

6

POLICY AND FINTECH, PART II: USE CASES

Oliver R. Goodenough, Mark Flood,
Matthew Reed, David L. Shrier,
Thomas Hardjono, and
Alex Pentland

6.1 SPECIFIC USE CASE I: TRANSACTION RECORDS AND TRADING MARKETS

Several efforts have been made to use new technologies to change how securities are issued and traded, but thus far these efforts have failed to deliver a sufficiently compelling value proposition to induce widespread change. With the rise of blockchain-based technologies, even the leaders of the world's largest incumbent stock exchanges are now acknowledging a threat to the status quo. This section examines a few critical issues related to technological innovation in the issuance and trading of securities and explores possible regulatory responses to these challenges.

6.1.1 Issuance and Trading in a Blockchain Context

Most stocks and bonds are issued as securities with known and recorded ownership. This is certainly the case since the Tax Equity and Fiscal Responsibility Act of 1982 (TEFRA) prohibited the issuance of US bearer bonds.¹ Security ownership recordation is, at its core, a process of recording a “fact” with distributed, shared agreement on its truth. Because a blockchain manages a consensus version of the truth, an appropriately designed blockchain could, in principle, be well suited to the tasks of securities transfer and ownership recording and

have the potential to make the process more accurate and efficient. In addition, the distributed nature of blockchain could create a greater sense of trust in the system, as any participant in the market can validate a transaction. Moreover, regulatory oversight would become easier because an irrevocable ledger is readily accessible. After issuance, much of the life of a security exists in secondary markets. Thus, demonstration of exclusive ownership and transfer becomes paramount, tasks for which cryptocurrency blockchains have shown capability. Later in the chapter we address the related issue of *identity*, which must also be solved for a recordation system to make sense.

Permissioned blockchains can solve for issues of identity of participants and exclusivity of ownership. The bitcoin blockchain (BCBC) protocol is less well suited to this purpose, as it strives to maintain the anonymity of participants in an effort to mimic old-fashioned cash (specie or paper currency) payments. Apropos of the TEFRA note above, ownership identification is required for numerous purposes, including property taxes and taxes on capital gains, so a permissioned/identified system will be needed. However, systems have already been proposed (by the Massachusetts Institute of Technology, among others) for privacy-protected, traceable transactions. Identity in these systems could be managed by a trusted third party, while the identity of a particular participant in a transaction could be cryptographically shielded. This system would allow for anonymous trading of beneficial ownership until the occurrence of an appropriately permissioned event (e.g., a warrant is issued by a duly recognized court of authority), at which point the guardian entity managing the identity could selectively release the required information. Similar anonymity requirements are commonplace in brokered financial markets, where the broker hides participants' identities (from each other) as a way to limit information asymmetries.

How do we decide who is "inside the wall"—that is, who gets to write blocks to the blockchain? Given that advance

knowledge of the index is valuable (tradable) information, who gets to read the blocks on the chain as the consensus is being formed? Some versions of permissioned blockchains allow a small set of trusted participants to trade with each other, akin to a private trading network. Yet, this would have an exclusionary effect on small investors. It is also possible to create a permissioned, public blockchain where only some have “write access” but anyone can “read” the transaction stream, and this may provide for the balance required between competing objectives.

What do we do in the case of errors of execution? The BCBC and most other blockchains do not have a convenient “undo” mechanism when mistakes are made. For example, a minor programming error forced Knight Capital to sell itself after losing \$440 million at a rate of \$10 million a minute.² To avoid such scenarios, one might impose stringent authentication of participants, but this too would deviate from the original BCBC protocol of user anonymity. In the case of error, it is possible to inject a “correcting entry,” but the counterparties would need to agree to this—if Snidely decided he liked his erroneous transfer, it would be difficult to undo absent a court order (and even then, that simply creates a legal claim). In cases where anonymity is less critical, other options are available. For example, institutional participants in wholesale funding markets typically have an “investigations” office for semiformal arbitration of mistakes and disagreements with regular counterparties. Blockchain technology does not preclude a similar approach.

The need for an “undo” tool is illuminated by a dramatic blockchain failure that occurred in June 2016. In this case, Ethereum was the context for an attack by a hacker using the *nom de fraude* “the Attacker.”³ This would-be bandit exploited a programming flaw in a digital currency fund—the decentralized autonomous organization (DAO)—to direct the transfer of 3.6 million ETH (then worth about \$53 million) into

his or her account. The cofounder of Ethereum countered by freezing the DAO tokens. The Attacker then added insult to injury by asserting, through a post on Pastebin, that he or she had a valid claim to the money, arguing that the record in the Ethereum chain was the only source for the title and that any attempt to change the record would be a breach of the rules. Even though the Attacker threatened to unleash lawyers on those seeking to correct the fraud, reverse hackers allied to the platform managed to recapture most of the funds.⁴

Resorting to the use of cybervigilantes to battle back from predatory exploitation of the system's architecture suggests that trusting in the technology is not a substitute for authoritative governance. Rather, creating some kind of reconciliation or correction capability looks increasingly to be a necessary element of a blockchain-based trading system. Such a capacity ultimately needs an adjudicator; in a traditional market or contract we look to choice of law and choice of forum provisions to set up the correction system. We can also specify *private* adjudication, through arbitration or some kind of market-specific committee. The choices are several; the need to have one in place is critical.

6.1.2 Settlement and Hypothecation

The settlement of a trade is an area currently burdened with several layers of a process. Much of this process predates the advent of electronic records and thus has the potential to be automated using blockchain technology. A good deal of securities settlement involves statements of ownership—of stocks, bonds, and so forth. This is broadly consistent with the original BCBC, which tracks uninterrupted ownership of specific coins through time. In part, this works because the individual coins are clearly defined and identified, and ownership is rivalrous. It makes sense that, at every instant, there is a one-to-one mapping between a coin and its owner, and that one should be able to track an individual coin's ownership relationships uninterrupted through time. Moreover, to the extent that a

registrar's blockchain uses a distributed ledger, the BCBC has a mining cost that can be calibrated to encourage truthful voting under a distributed consensus protocol. A number of blockchain variations would be capable of managing such a distributed ledger of ownership.

For some legal applications, a document's chain of custody is important. An analog is the chain of obligation for reused or rehypothecated collateral. Unmanaged rehypothecation chains—the Lehman collateral hairball—were an important factor in the September 2008 run. Collateral rehypothecation frequently occurs in bilateral over-the-counter (OTC) markets. A trusted, decentralized registration point for OTC collateral pledges could therefore be an especially valuable application for an appropriately managed blockchain. However, the ability for accurate identification and authentication (which permissioned blockchains can require but which are absent from the BCBC) would be crucial to make this work reliably.

