INSTITUT C.D. HOWE INSTITUTE

COMMENTARY
NO. 468

Blockchain Technology – What's in Store for Canada's Economy and Financial Markets?

Blockchain technology has the potential to transform the financial sector. This study identifies and advises on the challenges facing regulators and policymakers.

Thorsten Koeppl and Jeremy Kronick

About The Authors

Thorsten Koeppl
is Associate Professor and RBC Fellow, Queen's University.

Jeremy Kronick
is Senior Policy Analyst at the C.D. Howe Institute.

The Institute's Commitment to Quality

C.D. Howe Institute publications undergo rigorous external review by academics and independent experts drawn from the public and private sectors. The Institute's peer review ensures the quality, integrity and objectivity of its policy research. The Institute will not publish any study that, in its view, fails to meet these standards.

The Institute requires that its authors publicly disclose any actual or potential conflicts of interest of which they are aware.

In its mission to educate and foster debate on essential public policy issues, the C.D. Howe Institute provides nonpartisan policy advice to interested parties on a non-exclusive basis. The Institute will not endorse any political party, elected official, candidate for elected office, or interest group.

As a registered Canadian charity, the C.D. Howe Institute as a matter of course accepts donations from individuals, private and public organizations, charitable foundations and others, by way of general and project support. The Institute will not accept any donation that stipulates a predetermined result or policy stance or otherwise inhibits its independence, or that of its staff and authors, in pursuing scholarly activities or disseminating research results.

Commentary No. 468
February 2017
Financial Services and Regulation

Daniel Schwanen
Vice President, Research

$12.00
ISBN 978-1-987983-13-5
ISSN 0824-8001 (print);
ISSN 1703-0765 (online)

The Study In Brief

Blockchain technology has the potential to transform dramatically how a modern economy deals with maintaining and updating records. This innovation has already created lots of turbulence in financial markets and beyond. It will be a challenge to let markets figure out how to best use this technology while ensuring consumer safety and efficiency. Our goal in this paper is to unveil the potential of blockchain technology and guide regulators in how to approach the challenges this technology entails.

The most well-known examples of blockchains are found in the area of payments systems and, more generally, in financial markets. It is thus understandable that the financial industry is leading the charge to unearth the potential of this technology in order to find cost efficiencies, but also to recapture above normal profits. The potential application of this technology, however, reaches much further than merely being a currency like bitcoin or a record-keeping system. Early applications of this technology include smart contracts and attempts by governments to build universal online identification systems. Blockchain technology also introduces new concepts such as cryptographic communication protocols and distributed data storage that can increase the safety of electronic networks and offer potential cost efficiency.

We do not expect distributed ledgers to completely supplant traditional intermediaries, especially in areas where these intermediaries are of systemic importance or provide services that require a high degree of ad hoc coordination. Still, many elements of this new technology offer a unique opportunity for such intermediaries to modernize their infrastructures and offer their clients safer and cheaper systems. It is not clear, however, how to realize such benefits in a way that makes sure they are passed on to the economy as a whole.

This leads us to identify three major challenges and priorities for policymakers and regulators arising from blockchain technology:

1. Design a principle-based regulation regime that achieves high safety standards, legal certainty and a stable environment for transactions based on distributed ledger technology;

2. Ensure that this technology leads to appropriate end-user cost efficiencies rather than simply a redistribution of above-normal profits among intermediaries; and

3. Determine areas where government involvement is advisable, be it in the role of facilitator for a private or public distributed ledger, or as a direct central node that applies elements of the technology but retains the monopoly of managing the ledger entries.

C.D. Howe Institute Commentary© is a periodic analysis of, and commentary on, current public policy issues. Michael Benedict and James Fleming edited the manuscript; Yang Zhao prepared it for publication. As with all Institute publications, the views expressed here are those of the authors and do not necessarily reflect the opinions of the Institute's members or Board of Directors. Quotation with appropriate credit is permissible.

To order this publication please contact: the C.D. Howe Institute, 67 Yonge St., Suite 300, Toronto, Ontario M5E 1J8. The full text of this publication is also available on the Institute's website at www.cdhowe.org.

In 2009, the first bitcoin transaction took place. Until recently, many people viewed the idea of such an alternative currency, which existed only virtually, as a mere curiosity, another strange development of the computer age. But times have changed.

Many believe that distributed ledger technology – often also labelled blockchain technology, which is the main idea that powers bitcoin – is ushering in a new era for the organization of our economies, businesses and markets. The Bank of Canada's recent announcement to explore the potential of blockchain technology for the Canadian payment system is prime testimony of this trend.[1] So what's in store for the Canadian economy from this new technology?

A *distributed ledger* in its basic form is a digital record distributed among many participants connected by a network who have agreed on the rules for updating the ledger. This ledger is commonly maintained as a *blockchain* where records are collected into blocks of data and put into a chronological order with each block building on the previous one.

Record-keeping through a ledger – such as ownership records, financial accounts or the safe-keeping of securities – enables decentralized trading and contracting that are of crucial importance for a modern economy. Blockchain technology challenges the idea that these functions must rely on a centralized, public ledger or platform. Traditionally, neutral third parties have run public ledgers. Historically, the key challenge for successful economies was therefore: How can we ensure that a public ledger is safe and accurate, or in other words, how can the third party running the ledger itself be trusted?

Bitcoin's revolutionary idea is that a blockchain can be used to solve this problem. Based on cryptography and a peer-to-peer network, the blockchain creates an online ledger that, once distributed among the network's participants, is tamper proof and available to verify transactions without the need for a trusted third party such as a middleman or intermediary.

The most well-known examples of blockchains are found in the area of payments systems and, more generally, in financial markets. It is thus understandable that the financial industry is leading the charge to unearth the potential of this technology in order to find cost efficiencies, but also to recapture rents. The potential application of this technology, however, reaches much further than merely being a currency like bitcoin or a record-keeping system. Early applications of this technology include smart contracts and attempts by governments to build universal online identification systems.

Our goal in this *Commentary* is to describe the essential features of blockchain technology, outline the economic drivers behind it and show where questions and concerns for public policy arise.

The authors would like to thank Daniel Schwanen, David Andolfatto, Jonathan Chiu, Alex Ciappara, Gérald Cossette, Baiju Devani, Walter Engert, Blake Goldring, Phil Howell, Jim MacGee, Scott Wilkinson and anonymous reviewers for comments on an earlier draft. The authors retain responsibility for any errors and the views expressed here.

1 See https://www.bloomberg.com/news/articles/2016-06-16/bank-of-canada-experimenting-with-distributed-ledger-technology.

Blockchain technology also introduces new concepts such as cryptographic communication protocols and distributed data storage that can increase the safety of electronic networks and offer potential cost efficiency. We do not expect distributed ledgers to completely supplant traditional intermediaries, especially in areas where these intermediaries are of systemic importance or provide services that require a high degree of ad hoc coordination. Still, many elements of this new technology offer a unique opportunity for such intermediaries to modernize their infrastructures and offer their clients safer and cheaper systems. It is not clear, however, how to realize such benefits in a way that makes sure they are passed on to the economy as a whole.

This leads us to identify three major challenges for policymakers and regulators arising from blockchain technology.

- First, blockchains and ideas associated with them push the frontier of what is feasible. New applications such as smart contracts have the potential to revolutionize the corporate world. While regulation should not stifle business experimentation, it is indispensable for creating a basic legal framework and putting standards into place that offer safety and stability. As one cannot pinpoint the institutions that will drive this change, it is best to employ a principle-based approach that moves away from institution-based to activities-based regulation.
- Second, blockchain technology has already started to create turbulence in well-established areas where specialized intermediaries have performed critical functions for decades. Moving forward, policymakers have to be vigilant that blockchains are not used to reshuffle rents at the expense of users, but really do create cost efficiencies.[2] One way to achieve this goal is to engage in public-private partnerships to develop new systems that are stable, solve start-up problems associated with network externalities, and foster competition by ensuring fair access to blockchain-based systems.
- Third, many potential blockchain applications are in areas highly important for the economy such as payment systems, financial market infrastructure or government databases. We do not deem it feasible to move toward a truly distributed ledger based on blockchain technology in many of these areas. However, some ideas from the technology can be used to improve existing systems, but applications in areas of critical infrastructure will often necessitate direct government involvement. Policymakers will thus be forced to make decisions to what degree small private networks can provide services based on blockchains and how governments engage with these networks. Prime examples are the Bank of Canada's Project Jasper examining the feasibility of an interbank settlement engine along with several foreign government projects to harmonize online identities.

