Innovations in payment technologies and the emergence of digital currencies

By Robleh Ali of the Bank's Financial Market Infrastructure Directorate, John Barrdear of the Bank's Monetary Assessment and Strategy Division, and Roger Clews and James Southgate of the Bank's Markets Directorate.[1]

- Modern electronic payment systems rely on trusted, central third parties to process payments securely. Recent developments have seen the creation of digital currencies like Bitcoin, which combine new currencies with decentralised payment systems.
- Although the monetary aspects of digital currencies have attracted considerable attention, the distributed ledger underlying their payment systems is a significant innovation.
- As with money held as bank deposits, most financial assets today exist as purely digital records. This opens up the possibility for distributed ledgers to transform the financial system more generally.

Overview

Money and payment systems are intrinsically linked. In order for an asset to function as a medium of exchange, there needs to be a secure way of transferring that asset — a payment system. And for any system other than the exchange of physical banknotes or coins, a means of recording the values stored is also needed — a ledger. Modern payment systems are computerised and most money exists only as digital records on commercial banks' accounts.

This article considers recent innovations in payments technology, focusing on the emergence of privately developed, internet-based digital currencies such as Bitcoin. Digital currency schemes combine both new payment systems and new currencies. Users can trade digital currencies with each other in exchange for traditional currency or goods and services without the need for any third party (like a bank). And their creation is not controlled by any central bank. Bitcoin — currently the largest digital currency — was set up in 2009 and several thousand businesses worldwide currently accept bitcoins in payment for anything from pizza to webhosting. Most digital currencies, including Bitcoin, incorporate predetermined supply paths leading to fixed eventual supplies. An overview of how digital currencies work, including the creation of new currency, is included in this article.

Much of the media focus to date has been on the new currencies themselves (such as 'bitcoins') and the large price swings that these have experienced. **This article argues, however, that the key innovation of digital currencies is the 'distributed ledger' which allows a payment system to operate in an entirely decentralised way, without intermediaries such as banks.** This innovation draws on advances from a range of disciplines including cryptography (secure communication), game theory (strategic decision-making) and peer-to-peer networking (networks of connections formed without central co-ordination).

When payment systems were first computerised, the underlying processes were not significantly changed. Distributed ledger technology represents a fundamental change in how payment systems could work. And in principle, this decentralised approach is not limited to payments. For instance, the majority of financial assets such as shares or bonds already exist only as digital records, stored on centralised databases.

A companion piece to this article focuses in more detail on the economics of digital currencies. It considers the extent to which they serve the roles of money, the incentives embedded in the design of the schemes and touches on some of the risks they may pose to the monetary and financial stability of the United Kingdom if they reached significant scale.

Click here for a short video that discusses some of the key topics from this article.

[1] The authors would like to thank Victoria Cleland, Danny Eckloff and Tom Ludlow for their help in producing this article.

Money and payment systems are intrinsically linked. They evolved together and this connection remains evident in the responsibilities of many central banks, including the Bank of England's role of ensuring both the stability of the currency and the payment systems which support the UK economy. Recent innovations in payment technologies have prompted great interest, particularly those that also incorporate 'digital currencies'.

This article provides a brief overview of how payment technologies, and the principles that underpin secure and reliable payments, have evolved from the 16th century up to the present day. It considers the key risks that arise and need to be mitigated in modern payment systems. It then considers the motivation behind some of the more recent developments in payment systems and currencies, and to what extent these truly represent a new technological or economic model. In particular, it focuses on the key technological development that underpins digital currencies: the creation of a distributed ledger. It considers the extent to which this new technology eliminates some of the risks traditionally found in payment systems, as well as some of the new risks it poses. Finally, it considers to what extent this distributed ledger model could have other applications beyond payments. A short video explains some of the key topics covered in this article.[1]

A companion piece to this article, 'The economics of digital currencies', considers the extent to which digital currencies may be considered money; some of the challenges the existing schemes could face over the longer term; and provides an initial assessment of the risks that digital currency schemes may, in time, pose to the Bank's mission through their potential impact on UK monetary and financial stability. Other issues such as those concerning consumer protection, taxation and money laundering are beyond the scope of this article, but some publications from other institutions regarding some of these issues are cited in the companion article.

The evolution of payment technology

The payment technology used in most economies today evolved from the early banking system and still retains structural characteristics from those roots. Early payments were made by exchanging intrinsically valuable items such as gold coins. When goldsmith banks emerged in the 16th century they kept ledgers of their customers' deposits which enabled payments to be made by making changes in the ledgers rather than physically exchanging the assets. This only worked for customers who shared the same bank. Over time, the need to make payments between banks led to the emergence of a central 'clearing' bank at which all the member banks could hold accounts, making interbank payments much simpler. The box on page 3 traces the evolution of payment systems in more detail.

Modern payment systems

In modern payment systems, payments are made by reducing the balance in a customer's account and increasing the balance in the recipient's account by an equivalent amount — a process that has not changed since the 16th century. The difference lies in the technology employed to record the balances and transfer them between different banks.

Technological developments over the past 50 years have affected payment systems in two key ways. First, the records and ledgers have been converted from paper to electronic form, which has increased the speed of completing transactions and reduced operational risks. Second, the emergence of low-cost technology has allowed new payment schemes to emerge, such as mobile money schemes. These are discussed in **Figure 2** on page 4.