Scalability is also a critical issue. The BCBC protocol manages a distributed transaction ledger, so the current state of an individual's "account" must be calculated by rolling forward all historical transactions. Because current inventories of cash and securities are key variables in the settlement process, this calculation would need to be performed often. It is not clear that this process will scale adequately, especially in equities markets, where high-frequency trading is dominant.⁵ Some solutions to blockchain scalability have suggested creating "sub-consensus" nodes that aggregate to a larger consensus, but this only exacerbates the coordination issue noted above. Financial reform legislation moved most derivatives to clearinghouses, in part so that the regulators could have access to comprehensive information from the clearinghouses. It is not clear how this can be accomplished by a distributed system without some sort of coordinated reporting mechanism.

Similarly, the BCBC protocol does not directly support fungibility of cash and securities, relying instead on a relatively clumsy process of excessive lump-sum transfer, followed by

a mining and return of “change” in the appropriate amount. This introduces a coordination burden to ensure that these two messages are recognized as countermanding components of the same legal transaction. In principle, this should be straightforward and feasible, but *practical experience* shows that financial markets cannot always keep related transactions aligned.⁶ Disputes should be expected in practice, and some type of dispute-resolution mechanism will be needed. It would be hypothetically possible to train a machine-mediated dispute resolution system to facilitate efficiency, but it may not be feasible (at least in the near term) to eliminate human intervention entirely.

We have seen occasional flash crashes (precipitously steep declines in market prices) in a range of markets using high-frequency algorithmic trading. In many cases, the trading venue intervened to clamp trading and cancel executed transactions. This involves the unilateral suspension of trading to stem further acceleration of losses and instability and subsequent reversal by a third party (the exchange) of “completed” legal agreements. Clearly, this is not an optimal state of affairs, but conditional on a flash crash, unilateral intervention to cancel contracts is preferable to most alternatives. However, this requires a trusted relationship outside of the relationship between the transacting parties themselves, and some form of effective delegated authority permitting the trading venue to act pursuant to a set of predetermined rules or with the ex post involvement of authorities. Blockchain technologies also have the “undo” issues cited previously. Blockchain was designed as an irrevocable ledger, so unwinding errors becomes cumbersome, to say the least.

6.1.3 Transaction Monitoring

A blockchain defines a consensus version of the truth. In practice, we should expect to see an ecosystem of many blockchains, large and small, defining various “local truths” for

specific communities and purposes. The movement to give blockchains legal standing as evidence in contract enforcement is progressing, with the Vermont statutes discussed in chapter 5 standing as early beacons. In such a world, it is inevitable that two competing blockchain systems will eventually announce conflicting versions of the “truth.” It is possible that the blockchain consensus mechanism itself will step in to harmonize the differences. However, this may not happen automatically, because the consensus preference in each community may be to tolerate the inconsistency. Once again, this creates a need for a reconciliation mechanism. Industry coordination efforts such as those by Hyperledger and Interledger need to take into account the nuances of financial-securities-specific implications, and/or a new coordinating action will need to be taken around securities transactions to allow for reconciliation.

Suppose that industry blockchains successfully record much of the low-level data validation work that is currently handled by traditional double-entry bookkeeping and back-office confirmation and reconciliation processes. These blockchains could then be central staging points for supervision by regulatory bodies, archival recording, and enforcement of market manipulation laws. This raises issues such as the following: how much access should systemic supervisors have to these details without intruding on individual privacy, under what circumstances should supervisors be allowed to escalate their access, and who gets to decide whether escalation is permitted.

6.2 SPECIFIC USE CASE II: IDENTITY, TRUST, AND DATA SECURITY IN A BLOCKCHAIN ENVIRONMENT

As noted elsewhere in this chapter, the financial services industry must be able to provide its core services to the rest of the economy, and the stability of the system needed to provide these services relies on trust in the system, resilience and reliability, and robustness to manipulation. Blockchain

technology alone promises much, but many use cases will require integration with other existing and emerging fintech components, for example, for enhanced identity management and information security.

6.2.1 Current Infrastructure

The infrastructure of our financial system relies on an often-disparate network of networks and systems that have been built up over time. These networks must identify instruments and parties precisely and unambiguously in order to move money or settle securities. One such network is the payment system managed by the Society for Worldwide Interbank Financial Transfers (SWIFT), which moves massive sums of money among major financial institutions, including central banks. Like so many infrastructures underpinning our financial system—whether they are designed to facilitate trade or investment, payment systems, or risk transfer platforms—the SWIFT system relies on precise, unambiguous, and robust-to-fraud identification of counterparties and transactors.

SWIFT uses the Business Identifier Code (BIC) for this purpose. But the financial system more broadly lacks an unambiguous and ubiquitous means of identity management. For decades, disparate identification systems built up in our financial markets. Vendors provided proprietary partial solutions, such as the Committee on Uniform Securities Identification Procedures (CUSIP) number, which identifies individual securities; the Dunn and Bradstreet DUNS number; and the Markit Red Code, which identifies reference entities in credit default swaps. These vended solutions, each of which is different and covers a portion of the world's financial market participants, have proved costly and to be of limited use outside of internal systems, because of intellectual property limitations.

In other markets, a superior product may emerge as an effective market standard. For financial entity identification, no superior product emerged organically, as parochial financial

interests encouraged the various proprietary systems to remain proprietary, with the identifiers' owners capturing most of the economic rents. Although all end users would benefit from a common, open identification standard, no individual player examining the costs and benefits could unilaterally justify the costs of creating such a global system. Moreover, because a global system would be a natural monopoly with large network externalities, individual players would be unlikely to cooperate without external compulsion. A recent effort of authorities from around the globe has sought to solve this problem by creating the Legal Entity Identifier (LEI), which has now been adopted by over 1.5 million entities in 195 countries.⁷ This development is a good example of a productive use of the convening power of government.

In the case of these identity management systems, whether a closed, proprietary system (such as the BIC) or an open system (such as the LEI), the ability to rely on the identifier is paramount. Market participants and the authorities that oversee them must know "who is who" in our markets. Also critical is the need for identification systems that interoperate, especially as markets become more interconnected. This interoperability will aid oversight and risk management systems by supporting aggregation and netting of exposures, reducing opportunities for error or fraud, and generally improving confidence in markets and other infrastructures. But how does this confidence arise? How do we become confident that we are dealing with the intended "whom"?

Different solutions are available for this problem. One approach is self-identification. This may work well where the incentives to self-identify accurately are high and the costs of remedy low. Self-identification may also work where a counterparty has the opportunity to conduct due diligence to manage the risk of misidentification (or other counterparty attributes, such as credit risk). This approach may be sufficient where the costs of remedy (such as expensive litigation) are high. Likewise,

self-identification may also work in a closed system with repeat transactions, where the incentives to self-identify accurately and the costs of inaccurate identification are sufficiently high.