In summary, policymakers and regulators should focus on three priorities:

1. Design a principle-based regulation regime that achieves high safety standards, legal certainty and a stable environment for transactions based on distributed ledger technology;
2. Ensure that this technology leads to appropriate end-user cost efficiencies rather than simply a redistribution of above-normal profits among intermediaries; and
3. Determine areas where government involvement is advisable, be it in the role of facilitator for a private or public distributed ledger, or as a direct central node that applies elements of the technology but retains the monopoly of managing the ledger entries.

2 Economic rents occur when a payment ensues on a particular factor of production that is above the actual cost required to use that factor in the production of a good or service. Therefore, to reshuffle rents at the expense of the users would mean that the user is still forced to make the higher payment, but now the additional cost simply goes somewhere else.

What is Blockchain Technology?

A blockchain is a distributed ledger among the participants in a peer-to-peer network that allows one to keep records and execute contracts or agreements within the network. The ledger is distributed in the sense that all network members store an up-to-date copy. It is also decentralized since any peer has the right to update or maintain the ledger. Hence, the key feature of distributed ledger technology is that it can keep and update records without the use of a trusted third party.

How the ledger is maintained, how it is reliably updated and how actions based on the ledger are executed form the blockchain rules. For a successful regime, one needs to ensure that a distributed ledger is hard to forge, that the copies remain synchronized and that every member can trust the ledger. The blockchain thus becomes the trusted party itself (Box 1 illustrates the main elements of a blockchain through bitcoin, probably the most well-known implementation of the technology).

The blockchain can be seen as a book containing the ledger. Simply put, it is a list of signatures assigning ownership or expressing decisions and actions of its participants. New transactions or actions conducted within the network are essentially new ledger entries and are broadcast over the network to all participants (or nodes). The nodes can easily verify the authenticity of these transactions and whether they were conducted correctly.

A key blockchain feature is how it is updated with these new, verified transactions. In general, transactions are pooled into a block that can be seen as a new page in the overall ledger. The network members then compete for the right to update the ledger. Whoever wins this competition receives compensation and will update the ledger where it is again easy to confirm that the update was conducted in the correct fashion. After a node confirms an update, it will decide to work with what is likely to be the most accurate copy of the ledger across the network. This ensures that all nodes, once receiving and confirming updates, will work with the same ledger.

With this process, one is able to design a protocol that makes the ledger hard to forge while still giving participants access to the entries that they own and allowing them to conduct transactions that can be verified across the network. Similarly, participants have an incentive to work with what is likely to be the most accurate version of the ledger as long as they can trust that the blockchain is a correct representation of past transactions.

Up to this point, we have been discussing ledgers as being openly distributed or public, where everyone has the possibility to directly access and potentially update the ledger. An important alternative is the so-called private distributed ledger where only authorized participants have direct access and/or the possibility to update it. Such ledgers are often attractive to existing intermediaries since they allow them to maintain their central position while using some of the advantages offered by blockchain technology. In such systems, the main challenge is to provide fair access to the ledger while finding an efficient way of updating and maintaining it.

Furthermore, one can identify areas that will require the government to assume an important role in both updating and maintaining the ledger. In these situations, distributed networks would exist with the government acting as a centralized administrator while the ledger itself is distributed and can be accessed by the general public. We will discuss the functioning of these different forms of distributed ledgers and their trade-offs at a later stage.

In the context of open distributed ledgers, achieving trust is an essential element that can only be established through the ledger's design itself, as there are no entities that are charged with running the ledger. Such trust is established by making it difficult and somewhat expensive to update the ledger. Updating the ledger establishes

a new, additional path of transaction histories, or blocks, where each update is connected with the previous one – hence, the term blockchain. With this procedure, forging the ledger by changing transaction histories becomes more and more difficult and expensive. In a sense, the further a transaction in the ledger goes back, the more resources it would take to alter the record. Consequently, the trust among network peers concerning a particular transaction increases with the length of the block chain.[3]

But the problem remains that each update itself needs to reflect only valid transactions and is adopted throughout the network so that all members work off the same ledger. This result is achieved by having members engage in a contest for the right to update. The major innovation here is reliance on a combination of computer science and economics to get the balance right between participating in the contest and having little incentive to cheat when updating the ledger (for a thorough discussion of how to cheat within a blockchain see Box 2).

When blockchains were first introduced, the idea was to have members solve a computationally intense problem where the first member solving it gained the right to update it and earned a reward that was linked to the blockchain being used in the future. The contest required investment in computing power, while the reward made it attractive for members to participate. This implied that, on the one hand, cheating was expensive. A dishonest node needed to spend a lot of resources to have a fair chance of winning by fudging the ledger while honest members of the network competed to also solve the problem by investing resources. On the other hand, the expected gains from cheating were small, as the gain for a dishonest member tended to be limited as long as other, honest members eventually learned about a fraudulent update and potentially gave up on the blockchain as a valid ledger.

In more recent blockchain implementations, this so-called "proof-of-work" protocol, where members verify that the contest winner has spent a lot of resources in updating the ledger, has been supplanted by less costly and potentially more efficient protocols.[4] What these protocols have in common is that there is a cost for obtaining the right to update the ledger, which is most apparent in the permission-less, i.e., purely open, networks. Hence, blockchain technology is not a free lunch, as we will discuss in more detail in the next section.

A related problem is how to ensure that the updated ledger remains consistent for blockchain-system members. Inconsistencies arise if different parts of the network work with different versions or forks of the blockchain. Hence, the protocol needs to spell out rules regarding which chain to follow when they are in conflict. It is thus also clear that these rules are intimately linked to instil trust among members. One needs to be able to rely on one's copy of the blockchain to verify and confirm transactions. Sufficiently fast communication over the network is thus important to ensure that a common blockchain emerges consistently from

3 A distinction can be made between blockchain applications that keep track of all transactions and those that maintain only the state of a particular system. The latter form – sometimes referred to as a consensus ledger (see ECB 2016) – is usually not considered a proper chain as it does not need to keep track of the entire transaction history.

4 A prominent alternative concept is "proof-of-stake," where nodes have a probability of being able to update the ledger based on their stake in the chain. Nodes determine their stake by committing to give up temporarily the right to conduct some possible transaction or actions, thereby essentially pledging ring-fenced collateral. The probability of updating the ledger is then determined by a node's stake relative to the entire network's stake. Other concepts are "proof-of-activity," which combines the two protocols and "practical Byzantine fault tolerance," which is only feasible in a relatively small network (see Bank of Japan 2016).

updating. This tends to put limits on either the size of the network or the speed at which the chain can be updated. In a large, sufficiently distributed network, updating too frequently will lead to too many forks, thereby undermining the accuracy of and the trust in the blockchain.

The Economics of Blockchain Technology

Having presented the main building blocks of a distributed ledger, the question arises when and why this technology adds value to the economy, which relies to a large extent on record keeping. Financial markets are the best example, where double book-

Box 1: Bitcoin – The Original Example of a Blockchain-Based Distributed Ledger

Bitcoin was introduced in 2009 as a cryptocurrency, a currency that exists only virtually, and is regarded as inventing blockchain technology. The currency itself – bitcoins or its smaller denomination satoshi – is in the form of an "address," which is a sequence of bits that can be stored within a computer program called a "wallet." In principle, this makes the bitcoin holder anonymous unless an address can be associated with a wallet and a wallet with a person.

Bitcoin transactions, in general, take place on the bitcoin network, which is open to everyone and requires participants to simply install the appropriate free software.[5] When conducting a transaction, the bitcoin owner sends a message together with a signature over the network specifying that bitcoins are being sent or associated with a new address. All participants can easily verify that transactions are correct, since addresses associated with unspent bitcoins can be identified in the stored ledger or can at least be accessed by the entire network. This ledger is called the blockchain.

The blockchain contains a record of all transactions that have ever been conducted. It is based on a chain of blocks where each block groups together individual transactions once they have been verified as valid. A new block is added to the chain through a process of "mining." Network participants compete for the right to add such a new block of verified transactions. Once a new block has been added, the transactions within the block are considered confirmed. Since new blocks are confirmed based on the previous block's information, the more blocks that are added to an existing block, the more confirmations are given to a transaction within that block.

Mining is based on proof-of-work and is one of the key innovations associated with blockchain technology. It is costly to gain the right to add a new block, as one needs to solve a puzzle that can be done only by using brute computational force. There are fixed costs of buying computer equipment and energy costs associated with operating and cooling the equipment. Proof-of-work makes altering the transaction history prohibitively expensive, making the ledger tamper proof and, consequently, trustworthy. In order to induce mining, the bitcoin protocol relies on transaction fees and a reward in newly created bitcoins for winning the competition to update the ledger. The expected reward from mining amortizes the cost spent on computing power.

