Federal Reserve
Bank of Dallas

Embedded Supervision: How to Build Regulation into Blockchain Finance

Raphael Auer

Globalization Institute Working Paper 371

Research Department

https://doi.org/10.24149/gwp371

Embedded Supervision: How to Build Regulation into Blockchain Finance*

Raphael Auer†

October 2019

Abstract

The spread of distributed ledger technology (DLT) in finance could help to improve the efficiency and quality of supervision. This paper makes the case for *embedded supervision*, i.e., a regulatory framework that provides for compliance in tokenised markets to be automatically monitored by reading the market's ledger, thus reducing the need for firms to actively collect, verify and deliver data. After sketching out a design for such schemes, the paper explores the conditions under which distributed ledger data might be used to monitor compliance. To this end, a decentralised market is modelled that replaces today's intermediary-based verification of legal data with blockchain-enabled data credibility based on economic consensus. The key results set out the conditions under which the market's economic consensus would be strong enough to guarantee that transactions are economically final, so that supervisors can trust the distributed ledger's data. The paper concludes with a discussion of the legislative and operational requirements that would promote low-cost supervision and a level playing field for small and large firms.

JEL Classification: D40, D20, E42, E51, F31, G12, G18, G28, G32, G38, K22, K24, L10, L50, M40.

Keywords: tokenisation, stablecoins, asset-based tokens, cryptoassets, cryptocurrencies, regtech, suptech, regulation, supervision, Basel III, proportionality, blockchain, distributed ledger technology, central bank digital currencies, proof-of-work, proof-of-stake, permissioned DLT, economic consensus, economic finality, fintech, compliance, auditing, accounting, privacy, digitalisation, finance, banking.

---

*I thank David Archer, Ryan Banerjee, Morten Bech, Rainer Böhme, Dirk Broeders, Stijn Claessens, Johannes Ehrentraud, Marc Farag, Jon Frost, Leonardo Gambacorta, Marius Jurgilas, Thomas Leach, Henry Holden, Krista Hughes, Sébastien Kraenzlin, Joey Patel, Jermy Prenio, Joseph Noss, Tara Rice, Tarik Roukny, Hyun Song Shin, Philip Wooldridge, and an anonymous referee for the BIS working paper series; members of the Basel Committee for Banking Supervision's Task Force on Financial Technology, of the Financial Stability Board's Financial Innovation Network, and of the Financial Stability Institute's Suptech Network; as well as seminar participants at the Bank for International Settlements, the European Central Bank-University College London 2019 P2P Financial Systems conference, the Swiss Financial Markets Supervisory Authority, and the Swiss National Bank for comments. I further thank Alan Villegas and Giulio Cornelli for outstanding research support. The views expressed in this paper are those of the author and not necessarily those of the Bank for International Settlements, the Federal Reserve Bank of Dallas or the Federal Reserve System.

†Raphael Auer, Bank for International Settlements, Raphael.auer@bis.org.

Introduction

Authorities around the world today are grappling with the rise of distributed ledger technology (DLT) in finance. The challenge facing them is how best to apply technology-neutral regulation, so that similar risks are subject to the same regulation.[2]

This paper investigates how the "same risk, same regulation" principle might be applied to the financial supervision of DLT-based markets. It argues that, while regulation should remain technology-neutral, supervision should evolve in parallel with technology.[3] Although DLT may not change the underlying risks, it might open up new ways of supervising these risks.[4] So, instead of trying to fit cryptoassets into existing regulations, such as securities laws formulated long before the advent of DLT, it is worth asking how new technologies could serve to better monitor risks in financial markets.

Based on these characteristics, this paper puts forward the concept of "*embedded supervision*". This comprises a regulatory framework that provides for compliance to be automatically monitored by reading the market's ledger. As such, it reduces the need for firms to actively collect, verify and deliver data.

DLT makes possible the decentralised trading of asset-backed tokens, as well as decentralised financial engineering based on these tokens via self-executing ("smart") contracts. If such innovations take root, they will drive the development of financial markets via new forms of transparency and data credibility. The fundamental novelty is that DLT builds such credibility with a decentralised data structure based on economic consensus. Effectively, this harnesses the incentives of individual market participants to replace middleman-based data verification.

Compliance monitoring would then be automated, by relying on the trust-creating mechanism of decentralised markets for supervisory purposes. For example, for the case of a bank that holds asset-backed tokens, compliance with the Basel III capital standards could be automatically verified. This would be done by computing the ownership of (borrowing and lending) balances and the associated risk weights in the relevant distributed ownership ledgers.

Embedded supervision could ease the conflict between data availability, the cost of data collection and verification, and privacy. Compliance expenditure weighs heavily on financial institutions, and even more so on smaller firms. Supervisors thus face a trade-off between getting the data they need and keeping the costs of compliance within reasonable limits. Embedded supervision could further help

[2] The rise of so-called cryptocurrencies has also threatened to bypass existing legislation, in particularly with regard to anti-money laundering/know-your-customer (AML/KYC) legislation and facilitating illicit activity (see Möser et al (2013), Foley et al (2018) and Fanusie and Robinson (2018)), thus calling for a response to level the playing field (see Carstens (2018a,b,c), Landau and Genais (2018), Auer and Claessens (2018 and 2019), and FATF (2018)).

[3] Whereas "regulation" is the process of writing the rules that apply to the regulated entities, "supervision" is the enforcement of these rules.

[4] FIMNA (2018) and HM Treasury-Financial Conduct Authority-Bank of England Crypto-assets Taskforce (2018) apply existing regulatory frameworks to new DLT-based financial products according to underlying economic activity. They argue that, in most cases – such as the funding of a business via an initial coin offering (ICO) or a traditional initial public offering (IPO) – the choice of financial technology (ICO vs IPO) does not change the underlying risks.

maintain the confidentiality of firms and their customers, since cryptographic tools can be used to report an institution's aggregated financial exposures to the supervisor without disclosing the underlying individual transactions.

At this point, it should be noted that the concept of embedded supervision goes much further than simply reading a distributed ledger. The key issue is that data are not necessarily valid just because they are stored in multiple places. In today's compliance process, the data's trustworthiness is guaranteed by the legal system, the relevant authorities and the threat of legal penalties. In DLT-based markets, by contrast, data credibility is assured by economic incentives. In this world, the supervisors must primarily examine the conditions under which the market's economic consensus is strong enough to guarantee the quality of the data contained in the distributed ledger.

But what principles should govern a regulatory framework designed to use a market's distributed ledger for financial supervision? This paper discusses four principles for the deployment of embedded supervision (see Table 1).

Principles of embedded supervision — Table 1

Embedded supervision is a regulatory framework that provides for compliance with regulatory standards in DLT-based markets to be automatically monitored by reading the market's ledger. It would reduce the administrative burden for firms, while increasing the quality of data available to the supervisor. Four principles would guide their use:

- **Embedded supervision can only function as part of an overall regulatory framework that is backed up by an effective legal system and supporting institutions.**

  DLT-based exchange can evidence the transfer of ownership of asset-backed tokens from one known entity to another, but the connection between the underlying asset and the digital token must be guaranteed by the legal system. Additional institutions may also be required, for example, to guarantee the accuracy of external reference points that are relevant to payoffs of smart contracts.

- **Embedded supervision can be applied to decentralised markets that achieve economic finality.**

  If there is no central intermediary to guarantee that a transfer of funds or securities has become irrevocable, an economic one must be applied. Following Auer (2019), economic finality means that a transaction can be considered as final once it is certain that, from a specific moment, it will never be profitable to undo.

- **Embedded supervision needs to be designed within the context of economic market consensus, taking into account how the market will react to being automatically supervised.**

  Embedded supervision creates incentives for a regulated firm to cheat the supervisor by altering the transaction history in the blockchain. Supervisors thus need to ensure that the market's economic consensus is so strong that any attempt to deceive the supervisor will be unprofitable.

- **Embedded supervision should promote low-cost compliance and a level playing field for small and large firms.**

  Embedded supervision should be designed to keep the fixed costs of compliance low. The supervisor may need to monitor aspects of decentralised markets – such as the verification market and the governance of decentralised systems) to ensure a level playing field for entrants.

Source: Author's elaboration.

These applications run on "permissioned" DLT, in which peer-to-peer exchange is facilitated by decentralised economic consensus. At the same time, such systems retain an overarching coordination mechanism – tied to the legal system – that determines who can participate in the market and that guarantee the quality of the underlying assets.

Hence, the first principle of embedded supervision is that it must be part of an adequate entity-based regulatory framework, backed up by an effective legal system and supporting institutions. Foremost, this means that asset "tokenisation" – the process by which claims on real assets are digitally represented – is validated by the legal system. Although cryptography and distributed ledgers can prove the transfer of asset-backed tokens from one entity to another, the connection between the underlying asset and the digital token must ultimately be guaranteed by the legal system – which alone can underpin the ownership of assets such as real estate or shares in a brick-and-mortar business.

Summing up, the first principle of embedded supervision calls for a proper understanding of what DLT-based trading can achieve, and what it cannot. Just as in today's system, a decentralised financial system would need to be solidly rooted in both the legal system and supporting institutions such as land registries or rating agencies. What differs from today's system is the operational setup of how these entities trade with each other, how such trading is recorded, and how misbehaviour is deterred.[5]

The second and third principles – which constitute the paper's core theoretical results – concern the economic incentives at work to guarantee the finality of transactions in decentralised markets. For a supervisor to monitor compliance involving any set of transactions and ownerships, these transactions must be irrevocable and final (see CPMI-IOSCO (2012)). If there is no central counterparty capable of vouching with a legally binding signature, some different criterion for transaction finality must be established.

This paper focuses on the concept of economic finality proposed in Auer (2019), ie the notion that a transaction is final once it is no longer profitable to reverse it.[6] To this end, I develop a distributed and permissioned market in which "blocks" of financial contracts are verified by third parties. These verifiers stand to lose a given amount of verification capital should a blockchain reversal ever occur that voids existing transactions. From this setup, I develop the concept of "certain economic finality", meaning that a verifier's total skin in the game is so high that no market participant would ever find it profitable to bribe a verifier into reversing a transaction. I then argue that, if transactions are economically final, the supervisor can take them at face value.

The third principle is that, when designing embedded supervision, supervisors need to take into account the impact of their own actions on the regulated market. Regulated firms incur a cost in complying with regulation that they would not incur

[5] An additional key element is a watertight and ideally globally coordinated AML/KYC identity framework that keeps illicit activity out of this novel ecosystem.