But self-identification is also often coupled with a trusted source. That source could come from the government (a driver's license or Social Security card) or an institution (such as the SWIFT system, a credit provider with skin in the game, or perhaps a distributed system of verification). A trusted source of entity verification can be a proxy for the expensive due diligence regime described above. In practical cases, the level of assurance necessary, given the risks and costs of remediation, is measured against the level of proof of trust provided.

In a distributed system, what would be the source of proof, and would it be sufficient given the risks and costs of remediation? Again, different approaches are available, including limited blockchain networks (much like the SWIFT network), where some level of proof is a condition of entry into the system. Who should be the gatekeeper for such a system, and how can we establish trust in this gatekeeper? An analogy is the registration system our markets currently deploy, where exchanges and other market infrastructures must meet regulators' quality and fairness conditions before gaining registration authority. Oversight, via inspections, examinations, heightened sanctions regimes, and books and records requirements, ensures both ongoing compliance with the trust rules and confidence in the system.

In other regimes, economic forces and liability regimes support the trust infrastructure. Dark pools, where closed networks of large traders operate in an opaque-by-design framework that protects anonymity, provide an example. The dark pool's host has sufficient regulatory and legal exposure to give participants confidence in posted bids and offers, despite not knowing their counterparties' identities. Blockchain implementations could use similar mechanisms to manage counterparty identification.

What about more open systems that are not centrally managed or regulated but instead operate as peer-to-peer (P2P) networks—as many blockchain advocates envision? Such systems can rely on self-identification but counterparties will typically want additional assurances. The distributed network might provide proof of identity, perhaps using a hierarchical system of agreed-upon, standard, third-party, proof-checking sources. In such cases, self-identification is possible, but it presents problems ranging from simple fraud exposure to tactics for concealing beneficial ownership to avoid market shifts. These forces can significantly complicate the jobs of risk managers and regulators.

At a minimum, authorities and courts will require a blockchain platform operating in a regulated market to be able to reveal the identity of “anonymous” participants in certain circumstances. The technology must provide this flexibility of counterparties’ ensured legitimacy, even in an anonymous transaction, coupled with the ability to reveal the identity of participants. This naturally raises the question of who has permission to unlock this information, under what circumstances, and how authorities (courts, regulators, or self-regulatory institutions) can acquire both the keys and certainty that those keys will work. All of this also runs against certain expectations for “deputized” law enforcement agents—the banks and other intermediaries that are required to file “suspicious activity reports” with the Treasury—if they must rely on identity information from a third party instead of the agent itself, which bears the regulatory burden. Without careful design, the result could be further “derisking” by banks that find it cheaper to avoid certain clients than bear the risk that they might be associated with bad actors.

Understanding networks of counterparty connections is critical for counterparty and systemic risk management, for supervision and market monitoring, and to value investments. These interconnections can explain the propagation of risks

and reveal structural points of risk concentration and potential failure. A market crisis—for example, in the overnight funding market—might affect one firm directly but affect its counterparties indirectly. These networks of exposure are critical for financial stability oversight.

Network data—linking identified entities to their corporate sisters through regulatory filings, or linking them to counterparties through (traded or illiquid) instruments—can address questions of “who owns whom” and “who owns what.” Fintech platforms, such as blockchain storage or smart-contract representations of derivatives agreements, may be natural technology solutions in this context. Just as entity identification requires identification standards and governance frameworks, representations of counterparty networks will need common standards and governance mechanisms in order to provide certainty for both firms and supervisors that the network data reveals precise and actionable information.

These arrangements have numerous complexities. For example, ownership may derive from clear language in a regulatory filing or corporate ownership agreement, but it could also exist by virtue of a springing interest triggered, for example, by external factors or demands, or a minority shareholder interest that regulators deem to be a controlling interest.⁸ If a public blockchain is the storage platform for counterparty network data, that system would need to be dynamic and might also need to support selective revelation of network details to appropriate network participants or regulators. These complexities reveal difficult questions of governance, which authorities will likely translate into formal system requirements.

6.2.2 Security within Blockchains: Regulatory and Market Needs

Just as market participants and supervisors need assurance that the information in the financial system is accurate and actionable, they also require that the information technologies that

undergird the system be resilient and secure. US financial markets have a strong reputation for fairness and transparency, but even these markets sometimes suffer from concerns about unequal access and treatment. Blockchain-based systems will need to provide trustworthy security and resilience, but the built-in redundancy of distributed ledgers, properly managed, might offer improvements over existing infrastructures that suffer from single points of failure and “weakest link” problems.

The 2016 penetration of the SWIFT interbank payment system highlights the hazards of potential weak links. According to press reports, hackers penetrated SWIFT’s client software, diverting money from the Bangladesh Bank’s (the central bank) account at the New York Federal Reserve Bank to accounts in Sri Lanka and the Philippines. The heist attempted to transfer almost \$1 billion, but the New York Fed (alerted by a spelling error) thwarted most of these transfers. The hackers ended up stealing over \$100 million, much of which was subsequently laundered through casinos in the Philippines. SWIFT and Bangladesh Bank have since recovered much of the stolen money, but a significant amount is believed to be missing. In addition to stealing huge sums, cyberfraud also potentially threatens financial stability by endangering the credibility of the backbone of our financial system. Fortunately, both the three thousand institutions that own SWIFT and its eleven thousand users have strong motivation to prevent further damage. SWIFT is a closed system, designed to be secure. But the weak point was the Bangladesh Bank’s SWIFT software portal, which the fraudster(s) compromised with the use of stolen credentials.

Resilience to cyberattacks is particularly important for blockchain platforms, especially public blockchains, which must limit the ability of one node to disrupt the whole system, through penetration, malware, denials of service, and other means.

A registration regime, such as a permissioned blockchain, may provide assurance of cyberattack resilience to participants,

but regulators will also require assurance, as well as full access to the critical nodes of the system. Securities and Exchange Commission disclosure laws require that public companies disclose material risks. Annual reports are now full of discussions of the impact of cyberthreats, although a firm's need to address ongoing threats may legitimately weaken such disclosures. Nonpublic firms do not have such disclosure requirements.

Public blockchains may also face free-rider challenges. For a blockchain system under attack, who has the necessary authority, responsibility, and access to fight back? The law may authorize a regulator to intervene, but this authority may be inadequate if parts of the distributed network reside outside regulatory reach. Blockchain participants themselves may welcome supervision and a demonstration of adherence to understood rules, as a form of public assurance that the system is both safe and fair. This puts more pressure on authorities to understand fintech innovations and their proponents, and to provide clear guidance on security and resilience.

6.2.3 Identity within Blockchains: Current Limitations and Possible Solutions

Managing digital identities is a challenge, especially for public (permissionless) blockchains such as the bitcoin system and blockchain. Bitcoin uses *self-asserted* identities, meaning that any participant can simply create a public key pair and join the blockchain pseudonymously. Self-asserted identity is partly a way for individuals to assert control and maintain data privacy in an increasingly connected world.