Despite the application of new technology, the basic structure of centralised payment systems has remained unchanged. At the heart lies a central ledger, with settlement taking place across the books of a central authority, acting as a clearing bank (a service usually undertaken by the central bank of a given economy).[2] Each participant, typically a commercial financial institution, holds a balance at the central bank,[3] recorded in the ledger, but also reflected in the participant bank's own (internal) ledger. Individual customers, branches, or even other (typically smaller) banks would then hold balances at the participant bank, which would again be reflected in their own ledger. This 'tiered' arrangement is illustrated in **Figure 1**. The example traces a payment being made from one person to another via their commercial banks and the central bank.

Figure 1 A tiered payment system

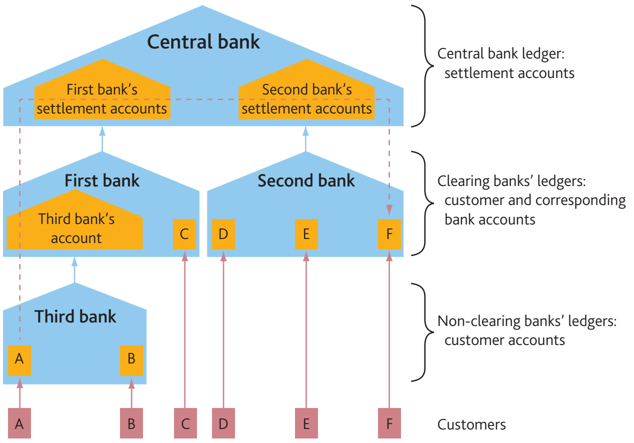
Central bank
First bank's settlement accounts
Second bank's settlement accounts
Central bank ledger: settlement accounts
First bank
Second bank
Third bank's account
C
D
E
F
Clearing banks' ledgers: customer and corresponding bank accounts
Third bank
A
B
Non-clearing banks' ledgers: customer accounts
A
B
C
D
E
F
Customers

Note: A payment from **A's** account to **F's** account passes through a number of intermediaries, which verify each step of the process. Participants only have sight of their own assets and liabilities. The solid lines indicate deposits and the dashed line payments.

[1] http://youtu.be/CxDKE_gQX_M.
[2] For more discussion of the role of the central bank in payment systems see Manning, Nier and Schanz (2009) and Norman, Shaw and Speight (2011).
[3] Settlement accounts may also serve as reserves accounts. And banks without settlement accounts may hold reserves accounts for other purposes. For the role of reserves see McLeay, Radia and Thomas (2014).

A brief history of money and payment systems

Throughout history there have been many different manifestations of money, both physical and electronic. Economists identify money through the roles that it serves in society. In particular, something may be considered money from the perspective of economic theory to the extent that it serves as a **medium of exchange** with which to make payments; a **store of value** with which to transfer 'purchasing power' (the ability to buy goods and services) from today to some future date; and a **unit of account** with which to measure the value of any particular item for sale.

In order for money to function as a medium of exchange, there needs to be a system to enable transfers of value — that is, a payment system — and for any system other than the exchange of physical banknotes or coins, a means of recording the values stored — a ledger.

Coins made of precious metals were one of the earliest methods of making payments in a number of regions of the world. Physical possession of the instrument denoted ownership, and the act of physical transfer acted as the payment system.

When goldsmith banks emerged in the 16th century, they issued notes — essentially IOUs — as receipts for gold deposits made with them. These IOUs could be transferred from one individual to another. Each bank kept its own ledger and in the earliest days there was no interbank settlement — that is, no way in which the ledgers of individual banks and branches could be 'connected' — so the notes could only be redeemed at the bank and the branch where they were issued. This meant that any payment requiring the transfer of money to a different bank would require the bearer of a note to first convert it into gold and then to physically transport it to the new bank, a cumbersome process.

The pressure to reduce these transaction costs led to banks starting to accept claims on each other. This innovation made trading more convenient as merchants could now deposit notes from other banks directly into their own bank, eliminating the burden of converting paper money into gold in order to transfer the funds. In accepting the note from a different bank, though, the payee's bank faced a new problem in that it was now exposed to the payer's bank until settlement in gold could be arranged. Where note acceptance was limited to a small number of banks this could be handled bilaterally. But as the number of banks in the system increased, interbank payments became more cumbersome and the incentive for banks to create a more efficient system increased.

The solution that eventually emerged was for one 'clearing' bank to sit at the centre of the system, with all member banks holding accounts with the clearing bank. The system worked by requiring all the member banks to hold balances against the risks they brought to the system. The bank operating the clearing system was, in effect, taking on some of the functions of a central bank (Goodhart (1988)).

New developments in payment systems and alternative currencies

A variety of developments in payment technologies and alternative currencies have emerged in recent years. Some of these innovations focus on making payments more accessible to a wider range of users — such as mobile phone payments — while still relying on a trusted central entity. More recent innovations have introduced a fundamentally different, decentralised structure to payment systems by relying on cryptography rather than a central authority. **Figure 2** describes four categories of recent innovations and some of their characteristics. They are split according to whether they establish a new payment system, a new currency, neither, or both (summarised in **Table A**).