[6] Auer (2019) examines economic finality for the case of proof-of-work-based consensus schemes. Bonneau (2016), Chiu and Koeppl (2017) and Budish (2018) offer related concepts and analyse the conditions under which blockchain transactions become prohibitively expensive to reverse via so-called 51% brute force attacks.

voluntarily.[7] By the same token, in the DLT world, this creates incentives for a regulated firm to cheat the supervisor by altering the transaction history in the blockchain. I thus also model the supervisor's impact on the market and show that, if a supervisor wishes to monitor compliance in real time, one strategy is to mandate a commensurate increase in the total amount of skin in the game for the verifiers.

The fourth and last principle concerns the broader societal goals when designing embedded supervision. The regulator's goal is neither a specific market structure nor a specific form of exchange. Rather, it is to create a stable financial system that offers high-quality services to consumers and business at the lowest possible cost. In this regard, a key puzzle is that, despite ample technological progress, financial services remain stubbornly expensive (see Graph 1). This might partly reflect the high barriers to entry created by the costs of complying with financial regulation.

Low-cost information technology has not yet brought down the price of financial services — Graph 1

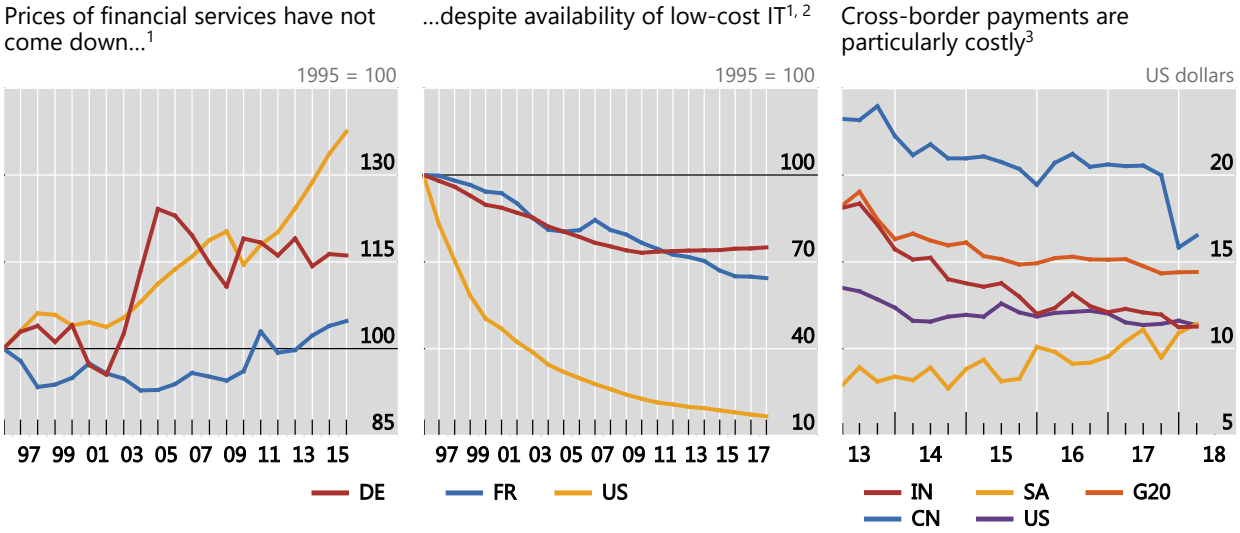

[1] Gross output price index normalised to equal 100 in 1995. [2] Simple average of the prices of computer components, software and communication equipment. For DE, price of software. [3] Average total cost for sending $200 with all remittance service providers worldwide. For CN and IN, receiving country average total cost; for G20, SA and US, sending country average total cost.

Sources: EU KLEMS; Eurostat; US Bureau of Economic Analysis (BEA); World Bank, *Remittance Prices Worldwide*, remittanceprices.worldbank.org; World Bank; BIS calculations; author's calculations.

Against this backdrop, I discuss how embedded supervision could be designed with a view to harnessing the "fintech opportunity" highlighted in Philippon (2016), as well as promoting low-cost financial service provision and a level playing field for both incumbents and potential entrants. In this aspect, the operational dimension is important. Public authorities can digitally sign and time-stamp relevant information – for example, the central bank's policy rate, data from national statistical offices or public land and firm registries – so that it can be fed directly into relevant market ledgers. Further, the fixed costs of compliance could be kept low by ensuring blockchain interoperability and developing an open-source suite of monitoring tools accessible to potential market entrants.

[7] Underlying regulation addresses issues related to limited liabilities, market contagion, and other externalities (see eg Allen and Gale (2000), Admati et al (2011), Admati and Hellwig (2013)).

I conclude by discussing challenges for legislators and regulators. The main legislative challenge is to provide for the concept of decentralised economic finality in legislation governing financial market infrastructure, ie to allow for ownership to be transferred without the involvement of a central registry. Regulators and supervisors would further need to develop auxiliary frameworks that govern distributed markets and their infrastructure, for example, when assigning the responsibility for dealing with crime in decentralised markets.[8] With this, the rise of DLT might lead to higher-quality compliance at a lower cost. This stands, of course, in stark contrast to the current situation where DLT investors lobby for light regulatory regimes while supervisory agencies struggle to apply AML/KYC standards to cryptocurrencies.

The structure of this paper is as follows. The next section outlines a potential future landscape for the financial industry in which regulated financial entities trade in decentralised marketplaces. It also discusses how these novel compliance processes could be organised. The next section develops a theoretical model of a decentralised market, and sets out the conditions under which embedded supervision could operate. The following section discusses what these considerations would entail on the part of regulators and supervisors and, in particular, how novel regulatory frameworks could harness a technological opportunity with a view to creating a stable and competitive financial sector.

Embedded supervision of token ecosystems: a primer

This section discusses the current compliance process and its costs, and how embedded supervision could improve the compliance process in DLT-based markets that allow for decentralised trading of asset-backed tokens, as well as, decentralised financial engineering based on these tokens.

The trade-off between costs and data gaps in today's compliance process

Today's compliance process involves compiling reports at multiple levels of data granularity. In both their retail and wholesale businesses, banks engage in millions of individual transactions. These data need to be collected, aggregated and delivered to a host of internal stakeholders (internal risk control, internal compliance, management, trading desks etc) as well as to supervisors.

These data not only need to be delivered, but they need to be continuously deliver*able*. For example, to ensure that account holders have access to their insured deposits in the event of a bank failure, 12 CFR Part 370 of the FDIC's Rules and Regulations requires larger insured depository institutions to identify all of their

[8] Quintenz (2018) offers a very useful discussion regarding under what conditions software developers, transaction validators, or users might be accountable for illegal activity on distributed platforms.

insured depositors (ie each individual account), so that their account information is always available in the event of a failure.[9]

Compliance is thus resource-intensive, confronting supervisors with a trade-off between getting the data they need and keeping the costs of compliance within reasonable limits. On the cost side, surveys indicate that compliance costs typically account for several percentage points of all operational costs at financial institutions (see Graph 2, left-hand panel), although not all of this is due to the administrative cost of complying with financial supervision. Costs are substantial for supervisors too (see Graph 2, right-hand panel).

The costs of compliance — Graph 2

Firm surveys point to high compliance costs[1, 2]

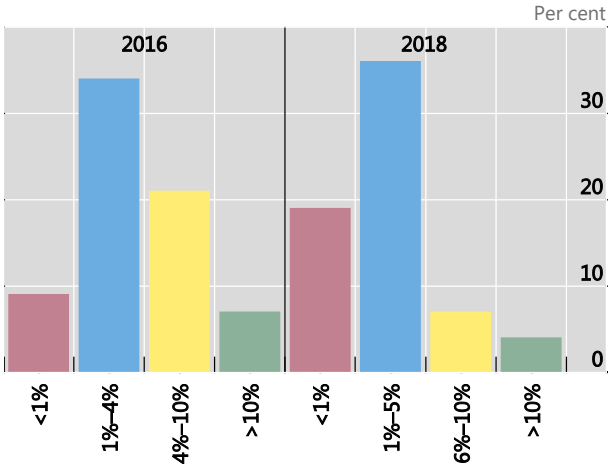

Costs are substantial for supervisors too

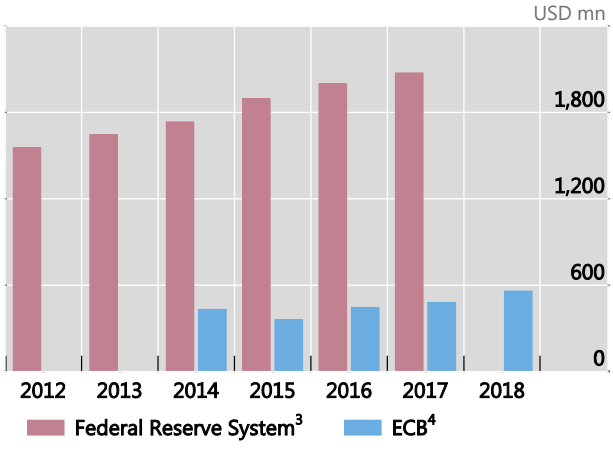

[1] Question: "As a percentage of annual revenue, how much do you believe your company spent or will spend on compliance?" Data do not specify the cost subcomponent of complying with financial supervision. [2] Remaining percentage of respondents answered "did not know". [3] Overall supervision and regulation and related operating expenses of the Federal Reserve System. [4] Supervisory fees.

Sources: Duff & Phelps "Global regulatory outlook", various years; ECB Banking Supervision, www.bankingsupervision.europa.eu; Board of Governors of the Federal Reserve System, Supervisory Assessment Fees Archive, www.federalreserve.gov/supervisionreg/supervisory-assessment-fees-archive.htm; national data; author's calculations.

Data gaps are the inevitable price paid by supervisors as they seek to keep the costs of compliance within limits. But the cost of such gaps can be devastating, as the collapse of Lehman Brothers in 2008 showed. At the time, the worlds' major financial institutions were not able to compute their consolidated exposure to the many subsidiaries of Lehman, so that "what would have been systemic risk morphed into systemic uncertainty" (see Haldane et al (2015)).

[9] See www.fdic.gov/regulations/resources/recordkeeping/index.html. Similarly, under too-big-to-fail regulation, large banks worldwide must have resolution plans ready, ie be able to separate their international and domestic financial businesses to insulate local credit conditions from turmoil abroad. This requires that banks can, at short notice, identify and classify all business activities as either domestic or foreign.