However, self-asserted identity has an inherent limitation in terms of scalability. This limitation is not unique to the BCBC; it is also present, for example, in the PGP (Pretty Good Privacy) system where users self-issue PGP key pairs.⁹ Key holders (PGP users) must provide their PGP public key directly to friends and colleagues, either in person or through a public "key ownership declaration" event, such as the "PGP

key-signing parties” at face-to-face meetings of the Internet Engineering Task Force.

Most current self-asserted digital identities (in the form of self-generated public key pairs) do not scale well, because they lack integration with existing digital and real-world infrastructures. A complete and scalable identity management system needs to ground identity in the physical world, and it should not rely solely and unconditionally on existing identity/service providers.¹⁰ We believe a new model is needed for “self-sourced identities” that would preserve privacy while ensuring scalability at the global internet level. Specifically, a scalable identity model must allow entities in the ecosystem to (1) verify the “quality” of an identity, (2) assess the relative “freedom” or independence of an identity from any given authority (e.g., government, businesses), and (3) assess the source of trust for a digital identity.

If anonymity is a requirement for self-asserted identities, then self-issuance of a public key pair (the case in the BCBC) is inadequate. True anonymity requires that identity be unlinkable across transactions, to prevent identity leakage through correlation attacks. Even if a digital identity is anonymous and unlinkable, counterparties must also still accept that identity. Parties relying on an anonymous, self-asserted identity will still need to assess its provenance and source of trust. To meet such requirements, a future self-asserted identity system would need to incorporate the notion of the varying degree of quality of the identity as a function of the veracity of the underlying provenance information (i.e., source of trust).

6.2.4 Digital Identities and Attributes for Future Blockchain Systems

Several avenues for scalable identity management and federation are available. These new approaches may require the introduction of new blockchain technologies and components, including new remuneration models for participants

in the ecosystem, as well as more efficient proof-of-work schemes. New identity technologies include the following:

1. *Verifiable pseudonymous identities and attributes*: Anonymous and verifiable identities have been a topic of research for over two decades.¹¹ Some of these schemes have been implemented in systems such as U-Prove and IdeMix, and some limited deployments have been carried out.¹² These proposed systems have not seen broad deployment on the internet because of a number of constraints (e.g., lack of use case or business model). The arrival of the bitcoin system and the potential of new forms of blockchain-based systems may provide use cases of the deployment of these existing anonymous and verifiable identities.
2. *Smart contracts for binding and revealing attributes*: A node on the blockchain P2P network can compute smart contracts, which are sequences of computations that map to legal agreements (e.g., between two transacting entities). The same computation model might bind attributes—regarding a pseudonymous identity—to a contract that names the pseudonymous identity. The computation could begin with attributes that are “blinded” and then subsequently release one or more of the attributes during the multiround smart-contract exchange protocol. The multiround negotiations would build toward a release of all the relevant attributes regarding both sides of the transaction. An example of such contracts is bidding at auctions, which could start with an anonymous or pseudonymous buyer/bidder accompanied by attributes of the buyer (e.g., buyer financial worth, history of bidding).

One innovation that could contribute to reliable identity management in future blockchain systems is data-driven distributed computation to derive attributes. In this approach, the P2P distributed nodes on a blockchain would each collect data regarding an identity and perform analytics based

on the data available to each node. Each node would first arrive at a “subattribute” value or parameter (e.g., single credit score) independently of the other nodes. Collectively the nodes would then contribute their respective subattributes to a group computation process, such as a multiparty computation (MPC) algorithm, which would result in a complete attribute. If a privacy-preserving algorithm is used for the MPC, this will have the added benefit of no single node knowing the subattributes of the other nodes.¹³ Proposed solutions, such as Enigma, can provide a foundation for deriving attributes.¹⁴ Attribute derivation, of course, would need to be reversible, for the reasons articulated above.

Another opportunity lies in the legal aspects of identities and attributes on the blockchain—namely, the introduction of a legal trust framework that uses automated contracts exchange (smart contracts or otherwise) to reduce friction in using digital identities through the blockchain. Such frameworks aim to reduce risks and liabilities of entities in the ecosystem through a set of agreed principles, operating rules, and mechanisms for legal recourse. Legal trust frameworks are crucial to the acceptance of digital identities and attributes in the real world. Some examples of legal trust frameworks for identity management and federation are the Federal Identity, Credential, and Access Management program, OpenID Exchange, and Safe-BioPharma.¹⁵

6.2.5 Scalable Digital Identities: Addressability, Source of Trust, and Verifiable Attributes

A number of desirable characteristics for digital identities can help guide future blockchain innovation around identity management. If satisfied, these characteristics could meet many of the regulatory needs described above.

1. *Addressability*: The notion of addressability refers not only to the uniqueness of an identity string at a global scale but also to any semantics embedded within the identity

structure that make it usable in practice. For example, the current “email identity” consists of a name (unique within the namespace of the domain) followed by a fully qualified domain name. These semantics enable the Simple Mail Transfer Protocol and Post Office Protocol v3 to interpret the identifier as a routable email address.

2. *Source of trust*: The source of trust of an identity (within a namespace) is derived or bequeathed from the authoritative entity that owns and/or manages that namespace. In essence, the source of trust vouches for the existence of the named identity and is associated with an individual or entity. For example, the Social Security number (SSN) of a US citizen is bequeathed by the government as the authoritative entity governing the SSN namespace. An email service provider vouches for an email address as the legal entity that owns the corresponding domain namespace. A public key infrastructure (PKI) service provider legally signs and issues a digital certificate to a user under the PKI provider’s authority as specified in the service contract, commonly referred to as the certificate practices statement. A source of trust could also be established through recurring interactions with other trusted sources, such as banks that have performed “know your customer” on clients and then issued credit cards. A source of trust is crucial for the legal acceptability of the digital identity within online transactions, especially when transactions cross boundaries between the digital world and the real world. A formal legal trust framework typically expresses the legal aspects of identity management within a given application ecosystem. The trust framework is a set of legal processes and operating rules for the issuance, management, and accreditation of identities and identity providers within that ecosystem.
3. *Verifiable attributes*: Related to an identity’s source of trust is the source of trust or *attribute authority* for attributes

associated with the digital identity. The attribute authority “binds” (often cryptographically) an attribute to a digital identity. In many instances, an attribute may have several authoritative sources, each making assertions with different degrees of veracity. The *relying party* in a transaction is the entity that ultimately decides whether to accept or reject a given assertion regarding an identity. For example, a state government in the United States may be an attribute authority for a person’s residency status (e.g., “Joe is a legal resident of Massachusetts”). A private banking consortium may be the attribute authority for some financial information regarding an identity (“Joe has a FICO score above 500 pts”).

