of unsophisticated competitors, they would be able to acquire information on their liquidity needs, which would allow them to practise front-running and thus increase the trading costs of these competitors. Some level of confidentiality would therefore be necessary in order to keep information on settled transactions inaccessible to traders in real time. Technical solutions to this problem are being developed, using zero-knowledge proof algorithms, i.e. validation methods that settle transactions across the network of participants without sharing their details with non-involved parties.

5 Conclusions

Distributed ledger technologies are of relevance beyond the Bitcoin and its unrestricted blockchain model. Consensus ledgers, restricted technologies and smart contracts all represent more viable and attractive options for financial institutions, as they draw on existing business models used in the post-trade industry for securities transactions and have the potential to create safer, more reliable and efficient post-trade processes.

Based on the analysis in this paper, it can be concluded that there is potential in distributed ledger technology. Furthermore, innovation is, in general, welcome in the European post-trade market for securities, wherever it can bring safety and efficiency. A number of factors could, however, pose potential barriers to the widespread uptake and use of DLTs. First, the technology is not yet mature; second, the clarification of critical legal, operational and governance issues will take time; and third, even were DLTs to be adopted widely, certain post-trade functions will continue to be performed by institutions. It is not yet, therefore, clear whether DLTs will cause a major revolution in mainstream financial markets or whether their use will remain limited to particular niches. It is possible that a DLT may find its way into the mainstream market, but should this happen, it is more likely to cause a gradual change in processes, rather than a revolution in the market.

References

Bayer, D., Haber, S. and Stornetta, W.S. (1993), “Improving the Efficiency and Reliability of Digital Time-Stamping”, in R. Capocelli, A. De Santis and U. Vaccaro (eds.), Sequences II: Methods in Communication, Security and Computer Science, Springer Verlag, New York, pp. 329-334.

Benaloh, J. and de Mare, M. (1991), “Efficient broadcast time-stamping”, Clarkson University, Department of Mathematics and Computer Science. TR 91-1, August.

Eyal, I. and Sirer, E.G. (2014), “Majority Is Not Enough: Bitcoin Mining Is Vulnerable”, in N. Cristin and R. Safavi-Neini (eds.), Financial Cryptography and Data Security, Vol. 8437 of Lecture Notes in Computer Science, Springer Verlag, Berlin and Heidelberg, pp. 436-454.

Haber, S. and Stornetta, W.S. (1991). “How to Time-Stamp a Digital Document”, in A.J. Menezes and S.A. Vanstone (eds.), Advances in Cryptology-CRYPT0’ 90, Vol. 537 of Lecture Notes in Computer Science, Springer Verlag, Berlin and Heidelberg, pp. 437-455.

Nakamoto, S. (2008) “Bitcoin: A Peer-to-Peer Electronic Cash System”, manuscript.

Sompolinsky, Y. and Zohar, A. (2013), “Accelerating Bitcoin’s Transaction Processing. Fast Money Grows on Trees, Not Chains”, IACR Cryptology ePrint Archive, Report 2013/881.

Swanson, T. (2015), “Watermarked tokens and pseudonymity on public blockchains”, R3 publication.

Acknowledgements
The authors are grateful to Helmut Wacket, Markus Mayers, Dirk Bullmann, Ad van Riet, and George Kalogeropoulos (European Central Bank) for their very helpful comments on earlier drafts. All errors remain with the authors.

Andrea Pinna
(Corresponding author)
European Central Bank; email: andrea.pinna@ecb.europa.eu

Wiebe Ruttenberg
European Central Bank; email: wiebe.ruttenberg@ecb.europa.eu

Postal address 60640 Frankfurt am Main, Germany
Telephone +49 69 1344 0
Website www.ecb.europa.eu

ISSN 1725-6534 (online)
ISBN 978-92-899-2335-4
DOI 10.2866/270533
EU catalogue No QB-AQ-16-005-EN-N