Table A Types of innovation

Category	New payment system	New currency
I: Wrappers	No	No
II: Mobile money	Yes	No
III: Credits and local currencies	No	Yes
IV: Digital currencies	Yes	Yes

There are some caveats to this simple categorisation. For example, while local currencies technically represent new currencies, any such scheme operating at a one-to-one fixed exchange rate and backed by national currencies bears a close relationship with an existing currency. It is also important, in the final category, to distinguish between digital currencies as candidate payment systems and digital currencies as potential forms of money. Although Bitcoin introduced a new currency and a new payment technology together, the distributed ledger technology could, in theory, be used without the creation of a new currency. As emphasised by Haldane and Qvigstad (2014), it would technically be possible for an existing central bank to issue digital-only liabilities in a distributed-ledger payment system equivalent to those deployed by recent digital currencies.

Figure 2 Recent innovations in payment technologies

Category I: Wrappers

The first category of innovation focuses on providing 'wrapper' services to improve the user interface and accessibility of existing payment systems architecture. These innovations therefore represent neither a new currency nor a new core payments system.

The core motivation can be either new entrants seeking to capture a segment of the market, or incumbents seeking to improve market share and reduce consumer use of other, more expensive payment systems. Examples include Google Wallet, Apple Pay and Paym, which builds on the existing infrastructure to make payments by linking users' mobile phone numbers to their bank accounts.

Category II: Mobile money

These schemes represent new payment systems, with money stored as credits on a smart card or a system-provider's books, but continue to use national currencies. One example is M-Pesa, a popular service in Kenya that grants access to financial services, including payments, to anybody with a mobile phone.

In areas where access to traditional banking infrastructure is limited, development and adoption of new payment systems serves to fulfil otherwise unmet demand. In more developed economies, new payment systems are probably developed in response to the high margins associated with incumbent systems and adopted on the basis of their ease of use.

Category III: Credits and local currencies

This category relies on users trusting a new currency as a unit of account and medium of exchange. Credits are schemes in which private companies accept money in exchange for an alternative unit of account which can be spent on a particular platform (such as within an online game). Nevertheless, they generally make use of existing payment systems, including use of 'wrapper' services, to make transfers. Local currencies are similar in that people exchange national currencies for a local equivalent which can be spent in a specific geographical area. UK local currencies such as the Bristol Pound are often backed by and remain on a fixed exchange rate with sterling. Naqvi and Southgate (2013) consider local currencies in more detail.

A key motivation for both the development and the adoption of local currencies surrounds a desire to promote spending at, and between, participants of the scheme in order to boost economic activity in a specific region, support local sustainability and shorten supply chains.

Category IV: Digital currencies

A digital currency scheme incorporates both a new decentralised payment system and a new currency. All the schemes exhibit a publicly visible ledger which is shared across a computing network. A key defining feature of each digital currency scheme is the process by which its users come to agree on changes to its ledger (that is, on which transactions to accept as valid).

Most digital currencies are 'cryptocurrencies', in that they seek consensus through means of techniques from the field of cryptography. There are also a small number of digital currencies, the most prominent of which is Ripple, that seek consensus through non-cryptographic means.[1]

[1] It is possible to have a digital currency with a centralised ledger. This is not discussed in this article because there is no recent example of a digital currency operating in this way.

The rest of this section provides an introduction to Bitcoin — currently the most prominent example of a digital currency — including a brief discussion of the motivation for setting up and using a digital currency.

What is Bitcoin?

Bitcoin was the first, and remains the largest, functioning digital currency. It was launched in January 2009 and is a privately developed, internet-based currency and payment system that requires no intermediaries (like banks) for the processing of payments. Furthermore, the supply of bitcoins is not controlled by a central bank.[1] It is commonly referred to as a 'cryptocurrency' as it relies on techniques from the field of cryptography to ensure the secure validation of transactions. There are currently several hundred cryptocurrencies in existence, such as Litecoin and Peercoin. Most of these were inspired by, or explicitly based on, Bitcoin.

Bitcoin users do not have to disclose who they are. They maintain a digital 'wallet' on their computers and, by use of special software, trade the currency among each other in exchange for traditional currency or goods and services. Several thousand businesses worldwide currently accept bitcoins in payment for anything from pizza to webhosting. Payments can be made at any time and between any two users worldwide. Users may acquire bitcoins as a reward for verifying earlier transactions (explained more below), by purchasing them from other users (in exchange for traditional currencies) or in exchange for goods and services.

A key innovation of digital currency systems is the use of a 'distributed ledger' that allows payments to be made in a decentralised way. How this works — and how it marks a key innovation in payment technology — is explained in the subsequent section of this article, but the basic process is as follows. A user, wishing to make a payment, issues payment instructions that are disseminated across the network of other users. Standard cryptographic techniques make it possible for users to verify that the transaction is valid — that the would-be payer owns the currency in question. Special users in the network, known as 'miners', gather together blocks of transactions and compete to verify them. In return for this service, miners that successfully verify a block of transactions receive both an allocation of newly created currency and any transaction fees offered by parties to the transactions under question. The box on pages 7–8 provides a step-by-step overview of how a transaction works using this payment system.