4. *Privacy preserving*: In many cases, a digital identity scheme should preserve the privacy of its owner. Currently, identities in the form of email addresses are designed for addressability at the expense of privacy. Self-asserted identities in a bitcoin or PGP system provide a degree of privacy, but at the expense of scalability. New paradigms, such as the ability to create an opt-in anonymous identity on a permissioned blockchain, with anonymous verification and the ability of a regulator following due process to reveal identity, help bridge these schemes.

6.3 SPECIFIC USE CASE III: BLOCKCHAIN, STABILITY, AND SYSTEMIC OVERSIGHT

Our discussions so far have focused on the needs of authorities and market participants in their direct activities, and on rules that can help facilitate the adoption of fintech platforms, such as blockchains, in a productive and protective way. We argued in the introduction above that fintech innovations are new technologies emerging to address the exponential big data challenges in financial services. These forces are especially relevant for systemic oversight to monitor financial stability,

where the full financial system is in scope and the implications of systemic externalities can be severe.

It is important that these technologies and the legal framework be adaptable to permit authorities, market participants, and the courts to address or avoid failures during periods of stress to the stability of the system. Big data challenges are impinging on systemic financial stability monitoring at every point of the data life cycle.¹⁶

For example, latency reductions in high-frequency trading (HFT) continue to increase both the velocity of trading messages generated and the ability of systems to process those messages. HFT algorithms are themselves a form of fintech, but their intensive focus on latency minimization seems to leave little room for time-consuming proof-of-work processes associated with many blockchains. Clearinghouses for HFT venues will need to balance complex trade-offs among competing priorities—for example, between the scope of access to market data and its staleness. The paperwork crisis of the 1960s led to a dramatic redesign of clearing and settlement processes, and a similar rethink may be in store now.

More generally, supervisory solutions include suitable points of entry for policy tools such as liquidity infusions, protocols to coordinate industry support in a crisis to backstop or restart market activities, and the possibility to unwind transactions and resolve failed institutions, markets, or processes.

We focus on two broad challenges: (1) the abstraction, evaluation, and analysis of data from a systemic perspective and (2) mechanisms for crisp and effective systemic intervention, particularly in times of stress or crisis. In traditional financial markets and transactions, these two functions have been the domain of a small group of regulatory bodies, such as central banks and analytical teams like the Office of Financial Research (OFR).

Particularly in the domain of cryptocurrencies, concerns are emerging about their potential to disrupt the traditional

monetary and financial stability policies of central banks. To date, blockchain-based digital currencies have operated as small-scale payments media without the backing of central banks. They thus lack the traditional monetary governance mechanisms to achieve inflation or unemployment targets or to stabilize financial systems. Perhaps because of this, together with the absence of adjustments such as legal tender recognition, no blockchain currency has thus far grown big enough to present stability concerns in a major economy. Nonetheless, central banks and others are paying attention to this issue.¹⁷

We address the intervention question first, before turning to the more complex issue of data.

6.3.1 Intervention Mechanisms

It is useful to consider financial crises, as these episodes are likely to expose important operational bottlenecks and gaps in supervisory authority. In a crisis, official intervention can be necessary, for example, to replace vanishing private-sector liquidity, to clamp the runaway feedback of a flash crash, or to otherwise mitigate the potential long-term harms of a short-term breakdown.

Liquidity intervention by central banks is a critical tool for crisis management, benefiting from many decades of hard-earned experience. In a liquidity crisis, central banks provide funding to critical nodes in the system, typically by lending against good collateral. It is unclear how best to translate such a protocol to a possible future world in which the payment system is centered on a cryptocurrency founded on a decentralized blockchain. For example, would a legal-tender fiat currency, managed by a central bank, be a useful recourse for a panicked flight to quality that might ensue in a crisis? Would authorities or market participants be able to restart private payment systems that are suffering from technical failures or a loss of confidence? Might a fully digital system, particularly one not linked to a fiat currency but nevertheless universally

accepted, provide an accelerant to a crisis? And could authorities intervene to tamp down any such accelerants? Would holders of a cryptocurrency further devalue that currency in a flight to a perceived more stable fiat currency, particularly where a government was injecting cash into the system to stabilize it? Answering such questions is an important (but open) research challenge.

In another context, flash crashes and associated trading halts can lead to voided transactions as trading venues triage their systems to limit the damage. But such transactions, particularly in derivatives markets, may have been intended as critical hedges for one or more counterparties. Voiding such transactions may thus expose large risks by removing a hedge in the midst of a very volatile market. Should markets coordinate transaction cancellations across trading venues? If so, who should make this decision, and how should coordination occur? For example, if each exchange clearinghouse uses a local blockchain as its system of record, what are the standards and procedures for integrating data across blockchains in a crisis? How should one resolve disputes regarding information recorded on the blockchain system(s) of record, noting that time for resolution may be very limited?

Market participants may require that the unwind protocol be incorporated in the same blockchain. But this may create a central point of failure. Technical issues (e.g., a denial-of-service attack or a mutating computer virus) may erode confidence in the blockchain and the market that relies on it. Again, what should be the protocol for restarting such a system and rebuilding confidence? What should be the entry points for human and/or supervisory intervention? It is important to clarify these issues in advance, thus avoiding paralysis or power grabs in a crisis.

Once again, answering such questions is an open research challenge. While these concerns exist for all aspects of business and regulation, they are particularly critical for financial

stability in the context of resolving firms, central counterparties, and financial market utilities. Should a critical node of the system fail, the existence of immutable blockchain transactions or contracts could hinder or thwart the ability of supervisors to resolve institutions, in much the same way that the International Swaps and Derivatives Association's master agreements cross-default provisions hindered the ability to transition swaps from failing or failed counterparties.¹⁸

6.3.2 Data: Extraction, Evaluation, and Analysis

Since 2008 there has been increased attention to the cumulative, systemic effects of individual actions in the financial markets. Individual actions that are acceptable assumptions of risk or failures of payment in a single transaction can, if taken in the aggregate, create the equivalent of a heart attack in the system as a whole. Any attempt to understand, anticipate, and ameliorate systemic effects must start with identifying and collecting the relevant data about transactions and markets. New data sets are emerging from many sources, ranging from official collections such as supervisory stress tests to informal access to public corpora of news reports. Culling this massive information resource for patterns of concern will require correspondingly massive efforts at data collection, integration, and analysis.

It is common to classify big data scalability challenges according to the predominant bottleneck that materializes. These are the “Vs” of big data: volume, velocity, variety, and veracity. All these dimensions are relevant to fintech.¹⁹ For example, both market participants and their regulators are turning to cloud storage systems to help address challenges associated with rapidly growing data volumes. We touched briefly on issues of velocity in our discussion of HFT above. Similarly, integration of identifiers—for example, in the LEI framework—addresses a fundamental problem with variety. In the remainder of this subsection, we will focus on blockchain

platforms, which can help address challenges with data quality (veracity). Although blockchain systems can improve data quality, they can also introduce new challenges.