5 Off-chain transactions are also possible. Under this scenario, a series of payments are made between different counterparties off the bitcoin network with more instantaneous transferring. Eventually, the transactions are added to the bitcoin network blockchain.

Box 1: Bitcoin (Cont'd)

To ensure that the ledger remains consistent, any newly mined block is broadcast across the network. Once the information is received, participants can easily verify that the node sending the new block has done the work and solved the puzzle. The bitcoin protocol specifies that anyone in the network needs to work with the longest chain. Hence, upon receiving a new block, which is controlled by an algorithm that adjusts the difficulty of the puzzle so that a new block is mined about every 10 minutes, the entire network updates the ledger and starts working on confirming a new block based on the just-added block.

The safety of bitcoin transactions depends on basic cryptography. Addresses and signatures are generated by using private/public key generation. They can only be produced by whoever holds the private key and verified by anyone seeing the public key, which can be sent with the signatures.[6] Hence, only the private key holder can spend a bitcoin just like currency. However, losing the private key or having it stolen from a wallet is like losing cash. It is nearly impossible to prove that one has owned a particular bill or key. Recent reports about security breaches at bitcoin exchanges or other platforms that operate blockchains underscore that individuals are responsible for the safety of their bitcoin storage. Notwithstanding, employing some basic rules can make storing bitcoins safer than cash.

(For more technical details on how bitcoin works see Antonopoulos (2014). A more accessible resource for understanding the basics of bitcoin is http://www.michaelnielsen.org/ddi/how-the-bitcoin-protocol-actually-works.)

6 Importantly, the receiver of a message with private/public key cryptography can also reply by using the public key that can be decrypted only by the private keyholder. This allows two parties to exchanges messages in a secure but economical fashion.

entry systems keep track of payments, ownership of assets and obligations from contracts. Indeed, one can argue that the introduction of formal record keeping and the development of supervising intermediaries was a cornerstone of developing a modern economy with decentralized transactions.

Blockchain technology offers a new opportunity for record keeping without relying on a trusted third party. In many of its applications, the technology also employs additional uncommon features. Examples are the use of cryptography for secure communication that does not need to rely on cumbersome encryption and data storage distributed among users. In fact, it is quite common to equate blockchain technology with these features, especially in the context of mere data management.[7]

In our discussion, we will concentrate on the trade offs that this new technology entails relative to an established record-keeping system that is based on the use of intermediaries who require incentives to act as trusted third parties. We will do

7 One can then speak of a mutualized data management system where security is based on cryptographic features compared to a secured central database with secure access (see for example Accenture 2015). It is easy to see that the former has attractive security features as hacking into a database without access to the cryptographic keys is pointless.

Box 2: Double Spending in Blockchains

Altering the history of transactions within blockchains in order to cheat can be nearly impossible, since one tries to undo confirmed blocks backwards while the network is working on adding additional blocks forwards, making the chain longer and, thus, accepted across the network.

Fraud within a blockchain technology arises instead as the problem of "double spending." Think for a moment about the bitcoin protocol. With cash, spending it twice is impossible. With bitcoin, however, one can in principle use an address to conduct two transactions simultaneously. The idea is to spend bitcoins for a real transaction, but then to send the bitcoin also to one's own address at the same time. For a blockchain to maintain its value, it must safeguard against these double- spending attacks.

For an attacker to be successful, double spending relies on two components. First, the attacker must conduct a regular transaction. The counterparty then has to receive this transaction and agree to it and settle the obligation. Second, the fake transaction must be included in the blockchain. This implies that the cheated counterparty cannot use what it has received at its address in the regular transaction, but the attacker still can. An easy way to protect against such an attack is to wait for at least one confirmation; i.e., the counterparty waits to see whether the transaction is included in one or more new blocks that are added to the chain before fulfilling its own part of the transaction.

A slightly more sophisticated attack is to create a fork in the chain. An attacker needs to mine blocks that include fraudulent transactions himself faster than other nodes in the network mine new blocks. He can then choose to not release blocks at first so that the real transaction or action is confirmed and thus settles. Afterwards, he releases a chain that has an equal number of blocks (or more) as the correct chain. To be successful with such a strategy, an attacker either needs to have a large share of the total computing power or get extremely lucky by outrunning the network in finding new blocks. For example, bitcoin amassing more than 50 percent of computing power leads to a 100 percent probability of eventually outrunning the chain.[8]

This implies that double spending attempts are not without a cost. Attackers need to spend significant amounts on computing power and on mining new blocks. Moreover, one can protect against these attempts by waiting for a sufficient number of confirmations of the transaction in the form of new blocks, which decreases the probability that eventually the current chain gets outrun by a fraudulent, different fork.

(For a discussion on double spending as it relates to the bitcoin protocol, see Rosenfeld (2014) and for cryptocurrencies more generally, see Chiu and Koeppl (work in progress)).

8 Recent research has challenged this fact. As Sirer et al. (2015) show, an attack has a can be successful with high success probability as long as the attacker amasses about one-third of the computation power. Further concerns have arisen in the context of bitcoin context as so-called mining pools have reached computation power close to this threshold.

so along four different, but connected dimensions: safety, efficiency, incentives and dynamic adjustment.

Safety

Clearly, keeping records safe and tamper proof is of critical importance. One can argue that a distributed ledger is safer than a centralized one: since an identical ledger is stored ideally on every network node, there is no single point of failure as exists in a centralized network where a third party is charged with maintaining and safeguarding the ledger. Furthermore, redundancy is built into a distributed system and, thus, makes it more failsafe.

With electronic communications and record keeping, a centralized system is not necessarily very robust. An attack on the central node can bring the entire network down or compromise critical data records stored on some dedicated data point. With a distributed system, losing a single node or multiple nodes is usually not a problem.[9]

Furthermore, since many blockchain elements are built directly on cryptography (see Box 1), there is no need to protect communications within the network. In a centralized network, however, messages between the centre and the nodes often need to be protected. Similarly, distributed networks work with individual, cryptographical security features that are much harder to compromise, provided the individual nodes or members are able to safeguard the information appropriately.

Notwithstanding, blockchains have to overcome some inherent disadvantages. A distributed ledger always faces the prospect of inconsistencies in data stored among different nodes. As we will discuss further below, this puts limits on its efficiency and is likely to lead to some extra costs when running the technology.

Similarly, any new blockchain technology application also will have to first build up trust among its potential users. The participants in a distributed ledger are typically not readily identified. This will create some perceived insecurity, at least during its introduction. It will also take time for the blockchain to be deep enough to become tamper proof. This is not the case with a third party introducing a new or running an existing record-keeping technology, especially when this third party has already gained a high degree of trust within an economy.

Importantly, some of these concerns can be mitigated in a blockchain where only a few existing intermediaries with a high reputation for trust have full updating privileges. In such permission-based or restricted blockchains, these intermediaries use the technology while continuing to play a central role by allowing end users indirect access to the information stored in the chain.

Efficiency

There has been much discussion about a distributed ledger's efficiency. For example, storage costs are larger since the ledger is being held by many or all nodes in the system. Communication can be more cumbersome and slower in a distributed system since nodes need to be connected, possibly in a hierarchical way compared to a centralized system that needs to have connections only between the centre and its nodes. On the other hand, blockchain technology offers a way to harmonize transaction protocols and to facilitate the processing and

9 The same principle applies to distributed computing where critical software applications run simultaneously on multiple nodes. Such redundancy can make the overall system more stable at the cost of ensuring that the applications stay coordinated.

sharing of information. These advantages could potentially lead to a substantial reduction in costs, especially in the area of financial markets.[10]

One particular concern is that many costs in a distributed ledger tend to be duplicated. An apparent major bitcoin flaw, as the first application of a blockchain, is that it duplicates the costs of updating the ledger (see Box 1). When network nodes compete with each other for the right of updating the ledger, resources are wasted, which is not necessarily the case when a neutral third party is charged with the task. Other costs arise whenever a proof-of-work protocol is being used, since confirmations in the form of ledger updates take some time to be produced.

Over the last few years, however, these costs have been partially alleviated in two ways. First, verification and confirmation times have been drastically reduced by alternative applications while maintaining the basic proof-of-work protocol. Second, alternative protocols have been introduced that experiment with avoiding altogether the duplication of effort. These have either improved the blockchain efficiency or clearly have the potential to do so. Notwithstanding, there remain four possible efficiency bottlenecks relative to a centralized, third-party, record-keeping model.