While important gaps have been filled since the Great Financial Crisis (GFC),[10] new ones are constantly emerging as the financial industry evolves.[11]

Compliance in DLT-based markets

How might the compliance process change in a DLT-based market? To be sure, DLT-based innovations have the potential to transform financial markets, in part as they offer radically new forms of transparency. The starting premise is that one needs to look beyond Bitcoin and other "permissionless" cryptocurrencies or cryptoassets[12] and instead focus on a "permissioned" version of the technology, which facilitates normal market functioning by decentralised consensus, yet retains, as a backup, an overarching (ie legal) coordination mechanism.

Such permissioned technology primarily enables the decentralised trading of asset-backed tokens, as well as decentralised financial engineering based on these tokens via self-executing ("smart") contracts.

The near-term potential of such tokenisation is highest in wholesale markets, ie trade between registered financial entities (see Mills et al (2016) and Benos et al (2017)). One example is the loan securitisation market, which has significant scale in the United States. In international markets too, banks' wholesale financial exposures are larger than their underlying business with non-financial customers. The dominance of wholesale financing and trading is most evident when it comes to payments: across the world, over 90% of all payments are of a wholesale nature. DLT could find widespread applications in such markets.

Exchanges or OTC markets could also be automated, as DLT can match demand and supply and automate price discovery. Similar developments could be envisaged for options and futures clearing houses. A principal application for DLT is to automate the flow of funds and the updating of security registers, which could reduce administrative costs and, most importantly, settlement risk (see Ruttenberg and Pinna (2016) and Chiu and Koeppl (2019)).[13]

[10] The cost of data gaps during the GFC was so evident that the G20 took the initiative in plugging them via a cooperation between the world's major international financial institutions, and national regulators.

[11] One contentious example, which demonstrates the trade-off between getting relevant data and keeping the costs of compliance within reasonable limits, is the supervision of the Basel III leverage ratio. The ratio sets a lower limit on banks' core capital as compared with their total assets. For banks in the European Union, the leverage ratio is implemented using quarter-end snapshots, while in other jurisdictions it is calculated on a daily average basis. This disparity creates incentives for banks to engage in regulatory arbitrage in the form of trading around quarter-end dates (see Munyan (2017), CGFS (2017), and Aldasoro et al (2018)). A consultative process proposing a globally uniform implementation of disclosure requirements has met with strong objections to the high process-related burden (see BCBS (2018)).

[12] Following FSB (2018a), a cryptoasset is a private asset that depends primarily on cryptography and distributed ledger or similar technology as part of its perceived or inherent value. An asset-backed token is a digital representation of an actual asset or revenue stream. The appendix contains a glossary of expressions used.

[13] See also Malinova and Park (2017), Lee (2015), and ECB (2019) for how DLT could further develop in financial markets.

For their part, smart contracts (as outlined in Szabo (1997)) could replace central securities depositories (CSDs). The latter are specialised financial organisations that hold securities so that ownership can be easily transferred through book entry. CSDs thus make electronic trading possible, by doing away with the transfer of physical certificates and by supporting trade automation. They also process dividend, interest and principal payments, as well as corporate actions including proxy voting. All these activities could, in principle, be automated via smart contracts.

Options and futures clearing houses could also be automated. These are financial organisations that clear payments and financial products (securities and derivatives), thus reducing the risk of non-payment or non-delivery of the financial products. In these instances, a DLT-based clearing system would have on its ledger both cash (for settlement) and the financial product, or operate via smart contracts that would connect a cash with a securities blockchain. As a smart contract can impose conditionality on both parties to a transaction (cash vs deliverables only), settlement risk would be eliminated.[14]

In the more distant future, exchanges or OTC markets for securities and derivatives might also become candidates for automation, and in particular less liquid OTC markets. In these markets, potential sellers are wary of disclosing their trading intentions in order to avoid driving prices against themselves. Instead of a standard open order book, they prefer to trade through a network of dealers/brokers, who in turn can rely on trusted contacts to execute trades with less effect on prices. A DLT-based version of such a market could automate the price discovery process via the demand and supply curves fed into the markets by the participants.

In all of these examples, it must be noted that the technology's primary advantage is automation (ie of the transaction process), thus reducing costs and settlement failure risks. But automation could also be provided through a centralised organisation, as currently. The advantage of DLT is that market participants could set up a market platform that would then function autonomously after the point of release.[15]

However, if such DLT-based markets were to develop, new ways of delivering data to financial supervisors and other stakeholders would open up. The key is that a DLT-based market already embodies all the relevant information, which supervisors could then readily access. Obviously, as financial firms will not want to reveal their trades, the ledgers would normally be encrypted. The compliance process would then consist essentially of determining which internal and external institutions could access which part of the underlying data and at what level of aggregation (see Graph 3).

As noted above, embedded supervision could monitor compliance with the Basel III capital standards.[16] The latter could be automatically verified by computing

[14] Of course, as is witnessed by recent flash-crashes, automation of financial trading can also create novel operational risks.

[15] When evaluating the economic case for DLT and its embedded supervision, therefore, the potentially cumbersome process of setting up a decentralised system must be compared with the cost of rent extraction by a monopolist providing a similar but centralised service, or the potential political economy dynamics if a government institution were to run the associated system. It is not a given that the decentralised system is more efficient (see Yermack (2017), Aste et al (2017), and Catalini and Gans (2017), and the formal analysis of Huberman et al (2017).

[16] See BCBS (2017) for a high-level overview of these standards.

the borrowing and lending balances and the associated risk weights in the relevant distributed ownership ledgers. Such calculations can apply not only to stock positions, eg end-of-reporting period compliance, but also be used for real-time sensitivity analysis of a balance sheet's exposure to market fluctuations, eg automated calculation of value-at-risk via simulation of ledger-based structured products and contractual obligations. In similar vein, also the full asset backing of an "on-chain" collateralised stablecoin could be automatically verified. For such a stablecoin, the value backing is provided by assets that are themselves traded on a distributed market. The coin itself is a smart contract that aims to create a stable value via financial engineering based on these underlying assets.[17]

Compliance process using embedded supervision — Graph 3

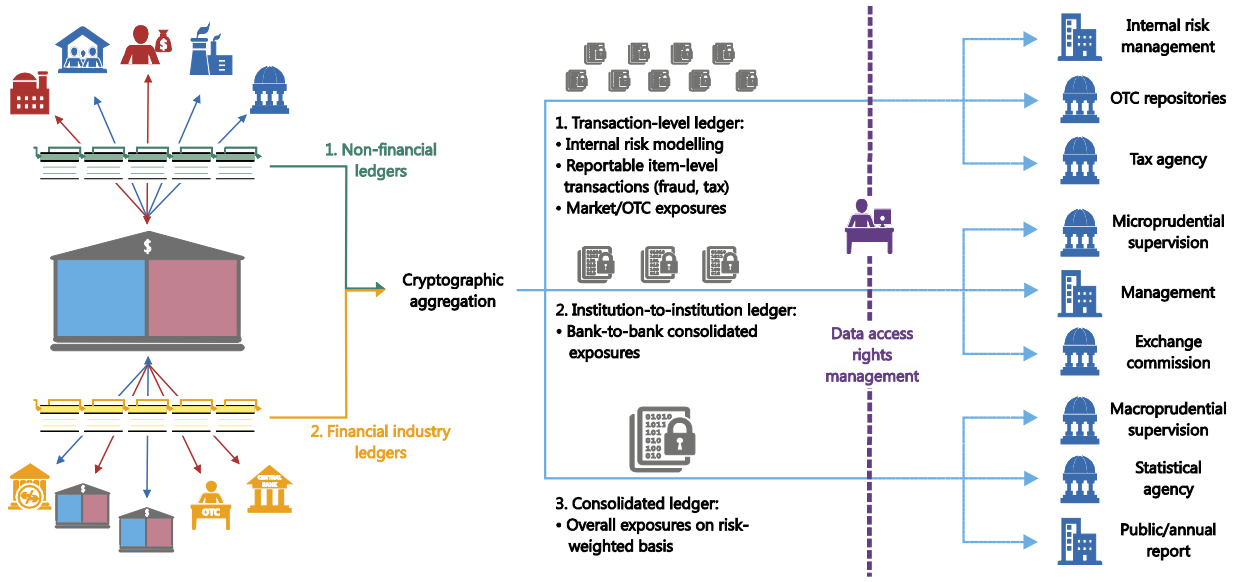

Embedded supervision can verify compliance with regulations by reading the distributed ledgers in both wholesale (symbolised by the green blockchain) and retail banking markets (symbolised by the yellow blockchain). Supervisors could access all transaction-level data. Alternatively, the use of smart contracts, Merkle trees, homomorphic encryption and other cryptographic tools might give supervisors verifiable access just to selected parts of such micro data, or relevant consolidated positions such as to institution-to-institution or sectoral exposures. Firms would only need to define the relevant access rights, obviating the need for them to collect, compile and deliver data.

Source: Author's elaboration.

Notably, this does not imply that data would need to be openly accessible,[18] nor does it mean that any supervisor would have access at an all-item level. Supervisors would gain access only to the relevant data, depending on whether they need transaction-level information or a more aggregated view. In this way, embedded supervision could help to maintain the confidentiality of firms and their customers, since cryptographic tools could be used to limit access only to selected parts of the underlying data or relevant aggregates. For example, information on aggregated

[17] See the stablecoin taxonomy of Bullmann et al (2019), and more generally the taxonomies of digital money of Bech and Garrat (2017) and Adrian and Mancini-Griffoli (2019).

[18] Due to coordination issues, it might not be socially desirable to make transaction data public (see in particular Morris and Shin (2018) for payments, Dang et al (2017) for banks, Goldstein and Leitner (2018) for stress tests and information disclosure, and the empirical treatise on bank transparency and deposit flows by Chen et al (2018)).

financial exposures could be disclosed, but without revealing the underlying micro data.[19]

Graph 4 summarises the elements of a blockchain adapted for embedded supervision. In the envisioned markets, the legal system would validate "oracles" – external reference points, such as ratings, on which the payoffs of certain ledger-based financial products may depend. Rating agencies, as well as land and other registries, and other external data providers would feed external data into a ledger, while decentralised exchanges would facilitate the trading of tokenised assets and financial contracts.

