



Top Ten Obstacles along Distributed Ledgers' Path to Adoption

Sarah Meiklejohn | University College London

This article presents the top ten obstacles toward the adoption of distributed ledgers, ranging from identifying the right ledger to use for the right use case to developing scalable consensus protocols that provide some meaningful notion of public verifiability.

In January 2009, Bitcoin was released into the world by its pseudonymous founder, Satoshi Nakamoto. In the ensuing years, this cryptocurrency and its underlying technology, called the *blockchain*, have gone on a rollercoaster ride that few could have predicted at the time of its deployment. It has been praised by governments around the world, and people have predicted that “the blockchain” will one day be like “the Internet.” It has been banned by governments around the world, and people have declared it “adrift” and “dead.” The price of Bitcoin skyrocketed in late 2013 up to \$1,200 per bitcoin, only to spend the entire next year languishing at anywhere from \$200 to \$500 per bitcoin, before beginning a steady climb in 2016 that now has Bitcoin’s price hovering around \$4,500 per bitcoin (a number that will no doubt be wildly out of date at the time of publication).

After years in which discussions focused entirely on Bitcoin, people began to realize the more abstract potential of the blockchain, and “next-generation” platforms such as Ethereum (www.ethereum.org), Steem (<http://steemit.com>), and Zcash (<http://z.cash>) were launched. More established companies also realized the value in the more abstract properties of the blockchain—resilience, integrity, and so forth—and repurposed it for their particular

industries to create an even wider class of technologies called *distributed ledgers* and to form industrial consortia such as R3 (www.r3.com) and Hyperledger (www.hyperledger.org). These more general distributed ledgers can look, to varying degrees, quite unlike blockchains and have a somewhat clearer (or at least different) path to adoption given their association with established partners in industry. As we describe below, however, they must nevertheless overcome many of the same obstacles to become truly productive and long-lasting solutions.

Amidst many unknowns, what is increasingly clear is that, even if they might not end up quite like the Internet, distributed ledgers—in one form or another—are here to stay. Nevertheless, a long path remains from where we are now to widespread adoption—that is, a point at which the blockchain is even within several orders of magnitude as ubiquitous as the industries it has the potential to disrupt, such as traditional databases or the Internet—and there are many important decisions to be made that will affect the security and usability of any final product.

In what follows, we present the top ten obstacles along this path as well as what we as a community can do (and have been doing) to address them. By necessity, many interesting aspects of distributed ledgers, both in

terms of problems and solutions, have been omitted, and the focus is largely technical in nature. Nevertheless, our hope is that this serves as an introduction to the topic and to potentially fruitful avenues of research.

10. Usability: Why Use Distributed Ledgers?

The problem. What does usability refer to, and how is it a problem? There are a few possibilities in this context; for example, one could say that currently deployed distributed ledgers do not have particularly nice interfaces and are difficult to use without expert knowledge. While this is certainly true, one can and should expect these problems with any new technology and further expect them to be resolved with time and increased usage. (Here relational databases and the development of graphical, if imperfect, user interfaces into them might serve as a useful history.)

Thus, a perhaps more interesting “usability” problem is: What do end users actually want from distributed ledgers? Most deployed cryptography that end users interact with today maps easily to processes and desires that have existed for centuries. If you want to hide the contents of your conversation with someone, use encryption. If you want to convince someone that it’s really you saying something, use a (digital) signature. But what is it that the full public verifiability (or accountability, immutability, and so on) of distributed ledgers really maps to in terms of what end users want?

Potential and developing solutions. While some research has explored users’ perceptions of Bitcoin,¹ there has been very little research into the broader usability of distributed ledgers or into the way they fit into existing mental models. It may be that in fact most end users have no interest in the properties of distributed ledgers (just as they also likely have no interest in the properties of regular databases). Given the growing number of users who already interact with platforms like Bitcoin and Ethereum, however, and the discussed potential for distributed ledgers to—among many other use cases—inform the choices of ethical consumers by providing supply chain transparency, it is still worth exploring which specific benefits people feel they can achieve by using these technologies.