The first is related to network latency, or how quickly information spreads through the network. A distributed ledger can be effective only if its network information is fairly homogenous across the individual nodes. There appears to be a minimum time required between updates to ensure that the ledger remains consistent without too much forking occurring that would undermine the accuracy and, hence, the usefulness of the ledger itself.

The second potential efficiency bottleneck is related to how much network members trust the information contained in the ledger. With a blockchain, every update confirms previous updates and ensures that older information remains within the blockchain. Hence, participants need to wait for a number of confirmations so they can be sure a transaction will be reflected appropriately and securely. Such a delay is costly whenever one needs a transaction to be in real time; i.e., when a transaction requires immediate finality.

The third relates to the cost of updating the ledger. The mere knowledge that it is difficult and costly to add new entries to the ledger generates faith among the members that the blockchain contains accurate information. This faith in the system allows it to operate efficiently without requiring additional internal checks at either the member or user level. Consequently, incentive considerations will determine how costly it will be to run a blockchain, as we will outline in more detail below.

The fourth and final possible bottleneck has to do with scalability. The throughput capacity of current applications is not compatible with the transaction volume in many financial and real markets. For example, bitcoin can support only a miniscule fraction of all payments that need to be made daily within even a small economy. Although many efforts are under way to increase scalability, it remains unclear at this point whether distributed ledger technology in its purest form can support trade to an extent that is required in many potential applications.

It is important to realize, however, that restricted or private distributed ledgers hold the promise to reduce these bottlenecks substantially. Putting some central nodes in charge of updating and maintaining the ledger combines existing infrastructure with the strengths of this new technology. For example, blockchain protocols could allow for a cost-efficient, straight-through processing of transactions that is currently often not possible with existing payments infrastructure.

10 For a detailed discussion, see ECB (2016).

Incentives

With any record-keeping system there are incentives to defraud it. A third party in charge will weigh the short-term costs and gains from changing the ledger to his advantage against the revenue it can generate in the long run from being trustworthy.

Similar considerations apply to blockchain applications. First, it must become difficult and costly to update the ledger. This implies that there should be few incentives to defraud the ledger since failed attempts carry a sunk cost. Second, updating the blockchain should usually come with a reward. This ensures that the network nodes have an incentive to gain the right to update the ledger. If many nodes seek this reward and thus act honestly, they will reduce the likelihood that a single fraudulent node will succeed in defrauding the chain. The rewards for malevolent actions should also be limited, as once a fraud has been detected, there is little incentive for maintaining the existing blockchain.

A blockchain becomes more secure as the responsibility for the ledger is distributed among more nodes. Attacks tend to be less successful as the number of nodes increases. This implies a trade off with many of the costs we have outlined in the previous section, as the number of nodes increases the costs of maintaining the blockchain.

Ultimately, it is thus open to debate whether the incentives to cheat are larger in a third-party ledger or with a blockchain. The actual application and context of a blockchain is likely to decide whether it is the better system or not. Importantly, one has to realize that in many circumstances a restricted chain can combine advantages from both alternatives. The designated central nodes in such a chain are likely to be existing, well-known trustworthy third parties that could be held accountable for maintaining the ledger. Even though the ledger would not be truly distributed and thus more vulnerable to outside attacks, one would need fewer incentives to prevent malevolent behaviour within the chain.

Dynamic Adjustment

This brings us to the final important challenge of this new technology: how can a blockchain evolve dynamically over time? As with any other network, having a distributed ledger raises the question of how to deal with externalities. Potential participants need to coordinate on its design and on who will be granted access.

Furthermore, as the needs of the network evolve, adjustments are likely to be necessary. Being a distributed system, there needs to be consensus among a critical number of participants to make changes or adjustments. Ensuring continuity of a chain is a vital issue here, especially in cases where the ledger's integrity has been called into question. Also, how to replace an existing chain and retain its information is currently not clear.

Finally, there is no guarantee that a ledger intended to be distributed will actually continue to do so. When some participants amass enough power within the network, they will have sufficient weight to determine how the ledger is run and updated. This either creates a public system where only a few members retain all the power or drives one back to a centralized structure where one node can act like a third party administering the ledger.

Our discussion thus points again to a case-by-case evaluation whether a blockchain solution is appropriate. A restricted blockchain where only a limited number, most likely very few, participants have direct access to the chain can alleviate many of the concerns we have raised here. Naturally, one can imagine even hybrid forms of a distributed ledger where access and participation is tiered among direct and indirect participants. The core of the network would then be in charge of the ledger's design features.

Applications of Blockchain Technology

While blockchain technology has been popularized by bitcoin, bitcoin in fact uses several concepts and ideas that reach even beyond the core concept of

a distributed ledger. A prime example is its use of cryptographic principles such as private-public key encryption that are not familiar to many outside the realm of computer specialists. Many such blockchain auxiliary ideas already have and will be applied to increase efficiencies, especially in the Internet-based economy.

Similarly, one can imagine that some distributed ledger approaches transcend a blockchain. As already pointed out, it is conceivable that some (or one) parties retain more control over updating a ledger, even when distributed among network members. Also, participation in the ledger could be restricted to a limited number of members that agree on specific updating rules that are not entirely based on the ones we have presented here or do not implement the idea of an irreversible blockchain.

Abstracting mostly from these issues, we ask here whether blockchain technology can supplant existing forms of record keeping. The answer to this question depends not only on the costs and benefits of the technology, but also on the interests of the institutions and economic actors that rely on the ledger.

Furthermore, as pointed out already, blockchain technology can be quite flexible and allows for different designs. These are likely to be driven by the exact application it is intended for. We will next look at four areas for which blockchain technology already has shown some potential and discuss its future for reshaping existing arrangements in Canada.

Payment Systems

Payment systems are at an economy's core as they have the purpose of settling transactions by transferring monetary value from one person or institution to another. There are two forms of payments systems, retail and wholesale.

Retail payments comprise the bulk of all payments and are linked to lower-value transactions such as the purchase of goods and services and paying one's bills. Here, the mere functionality and the ease in which to transfer payments between counterparties is the main concern.

Wholesale or large-value payment systems tend to involve financial institutions as intermediaries for settling high-value and often time-critical transactions such as financial market or real estate transactions. The volume in such payments systems tends to be lower, but the stakes for the counterparties to receive payments and on time are much higher. A critical feature of these systems is immediate finality, where payments become irreversible as soon as they are carried out, which is commonly ensured by having the central bank backing the transactions.

Retail payment transactions: Blockchain technology has already begun to reshape the payments landscape. With most retail transactions, counterparties are willing to take on some residual risk for accepting a transaction without receiving a guaranteed payment immediately. Similarly, for transactions that are not time critical, counterparties are willing to wait for confirmation of having received a payment before completion. Cryptocurrencies are already being used for such transactions since the costs of their usage can be much smaller than traditional payment systems. Leading examples are using bitcoin for person-to-person (P2P) international money transfers and the Ripple Transaction Protocol, which allows banks to move money across borders without the necessity of correspondent financial institutions.[11] In the

11 Some revenue-pool estimates associated with international correspondent banking are in the area of US$150 billion to $200 billion (see Williams, Gunn, Roma and Bansal 2016). One can assume that a significant share of this revenue is rents that are competed away by alternative, blockchain-based payments in the medium run.

Canadian context, some banks including National Bank, CIBC and ATB Financial, are already taking advantage of Ripple's flexibility in getting money across borders, essentially in real time.[12] Other applications like Litecoin have lowered the costs associated with cryptocurrencies further by being less restrictive and using a less secure protocol, which allows P2P transactions to be processed much faster than with the original bitcoin protocol.

These developments have created some turbulence in the financial sector. Intermediaries face pressure to make retail payments and simple cross-border payments easier, safer and less costly. Settlement and communication platforms such as Visa, MasterCard, Western Union or SWIFT face the risk that they will become obsolete once cryptocurrencies and blockchain applications for cross-border payments become more commonly accepted. What has so far kept blockchain technology from being adopted more broadly, in addition to trust and security concerns, are limits on transaction throughput, a relatively high confirmation latency and a high variability of confirmation times with high transaction volumes. Notwithstanding, many startup enterprises are currently working on resolving or mitigating these issues.

Interestingly, the main blockchain challenge is likely to come from other financial institutions such as banks and credit unions that see the potential to cut out such intermediaries and recapture some of the rents created by third parties in charge of a ledger. At a minimum, one can see that there is room for blockchain technology to provide additional competitiveness and to push for a more decentralized system with more choice for consumers and room for only the most efficient businesses. A major hurdle is coordination among payment-service providers to deliver these benefits in a widely used, uniform network based on blockchain technology.