Privacy and anonymity are entry-level concerns in this process. The decentralized ledger built into the first-generation blockchains presupposes all-or-nothing transparency in the record of transactions. This helps establish the consensus needed for a valid blockchain ledger, but it creates problems when participants wish to record a transaction privately. Even in more general implementations, the distributed ledger must be open at least to validators (and potentially others), presenting questions about who knows what about transactions—for example, for supervisory or risk management purposes—under what circumstances, and in what format. Conversely, many early adopters of cryptocurrency have valued anonymity, displaying a range of motivations from the innocent to the malign.

The BCBC, in particular, goes to great lengths to anonymize (or pseudonymize) the identities of blockchain participants. Although bitcoin's anonymity can be useful in certain cases, for many financial transactions the unambiguous disclosure of the true identities of obligors and obligees matters very much. Bitcoin's pronounced anonymity posture highlights the question of blockchain transparency. Again, who gets to see what? Particularly in a crisis, financial stability supervisors will have a sharply heightened need for information, which suggests the need for adjustable transparency in key blockchains.

Even outside of a crisis, there will be valid requirements for partial information revelation. For example, since the 2008–2009 global financial crisis, numerous financial stability indexes have emerged, such as the OFR's Financial Stress Index or the various SRISK indexes at New York University's V-Lab.²⁰ A key goal of these indexes is to provide aggregated signals of accumulating risks to market participants. In many cases, these signals may depend directly or indirectly on the identities of key participants in the financial system, putting

participant anonymity and financial stability in tension. In general, then, the law or its regulatory implementation should make credible levers available to supervisors to adjust policy in the face of changing risks. A key question is how the governance framework will intervene to stabilize the system during stress periods. The law must be capable of enforcing rights established and represented in blockchain.

Another attribute of a distributed (public and permissionless) ledger—whether to support a currency, a payment system, a clearing or settlement regime, or secure contracting—is the absence of a central node for collection of information about the system or its actors or transactions. The many benefits of such an approach have been celebrated, including (potentially) anonymity, reduction of costs, avoidance of reuse (rehypothecation) of collateral, and others.²¹ Removal of a common ledger, however, could have consequences for both access to information and the quality of that information. In all instances, it is important to understand the implications of the permission rules for the ledger in question.

The BCBC, for example, builds in two forms of latency that are central to its process for building a consensus version of the truth into the distributed ledger. First, there is a mining latency for the brute-force calculations needed to establish the right to add a new block to the chain. It is theoretically necessary that this process be computationally costly—when signaling is costless, arbitrary nonsense (e.g., spam) or malicious misinformation can swamp the ledger. Currently, the BCBC mining costs periodically increase relative to reward, by intrinsic design (the *halving*, a process where the mining incentive decreases by 50 percent). In theory, a different blockchain with a higher reward, or a blockchain that touches a financial instrument worth considerably more (say, a trading blockchain that deals with transactions in the hundreds of millions or billions of dollars), could incentivize pursuit of fraudulent activity.

Depending on the consensus mechanism, a “consensus latency” in addition to the mining latency may be needed for an adequate number of affirming votes to accumulate for a particular version of the blockchain to proceed as the consensus truth. Both forms of latency imply a temporal delay, which could create tension, for example, if the goal is to support HFT operations. Although early implementations, such as the BCBC, have strongly favored a decentralized ledger and distributed consensus in establishing the agreed-upon version of the truth, the possibilities under alternative centralization assumptions have not yet been thoroughly explored, either from a legal perspective or from a technical perspective.

Alternatively, consider the example of the US wholesale payment system, which handles trillions of dollars in transactions each day. Presently, access to information about that system, its operations, and the flow of dollars through it can be understood by surveying a relatively small number of clearing banks, including the Federal Reserve Bank of New York, and financial market utilities (FMUs) that provide the settlement infrastructure.²² A distributed ledger system could involve the same basic participants exchanging obligations, and thus result in the same basic access to information about the system as can be gathered now through the clearing banks and FMUs. As a result, the introduction of a blockchain-based system of record need not disrupt the ability of supervisors and market participants to monitor the system.

On the other hand, a decentralized ledger could diffuse this currently concentrated system with many implications. First, not enough is known about low-latency blockchain systems. Just as mining and consensus latency may recommend against a blockchain to support HFT, they may turn out to be an awkward choice for monitoring the financial system as a whole. Timely information gathering may be impossible in a highly diffuse financial system during rapidly changing, fluid situations, such as a market stress event. Capturing information

about millions of disaggregate transactions and attributing them to central nodes or participants from perhaps thousands of servers would take computing power, time, and specialized tools that do not yet exist. This could also complicate efforts to identify failures and subsequently restart the system in the event of a cyberattack such as a denial of service against multiple nodes of the distributed ledger.

Access to information in a distributed ledger could be further complicated by a lack of resources (reliable power grids and/or communications infrastructure), misaligned incentives (key miners who stand to lose from full information revelation amid a crisis), or incompatible legal frameworks (privacy and information sharing laws) in key jurisdictions. Any of these forces could motivate regulations that ensure cooperation and fair play.

Additionally, simple access to distributed information about the system does not imply the feasibility of integrating data sets from various venues and jurisdictions effectively. The technologies and standards for payment information exchange, particularly across borders, will require coordination. There are alternatives to coordinated information integration for purposes of monitoring, regulation, and resolution—an ad hoc patchwork of local systems or a set of isolated silos—but any solution will present implementation challenges.

As noted above, anonymity or pseudonymity of transactors in a blockchain could raise other supervisory (and law enforcement) concerns. For example, given the traditional role of nodes in the system, such as banks' use of nodes to file suspicious activity reports, information flow could be limited by a distributed ledger network. Anonymity was a low-level requirement for the BCBC, but it is not clear how crucial this feature is, or what other possibilities open up if this constraint is eased. Although law enforcement concerns are generally beyond the scope of this chapter, regulators will surely seek to have means to identify market actors and their actions both to protect its participants and to preserve trust in the system

itself. This should be balanced, however, against the legitimate need for some degree of anonymity to protect proprietary trading activity or avoid front running ahead of large market-moving activities (which created the need for dark pools).

On the other hand, pretrade anonymity is a crucial feature of many dealer markets, where knowledge of who is trading or whether they have come to buy or sell can, by itself, be sufficiently revelatory to drive trading away from the market altogether. The blockchain ledger would represent a new information source in these intricate environments, and the implications for existing microstructures are far from absolute. It is not clear that a blockchain would necessarily disrupt existing relationships, nor is it clear whether such disruption would be good or bad, on net. One should acknowledge, however, that the potential for disruption exists, and stay attuned to its ramifications as innovations are introduced. For example, the possible loss of anonymity might drive trading away from central nodes of the system and into bilateral markets, which in turn could have implications for firms' internal risk management systems, particularly if they lack the information necessary to monitor risk exposures.

