François R. Velde
Senior Economist and Research Advisor
Federal Reserve Bank of Chicago[1]

[1] *The views expressed here do not necessarily represent those of the Federal Reserve Bank of Chicago or the Federal Reserve System.*

Money and payments in the digital age: innovations and challenges

Information technology (IT) is having a growing impact on the financial industry. In some ways this is not new: computing power has been harnessed by banks and other financial intermediaries for decades. But now it is IT's ability to process information (the "I" part of IT) that is opening new avenues. Indeed, finance to a large extent is a matter of information, or lack thereof. In a frictionless world with no informational asymmetries and perfect record-keeping there is no place for financial intermediation. Now that increasing numbers of transactions are taking place in a realm (the Internet) where information can be acquired and exploited in novel ways, financial intermediation will be transformed.

This essay focuses on one early development, namely distributed ledger technology (DLT), starting from its use in creating a money-like asset, Bitcoin. I first briefly review the past of money, using the insights we can gain from this new technology into the nature of money. I then turn to DLT, exploring its basic features, asking what promise it really holds. I conclude with some thoughts on how central banks may have to react to these developments.

1 An overview of monetary evolution

Tokens

Why does money exist? We can start from the classic presentation of the barter problem, whose earliest known formulation is nearly 2000 years old. The setting is one of decentralized interactions. You and I meet and we each have some good. But I want what you have while you do not want what I have, which makes a direct, *quid pro quo* exchange impossible.

This problem arises only under certain assumptions. Some are natural: there is diversity in tastes and in goods, production and encounters occur at different times and cannot be synchronized. But the most critical assumption is the lack of information and record-keeping technology. If we had a ledger where we kept track of what everyone had exchanged, I could get what I want from you, inscribe it as a debit on my account, and later I give what I have to someone else and credit my account.

The classic solution, which allows many trades to take place that would not otherwise, is to use secure tokens that will embody your claim on society's resources arising from your act of giving me what I want. This is called money.

A new problem arises: how to make the tokens secure? Again, the classic solution is to make them costly to counterfeit, by using a costly material: gold or silver, which have been used to make coins since the 6th century BC.[2] While the exact origins of coined money are still obscure, coins have usually been produced in a standardized format and certified by political authorities, either directly or under license.

Tokens are now costly to counterfeit, but also costly to make as well. The real value of resources devoted to securing the tokens is roughly the value of the money stock's content, that is, the total real balances held by the economy, and this can be substantial.[3] There will naturally be a tendency to seek ways to economize on this resource cost.

[2] *Curiously, the very first coins were made of a mixture of the two metals, but within a hundred years pure gold and pure silver coinage came into use.*

[3] *The use of precious metal can also serve another purpose, which is to anchor the price level to the relative value of the metal used.*

One method is to use cheaper tokens, made of a cheaper material such as copper (from the 15th century) or paper (from the 17th century). But counterfeiting becomes increasingly profitable as the cost of making the tokens is lowered relative to their value in exchange, so a combination of technology and enforcement is required to keep counterfeiters at bay. In the 16th century a new technology for minting coins gave governments a temporary advantage in making recognizably better struck coins and allowed them to experiment with token copper, sometimes on a large scale. Later, the introduction of paper money was also accompanied by the use of various techniques (watermarks, counterfoils, high-quality engraving, secret points) to defeat counterfeiters.

Another method is to transfer private liabilities. Suppose that debtor B owes 10 to A and also 10 to C. A can make a payment of 5 to C by instructing B to decrease the first balance by 5 and increase the second by 5. This transforms the debtor B into a "bank", not in the sense of an intermediary between borrowers and lenders but in the sense of an agent whose liabilities are used in payments.[4] The prerequisites for such operations include an efficient technology for keeping debt records, which becomes readily available with the adoption of double-entry book-keeping in the 13th century. Also needed is a technology for sending secure messages. The simplest technology is walking over to the banker's office, but over time paper-based messages developed. The name we still use for them ("checks") indicates that security was a concern from the very beginning.

The earliest centralized payment systems emerged when political authorities set up their own public banks, oftentimes making settlement legal tender and final.

Recurring themes

From this terse overview of monetary history some recurring themes emerge.