Example of a DLT-based market ledger providing for embedded supervision — Graph 4

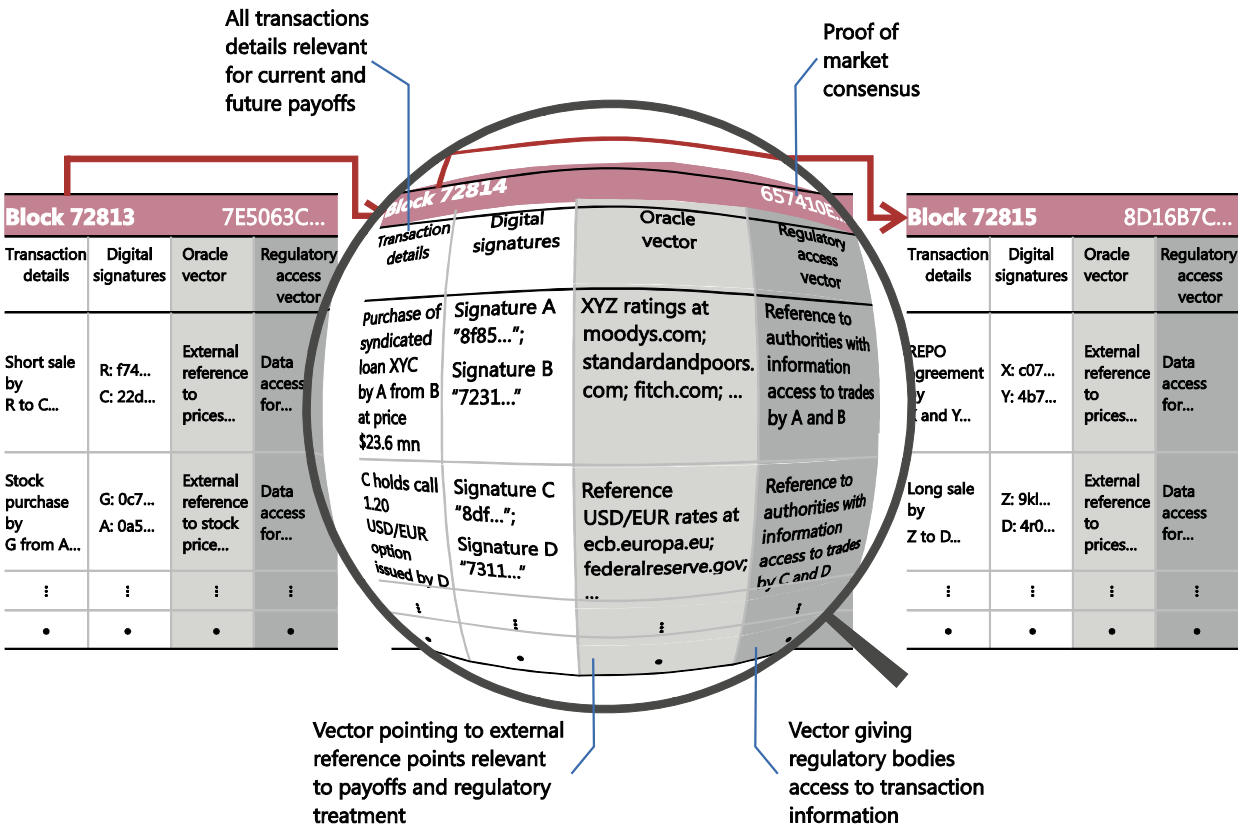

The blockchain records the history of transactions and contractual obligations, as well as links to relevant external information sources (oracles) and a vector giving regulatory bodies access to the information. Market participants transact on the blockchain, which records their transactions and obligations in the form of smart contracts. Payoffs of structured financial products, etc may depend on oracles, which are external reference points such as official interest rates, exchange rates, or market rates elsewhere. Supervisors in various jurisdictions have access to the (non-public) information in the ledger, can apply their regulatory model and may also specify circuit-breaker rules for the resultant payoffs.

Source: Author's elaboration.

It can also be surmised from Graph 4 that the first principle of embedded supervision calls for a proper understanding of what DLT-based trading can achieve, and what it cannot. As regulatory compliance only applies to regulated entities,

[19] The current state of technology supports this for transactions that encode pre-defined transfers of cryptoassets. In the field of computer science, it is an open question as to how far smart contracts can be used in such a setting. Of course, market participants must be obliged at all times to use transaction systems that allow access to supervisors.

embedded supervision can only be useful in the context of transactions involving regulated financial intermediators.

Embedded supervision would not relieve the boards and senior management of financial intermediaries of their responsibility to comply with regulation (see the United Kingdom's senior manager regime as detailed in Financial Conduct Authority (2017) for an example of current legislation). One aspect is that technology is fallible and management needs to know how adequately their institution is fulfilling the prudential requirements in order to correct any issues. But more fundamentally, even if a given institution is transacting on a DLT-based financial market, it might also be active on other non-DLT based markets, and reporting whether this is the case must be a management obligation.

Moreover, while DLT can evidence the transfer of ownership of asset-backed tokens from one entity to another, the connection between the underlying asset and the digital token must be underpinned by the legal system. For example, if there is a token-based bond of a specific company that itself does not participate in the distributed market, a legal institution must enforce the payment of interest and principal. Legal backing is also needed for the validity of oracles. For example, for the trading of a smart contract that reproduces the payoffs of an inflation-linked bond, the payouts depend on the inflation measure that is fed into the ledger from an external source. A final key element must be a watertight and potentially globally coordinated KYC identity framework that keeps illicit activity out of this novel ecosystem.

Embedding supervision in markets that achieve economic finality

Novel distributed markets can only be automatically supervised if transactions in these markets are final – the notion that the perceived balance of ownership that one communicates to the supervisor is immutable, or that "a transfer of funds [or] a transfer of securities that have become irrevocable and unconditional" (see CPSS (2003, p 496)). Traditional institution-based exchange is protected by the legal system: it is final by law and cannot be revoked.

Embedded compliance would replace this legal and institutionally based trust with a scheme by which the distributed market applies an economic incentive to achieve agreement (ie a consensus) on updates of the ledger (ie on transactions). The supervisor would then accept this consensus as valid if it can be proven to be irreversible. But what are the conditions for such irreversibility?

The remainder of this section exemplifies the conditions under which a supervisor can trust the information contained in a DLT-based market that lacks such a legal criterion. To this end, I examine under what conditions a distributed market functions, and derive the conditions under which the supervisor can trust the ledger's data. In doing so, I build on Auer (2019) and define a payment as final once it is certain that, from a given moment, it will never be profitable to undo the transaction via a double-spending attack.

Economic finality in a permissioned market with decentralised verification

In what follows below, I model a distributed market in which transactions are verified by third parties standing to lose a given amount (verifiers' "skin in the game", "stake", or "verification capital" in what follows below) should a blockchain reversal ever occur that voids existing transactions.

The model I describe is general. It could, from a technical perspective, be implemented in various ways. One is a permissioned DLT-based market with an overarching coordination mechanism that deters misbehaviour by ensuring that verifiers lose a given amount of deposited capital or pay a fine should they ever verify conflicting blocks.[20] Further in the future would be permissionless "proof-of-stake" consensus algorithms that tackle the so-called "nothing-at-stake" and "long-run attack" problems that plague current versions of this technology.[21]

The first theoretical result I derive is a sufficient condition for transactions in this market to be economically final. At each point in time, each block contains contracts that generate net transfers. This creates incentives for the party on the losing side of the contract to bribe verifiers into undoing the blockchain and voiding the contract. The total amount put up by the verifiers as skin in the game has to be high enough to deter this.

I show that economic finality in this market requires that the total amount of verification capital securing the block is higher than the maximum net transfers that could be generated by undoing the block in question. I derive this result by, first, showing that, while potential attackers could attempt to undo a chain of any length (ie undo only the last block, undo the last two blocks together, undo the last three blocks together etc), the most profitable attack strategy is to reverse only the last block. I then show that undoing the last block is never profitable if the maximum net transfers generated by the block in question are smaller than the total amount that is at stake for the verifiers, which is a sufficient condition such that there can be no coalition of losing parties who would find it profitable to bribe verifiers to undo the chain.

The second result concerns the supervisor's impact on the market. Embedded supervision is no free lunch, as the supervisor's actions might themselves strain market consensus. Axiomatically, regulatory compliance creates a financial burden for the regulated entities (for, if it did not, there would likely be no need to regulate the market). For example, if a market participant would like to finance the loss from a contract with debt, but minimum equity regulation were to bind, the cost of any loss

[20] In case of a permissioned model with fines, the market is similar to today's arrangements. One key difference is that fines would depend not on legal arguments but on cryptographic proofs of events or actions.

[21] See Kiayias et al (2017), Buterin and Griffith (2017) or Zamfir et al (2018)) for proposed implementations. The "nothing at stake problem" refers to the difficulty of deterring that verifiers can create multiple conflicting blockchain histories at no cost (see Poelstra (2014) for a description and Saleh (2018) for an economic analysis). "Long-range attacks" refers to early verifiers inventing alternative transaction histories long after having withdrawn their stake from the system. In particular, the latter problem might require some form of institutionalisation for successful proof-of-stake implementations.

created by the contract would be that of the marginal cost of equity, which many argue is higher than that of debt.

Thus, if market participants know that the market's data are being used to determine whether they are compliant with regulation, this would create incentives to fool the supervisor and undo the blockchain.

I show that there are two potential responses. If a supervisor wishes to monitor compliance in real time, they must mandate a commensurate increase in the total verification capital. However, an alternative strategy is to embed embedded supervision in the market equilibrium without requiring any additional verification capital, which is possible if the supervisor applies compliance with some time lag. The underlying intuition is that, while a competitive verification market will generically set a verification capital such that blockchain reversals are made marginally unprofitable at the time of signing, less capital is needed once a transaction is "buried" in the blockchain, ie once subsequent blocks have been verified and added to the chain. The supervisor can utilise the resultant residual verification capital in the deeper layers of the blockchain and read the market data with some lag without straining market consensus.

To establish these results, it is necessary to introduce some notation.

Notation: time t and block number b. In the environment described below, time is discrete and indexed by t, and one new block indexed by b is added to the blockchain in every time period.

Importantly, b is normalised so that it is equal to t. With this definition, block b is the one that was added to the chain at time t, and b-t corresponds to the time that has elapsed since block b was written into the blockchain.

Financial contracts and payoffs. In each block b, N_b market participants pay a fee π (solved for below), which gives them the right to sign a financial contract into the block. These contracts are indexed by i. b_i denotes the bock in which contract i is signed into the blockchain. After financial contract i is signed into block b_i, it generates a series of net payoffs for the involved parties. These financial contracts can be thought of as any type of financial transaction with uncertain future net payoffs. One example for such contracts is an American put option on a stock at strike price X issued by A and held by B. The net payoff to A is equal to the price of the put option when the contract is initially signed, equal to $min[0, -(X - Price)]$ when the stock matures, and 0 at any other point in time.

I assume that, before engaging in any transaction, market participants must hold on-ledger funds that are always sufficient to meet the contract's net payoffs directly on-chain (thus doing away with settlement risk).