This concern touches the regulatory system in several respects. Since 2002, the Sarbanes-Oxley law has required that corporate executives in the United States certify the adequacy of internal controls for risk management systems. Gathering the information necessary to meet this compliance requirement could be even more challenging in a distributed system if transactions are recorded and positions are reflected on multiple, dynamic platforms. The costs of maintaining internal risk management systems that ingest information from distributed ledgers could prove high. A similar concern arises in the context of the stress testing of banks now required for compliance with the Dodd-Frank Act. Risk management systems may need retooling to comply with the requirements of the law, and supervisors—in this case, the Federal Reserve and

the Federal Deposit Insurance Corporation—will likewise need to adapt to identify the source of and standards for the data necessary to perform their supervisory functions.

Supervisors have invested significant effort to align reporting and transaction standards to improve the quality of information about the financial system. A good example is the creation of swaps data repositories, central counterparties, and swap execution facilities, where the OFR and other regulators are working to align data standards for entities, products, transactions, and myriad other data fields. As previously discussed, a possibly transformative application of blockchain technology would be for messaging, settling, clearing, and reporting the transaction using a single decentralized ledger. Here again, blockchain technology is not a panacea, however. Without common standards to reflect the legal and economic terms of the transactions and their counterparties, the precrisis opacity overlaying our derivatives markets would persist, regardless of the promise of blockchain technology. What remains, then, are questions of who will develop these standards, and how. Will competitive forces drive narrow, proprietary blockchain systems? Or will early steps be taken to avoid a disjointed collection of specific, one-off blockchain implementations that suffer from the kinds of classic collective-action problems that have hamstrung financial markets for years?

Finally, blockchain technology presents a significant opportunity to improve on many aspects of our financial infrastructure. One critical area is access to information by appropriate authorities, and the ability to securely share that information. In making design decisions, technologists could consider the benefits of improving the ability of supervisors to govern such that they can perform their critical role in overseeing orderly markets, protecting investors, testing the soundness of institutions at the core of the financial system, and building trust and confidence in our financial markets. Hopefully this will be a taken opportunity rather than a lost one.

6.4 CONCLUSIONS AND STEPS FORWARD

We are only at the dawn of realizing the potential of disruptive financial technologies. The coming years will see new applications emerge, new kinds of organizations develop, and new consequences arise that will need to be dealt with by a diverse community of stakeholders.

The technological and application developments for financial innovation will inevitably be accompanied by developments in government rule setting and oversight: in short, *regulation*. Some of this will involve applying existing rules and oversight structures; some will involve creating new provisions that will better fit the particular needs of a technologically based application. In this chapter we have sought to provide a framework for imagining what this future interaction of financial innovations such as blockchain with government may look like, along with our admittedly speculative vision of some of the possible points of contact.

We hope that this combined exercise will be a useful starting point for further deliberation and consideration as the process goes forward. We also note the emergence of various toolkits for policy makers and encourage those toolkit creators to consider the architectural principles that we have outlined herein.

Our thinking so far has largely focused on the *substance* of the applicable rules. In closing, we also wish to speculate on the *process* by which they will be developed. We strongly urge technology advocates such as the blockchain user community to get involved, and in some cases seek partnerships, with the various points of contact in government as this goes forward. For the most part, this engagement will lead to better outcomes. The modern regulatory process often invokes a model of multistakeholder dialog, with the goal of eliciting approaches that best serve the industry as well as the public and the needs of government itself. The technical complexity

of blockchain applications makes this need for engagement even more critical. Indeed, regulatory and policy bodies such as the Organisation for Economic Co-operation and Development, the European Union, and the UK government have various formal and informal consultative engagements with private-sector actors to strive for policy that considers a variety of perspectives, including that of enterprise, alongside representatives of consumer advocacy, academia, and government.

Indeed, we note that the regulatory sophistication of technology market participants varies considerably. We admonish tech start-ups to engage in a proactive rather than reactive dialog with the applicable regulators for their respective businesses—the first contact with a regulator should not be the enforcement letter. They should also be open to hearing from those who have learned hard lessons from our financial markets even while they usefully disrupt. We also hope that those so-called entrenched interests use their positions and wisdom about financial markets carefully and avoid pushing out newcomers just to limit competition. Arguments that like activities should be regulated alike make sense if the regulation still makes sense, even if it means increased costs of entry.

A similar admonition can be aimed at governmental actors. As envisioned in such principles as those set out in Circular A-4, where possible, governmental intervention should consist of light-handed regulation, based on market and self-governing principles developed in consultation with those most affected. We favor the toolkit approach, where select actions can be chosen to address a given circumstance, as not all interventions are suitable for all domiciles, and every country must adapt its approach to optimize the benefits for its particular set of consumers, enterprise, and society.

Such a collaborative process can never be fully harmonious; there are simply too many diverging interests between and within the various classes of players. And the fractious and often tumultuous atmosphere of national government

coinciding with the Trump presidency adds a layer of uncertainty. But dialog can nonetheless be far more productive than contentious processes based on hostility, mutual suspicion, and avoidance. We hope that this chapter will help catalyze the next stages in the process.

ACKNOWLEDGMENTS

The authors are grateful for the research and conceptual contributions of Josh Jackson, which have informed portions of this chapter.

NOTES

The views of the authors are their own and not necessarily those of their institutions.

1. Federal Deposit Insurance Corporation (FDIC), "Role of the Transfer Agent," Section 11 in *Trust Examination Manual*, 2005, https://www.fdic.gov/regulations/examinations/trustmanual/section_11/rta_manualroleoftransferagent.html.
2. N. Popper, "Knight Capital Says Trading Glitch Cost It \$440 Million," *New York Times*, August 2, 2012, http://dealbook.nytimes.com/2012/08/02/knight-capital-says-trading-mishap-cost-it-440-million/?_r=0.
3. M. A. Russon, "The Curious Tale of Ethereum: How a Hacker Stole \$53m in Digital Currency and Could Legally Keep It," *International Business Times*, June 20, 2016, <http://www.ibtimes.co.uk/curious-tale-ethereum-how-hacker-stole-53m-digital-currency-could-legally-keep-it-1566524>.
4. A. Quantson, "Ethereum Devs Hack the Hacker, Price Skyrockets," *Crypto Coins News*, June 22, 2016, <https://www.cryptocoinsnews.com/ethereum-devs-hack-the-hacker-price-skyrockets/>.
5. The parable of Shlemiel the road painter is apt in this context. See, for example, <https://discuss.fogcreek.com/techinterview/default.asp?cmd=show&ixPost=153>.