One is that trust has always been at the core of money; and trust, ultimately, results from a lack of information. If I knew everything about your motives and your circumstances, I could predict your future actions and choose mine accordingly, without having to trust you. It is also lack of information that precludes nonmonetary (or credit-based) solutions to the lack of double coincidence and hence makes money useful. Understandably, then, technological advances that improve our ability to collect and use information can change both the degree to which money is useful and the form that useful money can take.

Another recurrent theme is that all monetary instruments have been subject to various forms of "attacks". As soon as the first coins appeared in the 7th century BC, counterfeits appeared as well in the form of coins with a veneer of precious metal over a cheap copper core, and the race has been on ever since. There have been various kinds of "attackers." Most were seeking to make a profit, namely the difference between the market value of a real token and the cost of production of a plausible-looking token: the wider the difference, the greater the incentive. Less common but occasionally occurring were what one might call "malicious" attackers, typically State actors intent on impairing an enemy's monetary system.[5] Finally, one might place in the category of "attackers" desperate governments impairing their own currency through debasement or inflation (either one seen as legal counterfeiting).

A final theme is the presence of a central authority. Historically, wherever there is money the State is not far away. That does not mean that privately issued currencies have not existed, but by and large they have either when there was no functioning State, or when the State tolerated these private issues. From the beginnings of coinage or soon after, coins were issued by political authorities (cities, kings). Roman law codified the notion that control of the currency was a regalian right, a prerogative of the sovereign, and this notion passed into both feudal law and Roman civil law of medieval Europe. It is true that weak sovereigns let feudal lords exercise the prerogative, but in most cases they regained control of the currency as their powers grew; or, if they did not, it was because they lost their sovereignty (as in the Holy Roman Empire). By the 19th century it became commonplace to think of currency as one of the marks of sovereignty, as symbolic as flags and anthems, even if States at times tolerated privately issued currencies, such as trade tokens and private bank-notes.

The State's involvement took various forms. At the simplest and the earliest, it was a mere form of certification: the central authority (or its designees) stamped its seal on standardized lumps of metal to certify the contents. Eventually it became the legal privilege to define what is, or isn't, money, in legal terms, which can be seen as a standard-setting function, forming an unambiguous consensus on what will discharge debts, effect payments in transactions, or serve as a unit of account. Modern constitutions (such as the U.S. constitution) and legal codes make clear that money remains a regalian right.

2 Bitcoin and distributed ledger technology

Bearing these themes in mind, let us turn to recent the technological changes that could affect the future of money.

Bitcoin

We may start with Bitcoin because, although it is not the first attempt at creating electronic forms of currency, it is the first to reach widespread recognition and (in some sense) use.

Bitcoin was designed under the pseudonym of Satoshi Nakamoto and launched in 2009. It is a protocol for communicating over the internet, but a highly specialized one. Whereas SMTP transmits e-mails with few limitations on their content, Bitcoin transmits formatted messages about transactions between pairs of agents (sender and receiver). The design problem that Bitcoin solves is to transfer value *over the Internet*, by issuing and managing a quantity of monetary tokens, *without any central authority* but rather letting anyone transfer value or even issue and manage the tokens.

Bitcoin is remarkable. Monetary history abounds in examples of monetary tokens that are not explicitly backed

[4] *In the Ripple protocol, the "rippling" feature potentially allows anyone to play that role in a given transaction.*

[5] *Examples include the British counterfeiting the French paper currency during the 1790s and the Germans counterfeiting British currency during World War II.*

or managed by a central authority, but they usually are tokens with an alternative use: a gold coin can always be melted down and turned into something else. Conversely, there are many examples of tokens that are intrinsically valueless, but there is always an issuing entity, private or public, that is supposed to provide some convertibility, guarantee, or acceptability. Bitcoin is not only the first completely dematerialized token: it is unique in monetary history in being intrinsically valueless (there are no alternative uses to a bitcoin if the protocol ceased to be used) yet it is no one's liability.

I hasten to add that I view Bitcoin as a proof of concept rather than a fully-fledged currency: eight years after its appearance, and five years after it gained worldwide notoriety, the aggregate value of its stock is tiny compared to existing monetary stocks, its use in ordinary transactions remains limited. Its value remains extremely volatile, and (as of writing) seems to be sought either for speculative reasons or as a way to evade capital controls. There seems little chance for Bitcoin to become much more, at least in advanced economies.