While they remain on the blockchain, contracts generate losses or gains. I denote the payoff generated by contract i at point in time $t > b_i$ by $c_{i,t}$.[22] I am assuming that the total cumulative payoff (ie the net present value of all the payouts that the contract has generated in its lifetime or is expected to generate) at point in time t is distributed i.i.d. over time:

[22] If contract i is between A and B, the payoffs to B are the exact mirror image of A's payoffs. Without loss of generality, we thus focus on A's payoff.

$$c_{i,t>b_i} \begin{cases} \in [-\underline{c}, \underline{c}] \; if \; contract \; remains \; on \; the \; ledger \\ 0 \; if \; the \; contract \; has \; been \; netted \end{cases} \quad (1)$$

Contract netting. I allow market participants to enter and leave the market. If they leave the market, participants can cash out of the market, ie they settle their contracts using off-ledger funds: the party with a negative balance transfers off-chain funds to the party with positive funds, and the two parties then void the contract on the blockchain.[23] I am assuming that this happens at least every L blocks (ie all contracts are taken off the ledger after L blocks) and also, that in each block there is a share $1-\beta$ $(0 < 1-\beta < 1)$ of contracts that are netted early (again, netting means to net on-chain positions via off-chain payments and then void the contract on the chain). All contracts are hence netted if $t - b_i \geq$L. Before that point in time, a share of $\beta^{(t+1)-b_i}$ is not yet netted.

Transaction verification. Blocks of new contracts are signed into the blockchain by verifiers, which are third parties who stand to lose should the block they have verified be reversed at some future time. Since contracts can generate net on-ledger payoffs, the losing party has an incentive to undo the blockchain and with it the transaction. There is thus a need for some actors to verify the contracts that are written into the blockchain. These verifiers could also perform other actions, such as KYC/AML or other legal background checks.

I assume that verification happens at the block level, and that this is done by verifiers indexed by $v \in V$. For each block b, the system randomly assigns a sufficient number of verifiers and a pre-determined order in which they can verify blocks. Each verifier has a verification capital of s,[24] which can be interpreted as the actual capital at stake or as the expected cost of legal fines in the case of misbehaviour. The latter amount is equal to the amount the verifier stands to lose should they verify a block that later turns out to be invalidated (ie is not included in the blockchain the market coordinates on).[25]

At a point in time/block b, the selected set of validators have two options: verify block b with their verification capital, earning the fees, or not to verify the block and invest their verification capital elsewhere to earn a return of δ.

The fee income is split among the validators, each receiving a fee income of $f = N_b\pi/v_b$. The latter fee income is paid with a delay of L, ie once all contracts have been netted. Once a validator has used its verification capital to verify a block, the verification capital is blocked for L blocks until when it is released and can be used to place as verification capital again.

If, in the meantime, a blockchain emerges in which v has verified any different block, v's verification capital is lost (and so is the fee f). I further assume that market

[23] Technically, one important condition is that the netting of a contract must not be subject to a double-spending attack. One implementation could be that participants issue each other "proof-of-netting" receipts that can be attached to the blockchain at any point in time, ie they could be reinserted to the blockchain by either party in case of a blockchain reversal.

[24] The assumption that verifiers are homogenous is made for ease of exposition. The model at hand with v verifiers each posting a stake s is isomorphic to a model in with heterogeneous competitive verifiers whose total verification capital sums to vs.

[25] Technically, this can happen. The original blocks created by the attackers can simply be imported into the main chain as "proof-of-malfeasance".

participants know the set of verifiers of each block before they sign a transaction into the blockchain, and that market participants follow two rules.

If two or more rival blockchains emerge, market participants sign their blocks only into blocks added to the blockchain with the highest cumulative amount of verification capital.

Market participants sign their contracts into a block only if the verification capital in the respective block is sufficient to ensure that the blockchain will never be reversed.

Given rule (b), no rival blockchains will ever come into existence, but equilibrium still requires a statement on the assumed off-equilibrium behaviour of market participants.[26]

Validators are assumed to act selfishly, ie they can be bribed and will take part in a "blockchain history reversion" attack if they receive a bribe marginally larger than their verification capital s plus the income f they lose on the voided chain.[27] Let b denote the most recent block and assume that an adversary wants to undo a contract that is contained in block b-x. The adversary needs to bribe an amount such that the resultant chain has more verification capital, ie an amount larger than $\sum_{k=b-x}^{k=b}(s + f)v_b$.

Rule (b) hence requires the incentives of a potential attack on this market to be analysed. The gain from a potential attack to the attacker is the value of the contracts that are being undone. At any point in time, since any contract only generates a transfer between two parties, many agents will have contracts with a losing value. These losing parties might form a coalition and jointly pay to undo the blockchain. We thus need to define $\overline{C}_{b,t}$, the maximum gain from voiding block b in the chain at time t:

$$\overline{C}_{b,t} \equiv \sum_{i\in b} \Pi_{i,t} max\left[\left|c_{i,t}\right|\right] = \begin{cases} \beta^{(t+1)-b} \sum_{i\in b} max\left[\left|c_{i,t}\right|\right], & t-b < L \\ 0, & t-b \geq L \end{cases} \quad (2)$$

Picking the maximum among the absolute values of payoff realisations ($max\left[\left|c_{i,t}\right|\right]$) reflects the fact that either A or B could be the losing party of contract i, and the highest loss has to be considered.[28]

$\Pi_{i,t}$ is the indicator function, taking a value of 1 if the contract is still active, and 0 otherwise. For $t - b < L$, the latter happens with probability $\beta^{(t+1)-b}$, ie one period after block b is added to the blockchain only $\beta < 1$ of the contracts are still live.

With the above-assumed support of potential payouts (see equation (1)), it holds that $\overline{C}_{b,t} = \beta^{(t+1)-b} N_b \underline{c}$ as long as $t - b < L$. Armed with the maximum value that can be gained by undoing block b at point in time t, it is possible to derive the amount of

[26] Rule (a) is an equilibrium strategy (as it coordinates market participants on a common chain and nobody would want to transact on a chain that is not transacted on in the future). The game theoretical analysis of Biais et al (2017) is likely also to apply in the setting of the present paper, meaning that there could be different strategies that are also an equilibrium.

[27] Such bribery would take place via off-chain payments.

[28] In summing over all maximum losses, I am allowing for the possibility that contractual payoffs are perfectly correlated across contracts in a block, so that idiosyncratic large payoffs do not even out. If instead contract transfers is idiosyncratic, the right-hand side summation in (1) is over the mean absolute loss rather than the maximum.

verification capital that is high enough to guarantee that it certainly will not be profitable to undo the blockchain.

One necessary (but alone not sufficient) condition needed for economic finality is that block b will not be reversed at period b+1:

$$\beta N_b \underline{c} \leq v_b(s+f)$$

Where v_b is the number of verifiers guaranteeing block b. The condition that it must be unprofitable to reverse the latest block is not sufficient, however, as it must also hold that it is not profitable at point b to undo both block b and the previous block b–1. Generally, the necessary and sufficient condition is that it will be unprofitable to undo any attack of length x:

Certain economic finality. Transactions on the market can be considered final if

$$\mathbf{E}\left[\max_{x<L} \sum_{k=o}^{x} \overline{\boldsymbol{C}}_{\boldsymbol{b-k,t}} - \boldsymbol{v_{b-k}}(\boldsymbol{s}+\boldsymbol{f})\right] \leq \mathbf{0} \tag{3}$$

Equation (3) says that no strategy to undo only the last block, or the last two blocks, or the last three blocks, and so forth, can ever be profitable, even under the most adverse realisation of payoffs.

However, it is noteworthy that an induction argument shows that $\beta^2 N_b \underline{c} < v_{b-1}(s+f)$ and that, therefore, if one can assume that market participants one period previously have set the total verification capital high enough to ensure that under no circumstances will it be profitable to undo block $b-1$, it also holds that with $\beta N_b \underline{c} < S_b$, it will not be profitable to undo a chain of length 2, as

$$\begin{aligned}\left(v_{b-1}(s+f) - \beta^2 N_b \underline{c}\right) + \left(v_b(s+f) - \beta N_b \underline{c}\right) \\ > \left(v_{b-1}(s+f) - \beta N_b \underline{c}\right) + \left(v_b(s+f) - \beta N_b \underline{c}\right) > 0\end{aligned}$$

Further iteration of this argument shows that $\beta N_b \underline{c} \leq v_b(s+f)$ is not only necessary for the equilibrium, but also sufficient.

The described market can hence be viewed as final if v_b, the number of verifiers of block b is equal or exceeds $\frac{\boldsymbol{\beta N_b \underline{c}}}{(s+f)}$.

Free entry of verifiers and market equilibrium. Entry into the pool of potential verifiers is open to anyone, but requires the verification capital s to be locked in, which could otherwise be invested at rate δ per unit of time/block.

In equilibrium, it holds that

$$\boldsymbol{v_b} = \frac{\boldsymbol{\beta N_b \underline{c}}}{\boldsymbol{s(1+\delta)^L}} \tag{4}$$

$$\boldsymbol{\pi} = \boldsymbol{\beta \underline{c}(1-(1+\delta)^{-L})} \tag{5}$$

Proof: In equilibrium, transactions must be final and validators must break even. Given that the stake needs to be deposited for L periods, a potential verifier compares investing the amount s for L periods to receive the compounded return $s(1+\delta)^L$ with the alternative of receiving stake s plus fee income f at the end of the period. The free entry condition of potential verifiers thus implies $f + s = s(1+\delta)^L$. Combining the latter free entry with the finality condition (3) yields $v_b = \frac{\beta N_b \underline{c}}{s(1+\delta)^L}$ and in turn also solves for (5).

The equilibrium user fee π is independent of s, the amount that is deposited by each verifier. This is so as lower s means that a proportionally higher number of verifiers is needed to ensure finality, leaving the total amount of verification capital

and thus implied opportunity costs unchanged.[29] The equilibrium fee is also independent of N_b, the number of contracts that are written into the chain. This is so because an increase in the number of contracts requires a proportionally higher number of verifiers to ensure finality, leaving average cost per contract unchanged.

The equilibrium user fee π is proportional in $\underline{c}$, the upper bound on losses that needs to be deterred: higher potential losses require more verification capital, in turn leading to higher costs per contract. Fees also increase with the opportunity cost of verification capital δ, the length L for which such capital has to be locked in, and decrease in $(1-\beta)$, the share of contracts that expire early.