9. Governance: Who Makes the Rules?

The problem. The beauty of distributed ledgers is that no one entity gets to control the decisions made by the network; in Bitcoin, for instance, coins are generated or transferred from one party to another only if a majority of the peers in the network agree on the validity of this action. While this process becomes threatened if any one peer becomes too powerful, there is a larger question looming over the operation of these decentralized networks: Who gets to decide which actions are valid in

the first place? The truth is that all these networks operate according to a defined set of rules, and that “who makes the rules matters at least as much as who enforces them.”²

In this process of making the rules, even the most decentralized networks turn out to be heavily centralized. In Bitcoin, the rate at which new blocks are added to the chain, the reward participants receive for sealing transactions into the ledger, and many other parameters were all handed down by Satoshi Nakamoto and have not changed since. The one parameter that users are proposing to change, the size of blocks, has led to a contentious debate that has been raging for years. In contrast, Ethereum has now carried out four so-called “hard forks” of its network, in which participants are essentially mandated to switch to a new version of the software to comply with new rules created by a core set of developers.

These increasingly common collapses in the governance structures underlying cryptocurrencies threaten to harm their value and reveal the issues associated with ad hoc forms of governance. Thus, the problem is not just that we don’t know how to govern these technologies, or that the governance structures are both centralized and relatively fractured, but that—somewhat ironically—we need more transparency around how these structures operate and who is responsible for which aspects of governance, so that users of these cryptocurrencies can be aware of them and make informed decisions about their own level of involvement.

Potential and developing solutions. Several research papers have focused on ways to maintain decentralization in terms of the entities that enforce the rules,^{3,4} in the form of a mining process that resists pooling, although in more general distributed ledgers, it is still very much an open question how to incentivize participants to help maintain the ledger. (For example, it is unclear who, other than Google or other large certificate authorities, would want to expend the resources necessary to maintain a log server in Certificate Transparency [www.certificate-transparency.org], which provides no explicit reward to these entities.)

In terms of the entities that make the rules, very few solutions have been proposed. The “Satoshi Oath” (http://ipfs.b9lab.com:8080/ipfs/QmXysWEAexXQqYZhTGpECvksnaBkSEWHdGhM7vNeHxu_e2g) outlines a pledge that the authors deem necessary for anyone setting out to create a blockchain-based application, but it still serves only as a guideline. Broader distributed ledger projects such as R3 and Hyperledger have opted for a consortium-based approach, and other projects such as Certificate Transparency have opted for a fairly centralized approach. As in Item 8, however, it is unclear how these approaches compare to one another or how the governance structures will evolve as these platforms gain popularity.

8. Meaningful Comparisons: Which Is Better?

The problem. Bitcoin was the first cryptocurrency to be based on the architecture we now refer to as the blockchain, but it certainly isn't the last; there are now thousands of alternative cryptocurrencies out there, each with its own unique selling point. Ethereum offers a more expressive scripting language and maintains state, Litecoin (<http://litecoin.org>) allows for faster block creation than Bitcoin, and each new ICO (http://en.wikipedia.org/wiki/Initial_coin_offering) promises a shiny feature of its own. Looking beyond blockchains, there are numerous proposals for cryptocurrencies based on consensus protocols other than proof of work (see Item 3) and proposals in non-currency-related settings, such as Certificate Transparency, R3 Corda, and Hyperledger Fabric, that still fit under the broad umbrella of distributed ledgers.

The natural issue that arises in such an increasingly crowded landscape is how to distinguish between these solutions and pick the one that is best for a given application. Do you need a blockchain or just a database? Maybe an Excel spreadsheet is sufficient for what you need to do? Related back to Item 10, what even are the properties that you want to satisfy? Do you need full public verifiability (and if so, why)? Even if someone could specify a list of necessary properties, it is still not clear which current platforms support which properties, and to what extent.

Potential and developing solutions. Some recent research⁵ has looked at the different parameters chosen by different cryptocurrencies (for example, Bitcoin's 10-minute block generation versus the 2.5-minute interval used by Litecoin) and the effect these parameters have on the security of the system, finding, for example, that the same level of security against so-called "selfish mining" attacks is achieved by 37 blocks in Ethereum as by 6 blocks in Bitcoin (due to the relative stale block rates of the two platforms). While insightful, this work focuses specifically on cryptocurrencies based on proof of work and thus does not extend to more general distributed ledgers or cryptocurrencies based on alternative consensus protocols (see Item 3 for examples).