Be that as it may, Bitcoin is at minimum a working prototype that has brought attention to its underlying technology: *distributed ledger technology,* also called *blockchain*. I now describe its mains characteristics, not just in terms of creating monetary tokens but also from a more general point of view.[6]

Chart 1

A blockchain

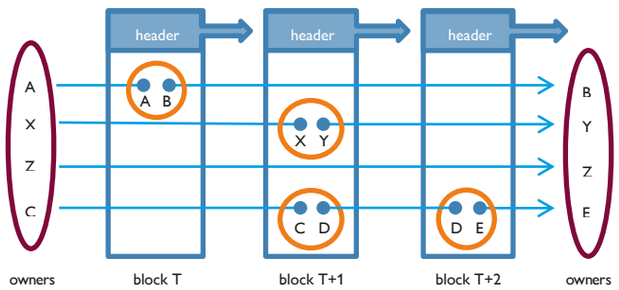

Source: Author's compilation.

Design elements

The original purpose of this technology is to ascertain and transfer property rights. These rights may be over assets that exist independently of the technology or not: the latter are called *native* or *on-blockchain* assets. Assume for now that the assets are well defined: what a native asset is will become clear shortly.

The concept of ownership is as follows: I own an asset X because everyone knows, and agrees, that I do. This differs from physical possession (e.g. cash) as well as from possession based on a registry (e.g., land). One way to formulate the process is recursive. Suppose that at some point in time T there is an agreed-upon state of the world, essentially a list of who owns what (a ledger). The technology provides a process for moving from T to T+1, which will consist in aggregating all valid changes of ownership. The result will be a new agreed-upon state of the world at T+1. The design correspondingly has three elements: a way to describe the state of the ledger, a language for transactions (changing ownership, in which cryptography will serve to verify identities of the previous and new owner, or sender and receiver), and finally a protocol for updating the ledger with validated transactions.

The third design element is the most difficult one, given the posited design problem. Updating the ledger means achieving a new consensus on who owns what. Each individual transaction can be easily evaluated by each actor to ascertain that it conforms to the rules and to the last known state of the ledger. But in a decentralized and asynchronous system, how do we reach the new consensus? The problem arises when actors disagree after the fact on the order in which transactions were made, because this allows me to send two mutually incompatible messages, each valid on its own ("I cede my coin to A", "I cede my coin to B") and each believed by a fraction of the network. This is the so-called "double-spending problem" and it is due to the combination of both features, decentralization and asynchronicity. There would be no conceptual difficulty if multiple actors could get together at fixed times to evaluate all new transactions jointly (synchronicity), or if a single actor evaluated all transactions as they came in (centralization).

There are two broad methods of establishing consensus, representing two conceptual extremes. The first is to appoint a dictator who chooses the block to be added to the blockchain. The first approach sounds like a terrible idea, but the trick is that the dictator is chosen at random for each block. Of course, in a decentralized environment where actors know little or nothing about each other, we have to be careful in how the selection takes place: effectively, there is no list of registered voters from which to select the dictator, and a malicious actor could create multiple fake identities to increase the chances of being selected. One method is to require the candidates to pay a cost, so that the one who adds the new block is the winner in a costly lottery: this is called "proof-of-work" and is the concept used by Bitcoin. Another method is to require candidates to post collateral: this is called "proof-of-stake" and may come into use for Ethereum.

The lottery requires the solution to a numerical problem that can only be found by random guessing, and guessing requires time and effort (a process called "mining"); but verifying that the solution is correct is easy.

In either case the choice of the dictator is embedded in the rule that the longest blockchain (more exactly, the one embodying the most proof of work) is the valid one. The lottery is running continuously, and every time a node wins it broadcasts the new block, with proof of work (the solution to the numerical problem) included in it. If two nodes find the solution nearly at the same time and the network does not agree (because part of the network received one new block and the rest the other new block) there is temporary disagreement and a fork in the chain. But the lottery keeps running, nodes will keep adding to both ends of the chain until one becomes longer than the other.