Residual verification capital buried in the ledger. An important insight is that, because the condition that the most recent block is not undone is the most stringent of the set of no-attack conditions in Equation (3), an excess of verification capital starts to build up in the ledger. The reason is that the verification capital is only freed after L periods, although a fraction of $1-\beta$ of the contracts are settled in each period. Therefore, if the current block is b, the free verification capital of block b–2 is equal to $N_b\beta\bar{c}(1-\beta)$. More generally, the residual verification capital in the chain from block b to b+t is equal to

$$\text{Residual verification capital}_{b,b+x} = \beta N_b \bar{c} \sum_{k=0}^{t}(1-\beta^k) \qquad (4)$$

Embedded supervision and economic finality

Equation (3) sets out the conditions for the market to be economically final if working on its own. But this does not automatically mean that the supervisor can trust the market's ledger: if market participants know that the supervisor will use information from the blockchain, this in itself might give market participants an incentive to report false information in order to fool the supervisor.

To model the supervisor's impact on the market, I assume that each contract, as long as it is live, generates an additional supervisory payoff $r_{i,t>b_i}$. The supervisory payoff can be thought of easily in the context of minimum equity regulation. If the contract is such that a negative payout for A is expected upon settlement, this reduces A's equity, which might bring it below the supervisor's mandatory target and necessitate a costly capital injection. If this is true, the regulatory cost is the marginal additional cost of equity over that of debt. For example, if the cost of debt is 2%, but the cost of equity is 6%, any loss c would cost $1.02c$ if the firm can finance itself with debt, but $1.06c$ if it is mandated to finance losses by raising additional equity. In this example, the net payoff would thus be $c_{i,t} = 1.02c$, while the regulatory payoff would be the additional cost $r_{i,t} = 0.04c$.

Going beyond the specific example, I assume that, much like the actual payoffs c, the regulatory payoffs r too are bounded by $\underline{r}$. Arguing along the above lines shows that, for the market to keep functioning (ie without any blockchain history reversals to a very high degree of certainty), with the supervisor applying compliance in this way, additional verification capital would be required equal to $\beta N_b \underline{r}$.

[29] This consideration also makes it clear that, while in the model at hand it is assumed that each verifier has to deposit a fixed amount s, as long as entry is contestable and no verifier has market power, the results are isomorphic to allowing for heterogeneous amounts of verification capital and fees that are proportional to a validator's verification capital.

Defining the maximum regulatory gain from voiding block b in the chain at time t by $\overline{R}_{b,t}$

$$\overline{R}_{b,t} \equiv \sum_{i \in b} \Pi_{i,t} max\left[\left|r_{i,t}\right|\right] = \begin{cases} \beta^{(t+1)-b} \sum_{i \in b} max\left[\left|r_{i,t}\right|\right], & t-b < L \\ 0, & t-b \geq L \end{cases}$$

Supervision-resistant economic finality. If the market is supervised via the distributed ledger's information, transactions in the market can be considered final if

$$\max_{x<L}\left[\sum_{k=o}^{x} \overline{C}_{b-k,t} + R_{b-k,t} - v_{b-k}s\right] \leq 0 \quad (5)$$

One solution for the supervisor is thus to mandate that the total verification capital satisfies $vs = \beta N_b(\underline{r} + \underline{c})$, but this necessitates a higher amount of verification capital, which is costly.

Equilibrium allowing for real-time embedded supervision: a market can be automatically supervised, with the regulator reading the ledger in real time, if the number of verifiers satisfies $v_b > \frac{\beta N_b(\underline{r}+\underline{c})}{s(1+\delta)^L}$. If the latter inequality binds, the fee π for a transaction is equal to $\beta(\underline{r} + \underline{c})(1 - (1+\delta)^{-L})$.

Another solution is possible, which is to lag the appliance of supervisory compliance. For example, assume that the supervisor sets equity requirements such that a firm's equity has to meet the block requirements one block removed. Then, it holds that residual capital equals:

$$\left(S_{b-1} - \beta^2 N_{b-1}(\underline{r} + \underline{c})\right) + \left(S_b - \beta N_b \underline{c}\right),$$

which can be positive if $\underline{r} < \underline{c}(\beta^{-1} - 1)$. More generally, consider a supervisor who allows regulatory requirements to be applied following an integer number of X blocks after the actual transaction block. The residual verification capital at lag x is equal to

$$\sum_{k=o}^{x} \beta^{b-k} N_{b-k} \underline{c} - S_{b-k}$$

The supervisor can thus apply embedded supervision, without mandating a higher verification capital, by applying compliance with a lag x, satisfying:

$$\sum_{k=o}^{x} \beta^{b-k} N_{b-k} \underline{c} - S_{b-k} > \beta N_b \underline{r}$$

Note that, in all existing regulation, data are delivered to the supervisor with a substantial lag. Here, the supervisor gets instant access to the data, but does not apply any supervisory measures until after a certain time lag has elapsed.

Operational aspects: harnessing the fintech opportunity

The above section explores the conditions under which a supervisor might take the data of a distributed ledger at face value. However, it is mute on why supervisors and regulators should actually invest in an infrastructure allowing for embedded supervision. Their goal is neither a specific market structure nor a specific form of exchange, but to assure high-quality, low-cost, and inclusive financial services, as well as a stable financial system. With these goals in mind, there are operational aspects to consider, aimed at bringing down the fixed and the marginal cost of doing business.

Bringing down the fixed cost of doing business is an important step towards levelling the playing field for small and large firms.[30] As a side effect of this focus on detailed regulation and supervision to tackle the risks of complex large financial intermediaries, supervisors may have created compliance costs that disproportionally affect smaller intermediaries (see Graph 5 and Philippon (2016)), thus favouring concentration.[31]

Smaller financial institutions are disproportionately affected by compliance costs

In per cent

Graph 5

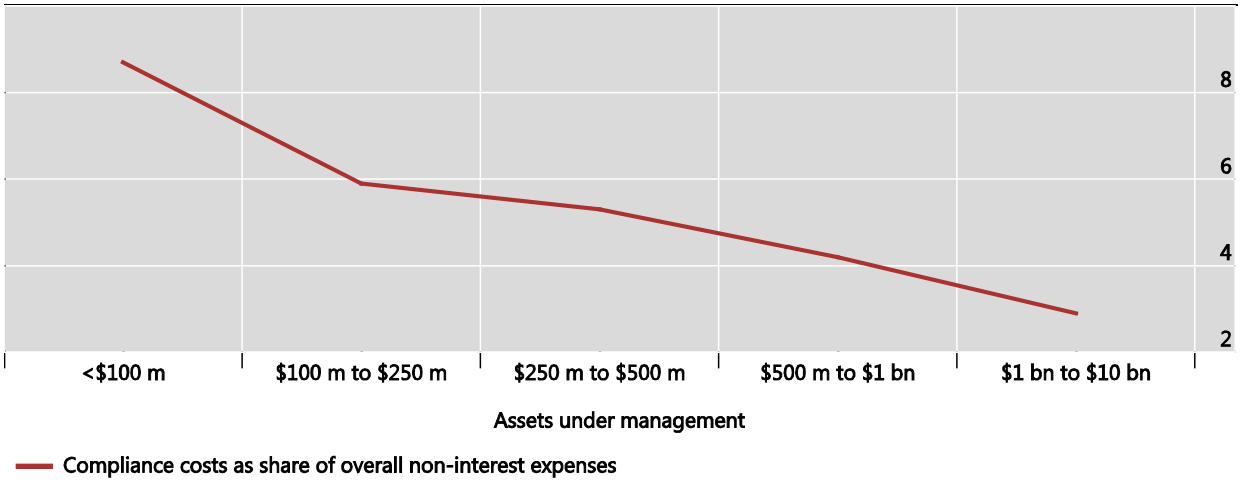

Source: Dahl et al (2016).

A first goal of embedded supervision should be to lower the fixed cost of compliance, thus levelling the playing field for large and small institutions.[32] One operational aspect is for regulators and supervisors to take an active role in the design of the market, in particular regarding standardisation of the database structure (by ensuring interoperability of various blockchains). Another one is developing a freely available open-source suite of monitoring tools to create clarity regarding how specific regulatory frameworks are applied in practice.

Efficient guidance of market standards to ensure contestability may also require adequate definitions of what it means to truly "decentralise" decision-making, risk-taking and system governance (see Buterin (2017) for a discussion and Walch (2017,

30 In particular, following the Great Financial Crisis, politicians, legislators and supervisors have focused on increasing the resilience of the financial system and, in particular, of the large banks that account for the bulk of total positions and thus aggregate risk (see G20 Leaders' Statement (2009)), an effort that is still ongoing (see Caruana (2016), BCBS (2017), FSB (2017), Carney (2017 and 2018), BIS (2018), and Carstens (2018e)).

31 It is well-established that despite ample technological progress, financial services remain expensive (see Philippon (2015), and Bazot (2018), as well as Graph 1 above).

32 Independent of technology, a regulator may want to offer differently tiered regulations that provide for lighter regimes for smaller ventures or for investors better equipped to analyse and deal with risk ("qualified investors"). For example, the US Securities and Exchange Commission (SEC) adopted Regulation Crowdfunding in 2015. ICO regulation could thus be seen as an "IPO light" for small, early-stage and risky ventures. Hohl et al (2018) offer a general review of how various jurisdictions apply the concept of proportionality to avoid excessive regulatory burden for smaller entities.

2019) for critical reviews).[33] Regulators and supervisors can steer some design elements of new decentralised markets, as they will set the market standards under which regulatory compliance can be automated.

A second operational goal is to reduce the marginal cost of doing business by facilitating access to trustworthy official information. One easily implementable aspect is for public authorities to directly offer digitally signed and time-stamped information that can be fed into relevant market ledgers. In many cases, financial contracts may reference data originating from the official sector, such as the central bank's policy rate or data releases by the national statistical officers. Moreover, in many jurisdictions, firm and land registries are operated by the government. Enabling low-cost tokenisation of the underlying firms and real estate will be facilitated if these registries make their information accessible in digitally signed, time-stamped and publicly available form.

A last operational aspect concerns the handling of disputes. Regulatory frameworks or standards could determine arbitration processes if information referenced in smart contracts turns out to be fraudulent. This could happen where the smart contract has a security vulnerability (which is frequently the case, see Luu et al (2016)) or in other unforeseen events such as if a smart contract is based on an interest rate benchmark that ceases to exist. Ultimately, the world is often too complex to be put into code, and the added value of decentralised automation has to be seen as simplifying the standard execution of a contract, while more complex cases might need to be handled via a legal procedure.

Conclusion

This paper has argued that supervisors might use DLT to efficiently supervise financial markets. The basic premise is that regulating blockchain-based finance should not require a departure from long-established principles on the regulation of specific economic activities. Rather, regulators and supervisors might consider investigating how their use of technology could evolve alongside that of the financial industry.

