Privacy implications of blockchain systems: a data management perspective

Heng Xu
Kogod School of Business, American University, Washington, District of Columbia, USA, and
Nan Zhang
Department of Management, Warrington College of Business, University of Florida, Gainesville, Florida, USA

Received 20 January 2023
Revised 16 March 2023
Accepted 16 March 2023

Abstract

Purpose – Privacy scholars appear to struggle in conceptualizing blockchain from a privacy perspective: is it a privacy-enhancing mechanism like differential privacy, a privacy-intruding tool like third-party cookies or a technology orthogonal to the issue of privacy? Blockchain does not seem to neatly fit into any of these buckets that we traditionally use to gauge the privacy implications of information technologies. In this article, the authors argue that blockchain transcends the extant conceptualization of privacy because it modifies the nature of data flow upon which the modern concept of privacy is based.

Design/methodology/approach – The authors introduce a conceptualization of blockchain as a new mechanism for data management. Then, following this conceptualization, the authors present a functional review of blockchain, summarizing the features it provides for the data it manages. This review sets up the discussion of how blockchain redefines data flow by separating the power of collection, access and query of data to different entities. After illustrating how this change regrounds privacy concerns in a blockchain system, the authors conclude with a discussion of the recommendations for future privacy research on blockchain.

Findings – The authors demonstrate that blockchain, by design, separates three core data-centric operations that are assumed to be inextricably linked in the canonical conceptualization of privacy: the collection, access and query of data. Collection means to capture and then store the data; access means to modify or augment the data and query means the ability to test or verify certain properties of the data (e.g. whether a bank account has a zero balance). Traditionally, any entities that collect data can evidently read, modify or query the same data as they wish. With blockchain, however, an entity that stores the data may not be able to modify the data, yet an entity that cannot even read the data may be able to verify certain properties of the data.

Originality/value – Privacy scholars appear to struggle in conceptualizing blockchain from a privacy perspective: is it a privacy-enhancing mechanism like differential privacy, a privacy-intruding tool like third-party cookies or a technology orthogonal to the issue of privacy? In this article, the authors aim to respond to this important question.

Keywords Privacy, Blockchain, Cryptocurrency

Paper type Conceptual paper

The authors were supported in part by the US National Science Foundation under Grants No. 1851637 and 2309853, and by gifts from Amazon and Meta. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the sponsors listed above.

Organizational Cybersecurity Journal: Practice, Process and People
Vol. 3 No. 1, 2023
pp. 71-79
Emerald Publishing Limited
e-ISSN: 2635-0289
p-ISSN: 2635-0270
DOI 10.1108/OCJ-01-2023-0003

Introduction

In the recent development of information technology, it is difficult to find anything more ambitious yet controversial than blockchain (Nakamoto, 2008) and the variety of products derived from it, such as cryptocurrency, non-fungible token (NFT), etc. Cryptocurrency is touted as revolutionary by some (Prasad, 2021) but "worthless" by others (Locke, 2021). Bitcoin, one of the most well-known blockchain systems, is recognized as the legal tender in two countries (BBC, 2022) but banned in nine others (Pérez, 2022).

The implications of blockchain technology on privacy are a similarly contentious topic. Some argue that the decentralized nature of blockchain strengthens privacy protection (Kshetri, 2017), while others question its compatibility with data privacy laws (Berberich and Steiner, 2016) and demonstrate the feasibility of de-anonymizing blockchain data (e.g. bitcoin transactions) with graph analytics algorithms (Ron and Shamir, 2013). While much technical research has been devoted to this topic (e.g. Kappos *et al.*, 2018), there is a paucity of studies on how the adoption of blockchain may affect the privacy concerns of individuals involved. Privacy scholars appear to struggle in conceptualizing blockchain from a privacy perspective: is it a privacy-enhancing mechanism like differential privacy, a privacy-intruding tool like third-party cookies or a technology orthogonal to the issue of privacy? Blockchain does not seem to neatly fit into any of these buckets that we traditionally use to gauge the privacy implications of information technologies.