The candidate blocks were 'empty' in the sense that they had no transactions in them other than the allocation of new bitcoins as a reward for solving the puzzle. This effectively served to create the initial endowment of bitcoins. The first blocks created 50 new bitcoins per block and the Bitcoin protocol calls for this reward to be halved every 210,000 blocks (roughly every four years).[2] The current

Chart 1 The projected supply of bitcoins in circulation

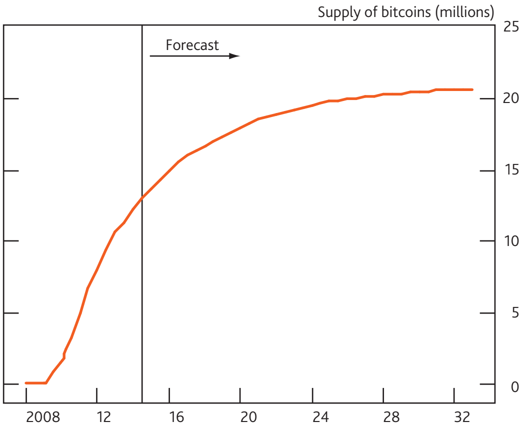

Source: http://blockchain.info.

Chart 2 The price of bitcoins (linear scale)

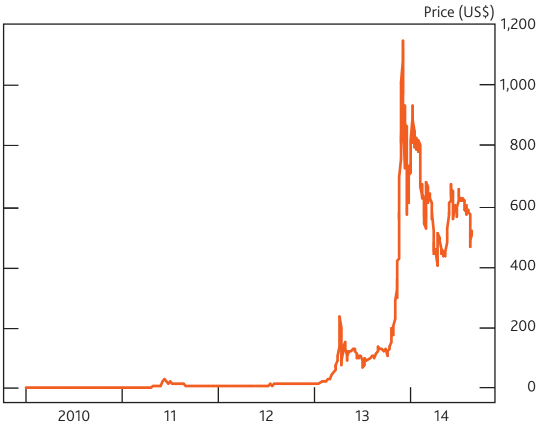

Source: http://blockchain.info.

Chart 3 The price of bitcoins (log scale)

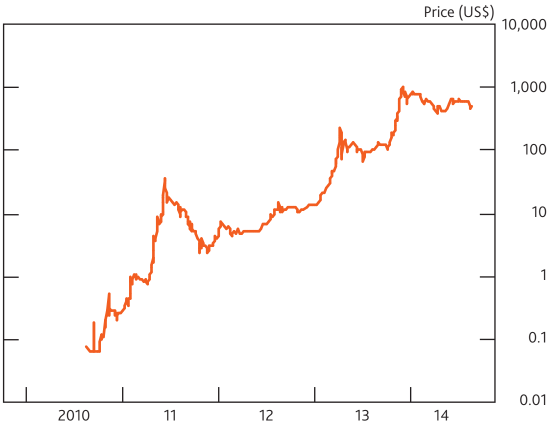

Source: http://blockchain.info.

[1] Note that throughout this article, 'Bitcoin' is used to refer to the system as a whole and 'bitcoin' to refer to individual units of the currency.
[2] The Bitcoin protocol seeks to maintain a roughly constant time of ten minutes between each successfully verified block. See the box on pages 7–8 and the annex for more detail.

reward is 25 bitcoins per block, and this is likely to be reduced to 12.5 bitcoins per block in 2017. The planned eventual total number of bitcoins is therefore 21 million, which will be mostly reached by 2040. There are currently a little over thirteen million bitcoins in circulation (**Chart 1**), distributed over perhaps one or two million users worldwide.

The price of bitcoins has increased markedly since the scheme was launched, rising roughly 5,000% over the past two years (**Charts 2** and **3**). It has also exhibited significant volatility, which has led to considerable debate and media attention.

Motivation for the development and adoption of digital currencies

Beyond a general increase in public willingness to use and trust computing technology, interest in and the adoption of digital currencies appears to be driven by three key factors: ideology, financial return and the pursuit of lower transaction fees.

The foundational motivations for Bitcoin appear to have been largely **ideological**. The digital currency was expressly designed to avoid any centralised control (of either the money supply or the payment system) and to minimise the degree of trust that participants need to place in any third party. The first block in Bitcoin's block chain (the 'genesis block') includes the text:

```
The Times 03/Jan/2009 Chancellor on brink
of second bailout for banks
```

in reference to a newspaper article from that day (Duncan and Elliott (2009)), presumably in order both to demonstrate that it could not have been created before that date and to highlight the conceptual distinction between Bitcoin and the structure of modern monetary economies. Complete adoption of Bitcoin by its users would allow them to exist economically almost entirely outside the prevailing monetary system, although this is not straightforward due to the relatively small number of businesses which accept it. In addition, some participants may be drawn to the near anonymity offered by such systems.

Second, digital currencies have come to be viewed by some as an **asset class** for financial investment, driven by an interaction between the schemes' planned fixed supplies and their increasing publicity. Since the future path of each such scheme's supply is predetermined and known with near certainty, movements in their price will essentially reflect only changes in demand.[1] Since digital currencies have no intrinsic demand (they are not used as a factor of production and are not sought out as a consumer good), expectations about medium and long-run future price growth will be predominantly driven by expectations relating to the future growth in the transactional use they support.

Advocates of digital currencies argue that they offer **lower transaction fees** on payments than existing electronic retail payment systems or international transfers. Based on this premise, a number of start-up businesses are seeking to offer payment facilities that use digital currencies as a bridge mechanism for settlement.[2] The sustainability of low transaction fees from digital currencies is discussed in more detail in the companion piece to this article.