Embedded supervision is distinct from other forms of "suptech" or "regtech", which aim to use machine learning or artificial intelligence to more efficiently monitor the financial industry (see FSB (2017) and Broeders and Prenio (2018)).

The key principle of embedded supervision is to rely on the trust-creating mechanism of decentralised markets for regulatory purposes too. If DLT-based markets were to develop, this would change the way assets are traded and how they are packaged into complex financial products. Since the information contained in the blockchain is verified by decentralised economic consensus, it could replace current processes for data delivery and verification. In today's compliance process, the data's trustworthiness is guaranteed by the legal system, the relevant authorities and the

[33] Even with the most decentralised systems, many aspects of centralisation remain, for example when it comes to the evolution of the code (core developers etc). Further to this, as shown by the concentration of the mining power of all of the world's major cryptocurrencies in the hands of only a few companies or mining pools, even systems that are intended to be decentralised have a tendency to centralise due to unforeseen returns to scale. Regulators and supervisors could counteract this, for example, by setting standards that guide or encourage entry into the verification market.

threat of legal penalties. In DLT-based markets, by contrast, data credibility is assured by economic incentives. In this world, the supervisor must examine the conditions under which the market's economic consensus is strong enough to guarantee the quality of the data contained in the distributed ledger.

These considerations highlight the main legal challenge facing legislators, regulators and global standard-setting bodies. This challenge goes deeper than current discussions on under what circumstances cryptoassets should be considered as commodities, securities or other asset classes.[34] Rather, it is how to embed the concept of economic finality in today's legal system, and the adjacent question of how to treat such assets on balance sheets.[35] In most jurisdictions, the legal setup is such that a single and regulated clearing and settlement provider is required to verify that an irreversible transfer of ownership has occurred. DLT, however, achieves such a transfer via the economic incentives of verifiers rather than by the authority of a central institution. Only if the principles of finality underlying the regulation and supervision of financial markets infrastructures are modified to recognise decentralised exchange could DLT ever gain traction in regulated finance.[36] Along with this, regulators and supervisors would also have to design rules regarding the assignment of responsibility in decentralised markets in the case of illegal activity.

To implement embedded supervision, regulators would also be required to acquire substantial technological know-how and the willingness to adjust their operational approach to the technology that is being developed by the financial sector.

Around the globe, many supervisors are open to this possibility and some are already developing the requisite sandboxes. One example is "LBchain", the Bank of Lithuania's blockchain-based sandbox that seeks to embed a regulatory infrastructure in a DLT-based market. Another one is the Federal Reserve Bank of Boston's supervisory node case study.[37] The benefits might include lower costs for both market participants and supervisors, real-time monitoring, deeper insights into the use of internal models, and improved detection of potential window-dressing and other abuses. In this way, contrary to the current situation where cryptocurrencies threaten to undermine AML/KYC standards, efficient supervision could become a key use case for DLT.

[34] Financial Market Supervisory Authority (2018) is an early contribution. It sets out how the existing regulatory frameworks will be applied to new DLT-based financial products according to underlying economic activity. Further, see US CFTC (2015), Financial Market Supervisory Authority (2018), HM Treasury-Financial Conduct Authority-Bank of England Crypto-assets Taskforce (2018), Gensler (2018), Clayton (2018), or Pierce (2018). FSB (2019) surveys regulators and their mandates for cryptoasset regulation.

[35] See European Banking Authority (2019) and BCBS (2019) for current guidance on the treatment of cryptoassets on balance sheets. These do not discuss the explicitly discuss the notion of finality, but focus more generally on the great risk such investments carry.

[36] Indeed, the *Principles of Financial Market Infrastructures* (see CPMI-IOSCO (2012)) are intended to be neutral to the organisation and function of financial market infrastructures (see paragraph 1.9).

[37] See Adamonis (2019) and Federal Reserve Bank of Boston (2019), respectively. Amstand (2019) argues for the general development of coded regulation to more effectively supervise fintechs.

References

Adamonis, A (2019): "LBChain platform-service", presentation at Bank of Lithuania, 6 February, https://www.lb.lt/uploads/documents/files/LBChain-blockchain-centre-vilnius(1).pdf.

Admati, A, P DeMarzo, M Hellwig and P Peiderer (2011): "Fallacies, irrelevant facts, and myths in the discussion of capital regulation: Why bank equity is not expensive?", working paper, Stanford University.

Admati, A and M Hellwig (2013): *The bankers' new clothes: What's wrong with banking and what to do about it*, Princeton University Press. https://doi.org/10.1515/9781400851195

Adrian, T and T Mancini-Griffoli (2019): "The rise of digital currency", *IMF Fintech Note*, no 19/01.

Allen, F and D Gale (2000): "Financial contagion", *Journal of Political Economy*, no 108, pp 1–33. https://doi.org/10.1086/262109

Amstad, M (2019): "Regulating fintech: Ignore, duck type, or code", in A Fatás (ed), *The Economics of Fintech and Digital Currencies*, VoxEU Books, March.

Aldasoro, I, T Ehlers and E Eren (2018): "Business models and dollar funding of global banks", *BIS Working Papers*, no 708, March.

Aste, T, P Tasca and T Di Matteo (2017): "Blockchain technologies: the foreseeable impact on society and industry", *Computer*, vol 50, no 9, pp 18–28. https://doi.org/10.1109/mc.2017.3571064

Auer, R (2019): "Beyond the doomsday economics of 'proof-of-work' in cryptocurrencies", *BIS Working Papers*, no 765.

Auer, R and S Claessens (2018): "Regulating cryptocurrencies: Assessing market reactions", *BIS Quarterly Review*, September.

——— (2019): "Cryptocurrencies: Why not (to) regulate?", in A Fatás (ed), *The Economics of Fintech and Digital Currencies*, VoxEU Books, March.

Bank for International Settlements (2018): *Annual Economic Report*, June.

Basel Committee on Banking Supervision (2017a): "High-level summary of Basel III reforms", December, www.bis.org/bcbs/publ/d424_hlsummary.pdf.

——— (2017b), "Implementation of Basel standards – A report to G20 Leaders on implementation of the Basel III regulatory reforms", July 2017

——— (2018): "Statement on leverage ratio window-dressing behaviour", October, see www.bis.org/publ/bcbs_nl20.htm.

——— (2019): "Statement on crypto-assets" 13 March 2019

Bazot, G (2018): "Financial consumption and the cost of finance: Measuring financial efficiency in Europe (1950–2007)", *Journal of the European Economic Association*, vol 16, no 1, February. https://doi.org/10.1093/jeea/jvx008

Bech, M and R Garratt (2017): "Central bank cryptocurrencies", *BIS Quarterly Review*, September 2017, pp 55–70.

Benos, E, R Garratt and and P Gurrola-Perez (2017): "The economics of distributed ledger technology for securities settlement", *Bank of England Staff Working Papers*, no 670.

Biais, B, C Bisière, M Bouvard and C Casamatta (2017): "The blockchain folk theorem", Toulouse School of Economics, *TSE Working Papers*, no 17-817.

Bonneau, J (2016): "Why buy when you can rent?", *International Conference on Financial Cryptography and Data Security*, Springer. https://doi.org/10.1007/978-3-662-53357-4_2

Borio, C (2018): "On money, debt, trust and central banking", keynote speech at 36th Annual Monetary Conference, Cato Institute, 15 November, Washington DC.

Broeders, D and J Prenio (2018): "Innovative technology in financial supervision (suptech) – the experience of early users", *FSI Insights*, no 9, July.

Budish, E (2018): "The economic limits of bitcoin and the blockchain", *NBER Working Papers*, no 24717, June. https://doi.org/10.3386/w24717

Buterin, V (2017): "The meaning of decentralization", Medium.com, 6 February.

Buterin, V and V Griffith (2017): "Casper the friendly finality gadget" (submitted on 25 Oct 2017 (v1), last revised 22 November 2018 (this version, v3)), arXiv:1710.09437.

Carney, M (2018): "FSB Chair's letter to G20 finance ministers and central bank governors", 13 March.

Carstens, A (2018a): "Money in the digital age: what role for central banks?", lecture at the House of Finance, Goethe University, Frankfurt, 6 February.

——— (2018b): "Central banks and cryptocurrencies: guarding trust in a digital age", remarks at Brookings Institution, Washington DC, 17 April.

——— (2018c): "Technology is no substitute for trust", *Börsen-Zeitung*, 23 May.

——— (2018d): "Big tech in finance and new challenges for public policy", keynote address at the FT Banking Summit, London, 4 December 2018.

——— (2018e): "Ten years after the Great Financial Crisis – where do we stand?", lecture at the People's Bank of China, Beijing, 19 November.

Bullmann, D, J Klemm and A Pinna (2019): "In search of stability in crypto-assets: are stablecoins the solution?", *ECB Occasional Paper Series*, no 230, August.

Caruana, J (2016): "Post-crisis financial safety net framework: lessons, responses and remaining challenges", keynote address at the FSI-IADI Conference on "Bank resolution, crisis management and deposit insurance issues", Basel, 6 December.

Catalini, C and J Gans (2017): "Some simple economics of the blockchain", *MIT Sloan Research Papers*, no 5191-16.

Chen, Q, I Goldstein, Z Huang and R Vashishtha (2018): "Bank transparency and deposit flows", Duke University, working paper, https://ssrn.com/abstract=3212873.

Chiu, J and T Koeppl (2017): "The economics of cryptocurrencies–bitcoin and beyond", Economics Department, Queen's University, working paper, no 1389.

——— (2019): "Blockchain-based settlement for asset trading", *Review of Financial Studies*, vol 32, no 5, pp 1716–53. https://doi.org/10.1093/rfs/hhy122

Clayton, J (2018): "Chairman's testimony on virtual currencies: the roles of the SEC and CFTC", US Securities and Exchange Commission, February.

Committee on Payments and Market Infrastructures (2012): "Payment, clearing and settlement systems in the CPSS countries – Volume 2", *CPMI Papers*, no 105, November.

——— (2015): *Digital currencies*, November.

——— (2017): *Statistics on payment, clearing and settlement systems in the CPMI countries*, December.

Committee on Payments and Market Infrastructures and Markets Committee (2018): *Central bank digital currencies*, March.

Committee on Payments and Market Infrastructures and International Organization of Securities Commissions (2012): *CPMI-IOSCO Principles for financial market infrastructures*, April.

Committee on Payment and Settlement Systems (2003): *Payment and settlement systems in selected countries*, April.

Dahl, D, A Meyer and M Neely (2016): "Scale matters: community banks and compliance costs", Federal Reserve Bank of St Louis, *The Regional Economist*, July, www.stlouisfed.org/~/media/publications/regional-economist/2016/july/scale_matters.pdf.