Consider, for example, bitcoin, one of the most popular blockchain-based cryptocurrencies today. Bitcoin is, on the one hand, privacy friendly by offering pseudonymity, meaning that it identifies users not through their real-world identities but using random identifications (IDs). Yet, on the other hand, bitcoin also undermines privacy by offering complete transparency, meaning that anyone can read or audit any transaction in history. From a privacy perspective, these contrastive stances raise novel questions that do not normally arise in classic settings (e.g. individual-to-business data collections), like whether an activity under a pseudonym could be linked back to an individual, whether such linkage becomes more likely when multiple activities are associated with the same pseudonym, how disclosing these activities affects an individual's privacy, etc. The presence of such novel questions makes the privacy implications of blockchain difficult to assess, both for the individuals involved and for privacy scholars interested in studying these individuals' attitudes, perceptions and beliefs towards privacy.

In this article, we argue that blockchain transcends the extant conceptualization of privacy because it modifies the nature of *data flow* upon which the modern concept of privacy is based. Specifically, we demonstrate that blockchain, by design, separates three core data-centric operations that are assumed to be inextricably linked in the canonical conceptualization of privacy: the *collection*, *access* and *query* of data. *Collection* means to capture and then store the data; *access* means to modify or augment the data; and *query* means the ability to test or verify certain properties of the data (e.g. whether a bank account has a zero balance). Traditionally, any entities that collect data can evidently read, modify or query the same data as they wish. With blockchain, however, an entity that stores the data may not be able to modify the data, yet an entity that cannot even read the data may be able to verify certain properties of the data.

This separation of power is unfortunately incompatible with the classic, intuitive yet simplistic, notion of how information "flows" from one entity to another. For example, now that the power to collect, access, and query the same data may belong to three different entities, which of the three should we mark as the destination of the "data flow" when data enters a blockchain system? If an entity only collects data but cannot access or query the collected data, does it count as a "recipient" of information? As we elaborate in the article, the separation of collection, access and query in a blockchain system shifts the focus of privacy concerns from the "flow" of data to the inference of knowledge about the data. In other words, to properly understand the privacy implications of blockchain, we need to burrow deeper into the technical design of blockchain systems to understand *what knowledge about the underlying data* may be learned by which entities in a system.

The rest of the article is organized as follows. We first introduce a conceptualization of blockchain as a new mechanism for data management. Then, following this conceptualization, we present a functional review of blockchain, summarizing the features it provides for the data it manages. This review sets up our discussion of how blockchain redefines data flow by separating the power of collection, access and query of data to different entities. After illustrating how this change regrounds privacy concerns in a blockchain system, we conclude with a discussion of our recommendations for future privacy research on blockchain.

Conceptualizing blockchain as data management

The modern business world is built on the premise of facilitating more and ever-speedier transactions between entities, i.e. individuals or organizations. To do so, there must be a fundamental mechanism that demarcates the boundaries between different entities, protects the assets of each entity and sets rules for the exchanges between them. Today, this mechanism is by and large provided by *structures* in the economic and legal systems, through instruments like contracts, bank accounts, consumer protection laws, etc. These structures govern individual behavior, managerial decisions and social actions, ensuring their orderly operation but also bringing along considerable bureaucracies and the associated overhead.

For example, as anyone who went through a real estate transaction can attest to, while a typical wire transfer can be executed automatically within minutes, the actual transfer of ownership could take days if not weeks in the USA (and even longer in some other countries, e.g. Hui and Png, 2021). This is because there is no legally robust method to automatically prove the ownership of a real estate asset and its eligibility to transfer (e.g. free of liens). As a result, costly manual efforts, easily amounting to thousands of dollars per transaction, have to be spent searching for legal records, establishing a chain of title (i.e. history of ownership) and compensating an issuer of title insurance to act as a guarantor for the real estate asset. For a generation used to instantaneous money transfers and stock trades with zero commission, the overhead incurred by the current mechanism of real estate transactions appears rather excessive and inefficient.