The distributed ledger as a key technological innovation

This section examines the concept of a distributed ledger — a key technical innovation of digital currencies — and how it is a feature that solves the problem of 'double spend' in a decentralised payment system. The distributed ledger (the 'block chain' in cryptocurrencies) was made possible by the emergence of several earlier innovations, including the internet. It rests on concepts from cryptography, game theory and peer-to-peer networking. Finally, this section also considers the risks in both centralised and decentralised payment systems.

The double-spend problem

A key problem for any electronic payment system is how to ensure that money cannot be 'double spent'. If Anne has a single £1 coin, it is not possible for her to give £1 to Bill and also £1 to Clare. The physical act of exchange prevents the payer from spending the same money twice. A payment system that relies on digital records must have a way of preventing double spending because it is simple to copy and edit digital records.

The approach used in the modern banking system, which emerged as a computerised replication of earlier paper-based records, is for specialised entities (usually banks) to maintain master ledgers that act as the definitive record of each individual's money holdings. In turn, they hold accounts recorded in the ledger of one central body (typically the central bank). Those holding the ledgers have the ability to prevent any transaction they deem to be invalid. In order to use the system, people must trust that these centralised ledgers will be maintained in a reliable, timely and honest manner.

An alternative approach is to implement a fully decentralised payment system, in which copies of the ledger are shared between all participants, and a process is established by which

[1] Markets that allow trading of digital currencies are also relatively illiquid, which may affect short-term price movements.

[2] For example, a payment provider might allow retailers to set their prices (and receive payment) in sterling, but allow consumers to pay with a digital currency. If consumers wished to pay with a different currency, such as the US dollar, then the payment provider might first convert the dollars to the digital currency before processing the payment.

Making payments securely with a distributed ledger

Any electronic payment system must have a reliable method of recording transactions that all participants can agree is accurate. For a decentralised system like Bitcoin this creates two challenges. The first is devising a secure and reliable method for updating a public ledger of which there are myriad copies distributed throughout the world. The second is, in the absence of a central authority to provide or co-ordinate resources, creating the necessary incentives for users to contribute resources to verifying transactions. This box describes how Bitcoin overcomes these challenges by explaining the mains steps in a transaction. The key concepts were first outlined by Nakamoto (2008).

Step 1 — Agreeing the transaction

__Anne__ is a Bitcoin miner who has previously verified a block of transactions successfully and received 25 new bitcoins as a reward. __Bill__ is a carpenter who sells furniture online and accepts bitcoin. __Anne__ decides to pay 1 bitcoin to __Bill__ for a chest of drawers and is prepared to pay 0.01 bitcoins as a transaction fee.

Bitcoin users are under no formal requirement to pay transaction fees and if they offer one, the size of that fee is at their discretion. However, Bitcoin miners are able to choose which transactions they process, so a higher fee offered gives them a greater incentive to validate **Anne's** transaction.

Step 2 — Creating the transaction message

__Anne__ creates a message with three basic elements: a reference to the previous transaction through which she acquired the bitcoins, the addresses to pay (including __Bill's__) and the amount to pay each one. The message also has other elements such as digital signatures and any conditions that __Anne__ may place on the payment.

The number of bitcoins at any address is derived from the output of earlier transactions that are all publicly available on the block chain for inspection. In this example there is a previous output of 25 bitcoins from **Anne's** mining activity which forms the input to the new transaction. Bitcoin transactions may have any number of inputs or outputs. The 'change' due to Anne is paid as an output of the transaction and any credit included in the input which is not accounted for in the output is accepted as a transaction fee.

Inputs:

- 25 bitcoins from **Anne** (the output from her previous transaction).

Outputs:

- 1 bitcoin to **Bill**.
- 23.99 bitcoins to **Anne** (her 'change' from the transaction).
- 0.01 bitcoins as a transaction fee to whichever miner successfully verifies the transaction.[1]

It is also possible for **Anne** to place some conditions on the payment, so that **Bill** cannot spend his proceeds unless they are met. Most payments do not impose any conditions, but more complex transactions may require multiple conditions to be met before any funds are released. This capability allows the technology to be expanded to support more complex transactions.

Step 3 — Signing the transaction message

Once the message has been created, __Anne__ digitally signs it to prove that she controls the payer address.

Similar to real signatures, digital signatures provide proof that the transaction message was created by the person who wants to make the payment.

Digital signatures are a form of public-key cryptography. They work by creating 'public' keys which can be used to decrypt messages encoded by a corresponding 'private' key. To create a digital signature, **Anne** encrypts the message she wishes to sign with her private key. This message can then only be decoded with the corresponding public key, which she also broadcasts in order that her transactions can be verified. Further information on public-key cryptography is contained in the technical annex.

Step 4 — Broadcasting the transaction message

__Anne__ broadcasts the signed message to the network for verification.

Bitcoin miners are arranged in a 'peer-to-peer network' — a network of connections that are formed informally with no central co-ordination. Although miners are under no obligation to do so, the Bitcoin protocol calls for all messages to be transmitted across the network on a 'best-efforts' basis, sharing the message with one's immediate peers. This means that **Anne's** transaction is not broadcast to the entire network at once, but instead goes to a random subset of her peers first, then to their peers and so on.

Peer-to-peer networks are commonly used to quickly and effectively share data between users in a number of other settings. Some video-streaming services, for example, make use of the technology.

Step 5 — Transaction verification ('mining')

Miners gather __Anne's__ new transaction and combine it with others into new candidate 'blocks'. They then compete to verify them in a way that other miners will accept.