Dang, T, G Gorton, B Holmström and G Ordoñez (2017): "Banks as secret keepers", *American Economic Review*, vol 107, no 4, pp 1005–29. https://doi.org/10.1257/aer.20140782

Darolles, S (2016): "The rise of fintechs and their regulation", Bank of France, *Financial Stability Review*, no 20, pp 85–92, April.

European Central Bank (2019): "Potential use cases for innovative technologies in securities post-trading", Advisory Group on Market Infrastructures for Securities and Collateral, January

Fanusie, Y and T Robinson (2018): "Bitcoin laundering: an analysis of illicit flows into digital currency services", Center on Sanctions and Illicit Finance memorandum, January.

Federal Reserve Bank of Boston (2019): *Beyond theory: getting practical with blockchain – Boston Fed learns by doing with blockchain technology*, February.

Financial Action Task Force (2015): *Guidance for a risk-based approach to virtual currencies*, June.

Financial Market Supervisory Authority (FINMA) (2018): Guidelines for enquiries regarding the regulatory framework for initial coin offerings (ICOs), 16 February.

Financial Stability Board (2017a): *Implementation and effects of the G20 Financial Regulatory reforms: 3rd Annual Report*, 3 July.

——— (2017b): *Artificial intelligence and machine learning in financial services*, 1 November.

——— (2018a): *Crypto-assets: report to the G20 on the work of the FSB and standard-setting bodies*, July.

——— (2018b): *Crypto-asset markets: potential channels for future financial stability implications*, October.

——— (2019): *Crypto-assets regulators directory*, April.

Foley, S, J Karlsen and T Putniņš (2018): "Sex, drugs, and bitcoin: how much illegal activity is financed through cryptocurrencies?", *Review of Financial Studies*, forthcoming. https://doi.org/10.1093/rfs/hhz015

Frost, J, L Gambacorta, Y Huang, H S Shin, and P Zbinden (2019): "BigTech and the changing structure of financial intermediation", *BIS Working Papers*, no 779.

Fröwis, M and R Böhme (2017): "In code we trust? Measuring the control flow immutability of all smart contracts deployed on Ethereum", in J Garcia-Alfaro, G Navarro-Arribas, H Hartenstein and J Herrera-Joancomartí (eds), *Data privacy management, cryptocurrencies and blockchain technology*, Springer, pp 357–72. https://doi.org/10.1007/978-3-319-67816-0_20

G20 Finance Ministers and Central Bank Governors (2018): Buenos Aires Summit communiqué, 19–20 March.

G20 Leaders' Statement (2009): Pittsburgh summit, 24–25 September, https://www.treasury.gov/resource-center/international/g7-g20/Documents/pittsburgh_summit_leaders_statement_250909.pdf.

Gensler, G (2018): "Remarks at the 2018 MIT Technology Review's Business of Blockchain conference", 23 April.

Goldstein, I and Y Leitner (2017): "Stress tests and information disclosure", *Federal Reserve Bank of Philadelphia Working Papers*, no 17-28.

Gorton, G and A Winton (2003): "Financial intermediation", in G Constantinides, M Harris and R Stulz (eds), *Handbook of the Economics of Finance: Corporate Finance*, Elsevier Science. https://doi.org/10.1016/S1574-0102(03)01012-4

Haber, S and S Stornetta (1991): "How to time stamp a digital document", in A Menezes and S Vanstone (eds), *Advances in Cryptology-CRYPTO' 90, Lecture Notes in Computer Science*, vol 537 Springer, pp 437–55. https://doi.org/10.1007/3-540-38424-3_32

Haldane, A, A Schubert and R Berner (2015): "Knowledge needed to prevent Lehman repeat", *Financial Times*, 14 January.

HM Treasury-Financial Conduct Authority-Bank of England Crypto-assets Taskforce (2018): "Crypto-assets Taskforce: final report", 30 July.

Hohl, S, M Sison, T Stastny and R Zamil (2018): "The Basel framework in 100 jurisdictions: implementation status and proportionality practices", FSI Insights, no 11, November.

Huberman, G, J Leshno and C Moellemi (2017): "Monopoly without a monopolist: an economic analysis of the Bitcoin payment system", *Columbia Business School Research Papers*, no 17-92.

Judmayer, A, N Stifter, P Schindler and E Weippl (2018): "Pitchforks in cryptocurrencies", in *Data privacy management, cryptocurrencies and blockchain technology*, Springer, pp 197–206. https://doi.org/10.1007/978-3-030-00305-0_15

Kiayias, A, A Russell, B David and R Oliynykov (2017): "Ouroboros: a provably secure proof-of-stake blockchain protocol", EPrint Archive.

Landau, J-P and A Genais (2018): *Les crypto-monnaies, rapport au Ministre de l'Économie et des Finances*, 4 July.

Lee, L (2016): "New kids on the blockchain: How bitcoin's technology could reinvent the stock market", *Hastings Business Law Journal*, vol 12, no 2.

Luu, L, D Chu, H Olickel, P Saxena and A Hobor, (2016): "Making smart contracts smarter", in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pp 254–69. https://doi.org/10.1145/2976749.2978309

Malinova, K and A Park (2017): "Market design with blockchain technology", https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2785626.

Mills, D, K Wang, B Malone, A Ravi, J Marquardt, C Chen, A Badev, T Brezinski, L Fahy, K Liao, V Kargenian, M Ellithorpe, W Ng and M Baird (2016): "Distributed ledger technology in payments, clearing, and settlement", Board of Governors of the Federal Reserve System, *Finance and Economics Discussion Series*, 2016-095. https://doi.org/10.17016/feds.2016.095

Morris, S and H S Shin (2018): "Distributed ledger technology and large value payments: a global game approach", mimeo, Princeton University, November.

Möser, M, R Böhme and D Breuker (2013): "An inquiry into money laundering tools in the Bitcoin ecosystem", in *Proceedings of the APWG eCrime Researchers Summit (ECRIME)*, San Francisco, pp 1–14. https://doi.org/10.1109/ecrs.2013.6805780

Munyan, B (2017): "Regulatory arbitrage in repo markets", *Office of Financial Research Working Papers*, no 15-22, June.

Nakamoto, S (2008): "Bitcoin: a peer-to-peer electronic cash system", white paper, https://bitcoin.org/bitcoin.pdf.

Open Banking Working Group (2018): "The Open Banking Standard", working paper, Open Data Institute, www.scribd.com/doc/298569302/The-Open-Banking-Standard.

Philippon, T (2015): "Has the US finance industry become less efficient? On the theory and measurement of financial intermediation", *American Economic Review*, vol 105, no 4, pp 1408–38. https://doi.org/10.1257/aer.20120578

——— (2016): "The fintech opportunity", mimeo, NYU Stern School of Business.

Peirce, H (2019): "Regulation: a view from inside the machine", remarks at the University of Missouri School of Law, 8 February.

Poelstra, A (2014): "Distributed consensus from proof of stake is impossible", May, https:// download.wpsoftware.net/bitcoin/old-pos.pdf.

Quintenz, B (2018): Remarks of Commissioner at the 38th Annual GITEX Technology Week Conference, 16 October, https://www.cftc.gov/PressRoom/SpeechesTestimony/opaquintenz16.

Ruttenberg, W and A Pinna (2016): "Distributed ledger technologies in securities post-trading – Revolution or evolution?", ECB Occasional Paper Series, no 172.

Saleh, F (2019): "Blockchain without waste: proof-of-stake", working paper, 27 February.

Szabo, N (1997): "Formalizing and securing relationships on public networks", *First Monday*, vol 2, no 9, September. https://doi.org/10.5210/fm.v2i9.548

Tapscott, D and A Tapscott (2017): "How blockchain will change organizations", *MIT Sloan Management Review*, vol 58, no 2.

Teutsch, J, S Jain and P Saxena (2016): "When cryptocurrencies mine their own business", *International Conference on Financial Cryptography and Data Security*, Springer. https://doi.org/10.1007/978-3-662-54970-4_29

US Commodity and Futures Trading Commission (2015): "CFTC orders Bitcoin options trading platform operator and its CEO to cease illegally offering Bitcoin options and to cease operating a facility for trading or processing of swaps without registering", 17 September.

Walch, A (2017): "Open-source operational risk: should public blockchains serve as financial market infrastructures?", in D Lee, K Chuen and R Deng (eds), *Handbook of blockchain, digital finance, and inclusion*, vol 2, Elsevier. https://doi.org/10.1016/b978-0-12-812282-2.00011-5

——— (2019): "Deconstructing 'decentralization': exploring the core claim of crypto systems", in C Brummer (ed), *Cryptoassets: legal and monetary perspectives*, Oxford University Press. https://doi.org/10.1093/oso/9780190077310.003.0003

Yermack, D (2015): "Corporate governance and blockchains", *NBER Working Papers*, no 21802. https://doi.org/10.3386/w21802

Zamfir, V, N Rush, A Asgaonkar and G Piliouras (2018): "Introducing the minimal CBC Casper family of consensus protocols", DRAFT v1.0, 5 November, Ethereum Research.

Appendix: glossary

Glossary — Table A1

Asset-backed token	A DLT-based digital representation of an actual real asset or revenue stream
Cryptoasset	A type of private digital asset that depends primarily on cryptography and distributed ledger or similar technology as part of their perceived or inherent value.
Cryptocurrency	A cryptoasset used exclusively/primarily for payments.
Double-spending	Strategy that consists of spending in one block and later undoing this by releasing a forged blockchain in which the transactions are erased. In blockchains based on proof-of-work, this requires short-term access to enough computational power to overwhelm the rest of a cryptocurrency's network of miners. In those based on proof-of-stake, this requires owning or bribing a majority of the staked resources.
Economic payment finality	Definition of payment finality in blockchain transactions developed in this paper. A cryptocurrency payment can be considered as final once it is certain that, from a certain moment of time onwards, it will never be profitable to undo the payment via a double-spending attack.
Miner	Class of agents, who update the blockchain via computational work, and in return receive block rewards and transaction fees when they add batches of valid transactions to the blockchain.
Proof-of-work	Mathematical evidence that a certain amount of computational work has been done, in turn calling for costly equipment and electricity use.
Proof-of-stake	A system in which coordination on blockchain updates is enforced by ensuring that transaction verifiers pledge their coin holdings as guarantees that their payment confirmations are accurate.
Protocol	The coded "laws" of a cryptocurrency. Set of rules that governs what constitutes a blockchain that is accepted by the network of users.

Sources: Financial Stability Board (2018a,b); Auer (2019).