In considering distributed ledgers as a whole, recent research⁶ explored the connections between the security properties provided by Certificate Transparency and Bitcoin, finding that the tradeoff between the two seemed to be between trust in a distributed set of authorities on the one hand and the need to (inefficiently) broadcast messages and engage in consensus protocols like proof of work on the other hand. This research again focuses on a specific underlying structure, so significant work remains to allow comparisons between other platforms, or more generally to

understand the set of tradeoffs that one should consider in evaluating any particular choice.

7. Key Management: How to Transact?

The problem. Over the years, many Bitcoin users have reported incidents in which they have lost all their bitcoins because they threw away a hard drive, forgot a password, or didn't take appropriate measures to secure the wallets on their computers; in the worst publicly reported incident, a user lost 7,500 bitcoins. Indeed, once the hard drive containing the wallet is lost or your money is stolen, the irreversibility of both the underlying cryptography and the transactions means that there's no way to recover. In one very public example, an Ethereum smart contract called The DAO (short for Decentralized Autonomous Organization), designed to act as a collective investment vehicle, managed to attract over \$150 million in what is considered the largest crowdfunding campaign ever (<https://www.nytimes.com/2016/05/22/business/dealbook/crypto-ether-bitcoincurrency.html>). After an attacker exploited a loophole in the code of the smart contract behind The DAO, the community was essentially powerless to do anything except watch them steal all the funds stored inside (until, that is, some of the governing Ethereum developers discussed in Item 9 created and advocated for a hard fork to undo the damage).

While these incidents obviously have serious financial implications for the individuals involved in them, they should in some sense not be particularly surprising to these individuals given the unregulated "Wild West" world of cryptocurrencies. People are proposing more mainstream uses of cryptocurrencies—for instance, integrating Bitcoin wallets into Linux distributions—that will put these wallets in the hands of less experienced users who may be less prepared for these types of risks; however, this will lead to wider and more serious consequences. We thus need robust solutions that can better tolerate both the loss and the theft of keys.

Potential and developing solutions. A commonly used Bitcoin technique that has the potential to mitigate this issue is a *multisignature*, in which multiple parties join their public keys together to create an address for which some subset of their signatures is sufficient to spend its contents. For example, in a 3-of-3 multisignature address, all three parties would need to contribute signatures in order to spend coins from the address, and in a 1-of-3 address, any one of the signatures would suffice. A 2-of-3 multisignature address arguably provides the best defense against key loss, as funds stored in the address are still accessible even if one key is lost, but an attacker would need to access two separate keys to steal from it.

One could similarly argue that secret sharing, in which the key is split among some number of friends (or

devices) who can be trusted to not collude but to provide their shares if and when key loss occurs, could also help. With both these solutions, however, we must understand if the threat model in which they work is realistic for the scenario (going back to Item 10) in which we expect to be distributing, storing, and using these keys. It is also still an open problem to identify solutions in which one could regain access to lost funds with the same ease with which users currently regain access to accounts for which they have forgotten the password.

6. Agility: Which Algorithms Do We Use?

The problem. As discussed in Item 9, in many cryptocurrencies, the rules seem to have been handed down from on high: for example, Bitcoin addresses are computed by taking the ECDSA public key, performing SHA-256, performing extended RIPEMD-160, performing SHA-256 again, and again, rearranging the bytes of this output and the output of the extended RIPEMD-160 hash, and converting the result into a base58 string. Similarly, once the governing developers have decided on a proof-of-stake protocol that they approve of (more on that in Item 3), it will replace Ethash as the consensus protocol in Ethereum, and everyone will have no choice but to agree if they wish to keep using the cryptocurrency.