Blockchain promises to provide such mechanisms in a far more efficient way, not through a societal structure but by technological means. A key premise of blockchain is an observation that all the structures necessary to facilitate modern business transactions, e.g. assets, rules of exchange, etc., can be expressed as and therefore captured by *data*. With this conceptualization, the fundamental mechanism to ensure orderly transactions is, in essence, a way to manage such data and to provide *trustworthy* answers when someone needs to query the data, e.g. when a buyer wants to verify whether a property is indeed owned by the seller or when a seller wants to verify whether a buyer has the financial means to purchase the property.

Consider what a closing agent does today for a real estate transaction. The agent's main job is exactly to collect data, e.g. a chain of title, a list of outstanding liens and a commitment from a mortgage lender, before vouching for the trustworthiness of such data, e.g. by issuing a title insurance to the buyer and disbursing financial assets to the seller. From this example arise two notable characteristics of the extant mechanism for data management: first, all the tasks of data collection, management and query answering are done by a third party, i.e. an entity that is neither the buyer nor the seller, often established to serve the very purpose of facilitating real-estate transactions and second, the trust in the data is established through economic and legal structures governing the third party, such as the legally enforceable nature of a title insurance policy. Both the operation of the third party and the enforcement of trust (e.g. through litigation) is costly, contributing to the aforementioned overhead of real estate transactions.

Blockchain drastically reduces cost by removing the need of a third party and the corresponding reliance on the costly economic/legal structures. Instead, it implements a *data*

management system to handle the tasks of data collection, management and query answering, while using cryptographic techniques to establish the trustworthiness of query answers. For publicly accessible blockchain systems, such a system is usually distributed across numerous *nodes*, i.e. computers from all over the world that voluntarily join the blockchain system. For example, bitcoin has its core database functions handled by tens of thousands of voluntary nodes as of 2022 (Bitnodes, 2022).

Consider again the real estate example. Today, in the USA, transactions such as deed transfers, mortgage initiations and satisfactions, liens, etc., are usually recorded in the local county clerk's office. Imagine a future where such transactions are stored into a blockchain system, which then performs cryptographic operations to protect against any future tampering of stored data. When a buyer wants to purchase a property, instead of hiring a closing agent to search local county's records and certify the results, it can simply issue queries to the blockchain system and obtain trustworthy answers on who the current owner is, whether there is any outstanding liens against the property, etc. The *trust* here is not established through an insurance policy, but in the form of a cryptographic proof showing that, if a malicious entity were to manipulate or fake a query answer, it would have to solve notoriously hard mathematical problems that (are currently believed to) take even the fastest supercomputer today hundreds of years to solve. As can be seen from this example, the reason why a blockchain system can remove the costly manual efforts associated with real estate transactions is because it *automates* data management and trustworthy query answering, replacing the costly economic/legal underpinning of these operations with cryptographic guarantees.

How blockchain redefines data flow and impacts privacy

Given the conceptualization of bitcoin as a new mechanism for data management, we now illustrate why the technical design of blockchain modifies the conventional notion of a data flow, thereby altering the ground on which we conceptualize people's privacy concerns.

Overview

There are three basic functions essential to any data management system: *collection*, i.e. the ability to store data in the system; *access*, i.e. the ability to specify what to store in the system (e.g. by modifying or augmenting data) and *query*, the ability to retrieve or verify certain properties of data in the system (Ramakrishnan and Gehrke, 2003). Traditionally, the organization that operates a data management system has full control over all three functions. By virtue of storing the data (physically or on the cloud), the organization can revise the data at ease and determine what queries may or may not be asked about the data. As the organization has full control, any trust (or the lack thereof) we may have on the functioning of data management (e.g. how the data may be used) is derived entirely from our trust of the organization which, in turn, is enabled by the costly economic and legal structures discussed earlier. From this perspective, it is easy to see why, in the context of private information being entered into a data management system, the extant notion of privacy concern centers around the data flow from individuals to the organization and why the current efforts on regulating privacy, like the European Union's General Data Protection Regulation (GDPR), mostly focus on setting rules for such data flows using levers in the economic and legal structures.