In Canada, retail payment systems (with the exception of credit card transactions) are built on central clearing arrangements, with Payments Canada providing and operating the necessary infrastructure and the Bank of Canada providing settlement services. Naturally, Payments Canada should thus take an active role in facilitating and coordinating private efforts for using the new technology. A problem here is that any alternative, blockchain-based, private-sector-payments solution on the retail side would likely compete with the engine that Payments Canada is designing for the major financial institutions."[13]

Ultimately, this raises the question whether blockchain technology should be an integral part of the retail-payment infrastructure that Payments Canada is building for the near future. From the perspective of our four major economic considerations discussed above, retail payments seem to be a strong candidate for successful implementation of some blockchain-technology elements.

In particular, a blockchain-based retail payment system could offer a safe implementation of P2P payments. The associated costs and risks seem to be low, as short delays in the settlement and finality of such transactions are acceptable, and the transactions' low values give little incentive to forge the ledger.

Still, the large volume of transactions should provide enough incentives for decentralized verification at relatively low cost. The key

12 See http://www.theglobeandmail.com/report-on-business/banks-to-use-blockchain-technology-to-speed-up-cross-border-transactions/article30545656.

13 Currently, there does not seem to be any effort to incorporate distributed ledger ideas into the main engine for retail payments, the Automated Clearing Settlement System.

challenge for introducing such a system is to avoid unnecessary cost duplication and creating a universal, broadly based network.

Here Payments Canada is in a unique position to spearhead such an implementation of a modern payments infrastructure based on blockchain technology. It is not clear at the moment whether this involves only building the necessary communication protocols or moving to a separate, new system where accounts are also distributed and updated accordingly. One could imagine that within a more complete blockchain based system, banks and credit unions maintain their central role in updating the ledger and providing access for their customers while losing individual control over transaction accounts. Indeed, a restricted, tiered ledger might be necessary to ensure that such a payment solution is not only cost effective, but also can be scaled up to handle a large volume of day-to-day transactions. Ultimately, it is pivotal to get all stakeholders to buy into such a radically new system that will affect bank business models in a dramatic fashion.[14] Clearly, Payments Canada would have to play a strong coordinating role in designing and building such a next-generation distributed ledger.[15]

Wholesale payment transactions: It is less clear, however, how blockchain technology can change the large-value payments sector. Under current blockchain technology, payments cannot be real-time with immediate and ultimate finality and essentially zero risk for the counterparties. The reason is twofold. First, there is a period of time, however short, between processing the transaction and confirming its finality. Second, with blockchain technology, so-called forks can occur, which implies that previously confirmed transactions might not be valid anymore. In state-of-the-art, large-value payment systems, the involvement of a third party such as a central bank guarantees immediate and ultimate finality of payments.

So-called permission-based consensus ledgers are seen as an alternative to a traditional blockchain. In such a system, a small network of nodes maintains the ledger so that alternative, faster consensus protocols can be used that reduce latency and allow for greater scalability with less room for inconsistencies. Settlement finality in such a system might be achievable since the central nodes are known and inconsistencies can be quickly resolved.

The Bank of Canada oversees a large-value transfer system (LVTS) that links the main financial counterparties in order for them to carry out critical payments with immediate finality. Together with Payments Canada, the Bank is currently in the process of modernizing its infrastructure. Interestingly, part of these efforts is the Bank's Project Jasper, which is looking into blockchain technology for developing a new system of settling interbank payments based on a cryptocurrency for large-value transactions called CAD-Coin.

This system is promising to be an innovative blockchain application. As a private network among only a few participants, it could offer the potential to be safer and more stable than a centralized system like LVTS. While such benefits are currently open to debate, several challenges will need to be solved.

14 See Williams, Gunn, Roma and Bansal (2016).

15 Payments Canada has recently announced that it will be rolling out their five-year plan to modernize payments, clearing and settlements. Its major focus is to create a faster system that is safe and secure and allows for the use of greater information within transactions. A lengthy consultation period preceded crafting the five-year plan with consumers and businesses demanding around-the-clock and speedy transaction capabilities where security concerns are minimized and constraints on data transmission are limited.

First, the Bank of Canada uses LVTS as a tool in supporting monetary policy. The Bank sets the overnight rate and generally uses different forms of repurchase agreements to enforce the target rate. However, as part of either liquidity management or exceptional monetary policy circumstances, it needs the ability to control settlement balances, which it does through the LVTS. Hence, within the blockchain, the Bank will continue to need to be able to create and remove balances at its discretion.

Second, the project's success will hinge on creating cost efficiencies. Having a private network will certainly help to contain duplication costs for updating the ledger, but a key design question will be how to structure the protocol to avoid proof-of-work or proof-of-stake concepts that are too costly.[16] Current systems already are quite cost efficient so the bar is relatively high to achieve additional cost efficiencies.

Third, an unresolved issue seems to be how to achieve immediate finality and real-time gross settlement (RTGS) with a blockchain-based, large-value payment system. The central issue here is to have a settlement asset provided within the system that is already part of the ongoing efforts to design a new system. Still, more thought is required for achieving a design using this new asset since it is not clear whether such an RTGS system would, indeed, be beneficial.[17]

Finally, a concern with having high-value transactions is that there is a far greater incentive to attempt to breach the system as the value of one breach can outweigh the value of remaining trustworthy. Hence, a new system will have to be carefully designed to minimize such adverse incentives.

Having a small private network with trusted third parties like the Bank of Canada and Payments Canada overseeing and operating respectively, and perhaps guaranteeing the integrity of the system, will mitigate some of these concerns while still having the potential to generate efficiency gains. However, such a design principle would be somewhat contrary to the very idea of a genuinely distributed ledger. Granted that these are difficult challenges, Project Jasper still has the potential to create a new standard for large-value payments systems in terms of liquidity management, credit risk and operational efficiency.

Smart Contracts

Smart contracts are one of the more interesting areas for the use of blockchains. According to Kiviat (2015), these digital contracts are "computer protocols that facilitate, verify, execute, and enforce the terms of a commercial agreement." Such contracts have already been successfully used by software and data providers within centralized computer networks. Examples include the digital rights management market, in which US copyright law has been embedded into digital files to prevent a user's ability to see, play, print or change the work itself, or Apple's servers that are programmed to enforce the terms and conditions of its iTunes store.

Instead of using a centralized system where a central node is in charge of managing the contracts, blockchain technology allows users to design

16 Proof of stake is similar to proof of work in that they both are designed to generate consensus. In proof of work, the probability that a miner will mine a particular block depends on the previous work the miner has done. With proof of stake, it depends on how much bitcoin a miner owns. The hope is that by forcing miners to have large bitcoin holdings, the incentive to attack the system and jeopardize their wealth will be lower. An innovative idea is to use so-called "practical Byzantine fault tolerance," where a two-thirds majority of participants decides on an update precisely when they learn that at least a two-thirds majority of participants has agreed on the new information (see Castro and Liskow 1999).

17 LVTS currently gives participants two options, a pre-funded real-time gross settlement or a collateral-based net settlement. It is not officially known whether a new system will maintain both options or establish an RTGS system.

contracts that are automatically executed following a trigger event, without having to rely on some form of costly third-party monitoring and enforcement mechanism.

Once again, one can see the potential disruption arising from this new technology. Decentralized smart contracts allow people who may not know each other to transact in a trustworthy way without intermediaries. From an efficiency standpoint, contracts such as these do not need to be executed in real time, so verification delays are often not a concern. And from a safety perspective, such contracts reduce the vulnerability arising from having a dedicated, centralized server for managing the contracts.

A prime example of how a blockchain can be used for these smart contracts is the Canadian-founded Ethereum platform[18] for distributed computing. It allows users to develop and run their own applications in a distributed fashion across network nodes. One of the Ethereum platform's key innovations is to base crowdfunding on a cryptocurrency.

Here's how it can work. Crowdfunding begins with a smart contract being added as a block in the chain. People interested in funding a particular project will contribute funds in the cryptocurrency directly through the smart contract. Then, when the contract's stated time period expires and the overall funding target has been met, the amount contributed will be paid to the project account.

If the target is not met, the smart contract returns the contributed funds to the donors. This mechanism eliminates the need for a third party to design, implement and control the entire crowdfunding process. And once again, having multiple records stored in a decentralized fashion increases computing security, while incentives to fudge the ledger tend to be limited due to the fact that individuals are likely to hold only small stakes in any particular project.