Aside from the governance issues raised by these rigid specifications, cryptographic primitives break, and they do so frequently. Eventually, computers may become powerful enough that SHA-256 will be broken, which would allow anyone to rewrite the entire history of a blockchain secured using only hashes. While there would likely be sufficient warning before this happened to give the developers time to switch to more secure alternatives (and anyway the problem would go far beyond distributed ledgers!), there is an argument to be made for providing users with multiple options even today, just as is done with systems such as TLS. As with TLS, however, agility can enable dangerous attacks, so significant caution is needed in both how it is implemented and in which cryptographic primitives are supported. The arguably safer option is not to enable users to engage with multiple cryptographic primitives, but instead allow them to choose the consensus required for their transactions (see Item 3 for more on different consensus protocols), depending on their trust assumptions about both the people with whom they transact and the maintainers of the system.

Potential and developing solutions. Almost all cryptocurrencies achieve no agility whatsoever: for each cryptographic primitive, there is one valid instantiation, and there is only one way to achieve consensus. One arguable exception is Ethereum, in which one could technically encode any cryptographic primitive—even

something that is completely insecure—into a smart contract by including a custom library within the contract code. Nevertheless, the underlying cryptography (that is, the checks that peers in the network make to verify individual transactions) and consensus protocol are just as rigid as in other cryptocurrencies.

Perhaps because they are unencumbered by the ideological mindset frequently encountered in cryptocurrencies, industry-led distributed ledgers are much more agile. In Corda, for example, an individual user can pick from a selection of available algorithms when forming a contract. In Hyperledger Fabric, participants can essentially plug in their own consensus protocol, leading to many different possible instantiations of the abstract design. Here then, the question is not necessarily whether or not it can be done at all, but how these different options can compose (for instance, what does it mean if two parts of the ledger have been agreed on by two distinct consensus protocols?) and what effect these options have on the overall security of the system.

5. Interoperability: How to Talk to Each Other?

The problem. Some people believe that the future will contain one single ledger (like the Internet), and in fact general-purpose platforms could in theory—that is, if the other issues on this list were solved—support most known applications of distributed ledgers. The more likely scenario, however, is that different companies will gravitate toward different ledgers based on their particular requirements. For example, financial applications requiring consensus between a fixed set of banks are more likely to adopt a platform such as Corda or Hyperledger Fabric, whereas applications that require fully open participation are more likely to adopt a platform such as Bitcoin or Ethereum.

To achieve the much-discussed potential of distributed ledgers to eliminate (or at least significantly open) the existing landscape of proprietary data silos, it is thus essential to achieve some kind of interoperability or a set of methods that allows these disparate ledgers to talk to each other.

Potential and developing solutions. Within the realm of platforms based on blockchains, the idea of a sidechain provides a simple way to translate actions from one blockchain into another—that is, to have transactions published in one ledger have an effect on another ledger. The security of these sidechains has been relatively unstudied to date, however, and while they have attracted significant attention, they have not yet seen much adoption. There are also several other methods proposed to transact across the blockchains of different cryptocurrencies, such as *atomic cross-chain trading* (http://en.bitcoin.it/wiki/Atomic_cross-chain_trading), but

again these solutions have been neither particularly well studied nor adopted in any meaningful way.

Beyond this, there have been a number of attempts to provide interfaces between distributed ledgers and real-world data, such as Town Crier⁷ and Oraclize (www.oraclize.it). Otherwise, both Corda and Hyperledger Fabric attempt to be modular in their choice of protocols (see Item 6), which means that they are designed to interoperate across different usages of the overall system, but to the best of our knowledge, there has been little attempt among existing solutions to interface with one another.

4. Scalability: Why Store Every Transaction?

The problem. Scalability can mean a lot of things. Item 1 discusses the need to scale the time it takes to process individual transactions with the computational power that is added to the network. Another aspect of scalability is that, as the system becomes more popular, it should not significantly increase the storage load placed on users of the system, as this deters participation and increases the barrier to entry. For integrity purposes, however, it is important in systems in which we expect full public verifiability (discussed further in Item 1) to not delete entries from the ledger. This results in, for instance, the Bitcoin blockchain requiring currently over 130 GB to store, as it contains every transaction that has ever taken place, and growing at a rate of roughly 144 MB per day (<http://blockchain.info/charts/blocks-size>). The main problem thus lies in how to provide a balance between these two (seemingly contradictory) requirements.