Blockchain fundamentally alters this arrangement. As illustrated in Figure 1, there may not be any centralized "organization" that operates the system and controls the functions. Instead, data management can be enabled by a network of distributed nodes (i.e. voluntarily participating computers) running the software that implements the blockchain system. As we

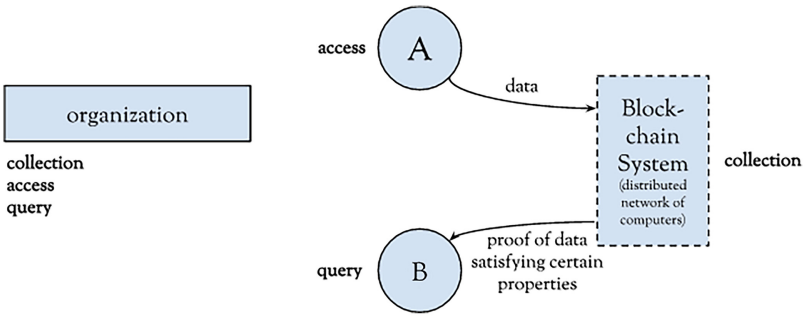

Centralized data management
(a)

Blockchain data management
(b)

Source(s): Author's work

Figure 1.
Data management functions

will elaborate later in the section, this distributed arrangement leads to a separation of the three functions, as a node that stores a data record may not be able to modify it (i.e. a separation of collection and access), while an entity that can neither read nor write a record is nevertheless able to query certain properties of it (i.e. a separation of access and query).

With this new arrangement, the proper performance of the three functions is not dependent upon the goodwill of any organization (or any participating node), but instead derives from the *functional guarantees* offered by the technical (specifically, cryptographic) design of the blockchain system itself. In the context of privacy, this means that it is now the technical design of a blockchain system, rather than any specific organization, that controls how private data may be collected, accessed and used in the future. It is thus quintessential to examine the functional guarantees provided by blockchain systems in order to understand the potential privacy concerns that could arise from their use.

In the rest of this section, we examine the prevailing guarantees provided by existing blockchain systems for each of the three core data management functions and discuss the privacy implications of such guarantees. We offer a caveat at the outset that our review assumes the functioning of a blockchain system to remain within its originally designed operating envelope. This means that we do not consider edge cases where a blockchain is compromised by malicious entities, such as through the 51% attack (Aponte-Novoa *et al.*, 2021) which requires adversaries to take over more than half of the computational power present in the system. We also do not consider the so-called layer-two protocols (Sguanci *et al.*, 2021) which repurpose a blockchain system for a different mechanism of data management.

Collection

In terms of *collection*, the most important guarantee offered by blockchain is tamper-resistant, meaning that once a record has been written into the system, it will never be modified in any way. Blockchain derives its name from the design that offers such tamper-resistance. Specifically, the design of blockchain requires a cryptographic hash of the previous record (or block of records) to be included in the next record. This ensures that, should a malicious entity attempt to modify an existing record in any way, it would have to first find a modified record that has the exact same hash as the existing one. The so-called collision-resistance property of cryptographic hash functions (Goldreich, 2001) guards against this possibility, therefore providing the tamper-resistance guarantee for data managed in a blockchain system.

In terms of privacy, there is an obvious conflict between the tamper-resistant property and the *right to be forgotten*, a concept enshrined in privacy laws such as GDPR. While there have

been extensive technical efforts on making blockchain GDPR-compliant (Politou *et al.*, 2019), they rely on reverting back to the traditional way of data management to various extent. For example, some shift the management of private data to outside the blockchain system (Eberhardt and Tai, 2017), while others reintroduce a trusted organization that is given the secret key required to revise historic records (Ateniese *et al.*, 2017). In essence, since any trust on the data managed by a blockchain system relies on its tamper-resistant property, supporting the right to be forgotten would inevitably undermine such trust, requiring it to be shifted elsewhere, e.g. to an organization (and its associated economic/legal structures) or to another technical solution.