Verification of a transaction block has two elements: validation and achieving consensus. Validating a block of transactions — which includes checking that the digital signatures are correct — takes a very short amount of time. Establishing consensus is purposefully more difficult and requires each miner to demonstrate the investment of computing resources known as a 'proof of work'. The proof of work scheme used by Bitcoin is explained in detail in the annex.

Proof of work schemes need to be difficult to achieve but simple to check. This allows the incentives of the system to be balanced in favour of transaction verification by making it very easy to spot a fraudulent transaction. The only way the system can be attacked is by assembling sufficient computing power on the network to 'verify' fraudulent transactions. This would undermine trust in the system as a whole and the value of any bitcoins the attacker could steal. It therefore makes more sense for anyone capable of assembling the necessary computing power to contribute to the continuation of the system, rather than attacking it.

The proof of work scheme used by Bitcoin means that the time taken for a miner to successfully verify a block of transactions is random. But as new miners join the network, or existing miners invest in faster computers, the time taken for a successful verification can fall. In order to allow time for news of each success to pass across the entire network, the difficulty of the proof of work problem is periodically adjusted so that the average time between blocks remains broadly constant at ten minutes for Bitcoin, meaning that payments are not instantaneous.

Step 6 — Success

__Clare__ is a miner and successful at verifying a block with __Anne's__ transaction in it, so she will receive both a reward of new bitcoins, as well as the transaction fee from __Anne's__ transaction. __Clare__ broadcasts this result and other miners add the block to the end of their copies of the block chain and return to step 5. __Bill__ receives the 1 bitcoin sent to him and delivers the chest of drawers to __Anne__.

Coinbase transactions

The first transaction in each block is a special 'coinbase' transaction which (i) grants the miner new bitcoins as a reward and (ii) pays the miner any transaction fees offered by transactions within the block.[2] The allocation of new bitcoins to each coinbase transaction is halved every 210,000 blocks (which, at ten minutes per block on average, equates to roughly once every four years). The current allocation is 25 bitcoins per block, which should halve to 12.5 bitcoins per block in 2017. The motivation behind such a money supply rule — and some issues associated with it — are discussed further in the companion article.

Orphaned blocks

The nature of a distributed system means that it is possible — albeit fairly infrequently — for two miners to successfully verify two different candidates for the next block at essentially the same time. When this happens, both copies are initially retained by the network as branches of the main chain, but miners will proceed to work on candidate blocks that follow on from whichever one they first receive.

The chain of blocks representing the greatest sum of work done is the accepted truth within the Bitcoin network (sometimes referred to as the 'longest chain'). Whichever branch is received by the majority of the network will initially be selected. However the branch with the most computation resources should ultimately take the lead. This branch will be most likely to have a subsequent block built on top of it and is therefore more likely to eventually 'win' the race. Miners that were working off blocks in the 'shorter' branch (that is, the branch with less demonstrated work done) then have a significant incentive to switch to the longer branch, as any work they contribute to the shorter branch will never be accepted by the majority of the network.

In this scenario, blocks within the abandoned shorter branch are referred to as 'orphans', such as the blocks in red shown in **Figure A**. Any transactions listed in an orphan block will need to be verified again. No reward a miner claims from an orphan block is recognised, as it is not part of the longest block chain.

Figure A Orphaned blocks

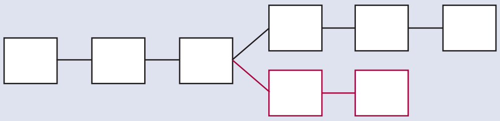

The rule that the chain with the greatest sum of work done wins is an important element in combating fraud in the Bitcoin network. Any attacker attempting to modify earlier blocks (so that bitcoins could be spent twice) would have to control enough computing power for them to both catch up with and then overtake the genuine block chain as the 'longest'. To be assured of success, the would-be attacker would need to obtain, and retain, a majority of all computing resources on the network. For this reason, the attack is known as a '50%+1' attack.

[1] Strictly, transaction fees are defined implicitly as the difference between the inputs and the explicitly listed outputs for each transaction. They are paid to miners as part of the 'coinbase' transaction in each block — see below for more detail.

[2] For example, in block number 310,000, the coinbase transaction was for a total of 25.15638661 bitcoins, comprising 25 new bitcoins and 0.15638661 bitcoins that were offered as transaction fees from the other 711 transactions in that block.

users agree on changes to the ledger (that is, on which transactions are valid). **Since anybody can check any proposed transaction against the ledger, this approach removes the need for a central authority and thus for participants to have confidence in the integrity of any single entity**.

Achieving consensus

The defining feature of a distributed payment system is the manner by which consensus is reached about any proposed changes to the ledger. How to achieve consensus between people in a network when nobody can be completely sure who can be trusted has long been recognised as a problem in the field of computer science.[1] It is not sufficient to offer blanket acceptance to all statements, for example, because this creates an incentive to lie in order to gain an advantage.

It is also not sufficient to have users vote on whether to accept a proposed change. This is because it is generally very easy for a single person to create many nodes on a computer network in order to distort the vote. Instead, digital currencies make use of game theory and recognise that, on its own, any proposed change to the ledger is 'cheap talk' — a statement that, since it was effectively free to issue, should receive very little weight. In order for a proposed change to the ledger to be accepted by others as true, those proposing the change — the 'miners' that serve as transaction verifiers — must demonstrate that it was costly for them to issue the proposal.