In certain applications, it may be sufficiently important to provide full auditability (for example, to satisfy regulatory requirements) that we cannot help but increase the storage load of participants. Especially in consumer applications, however, a relatively compelling case can be made that it is not actually necessary to store every transaction, as—for example—we are unlikely today to want or need to meaningfully examine someone's coffee purchases from early 2010. In fact, retaining such information over a long period of time is clearly at odds with the privacy of users (Item 2).

In addition, in a slightly altered trust model, more “casual” users may be willing to offload the work of auditing the system to more “archival” users; thus, while archival users would be required to store the entirety of the ledger, casual users could store only the information needed to check the validity of their own personal transactions.

Potential and developing solutions. The main innovation that allows Bitcoin wallets to run on smartphones—crucially, without storing the full Bitcoin ledger—is the idea of an SPV (Simplified Payment Verification) client,

which is analogous to the casual user mentioned above. These clients retrieve from archival peers and store only the headers of blocks along the blockchain, rather than their full contents, and rely on these headers to check for double spending. More specifically, with the help of an archival (or “full”) node, an SPV client checks only that a given transaction is included in the blockchain; it thus reveals the transactions in which it is interested to archival nodes and trusts the miners to prevent double spending from being included. While such clients have gained wide adoption, their security—and in particular the privacy and trust issues mentioned—is imperfect and somewhat poorly understood.

Another major development that would help prevent excessive transactions from being included in the ledger is the Lightning network, which is an example of a more general (and increasingly active) line of research on payment channels (http://en.bitcoin.it/wiki/Payment_channels). To open such a channel, two parties can create a transaction that effectively serves to jointly lock up a certain amount of funds. The parties can then transact an arbitrary number of times between themselves (crucially, without placing any transactions on the ledger), and when they are ready to close the channel—either because they have depleted their funds or no longer expect to transact—they can place another transaction on the blockchain that serves this purpose. The storage load on the blockchain is thus significantly reduced, as a single pair of transactions represents an arbitrary number of transfers. Furthermore, because this process can be designed to handle disputes, it still doesn't require the two parties to trust each other. Again, while this approach seems promising and has attracted significant attention, the security and privacy aspects of using such channels are still not well understood.

Finally, another promising approach is the idea of *sharding* the ledger, as discussed in Item 1. If implemented in a certain way, participants do not need to store (or even hear about) transactions that are irrelevant to them; for example, someone doesn't need to store everyone's morning coffee purchases, just their own. While this significantly reduces the storage requirements of an individual user, it is still important to understand the practical costs of storing what is ultimately a monotonically growing ledger.

3. Cost-Effectiveness: What Is the Cheapest Way?

The problem. People like to complain a lot about Bitcoin and its proof-of-work-based consensus protocol. They compare the electricity it uses to the electricity used by whole nations, or at least by large power plants. They call it an “environmental disaster.” While many of these arguments are quite overblown, the fact remains that proof of work is indeed very expensive,

and it would of course be a good thing if the same result could be achieved in a cheaper way.

To this end, there have been a huge number of alternative consensus protocols proposed, and even used in alternative cryptocurrencies. For example, proof of stake is a protocol in which, rather than preventing Sybil attacks on the mining process by requiring participants to consume some expensive resource, as is done in proof of work, so-called validators prove their “stake” in the system, in the form of an investment they have made that provides an economic disincentive for them to misbehave by forming bad blocks. Peercoin (<http://peercoin.net>), Nxt (<http://nxt.org>), and BlackCoin (<http://blackcoin.co>) use a version of proof of stake in which the investment is (roughly) coins that have been accumulated over time, and Ethereum has a planned transition (that is, hard fork) to its own version of proof of stake, Casper, within the next two years, in which the investment is a security deposit that locks up some of the validator’s coins for a certain period of time. Intel’s Sawtooth Lake platform uses proof of elapsed time (PoET), which relies on its SGX architecture. One of the most popular Bitcoin-based proposals, Bitcoin-NG,⁸ proposes isolating the usage of proof of work to elect a periodic “leader” who can then quickly (that is, without performing large amounts of computation) bear witness to individual transactions.