Access

For a blockchain system to facilitate real-world transactions, it cannot be a free-for-all data management system into which everyone has access to store or modify any data. Consider the real estate example discussed earlier. If anyone could enter a transaction into the system to claim ownership on a property, then the data being stored would become meaningless. In order to make a blockchain useful in practice, there have to be *rules* governing data access that are algorithmically enforceable without human intervention.

These rules vary between different implementations of blockchain. With bitcoin, anyone who wants to write a record into the system must prove, using a cryptographic digital signature (Goldreich, 2001), that it is the only entity with the right to access this record. For example, in order to spend the balance of a bitcoin account (i.e. *address* in bitcoin terminology), one needs to enter into the system a new data record representing the spending transaction. Bitcoin requires this new record to include a digital signature that could only have been generated by someone who possesses the secret key associated with the bitcoin account. In other words, with bitcoin, the access to each data record is tied to a specific entity (or everyone in a group of entities when the transaction spends bitcoins from multiple addresses).

From a privacy perspective, while the bitcoin implementation ensures proper access to records in the system, it inevitably discloses *some* information about the entities that are authorized to access the records. This is because the right to access is assured by the correspondence between digital signatures in different records, e.g. spending and deposit transactions for bitcoin, or deed transfers for the same real estate property. This assurance effectively links the pair of records together. By following such links and tracking the digital signatures in different records, one could then identify the sequence of all records associated with an individual, such as all transactions associated with the same bitcoin account. Over time, such sequences would inevitably reveal patterns of records and transactions that can be used to unveil the real-world identity of the individual.

As can be seen from this example, blockchain's separation of two data management functions, collection and access, introduces an intriguing new privacy issue. Traditionally, the anonymization of data records is frequently touted as an effective mechanism to protect information privacy (Article 29 Working Party, 2014). Yet, in a blockchain system, there is a *trade-off* between offering anonymity and ensuring proper data access. In the case of bitcoin (or any currency system), a spending transaction has to be linked somehow to a previous transaction vouching for the available asset of the spender. Such a link is necessary to ensure proper data access, yet by nature reduces the anonymity of individuals involved in the transactions.

There have been considerable technical efforts attempting to provide better anonymity while ensuring proper data access. For example, "private" cryptocurrencies like Monero enforce access rules with a special type of cryptographic digital signature called *ring signature* (Noether, 2015), which links a spending transaction to not one account but a group of accounts sharing the same balance. What ring signature offers is a digital signature

proving that the signer knows one of a ring (i.e. set) of secret keys, without divulging *which* of the keys the signer actually knows. Monero uses a variation of the ring signature with an additional guard against double-spending, which ensures that no one can use the same secret key to generate more than one ring signatures. As can be seen from Monero's design, its key idea is to enforce access rules *not* on the level of each account, but on a higher level of *balance* (i.e. allowing one access to all counts with the same balance). As access rules are now enforced at a higher level, such private cryptocurrencies can clearly offer better anonymity than bitcoin. Yet the fundamental trade-off between anonymity and access remains, as evidenced by recent studies on the possibility of deanonymizing transactions on private cryptocurrency networks (Kappos *et al.*, 2018).

Query

Like access rules, the rules governing queries, i.e. what information may or may not be asked about a data record, also vary considerably between different implementations of blockchain. With bitcoin, all information of a data record is public, meaning that anyone would be able to query the timestamp and amount of any bitcoin transaction. It is important to understand why bitcoin was designed this way. A key rationale here is to ensure the validity of spending transactions, meaning that the amount being spent does not exceed the balance of the spender's account. Since bitcoin already provides tamper resistance for data collection and storage, the easiest way to ensure a non-negative account balance is to publicize all transaction amounts associated with the account, so anyone could easily verify its current balance.