Corporate Governance

More generally, blockchains have the potential to revolutionize corporate governance and corporate control. The standard has been for governance to take place in the form of annual shareholder voting, often via proxies. Maintaining a distributed ledger makes it possible to have votes on corporate issues more often with direct participation of shareholders. It is even possible to pre-program corporate decisions and have them automatically executed via a voting scheme. A blockchain implementation would provide users with tokens that they could transmit to specific addresses by a deadline to cast a vote on a particular issue.[19] As no centralized system now exists for this type of operation, blockchain technology thus offers possibilities for establishing businesses that operate in the digital world and often require individualized, innovative governance solutions.

Existing firms could also use blockchains to perform real-time accounting or create contracts with automatic execution. Since time stamps are attached to each block in a chain and cannot be changed afterwards, a company's ledger can be made available for anyone to see, is up to date and allows an interested party to aggregate all of the firm's transactions into a real-time income statement or balance sheet. This is a vast improvement over today's situation in which current or prospective shareholders must wait for quarterly statements to be made publicly available.

Similarly, contract execution such as bill payments or delivery of services could be triggered

18 See https://www.ethereum.org.

19 See, for example, Yermack (2015).

by actions within a blockchain. This could facilitate the internal organization of firms, once again saving costs and increasing efficiency with potentially less error-prone processes.

Financial Markets

There are also first attempts under way to use distributed ledgers to enable the exchange of financial securities. At a minimum, one can imagine that cryptographic keys being used as a messaging system to sign trade agreements within a ledger for transactions that are later confirmed by central nodes. These nodes could be existing intermediaries; for example, swap dealers or main swap participants in derivatives trading.[20]

One particularly interesting application relates to private equity for which no central registries exist and where exchange of shares is cumbersome. NASDAQ has started to build a blockchain-based platform to improve private equity trading. Given the modest scale of the Canadian market, there is potential to give smaller, privately held Canadian companies access to a liquid financial market. Once again, traditional intermediaries such as banks or specialized lenders will face competition from such new solutions that are based on distributed record keeping and intermediation.

Similarly, one can imagine that blockchains could be used in over-the-counter markets where a central authority is missing for keeping records and providing information on counterparties' exposure. This would increase transparency in financial markets, making it easier for regulators and supervisors to obtain real-time information. As a consequence, regulatory concerns about transparency and value-increasing customized transactions would be muted, ultimately allowing more of these transactions to take place.[21]

Beyond the mere trading of financial instruments, blockchain technology holds considerable promise for achieving true, real-time, straight-through processing of financial transactions. Currently, the post-trading requirements of clearing and settling transactions often take time and involve cumbersome procedures for financial market participants. Using smart contracts for trading, one could imagine that traders could calculate exposures and margin calls right up to the automatic transfer of securities and cash in what is called a delivery-versus-payment mechanism. This could enable a near real-time settlement in cash markets on a gross basis.

Naturally, this means that one would need to integrate or link several ledgers that keep information on asset, collateral and cash positions of market participants. Recently, several blockchain solutions have been tested in this respect. In the US, the Depository Trust & Clearance Corporation in a joint venture with blockchain company Digital Asset has demonstrated the potential for clearing and settling repurchase agreements via a distributed ledger while European blockchain developer Clearmatics is working on a decentralized clearing network.

It is unclear how much blockchain technology will change trading in financial markets. The possibilities range from one integrated ledger for all processes to just some isolated ledgers being used to streamline specific post-trading functions (see ECB 2016 for an excellent, detailed discussion). The potential for simplifying trade protocols and risk management and the associated cost savings

20 One example is Corda, a blockchain-based derivatives trading platform developed by Barclays in a consortium with the blockchain consultancy R3.

21 An open question is whether blockchain technology requires standardization of transactions or contracts. One of the prime benefits of over-the-counter trading is that transactions can be customized to the need of the counterparties.

are beyond doubt. The ultimate implication is that many intermediaries such as clearinghouses and settlement agents could lose their unique positions of being necessary third parties for some services associated with the clearing and settling of transactions in financial markets. Their role could change to mere gatekeepers for accessing the ledgers or to specialized providers of specific services. Indeed, some institutions might disappear altogether, possibly being supplanted by intermediaries such as banks or brokers that give access to an integrated trading and post-trading system.[22]

Finally, distributed ledger technology could also be used for keeping track of real, physical assets. To utilize blockchain technology in this context, information pertaining to the asset has to be embedded in a typical transaction that is part of a block in the chain.[23]

Already, existing record-keeping systems could be moved to blockchains for reasons related to cost effectiveness and cybersecurity. One prominent example is land registries that keep track of real estate titles and transactions. But blockchains also offer an opportunity to record ownership where no centralized system has previously existed. This is the case, for example, for precious metals, gems or artworks. One current application in this area is provided by Everledger, which uses blockchain technology to look after diamonds and other luxury goods with significant financial value. The data in a block is irrefutable information regarding the ownership of the good in question, making it easy to track should it be stolen.

Government Services

A recent UK government study suggests that countries around the world should be able to use blockchain technology to "collect taxes, deliver benefits, issue passports, record land registries, assure supply chain of goods and generally ensure the integrity of government records and services."[24] Many of the gains from blockchains that arise within the private sector should be realistic possibilities for government services as well. Given the sensitivity of information used by governments, however, the advantages we discuss are likely to occur within a permissioned blockchain network.

In this regard, one benefit of blockchain technology is that it could be used to assure compliance with existing laws and regulations. This has the potential of lowering compliance and enforcement costs. It also puts forward an entirely new infrastructure based on so-called public key infrastructure popularized by the bitcoin protocol. This technology also allows one to deliver services and share information over unsecured networks simply by using private/public key cryptography.

For example, the government of Estonia has been experimenting with a distributed ledger technology entitled Keyless Signature Infrastructure that allows citizens to check the validity of their own personal records on different government databases. One can imagine using a similar scheme in Canada to provide a single, secure access point for all services ranging from tax filings to voting in elections, as well as accessing the healthcare system.

Ultimately, a universal personal identification system, where one's ID is stored and protected by the blockchain, is possible. Within our world

22 Discussing specific areas goes beyond the scope of this *Commentary.* Craig Pirrong provides an excellent discussion on central counterparties and blockchain technology at http//streetwiseprofessor.com/?p=10201.

23 For example, Colored Coins is a meta protocol that allows information about asset ownership to be attached to transactions involving tiny amounts of bitcoin. In a sense, bitcoin is merely a vehicle to store other information in a non-dedicated ledger.

24 See UK (2015).

that relies so much on electronic communication, e-commerce and remote administration, a secure and unified ID system with a tamper-proof signature holds the promise of immense cost savings. But even more important, the decentralized nature of blockchains offers an additional degree of cybersecurity that could make such a development a priority. One concern here involves how our services are run through different levels of government. With passports and taxes being managed at the federal level, and health and driver's licenses at the provincial level, managing this collaboration will be essential if the benefits of a unique ID are to be realized.

Policy Challenges

To this point, we have discussed the major economic factors that determine whether it would be useful to employ blockchain technology. Based on these factors, we have also given a few examples of where this technology is likely to succeed. What arises from this analysis is a set of challenges and questions for policymakers and regulators including who gets regulated, who is responsible for compliance with those regulations and which activities related to the technology should be regulated.[25]

Before discussing these challenges, we make two acknowledgements. First, it is important to recognize that new technologies are unique in their ability to solve challenges where and when we least expect them. When policymakers think about how to address these challenges, they must recognize the need to maintain flexibility as new challenges arise and that the new technology provides myriad ways in which to solve them.

Second, what follows is by no means an exhaustive list of challenges and, therefore, of policy responses. Our list merely reflects what we consider to be top priorities. Policymakers in Canada will be wise to review actions being taken internationally both to learn best practices but also to ensure that Canada remains competitive in this fast-growing area.

Safety and Stability of the Ledger

The foremost concern with blockchains is that the new ledger is safe and accessible. This means that users must be able to rely on the accuracy of the information in the distributed ledger and can, therefore, use it reliably to engage in transactions. A similar critical aspect is that users can easily and cheaply access the information whenever necessary. For example, in order to be reliable, a blockchain-based payment system would need to ensure that the ledger itself is accurate, the communication infrastructure is stable and the ledger can be accessed easily so that transactions can be conducted at any time, with only short delays, and in a wide range of situations.

It is often hard to assess in its early stages whether a blockchain application offers such features. While sophisticated users will be able to do so, it is far less clear whether the broad public can do so as well. A prime example that comes to mind is the Mount Gox Exchange where many bitcoin users lost their coin holdings due to a safekeeping loophole. Even more problematic would be incidents where newly created ledgers get compromised due to a design flaw. One leading example is the Ethereum platform that has been recently compromised so that a significant part of its currency could be stolen. Such experiences can lead to a quick loss of confidence in the new technology and problems in the early adoption stages.