Potential and developing solutions. In the distributed ledger research community, this seems to be the issue being explored most actively. In addition to the protocols mentioned above, there are dozens more articles with their own proposals, both for the protocols mentioned above (especially proof of stake, which seems to be a sort of “holy grail” as of this writing) and for new and increasingly exotic ones. For now, however, the vast majority of cryptocurrencies use some form of proof of work.

If one expands outward to general distributed ledgers, the list of alternative consensus protocols grows and begins to include more classical consensus protocols: two-phase commit, PBFT,⁹ and so on. These tend to be significantly more efficient than the “proof-of-X” protocols used in cryptocurrencies, but with the tradeoff that they require a fixed set of known participants. As in Item 8, what is needed is a way to understand this tradeoff, provide meaningful comparisons within this growing landscape of consensus protocols, and understand the benefits that each provides in a given setting.

2. Privacy: How to Protect Data?

The problem. While significant research has focused on the anonymity of users in cryptocurrencies (that is, protecting the identity of participants), little research has focused on their privacy. For example, Bitcoin users are technically “pseudonymous,” as their on-chain identities

are not inherently linked to their real-world identity, but the details of each transaction—for example, the amount of bitcoins being sent—are still completely transparent. If we consider more expressive platforms and more exotic use cases, the desire to store health records on a distributed ledger means transactions would contain information about a patient’s name (that is, their real-world identity), their age, the nature and justification of medical procedures, and so on. While a naive solution would be to encrypt this data and provide decryption keys only to the necessary parties, the fact is that encryption schemes, as with all cryptographic primitives, break or are compromised (see Item 6), so this does not provide a long-term solution for privacy.

Even with respect to anonymity, there’s still a lot of work to do. A long line of research has demonstrated the limitations of Bitcoin’s anonymity,¹⁰ and while emerging platforms such as Monero¹¹ and Zcash¹² provide strong theoretical guarantees of improved anonymity, they have not yet been analyzed empirically. It is thus possible that they have their own set of weaknesses that fall outside the abstract protocol specification but could emerge with increased adoption and patterns of usage.

Potential and developing solutions. One of the simplest solutions is to send transactions to only those participants with whom you trust the associated details; this is akin to the sharding-based solutions we explore in Item 1. Even here, however, it is still important to consider privacy, as transactions may be shared beyond the initial set of participants, or it may be necessary to hide certain details from one participant but reveal them to another.

One general solution that has been proposed is Hawk,¹³ in which users can hide the details of their transactions but still convince the rest of the network that the transactions are valid. While useful, Hawk is specific to Ethereum and makes use of fairly advanced cryptography. Thus, it is still very much an open question how to achieve privacy in a lightweight and flexible manner for general distributed ledgers.

1. Scalability: Do We Need Full Agreement?

The problem. Arguably the biggest hurdle for fully distributed ledgers is the insistence that every node in the network needs to agree on the full state of the entire ledger. Aside from the issues with this approach raised earlier (for example, in Item 4), this means that distributed ledgers cannot and do not scale in terms of their ability to process growing numbers of transactions (throughput) while still ensuring that users do not have to wait for their transactions to be included (latency). In other words, the more computational power that joins the network, the worse it will perform in terms of throughput and latency. It is precisely because of this requirement

that every node must agree on every transaction, as it means the more transactions are in the system, the longer the nodes must wait for them to flood the network.

Because this insistence on full replication thus violates one of the most basic properties of distributed systems, we might naturally ask ourselves: Why do it in the first place? One of the primary benefits of cryptocurrencies, enabled by this requirement, is the idea of full public verifiability: any participant can verify for themselves the correct functioning of the system by, for example, replaying all transactions and ensuring that any agreed-on rules haven't been violated. If only certain nodes agree on certain parts of the ledger, then there is no one participant that can satisfy this notion of verifiability. To allow for increased throughput and decreased latency, one must therefore provide a balance between avoiding full replication—which is often accomplished using sharding, in which each participant processes transactions only within a given shard—and enabling at least some degree of openness and verifiability.