This design exacerbates the tracking problem discussed earlier, as the more information a transaction reveals, the more likely one would be able to unveil the real-world identity of a bitcoin account owner based on the sequence of transactions associated with the account. To address this concern, cryptocurrencies like Monero implemented a different design known as a *zero-knowledge proof* (e.g. Bulletproofs; Bünz *et al.*, 2018) to guarantee the validity of a spending transaction without revealing its amount. The key observation here is that to ensure the validity of a spending transaction, one only needs to verify two inequalities: first, the amount flowing out of the spender's account is greater than or equal to the amount deposited into the receiver's account. Second, the amount being deposited is at least zero. What Monero does is to include in each transaction a cryptographic proof for both inequalities that can be publicly verified. Since neither inequality directly reveals the amount of the transaction, Monero ensures the validity of a transaction without publicly disclosing its amount.

From a privacy perspective, this zero-knowledge proof-based design further complicates what "disclosure" means for private information as it changes what can be queried about a data record from an all-or-nothing dichotomy to a much more fine-grained choice, allowing the system designer to support an arbitrary set of questions but deny others. Consider the amount of a transaction as the private information of interest. It is clearly disclosed to the blockchain system as it is stored *within the transaction record*, albeit in an encrypted form that can only be read by entities involved in the transaction. To others, the encrypted amount is nothing but a random number, and what they can learn from the transaction record is limited to the proof that both inequalities check out. In other words, *"what is in the data"* now becomes different from *"what can be learned from the data"*, further complicating the picture of how data, information or knowledge flows from one entity to another.

Implications on future privacy research

As the previous section shows how blockchain separates the collection, access and query functions of data management, we now conclude the article with a discussion of how such a

separation could affect future privacy research for blockchain. Specifically, we submit that there are at least two important implications:

First, the separation of data management functions requires privacy researchers to attend to the *partial* disclosure of data from one entity to another. For example, the separation of collection and access means that the identity of someone involved in a transaction is not disclosed as is. Instead, *some* (i.e. partial) information about the identity could be inferred from the chain of transactions, timestamps, etc. Similarly, the separation of access and query means that *some* information about the transaction amount (e.g. the fact that it satisfies the two aforementioned inequalities) has to be disclosed while the exact amount could remain hidden. Given the prevalence of such partial disclosure in blockchain systems, it becomes important for privacy researchers to delineate what knowledge about an individual may or may not be learned from the partial disclosure that occurs in a blockchain system. To this end, future privacy research may be able to lean on insights from a technical research field of *inference* (e.g. Naveed *et al.*, 2015; Jegorova *et al.*, Forthcoming), which focuses on how certain knowledge about data could be inferred from various sources such as query answers over data, statistics or machine learning models built over data, etc.

Second, the complex design of the cryptographic techniques used to separate the three data management functions also brings about the question of whether, and to what degree, the implications of such techniques on privacy could be understood by individuals involved in a blockchain system. As discussed in the beginning of the paper, even experts have divergent opinions, and sometimes misconceptions, about the privacy implications of blockchain. With the growing exposure of general population to blockchain technology, e.g. a survey of New York Digital Investment Group (2021) shows that 22% of Americans own bitcoin, it becomes increasingly important for privacy researchers to carefully study people's attitudes, perceptions and beliefs about their privacy in the context of blockchain. Yet, with separation of data management functions, the instruments used to elicit such attitudes, perceptions and beliefs may well need to be updated in order to properly separate people's true opinions from the considerable noise that stems from their lack of understanding of subtleties in the cryptographic tools that imbue blockchain systems. From this perspective, there is a massive amount of open research space for privacy scholars to explore in the dynamic and rapidly evolving landscape of blockchain technology.

References

Aponte-Novoa, F.A., Orozco, A.L.S., Villanueva-Polanco, R. and Wightman, P. (2021), "The 51% attack on blockchains: a mining behavior study", *IEEE Access*, Vol. 9, pp. 140549-140564.

Article 29 Working Party (2014), "Opinion 05/2014 on anonymisation techniques", available at: https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2014/wp216_en.pdf

Ateniese, G., Magri, B., Venturi, D. and Andrade, E. (2017), "Redactable blockchain–or–rewriting history in bitcoin and friends", *2017 IEEE European Symposium on Security and Privacy (EuroS&P)*, IEEE, pp. 111-126.