6. “How to Clean TRACE Data,” Copenhagen Business School Department of Finance, 2016, http://sf.cbs.dk/jdnielsen/how_to_clean_trace_data.
7. Global Legal Entity Identifier Foundation website, as of October 30, 2019, <https://www.gleif.org/en/lei-data/global-lei-index/lei-statistics>.
8. A *springing interest* shifts ownership in the event that a particular event occurs.
9. D. Atkins, W. Stallings, and P. Zimmerman, PGP Message Exchange Formats, IETF RFC1991, Internet Engineering Task Force, August 1996.
10. A. Pentland, D. Shrier, T. Hardjono, and I. Wladawsky-Berger, “Towards the Internet of Trusted Data,” in *Trusted Data*, 2nd ed., ed. T. Hardjono, D. Shrier, and A. Pentland (Cambridge, MA: MIT Press), 15–40.
11. See, for example, D. L. Chaum, “Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms,” *Communications of the ACM* 24, no. 2 (February 1981): 84–88; S. Brands, “Untraceable Offline Cash in Wallets with Observers,” in *Advances in Cryptology—CRYPTO’93: 13th Annual International Cryptology Conference, Santa Barbara, California, USA, August 22–26, 1993, Proceedings*, ed. D. Stinson (Berlin: Springer-Verlag, 1993), 302–318; S. Brands, *Rethinking Public Key Infrastructures and Digital Certificates* (Cambridge, MA: MIT Press, 2000); J. Camenisch and E. Van Herreweghen, “Design and Implementation of the Idemix Anonymous Credential System,” in *CCS ’02: Proceedings of the 9th ACM Conference on Computer and Communications Security*, ed. V. Athuri (New York: Association for Computing Machinery, 2002), 21–30; A. Lysyanskaya, R. L. Rivest, A. Sahai, and S. Wolf, “Pseudonym Systems,” in *Selected Areas in Cryptography: 6th Annual International Workshop, SAC’99, Kingston, Ontario, Canada, August 9–10, 1999, Proceedings*, ed. H. Heys and C. Adams (Berlin: Springer, 1999), 184–199; B. Rosenberg, ed., *Handbook of Financial Cryptography and Security* (Boca Raton, FL: CRC Press, 2011).
12. See “Overview,” U-Prove, February 25, 2012, <https://www.microsoft.com/en-us/research/project/u-prove/>; J. Camenisch, M. Dubovitskaya, P. Kalambet, A. Lehmann, G. Neven, F.-S. Preiss, and T. Usatiy, *IBM Identity Mixer: Authentication without Identification*, 2015, <https://www.zurich>

.ibm.com/pdf/csc/2015-11-12-Idemix-Presentation.pdf; ABC4Trust, “ABC4Trust: Attribute-Based Credentials for Trust,” March 28, 2012, <https://abc4trust.eu>.

13. For an overview of privacy-preserving MPC, see M. Flood, J. Katz, S. Ong, and A. Smith, “Cryptography and the Economics of Supervisory Information: Balancing Transparency and Confidentiality” (OFR Working Paper No. 0011, September 2013), http://financialresearch.gov/working-papers/files/OFRwp0011_FloodKatzOngSmith_CryptographyAndTheEconomicsOfSupervisoryInformation.pdf.

14. G. Zyskind, O. Nathan, and A. Pentland, “Enigma: Decentralized Computation Platform with Guaranteed Privacy,” submitted June 10, 2015, <https://arxiv.org/abs/1506.03471>.

15. See Federal Identity, Credential and Access Management (FICAM) Program, <https://www.idmanagement.gov/manage/>; OIX, OpenID Exchange, <http://openidentityexchange.org>; SAFE-BioPharma Association, Trust Framework Provider Services, https://www.safe-biopharma.org/infocenter/Trust_Framework_Provder_Services.pdf.

16. See M. Flood, H. V. Jagadish, and L. Raschid, “Big Data Challenges and Opportunities in Financial Stability Monitoring,” *Financial Stability Review of the Banque de France*, April 2016, 129–142, https://publications.banque-france.fr/sites/default/files/medias/documents/financial-stability-review-20_2016-04.pdf.

17. See, for example, R. Ali, “Innovations in Payment Technologies and the Emergence of Digital Currencies,” *Bank of England Quarterly Bulletin*, 2014, <http://www.bankofengland.co.uk/publications/Documents/quarterlybulletin/2014/qb14q3digitalcurrenciesbitcoin1.pdf>; T. Adrian and T. Mancini-Griffoli, *The Rise of Digital Money*, Fintech Notes, Note/19/01, International Monetary Fund, July 2019, <https://www.imf.org/~media/Files/Publications/FTN063/2019/English/FTNEA2019001.ashx>; International Monetary Fund, “Fintech: The Experience So Far” (IMF Policy Paper, June 2019), <https://www.imf.org/en/Publications/Policy-Papers/Issues/2019/06/27/Fintech-The-Experience-So-Far-47056/>.

18. ISDA, “Major Banks Agree to Sign ISDA Resolution Stay Protocol,” press release, October 11, 2014, <http://www2.isda.org/news/major-banks-agree-to-sign-isd-resolution-stay-protocol>.

19. See Flood, Jagadish, and Raschid, “Big Data Challenges and Opportunities in Financial Stability Monitoring,” for a fuller discussion of these issues.

20. Office of Financial Research, “OFR Financial Stress Index,” <https://www.financialresearch.gov/financial-stress-index>, accessed July 13, 2021; V-Lab, “Global Volatility,” <https://vlab.stern.nyu.edu>, accessed July 13, 2021.

21. McKinsey & Company, “Beyond the Hype: Blockchain in the Capital Markets” (McKinsey Working Paper No. 12, December 2015), <http://www.mckinsey.com/industries/financial-services/our-insights/beyond-the-hype-blockchains-in-capital-markets>; M. Mainelli and A. Milne, “The Impact and Potential of Blockchain on the Securities Transaction Lifecycle” (SWIFT Institute, Working Paper No. 2015-007, May 2016), http://www.swiftinstitute.org/wp-content/uploads/2016/05/The-Impact-and-Potential-of-Blockchain-on-the-Securities-Transaction-Lifecycle_Mainelli-and-Milne-FINAL.pdf; D. He et al., “Virtual Currencies and Beyond: Initial Considerations” (International Monetary Fund SDN/16/03, January 2016), <https://www.imf.org/external/pubs/ft/sdn/2016/sdn1603.pdf>.

22. Payments Risk Committee, *Intraday Liquidity Flows*, technical report, Federal Reserve Bank of New York, March 30, 2012, https://www.newyorkfed.org/medialibrary/microsites/prc/files/prc_120329.pdf.