25 See http://www.nortonrosefulbright.com/knowledge/publications/138038/financial-institutions-and-blockchain-technology.

Hence, over time, some standards will likely have to be developed to protect consumers from unsafe implementation. This requires a continuous dialogue among regulators, developers and users during the creation of applications. But it also needs a certain degree of monitoring to ensure basic cybersecurity and stability protocols are being followed.

Another concern is the legal certainty of transactions based on the distributed ledger. For users to have trust in the ledger, they must be confident that its information can be used in legal disputes. At the moment, to our knowledge it is not clear how to treat ledger discrepancies or how to change a ledger in response to legal decisions about its entries.[26] This is a particular problem for unrestricted ledgers. One of the design principles in such ledgers often is that everyone has direct access to the ledger, but can remain anonymous. This makes participants potentially unaccountable for their actions. Of course, with restricted ledgers that are permission based, all the nodes can easily be identified and held accountable for their actions. Consequently, in many applications related to financial markets, such ledgers hold a distinctive advantage.

Another particular concern is that any change in the ledger needs to be coordinated and accepted across the network nodes. It is currently not clear how a blockchain would handle mistakes in transactions that have been undone retrospectively in the ledger. Standard blockchain applications like bitcoin, for example, are irreversible and do not have a mechanism that would allow past transactions or records to be adjusted. Similarly, any well-functioning payment system needs to provide a mechanism to undo transactions if requested by the original counterparties.[27] Without a central settlement engine, it is not clear how such a change could be reflected quickly in the ledger.[28]

Anonymity and Regulatory Entry

The very design of a distributed ledger where nodes remain anonymous implies that no party can be readily identified that would be responsible for implementing or complying with regulations.[29] Indeed, the members of a truly distributed system that actively maintains the ledger are often hard to identify and possibly shift constantly over time. One particular concern is that members could avoid regulation by shifting their activity across borders to escape the reach of any particular regulator.

The implication then is a focus on activity-based regulation. Here, rules focus on the activity rather than the institution carrying it out, thus bypassing the need to identify constantly changing "institutions." Activity-based regulation, however, often suffers from regulatory arbitrage. Once an activity has been singled out to be regulated, a network can fairly quickly shift to a slightly altered activity, thereby avoiding such regulation. One can see that this puts enormous strain on current regulatory bodies to keep up with a fast-evolving technology.

26 Accenture is looking to patent an editable blockchain (Accenture 2016). However, it is not clear how a blockchain that can go back and make amendments will fit in the larger regulatory scheme and assist in resolving legal disputes.

27 All LVTS payments are final and irrevocable. Any adjustment can be done only through a reversing transaction. It is not clear how that would work under a blockchain-payments system.

28 Such concerns seem to be mute in consensus ledgers that simply agree upon and update the state of records without maintaining a continuous record of transaction histories. Such ledgers are often permission based and have only a few participants with direct access.

29 In principle, it is possible to trace participants and transactions in distributed ledgers. This is part of the reason why bitcoin – although heralded as granting anonymity – is not seen as a concern for law enforcement. Recent developments, however, such as Zcash claim to have achieved full anonymity rather than what some experts describe as "pseudo-anonymity."

A principle-based approach to regulation has worked in the past to contain such problems in the area of payments and settlements. However, one needs to recognize that such success was tied to having intermediaries in charge of critical infrastructure. With a blockchain-based payment system, however, there might not be a direct counterpart for regulators to monitor to ensure that principles are being followed.

As we have pointed out, however, blockchain applications within financial markets are likely to be permission-based networks with only a limited number of participants. This facilitates regulation and supervision in the sense that it might be possible to still hold institutions accountable for their trades and actions.

In this context, however, two different challenges arise. First, participants in financial markets need to stay anonymous as far as their trading strategies are concerned. Hence, distributed ledgers will have to be designed to respect such anonymity among members, but still have to ensure transparency for regulators.[30] This points to a more sophisticated ledger where different entities have different rights to access information. Some new efforts have been undertaken where regulators hold special private/public keys to have unique access to specific information stored in a ledger.

Second, even in permission-based chains, one has to update the ledger with information that parties might want to keep private from other members. While this seems to inhibit confirming and updating transactions within the ledger with regular protocols, new methods are being designed based on so-called "zero-knowledge proofs" to make distributed ledgers a possibility for trading in financial markets. Once again, regulators and clearing and settlement agents, including a central bank, would need to have access to the information even within an environment where the network members have no access to this information.

Coordination and Network Dynamics

Distributed ledgers are network based and, hence, their viability relies on being accepted among a sufficiently large number of users. In other words, blockchain technology involves network externalities where the benefits of an application increase with the number of participants. Adopting a new blockchain application thus suffers from the requirement that one needs to gain a critical mass of participants to reap the benefits of the technology.

The adoption of any new ledger will take time. For example, if a new blockchain-based payment system is to be successful, one needs to ensure that there are sufficiently many potential payees and payers for any individual person or business to accept payments made through this technology. Similarly, too many competing applications that are not compatible may cause a barrier to adopting any payments solution based on this technology.

Hence, one needs to have some degree of coordination when introducing a blockchain alternative or a new application of this technology. Furthermore, the immediate gains from the technology might not be large enough initially for a single private entity to introduce a new application that is beneficial in the long run for itself and an efficient solution for the economy.

This implies that existing record-keeping systems will not necessarily be transformed into newer blockchain-based ones, even if it would be efficient to do so. While this does not suggest direct involvement of government, it points to a role for leading industry participants and government to act

30 Indeed, after the 2008 financial crisis, regulatory efforts such as the *Dodd-Frank Act* and the introduction of ESMA (European Security Markets Authority) in the European Union make such transparency a cornerstone of financial regulation and supervision.

as facilitators for modernizing such infrastructure and achieving sufficient harmonization across different national or regional regulators.

Fortunately, in the payments area, Payments Canada as a private-public partnership is well situated to take a lead role in facilitating the design of an appropriate infrastructure, setting principles such as fair access for payment providers and determining who can update and modify the distributed ledger. It could also spearhead efforts to achieve sufficient standardization that allows one to operate a new blockchain-based payments technology nationally and possibly integrate it into cross-border payment applications.

Unfortunately, it is also conceivable that some of the technological change is being driven by rent-seeking and not by efficiency considerations. Rent-seeking behaviour seems especially likely to materialize in the case of financial markets, where blockchain technology offers a tremendous opportunity to reshape market infrastructure.

For many intermediaries, there is an imminent risk of being made redundant by the principal users of the infrastructure itself. At the same time, blockchain technology offers the most benefit in permission-based ledgers with only few direct participants. Consequently, intermediaries at every layer of the trading and post-trading landscape are bound to exploit new market opportunities. In this scenario, if traditional intermediaries are simply replaced by other ones that branch into different areas of financial markets, then efficiency gains and perhaps, more importantly, cost cutting for end users will not necessarily be realized with existing rents simply redistributed.

Therefore, policymakers should be wary of blockchain applications that are designed to purely cut out intermediaries for the sake of recapturing rents rather than to realize efficiency gains that ultimately benefit consumers. Now, there can be positive economic externalities that emerge from a redistribution of rents that do not completely get passed down to end users. However, policymakers need to be aware of how large and truly positive those redistribution externalities are to ensure that regulation and the spread of this technology are appropriate, given incentives to maintain or obtain rents.

Public, Private or Administered Networks

Coordination problems and network considerations raise the question of whether distributed ledgers can be arranged in a purely decentralized fashion for critical infrastructure. As we have pointed out repeatedly, most applications are also likely to be driven by new or existing businesses. Thus, it seems reasonable to expect that future blockchain applications will rarely be fully public networks where every user has unrestricted access to the ledger.

Consequently, private networks that retain the distributed nature of the ledger, but restrict the right to update and modify it, are thus the most likely outcome. Of course, traditional intermediaries -- for example, banks and settlement providers in payment systems -- would still play a central role within the blockchain. How users can access and use the ledger then become important design principles. This points once again to private-public partnerships where public involvement guarantees fair access for payment providers and safety for users, while designated parties maintain and update the ledger. Such solutions have the potential to most efficiently use parts of the blockchain idea such as cryptographic communication and the distributed nature of the ledger.

In certain areas, the government would clearly need to take on a more active role in supplying a blockchain application. Any area that requires a very high degree of security should see the government taking a special role in updating and maintaining the ledger. One solution would be distributed networks where the government serves as the network's centralized administrator.