The other general method of establishing consensus is to hold a vote on the contents of the new block. This may sound better than random dictators, but, as with political elections, it raises the preliminary question of who is entitled to vote. In general, it is difficult

[6] *I will therefore not describe the specifics of Bitcoin any further, except by way of illustration.*

to maintain the free access and anonymity of the bitcoin model in this case. The method will therefore be mainly implemented in "permissioned" networks, where access is restricted. This of course raises the further question of who "permissions" the network, but I will set that aside. Assuming, then, that potential voters are identified and vetted to some degree, we still have to deal with the asynchronicity of the network (not everyone is voting on the same thing, or at the same time) and with the possibility of malicious users, either because vetted users are impersonated, or because users are not willing to trust all other users blindly.

The classic solution to this type of problem, called the "Byzantine generals problem", relies on multiple rounds of voting and an algorithm that is expected to reach consensus with very high probability. A high-profile example is the Ripple protocol, in which users do not trust all other users but a subset, specific to each user.

Whatever the method used, it is important to remember that trust is never eliminated, but rather displaced. Parties to a transaction need not trust each other, but, as in all monetary exchange, they have to trust the token exchanged: in the world of DLT, trust in the protocol is what dispenses with trust between parties.

With the three design elements in place, we see that ownership ultimately rests on a sequence of valid transfers, starting from some point and ending with the current owner. This sequence is what uniquely defines each bitcoin, and the blockchain is the collection of all such sequences. When the creation of these starting points is part of the protocol these chains constitute assets "native" to the technology, whose essential property is that they can be transferred with the technology. As the protocol regulates the creation of new starting points, the chains of ownership are in restricted supply. Thus, a bitcoin is useful because it can be transferred and is in limited supply.

Bitcoins are monetary tokens, but of a new kind: just as medieval bankers received coins in deposit and replaced the hand-to-hand exchange of the coins with written operations on ledgers, Bitcoin dispenses with the physical token, but contrary to medieval bankers there is no single ledger. Also, Bitcoin replaces the natural scarcity of precious metals with the artificial scarcity of controlled money creation.[7]

DLT: a solution in search of a problem?

About two years ago the notoriety of Bitcoin generated growing interest in its underlying technology. But Bitcoin was designed to solve a particular problem: a monetary token on the Internet that does not rely on any central authority. Much of the interest in DLT does not come from anyone genuinely interested in that problem. Instead, many have been taken with the attractive properties of DLT, which include resilience, speed, decentralization, immutability. But these properties are those of the solution to a particular problem. In many proposed applications, it is far from clear that decentralization and lack of trust are essential features of the relevant problem, leading to the nagging suspicion that DLT is a solution in search of a problem.

Furthermore, these desirable properties are not absolute or immutable: they arise from tradeoffs that might be resolved differently in other applications. I will cite a few. First, there is a tradeoff between scale and speed (or its inverse, latency). A truly decentralized or distributed system has to confront the problem of latency, because DLT requires full histories and multilateral transmissions of information, which implies amounts of data far beyond the needs of an equivalent centralized system (one hub talking bilaterally with many spokes). Second, DLT bases ownership on public information, that is, massive public disclosure. It is true that pseudonymity in Bitcoin mitigates the publicity,[8] but the tradeoff between information and privacy might be resolved differently in more restricted networks where participants are more readily identified. Third, Bitcoin's recent history illustrates the tension between openness and governance. Bitcoin is open-source protocol that Bitcoin users use, nothing more. Changes to the protocol are in principle endorsed by consensus, but difficult or strategic decisions are difficult to coordinate and can lead to fragmentation and incompatible splits.[9] Finally, several properties are more limited than is often recognized. For example, a decentralized network may be more resilient to shocks that might affect individual nodes, but the protocol itself becomes a single point of failure. Likewise the record of ownership is immutable only to the extent that the consensus does not change. The events surrounding Ethereum in August 2016 are a case in point.[10]

Finally, it always bears repeating that even if arrangement B is better than the status quo A, the costs of moving from A to B might be greater than the gains. Many of the touted advantages of DLT (e.g., transaction, settlement, and reconciliation will be faster and more reliable) come from multiple actors using a common language to store and update information, not from DLT per se. These advantages have been obvious to industry participants for decades, and one might wonder if high transition costs are not the reason why they have not been reaped earlier.[11]

3 Central Banks and DLT

Today central banks have a number of functions that could be affected by the development of DLT: among other things, they regulate financial actors, they often manage large payments systems, and they issue and manage outside money which is also the legal tender and unit of account.