Cryptocurrencies require that users contributing to the verification process must demonstrate a cryptographic 'proof of work' to show that they have paid a cost in computation time before their proposals are accepted. The box on pages 7–8 and the technical annex describe a proof of work scheme in more detail. Some other digital currencies impose a cost in the form of a small amount of currency that is destroyed as part of the transaction.[2] **Figure 3** shows an example of a distributed payment system.

Figure 3 A distributed payment system

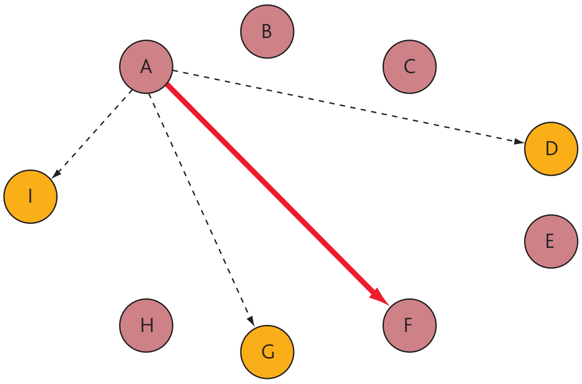

Note: All participants have sight of all accounts (and their entire history). Payments pass directly between users — shown here by the red arrow from **A** to **F** — but are verified by other users: in particular, new transactions are broadcast to 'miners' (shown here as participants **D**, **G** and **I**). When verified, the transactions are added to the history of the ledger.

Risks in payment systems

Centralised systems

There are certain risks that are common to all existing tiered payment systems. Finan, Lasaosa and Sunderland (2013) identify the three greatest risks as:

- **Credit risk**, in that a paying bank may become insolvent with a large amount of money owed to other members of the system.
- **Liquidity risk**, in that a member bank that is fundamentally solvent may not have the funds to settle a required payment at a particular moment in time.
- **Operational risk**, in that one of the banks involved in a payment transaction may cease to function (either temporarily or permanently) because of some event, such as an IT failure.

These risks are inherent to any intermediated banking system. As discussed in the box on page 3, this structure evolved in response to the need to make payments more efficiently and when payment systems were computerised, this intermediated structure remained — along with the main credit and liquidity risks present in those original systems. Prudential regulation of systemically important payment systems has led to the introduction of several measures which significantly reduce or remove these systemic risks.[3] When making these decisions regulators face a trade-off: on the one hand, prudential regulation of systemically important payment systems contributes to stable and efficient payments which promotes economic activity by reducing risk and uncertainty in the economy; but, on the other hand, some of the measures needed to reduce systemic risks in payment systems require participants to contribute money up front to cover these risks. Economic theory would suggest that such 'barriers to entry' may serve to weaken competition between existing members which may, in turn, lead to increased transaction fees and reduced levels of economic activity. But when constrained to the existing payment system architecture, these requirements are necessary in order to protect the broader financial stability of the United Kingdom.

[1] This is known as the 'Byzantine Generals Problem'. See Lamport, Shostak and Pease (1982).

[2] For example, Peercoin imposes that transaction fees paid by parties to the transaction (which are mandatory and set by the Peercoin protocol) are destroyed, rather than paid to transaction verifiers (miners). To ensure that this does not lead to an overall reduction in the supply of the currency, Peercoin also implements a 1% per annum increase in the supply that is paid to miners in a 'proof of stake' system broadly analogous to the payment of interest.

[3] For example CHAPS de-tiering referred to in Finan, Lasaosa and Sunderland (2013) and Bacs (the United Kingdom's automated clearing house, through which Direct Debits are processed) Liquidity Funding and Collateralisation Agreement in order to reduce credit risk in that system. Other examples can be found in *The Bank of England's supervision of financial market infrastructures — Annual Report*, formerly the *Payment Systems Oversight Report*.

Another important but generally non-systemic risk is **fraud**. For example a credit card user wishing to make a purchase over the internet must disclose their card details to the retailer. If these card details are stolen, the thief is then able to fraudulently make payments from the account of the card holder.

Decentralised systems

Existing distributed payment systems remove the credit and liquidity risks discussed above by eliminating intermediaries: payments are made directly between payer and payee. To be sure of this, users need to have confidence that for any distributed system they use, the cryptography employed has been implemented correctly.

In general, distributed systems designed in this way should also be more resilient to systemic operational risk because the whole system is not dependent on a centralised third party. A distributed system effectively has as many redundant backups as there are contributors to the network (which can easily number in the thousands, many more than centralised payment systems typically operate).

The nature of fraud risk — and other ways that customers may be susceptible to lose money — changes significantly between centralised and decentralised payment systems. In a decentralised system there is no need for users to disclose their complete payment details when making a payment, thus removing the risk of payment details being stolen from a retailer. However, the risk of direct loss of digital currencies is higher than that for deposits held (electronically) with commercial banks: if a user's private key is lost — because of a corrupted hard drive, say — then their digital currency will not be recoverable. This contrasts to a lost password used for internet banking with a commercial bank, say, which could be recovered or reset by contacting the bank in question. In this sense, a digital wallet is more analogous to a physical wallet containing physical currency than a bank account accessed online.