Of course, it is then unclear whether the ledger has to be operated as a blockchain per se. It seems

most likely that only features such as public/private key security and distributed storage of the ledger would be essential in such administered networks.

One example would be an integrated national or international system that allows people to access all kinds of online services, Internet sites and computer applications with a unique ID stored and maintained in a single, but distributed ledger. While the records might be distributed among many servers, the updating procedures would probably remain under centralized government control.

Indeed, some recent developments like Jasper see a public sector involvement to integrate private efforts to design a new blockchain application. Here, the Bank of Canada and Payments Canada would have to play a special role in administering a distributed ledger. While the ledger is distributed and possibly jointly maintained by private participants, there are third parties that can alter and change it.

One could imagine a similar setup on the retail payments side as well. It seems technologically possible to build a genuinely distributed payment system where Canadians could have transaction accounts that are distributed across the network with payment providers facilitating access to the network, and financial intermediaries providing auxiliary services related to these accounts, but not owning them any longer.

How to update such a distributed ledger seems then to be a rather minor design issue. A small number of designated settlement providers is likely to be more efficient than running a fully distributed ledger where all users in principle also have this possibility. A key impediment to achieving such a payment system is that legacy systems offer rents to intermediaries. As well, network externalities will slow the adoption of such a new system. To overcome this inertia, it will take some effort and leadership from the public side, possibly directed through Payments Canada, which can be a forum to initiate and guide the transition.

Concluding Remarks

Blockchain technology has the potential to transform dramatically how a modern economy deals with maintaining and updating records. This innovation has already created lots of turbulence in financial markets and beyond. It will be a challenge to let markets figure out how to best use this technology while ensuring consumer safety and efficiency. Our goal has been to unveil the potential of blockchain technology and guide regulators in how to approach the challenges this technology entails.

It will also be crucial to achieve safe and secure applications of this new technology without stifling innovation. In determining this balance, policymakers and regulators will have to decide whether to design rules and regulations along a principle-based methodology like was done with the Internet in the 1990s or whether to operate on a case-by-case basis. Given the uncertainty of how blockchain technology will evolve, it seems reasonable, however, to rely only on a relatively narrow set of guiding principles with a view that allows the technology to develop flexibly in different directions over time.[31]

The key challenge will be to define public involvement in blockchain applications. Many of its most promising uses lie in areas where critical infrastructure is concerned. Governments do not usually have the same expertise and incentives as private businesses to pursue new technologies. Hence, blockchains will test whether public-private partnerships can really implement frontier technology in a cost efficient and safe way.

31 See http://www.brookings.edu/blogs/up-front/posts/2016/02/25-bitcoin-explainer.

References

Accenture. 2015. *Blockchain-enabled Distributed Ledgers: Are Investment Banks Ready?*

Accenture. 2016. *Editing the Uneditable Blockchain: Why Distributed Ledger Technology Must Adapt to an Imperfect World.*

Antonopoulos, A. 2014. *Mastering Bitcoin*. O'Reilly Publishing.

Castro, M., and Liskov, B. 1999. *Practical Byzantine Fault Tolerance, Proceedings of the Third Symposium on Operating Systems Design and Implementation*. New Orleans.

Chiu, J., and Koeppl, T. 2016. *The Economics of Cryptocurrency.* Mimeo, Queen's University.

Eyal, I., Gencer, A., Sirer, E., van Renesse, R. 2015. *Bitcoin-NG: A Scalable Blockchain Protocol*. Mimeo, arXiv:1510.02037v2ECB. 2016.

Kiviat, T. 2015. "Beyond Bitcoin: Issues in Regulating Blockchain Transactions." *Duke Law Journal*. Vol 65: 569.

Nielsen, M. 2013. *How the Bitcoin protocol actually works*. Blog post: http://www.michaelnielsen.org/ddi/how-the-bitcoin-protocol-actually-works/.

Norton Rose Fulbright. 2016. *Financial institutions and blockchain technology*. Internal Publication. http://www.nortonrosefulbright.com/knowledge/publications/138038/financial-institutions-and-blockchain-technology.

Olson, P., and D. Wessel. 2016. *The Hutchins Center Explains: How blockchain could change the financial system, Part 2*. Brookings Institution Up Front blog post. http://www.brookings.edu/blogs/up-front/posts/2016/02/25-bitcoin-explainer.

Rosenfeld, M. 2014. Analysis of Hash-rate Based Double Spending, mimeo, arXiv:1402.2009v1.

Santo, S., Minowa, I., Hosaka, G, Hayakawa, S., Kondo, M., Ichiki, S., Kaneko, Y. 2016. "Applicability of Distributed Ledger Technology to Capital Market Infrastructure." JPX Working Paper.

Sirer, E., and Eyal, I. 2013. *Majority is not Enough: Bitcoin Mining is Vulnerable*. Mimeo, arXiv:1311.0243v5.

Stafford, P. 2016. "Canada experiments with digital dollar on blockchain." *Financial Times*, June 16. http://www.ft.com/cms/s/0/1117c780-3397-11e6-bda0-04585c31b153.html#axzz4EsWZJXSh

UK Government Chief Scientific Adviser. 2015. *Distributed Ledger Technology: Beyond Block Chain*.

Williams, G., D. Gunn, E. Roma, and B. Bansal. 2016. *Distributed Ledgers in Payments: Beyond the Bitcoin Hype*. Bain & Company.

Yermack, D., 2016. "Corporate Governance and Blockchains." SSRN Working Paper.

Recent C.D. Howe Institute Publications

January 2017	Robson, William B.P. *Enduring Virtues: Saving and Investing as National Priorities for Canada in 2017.* C.D. Howe Institute Commentary 467.
January 2017	Ciuriak, Dan, and Jingliang Xiao. *Protectionism and Retaliation.* C.D. Howe Institute Working Paper.
January 2017	Rémillard, Richard. *Government Intervention in Venture Capital in Canada: Toward Greater Transparency and Accountability.* C.D. Howe Institute Commentary 466.
January 2017	Hicks, Peter. *Toward a New Balance in Social Policy: The Future Role of Guaranteed Annual Income within the Safety Net.* C.D. Howe Institute Commentary 465.
January 2017	Johnson, Jon R. *The Art of Breaking the Deal: What President Trump Can and Can't Do About NAFTA.* C.D. Howe Institute Commentary 464.
January 2017	Sutherland, Jason M., and Erik Hellsten. *Integrated Funding: Connecting the Silos for the Healthcare We Need.* C.D. Howe Institute Commentary 463.
January 2017	Guo, Huijie, and David R. Johnson. "Unfair Advantage? School Fundraising Capabilities and Student Results." C.D. Howe Institute E-Brief.
December 2016	Found, Adam, and Peter Tomlinson. "Business Tax Burdens in Canada's Major Cities: The 2016 Report Card." C.D. Howe Institute E-Brief.
December 2016	Cross, Philip. *What do the Different Measures of GDP Tell Us?* C.D. Howe Institute Working Paper.
December 2016	Dodge, David A. "The Role of Macro-Economic Policies in an Era of Global Economic Stagnation." C.D. Howe Institute Verbatim.
December 2016	Busby, Colin, and Ramya Muthukumaran. *Precarious Positions: Policy Options to Mitigate Risks in Non-standard Employment.* C.D. Howe Institute Commentary 462.
November 2016	Schwanen, Daniel, Jeremy Kronick, and Ramya Muthukumaran. *Playing from Strength: Canada's Trade Deal Priorities for Financial Services.* C.D. Howe Institute Commentary 461.
November 2016	Dachis, Benjamin, William B.P. Robson, and Jennifer Y. Tsao. *Two Sets of Books at City Hall? Grading the Financial Reports of Canada's Cities.* C.D. Howe Institute Commentary 460.

Support the Institute

For more information on supporting the C.D. Howe Institute's vital policy work, through charitable giving or membership, please go to www.cdhowe.org or call 416-865-1904. Learn more about the Institute's activities and how to make a donation at the same time. You will receive a tax receipt for your gift.

A Reputation for Independent, Nonpartisan Research

The C.D. Howe Institute's reputation for independent, reasoned and relevant public policy research of the highest quality is its chief asset, and underpins the credibility and effectiveness of its work. Independence and nonpartisanship are core Institute values that inform its approach to research, guide the actions of its professional staff and limit the types of financial contributions that the Institute will accept.

For our full Independence and Nonpartisanship Policy go to www.cdhowe.org.

C.D. HOWE
INSTITUTE

67 Yonge Street, Suite 300,
Toronto, Ontario
M5E 1J8

Canadian Publication Mail Sales
Product Agreement #40008848