BBC (2022), "Bitcoin becomes official currency in Central African Republic", *BBC*, available at: https://www.bbc.com/news/world-africa-61248809

Berberich, M. and Steiner, M. (2016), "Blockchain technology and the GDPR-how to reconcile privacy and distributed ledgers?", *European Data Protection Law Review*, Vol. 2, p. 422.

Bitnodes (2022), "Global bitcoin nodes", available at: https://bitnodes.io/nodes/all/ (accessed 14 December 2022).

Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P. and Maxwell, G. (2018), "Bulletproofs: short proofs for confidential transactions and more", *Proceedings of the 2018 IEEE Symposium on Security and Privacy (SP)*, pp. 315-334.

Eberhardt, J. and Tai, S. (2017), "On or off the blockchain? Insights on off-chaining computation and data", *European Conference on Service-Oriented and Cloud Computing*, Cham, Springer, pp. 3-15.

Goldreich, O. (2001), *Foundations of Cryptography*, Cambridge University Press, Cambridge, Vol. 1.

Hui, K.L. and Png, I. (2021), "Why hasn't Hong Kong updated its antiquated property title system yet? South China Morning Post", March 27, 2021, available at: https://www.scmp.com/comment/opinion/article/3127003/why-hasnt-hong-kong-updated-its-antiquated-property-title-system

Jegorova, M., Kaul, C., Mayor, C., O'Neil, A.Q., Weir, A., Murray-Smith, R. and Tsaftaris, S.A. (Forthcoming), "Survey: leakage and privacy at inference time", *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Kappos, G., Yousaf, H., Maller, M. and Meiklejohn, S. (2018), "An empirical analysis of anonymity in zcash", *Proceedings of the 27th USENIX Security Symposium*, pp. 463-477.

Kshetri, N. (2017), "Blockchain's roles in strengthening cybersecurity and protecting privacy", *Telecommunications Policy*, Vol. 41 No. 10, pp. 1027-1038.

Locke, T. (2021), "Jamie Dimon says bitcoin is 'worthless'", *CNBC*, October 11, 2021, available at: https://www.cnbc.com/2021/10/11/jpmorgan-chase-ceo-jamie-dimon-says-bitcoin-is-worthless.html

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

Naveed, M., Kamara, S. and Wright, C.V. (2015), "Inference attacks on property-preserving encrypted databases", *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pp. 644-655.

New York Digital Investment Group (2021), "Banking + bitcoin survey", available at: https://nydig.com/research/nydig-bitcoin-banking-survey

Noether, S. (2015), "Ring signature confidential transactions for Monero", *IACR Cryptology ePrint Archive*.

Pérez, J.V. (2022), "Countries where cryptocurrency is legal and illegal", *Money.com*, October 21, 2022, available at: https://money.com/cryptocurrency-legal-status-by-country/

Politou, E., Casino, F., Alepis, E. and Patsakis, C. (2019), "Blockchain mutability: challenges and proposed solutions", *IEEE Transactions on Emerging Topics in Computing*, Vol. 9 No. 4, pp. 1972-1986.

Prasad, E.S. (2021), *The Future of Money: How the Digital Revolution Is Transforming Currencies and Finance*, Harvard University Press.

Ramakrishnan, R. and Gehrke, J. (2003), *Database Management Systems*, McGraw-Hill, New York, Vol. 3.

Ron, D. and Shamir, A. (2013), "Quantitative analysis of the full bitcoin transaction graph. Financial Cryptography and Data Security", *Lecture Notes in Computer Science*, Vol. 10323, pp. 478-493, Springer.

Sguanci, C., Spatafora, R. and Vergani, A.M. (2021), "Layer 2 blockchain scaling: a survey", available at: https://arxiv.org/pdf/2107.10881

Corresponding author
Heng Xu can be contacted at: xu@american.edu