Potential and developing solutions. This topic has received significant attention, and a number of academic proposals adopt some form of sharding.^{14,15} Many industrial proposals similarly adopt a type of sharding; for example, in Corda, participants need to achieve consensus only on transactions that are directly relevant to them, and in Certificate Transparency there is no global consensus on the contents of the ledger. While these approaches achieve significantly better scalability, they raise questions about verifiability that fully decentralized solutions do not. For example, if only certain participants see certain transactions, how can other participants tell that their transactions obey the global set of rules? In the absence of such a global set of rules, what meaningful notions of integrity can we even satisfy?

In this last item, as with all the items on the list, we again see that each platform does not provide one uniquely perfect solution, but rather a set of tradeoffs that—to come full circle back to Item 10—must ultimately be balanced according to the individual use case in which the technology will be used. ■

Acknowledgments

This work was generously supported in part by R3 and by EPSRC Grant EP/N028104/1. Thanks are also due to the anonymous reviewers for their helpful feedback.

References

1. S. Abramova and R. Böhme, "Perceived Benefit and Risk as Multidimensional Determinants of Bitcoin Use: A Quantitative Exploratory Study," *Proceedings of the 37th International Conference on Information Systems (ICIS 16)*, 2016.
2. V. Lehdonvirta, "The Blockchain Paradox: Why Distributed Ledger Technologies May Do Little to Transform the Economy," 2016; <http://blogs.oii.ox.ac.uk/policy/the-blockchain-paradox-why-distributed-ledger-technologies-may-do-little-to-transform-the-economy>.
3. A. Miller et al., "Nonoutsourcable Scratch-Off Puzzles to Discourage Bitcoin Mining Coalitions," *ACM CCS*, 2015, pp. 680–691.
4. L. Luu et al., "Smart Pool: Practical Decentralized Pooled Mining," *Proceedings of the 26th USENIX Security Symposium*, 2017.
5. A. Gervais et al., "On the Security and Performance of Proof of Work Blockchains," *ACM CCS*, 2016, pp. 3–16.
6. M. Chase and S. Meiklejohn, "Transparency Overlays and Applications," *ACM CCS*, 2016, pp. 168–179.
7. F. Zhang et al., "Town Crier: An Authenticated Data Feed for Smart Contracts," *ACM CCS*, 2016, pp. 270–282.
8. I. Eyal et al., "Bitcoin-NG: A Scalable Blockchain Protocol," *Proceedings of the 13th Symposium on Networked Systems Design and Implementation (NSDI 16)*, 2016.
9. M. Castro and B. Liskov, "Practical Byzantine Fault Tolerance," *Proceedings of the Third Symposium on Operating Systems Design and Implementation (OSDI 99)*, 1999, pp. 179–186.
10. J. Bonneau et al., "SoK: Research Perspectives and Challenges for Bitcoin and Cryptocurrencies," *IEEE Symposium on Security and Privacy*, 2015, pp. 104–121.
11. S. Noether, A. Mackenzie, and the Monero Research Lab, "Ring Confidential Transactions," *Ledger*, vol. 1, 2016, pp. 1–18.
12. E. Ben-Sasson et al., "Zerocash: Decentralized Anonymous Payments from Bitcoin," *IEEE Symposium on Security and Privacy*, 2014, pp. 459–474.
13. A.E. Kosba et al., "Hawk: The Blockchain Model of Cryptography and Privacy-Preserving Smart Contracts," *IEEE Symposium on Security and Privacy*, 2016, pp. 839–858.
14. G. Danezis and S. Meiklejohn, "Centrally Banked Cryptocurrencies," *NDSS*, 2016.
15. L. Luu et al., "A Secure Sharding Protocol for Open Blockchains," *ACM CCS*, 2016, pp. 17–30.

Sarah Meiklejohn is a Reader in Cryptography and Security at University College London. She has broad research interests in computer security and cryptography and has worked on topics such as anonymity and criminal abuses in virtual currencies, anonymous credentials, and understanding the interface between cryptographic primitives and their mathematical underpinnings. Previously, Meiklejohn received a PhD from University of California, San Diego, in May 2014 under the guidance of Mihir Bellare and Stefan Savage, as well as an ScB in mathematics in 2008 and an ScM in computer science in 2009, both from Brown University. Contact at s.meiklejohn@ucl.ac.uk.