As regulators, central banks are challenged in many ways by DLT. In the most extreme form of open consensus-based protocols, it is difficult to say who or what could or should be regulated. Here again Bitcoin's experience, whether or not one counts is as a form

[7] *Cleverly, Bitcoin's money creation mechanism is used to compensate the costs of proof-of-work.*

[8] *On Bitcoin's blockchain, parties to transactions are only identified by addresses (randomly generated strings of letters and numbers). The link to an individual exists only through the individual's knowledge of the password associated with the address.*

[9] *I refer here to recent debates on raising the block size limit in the Bitcoin protocol.*

[10] *To erase the consequences of a faulty smart contract, part of the Ethereum community agreed to go back in time to an earlier version of the blockchain. This was not accepted by all and two versions of Ethereum now coexist.*

[11] *The failure of the Global Straight Through Processing Association in the 1990s is instructive.*

of money, has blazed trails and brought DLT generally to the front of regulators' minds. It has also shown that a good part of the Bitcoin community, especially those who provide additional services around the use of Bitcoin, have themselves sought legal clarity. More broadly DLT can change the risk profile of existing regulated actors, allow the emergence of new actors who perhaps ought to be regulated, and create new systemic risks. On this last point word on smart contracts is in order. Once transfers of ownership can be effected by lines of code, it is conceptually a small steps to making transfers contingent on any variety of events either on-chain or (less easily) off-chain. The simplest example would be a transfer from A to B contingent on a prior transfer from C to D. Ethereum is to a large extent an attempt to extend Bitcoin's design to exploit the full potential of smart contracts. The recent history of algorithmic trading, high-speed trading, and occasional "flash crashes" make clear, however, the possibility for unintended consequences in a complex and partly automated system.

The other two functions of central banks, providing payments and legal tender, are possible avenues through which they might find themselves involved in DLT, either because of payments or because of legal tender. Even central banks that do not directly manage large value payments systems (LVP) do manage a "small value" payment system (SVP), namely physical currency. Payments systems, like any other way of doing things, should be improved when a more efficient technology comes along. The initial promise of Bitcoin was to offer a currency for the age of the Internet. Perhaps metallic coins and paper notes, relics or vestiges of monetary history, are destined to be replaced, just as paper as physical support to convey information.[12] Some aspects of a digital currency, like the vast amounts of information it could generate, would be attractive to some (regulators, law enforcement, and economists!) but not others (those who see value in the anonymity that cash offers). Similarly, the ability to pay interest on currency could be attractive to policy makers because it could make monetary policy more potent; this ability might be welcome or not, depending on whether the interest paid is positive or negative. Of course, improvements to SVPs could come from the private sector, and if physical cash is destined to be replaced then central banks may only need to manage its gradual disappearance. But is physical cash doomed? It has resisted well to the emergence over recent decades of electronic means of payment, in part because the use of cash protects the privacy of individuals' transactions in ways that even the pseudonymity of Bitcoin cannot fully duplicate.

But other considerations might lead some central banks to investigate their own version of digital currency, whether for small or large payments. First, there might be a reluctance on the part of the public to deal only with inside money (liabilities of the private sector). Second, suppose that DLT becomes widely adopted in a variety of contexts to record and transfer ownership of assets. In such transfers the other leg will often be a settlement in cash. This "cash" will often be inside money, but as legal tender, the currency provided by central banks is the ultimate final means of payment, and participants may well demand the option of settling in central bank money. But here again the key question remains: is a decentralized system needed to implement a digital version of a central bank's currency? The Bank of Canada recently concluded that the answer is not obviously positive.

4 Conclusion

In summary, Bitcoin is in my view a remarkable achievement, although unlikely by itself to replace monies in well-functioning monetary systems. It has nevertheless offered an interesting prototype and has generated interest in its underlying principles with potential applications to asset transfers. Central banks and regulators face a host of potential challenges, and the time may come soon when they will have to become involved in blockchains.

[12] *The scroll was replaced by the codex, or book, in the early centuries AD; parchment was replaced by rag-based paper in the 16th century and wood-based paper in the 19th century. No doubt the emergence of paper money in the 17th century was facilitated by the availability of a support that was both cheaper than parchment and still offered devices (such as watermark) to hinder counterfeiting.*