More substantially, distributed systems are subject to a risk of **system-wide fraud** if the process of achieving consensus is compromised. Cryptocurrency schemes, for example, are currently designed such that a would-be attacker would require sustained control of a majority of the total computer power across the entire network of miners. Some loosely co-ordinated pools of miners have, on occasion, represented a majority of computing power in the Bitcoin network.[1] Some researchers have also suggested that the necessary threshold for a successful attack may be less than 50%. This issue is examined in more depth in the annex.

Applications of the distributed ledger beyond payment systems

The introduction of any new technology enables the rethinking of business processes associated with the former technology. In the case of payments, when paper ledgers were first computerised, the underlying processes were not significantly changed.

It is often the case that the bulk of the gains from the introduction of a new technology do not arise immediately because processes that make use of the technology also need to be rethought. For example, Brynjolfsson and McAfee (2014) observe that when the electric motor was first introduced to factories, the productivity improvements it enabled only emerged after a lag of 30 years. This was approximately the time it took for a new cohort of factory managers to emerge who realised that instead of merely electrifying the single steam engine powering all the machinery in a factory, small electric motors could be fitted to each machine. While the initial installation did reduce costs, the authors argue that the greatest gains came from factories being rearranged according to the most efficient flow of materials, rather than the limitations of the machinery. It was not the electrification itself which produced the gains but the changes in processes which it made possible.

In a similar way, the potential impact of the distributed ledger may be much broader than on payment systems alone. The majority of financial assets — such as loans, bonds, stocks and derivatives — now exist only in electronic form, meaning that the financial system itself is already simply a set of digital records. These records are currently held in a tiered structure (that is, with records of individuals' accounts stored centrally at their bank, and banks' reserves accounts held centrally at the central bank), but it may be possible in the future — in theory, at least — for the existing infrastructure of the financial system to be gradually replaced by a variety of distributed systems (although this article makes no prediction in this regard). Some developers have already implemented so-called 'coloured coins' which means using digital currencies as tokens for other assets by attaching additional information. This development could allow any type of financial asset, for example shares in a company, to be recorded on a distributed ledger. Distributed ledger technology could also be applied to physical assets where no centralised register exists, such as gold or silver.[2]

[1] The proof of work scheme used by Bitcoin means that the time taken for any given miner to successfully verify a block of transactions is random. In order to smooth out the consequent volatility of earnings, miners often pool their resources and agree to share their earnings in proportion to the computing resources contributed.

[2] For an asset such as gold there is a necessary link to physical custody which is different for most financial assets which are already purely digital.

Some commentators (Wenger (2013)) have suggested that the key to understanding Bitcoin is to think of it as a protocol, akin to those that underpin the internet. Others have extended this analogy further, suggesting that digital currencies may be thought of as an 'internet of money'. But since the potential applications are, in principle, broader than just payments, the distributed ledger technology may perhaps be better described as a first attempt at an 'internet of finance'.

Conclusion

Digital currencies, as presently designed, carry both risks and benefits. As explained in the companion piece to this article, digital currencies do not currently pose a material risk to monetary or financial stability in the United Kingdom, but it is conceivable that potential risks could develop over time. The distributed ledger is a genuine technological innovation which demonstrates that digital records can be held securely without any central authority.

The total stock of digital currencies is at present too small to pose a threat to financial stability, but further increases cannot be ruled out and it is conceivable in time that there could be an asset price crash among free-floating digital currencies that had the potential to affect financial stability. Potential risks to monetary stability would only be likely to emerge once digital currencies had achieved substantial usage across the economy. If a subset of people transacted exclusively in a digital currency, then the Bank's ability to influence demand for this group may potentially be impaired. The incentives of existing digital currency schemes pose considerable obstacles to their widespread adoption, however. This is discussed in more detail in the companion article.

Ultimately every transaction involving a financial asset must be recorded and most of these records are digital. The structure of the broader financial system is similar to payments in that these records are held by centralised third parties. The application of decentralised technology to this platform of digital information could have far-reaching implications, other industries whose products were digitised have been reshaped by new technology. The impact of the distributed ledger on the financial industry could be much wider than payments.

Annex
Technical issues

This technical annex provides further details on digital signatures and cryptographic hash functions. It also discusses whether digital currencies are fraud-proof.

Digital signatures and public-key cryptography

Digital signatures provide a mathematical proof that a particular message was approved by a particular person. They are an application of public-key cryptography, which relies on two separate, but mathematically interrelated keys: one private and one public.

Bitcoin addresses are a version of the public key, which can be made widely available and published. Addresses and their private keys are random strings of alphanumeric characters. An address is typically 34 characters long (for example 1FfmbHfnpaZjKFvyi1okTjJJusN455paPH), while a private key is typically 51 characters long.

Each Bitcoin address is paired with a corresponding private key, which is kept secret by the owner of the address, and needed to sign transactions from — and, hence, prove ownership of — the address. It is also possible to create addresses that are linked to multiple private keys. These may be set up such that any of the private keys may be used to sign a transaction, or all of them must be used together.

Figure A depicts the process of signing a transaction in Bitcoin. **Anne** encrypts a copy of the transaction with her private key and then broadcasts both the plain and the encrypted versions of the transaction details. Anybody can combine the encrypted version with **Anne's** public key to obtain another plain version. If it is the same as the plain version that **Anne** broadcast, then it proves that **Anne's** private key must have been used.

Figure A Digital signatures

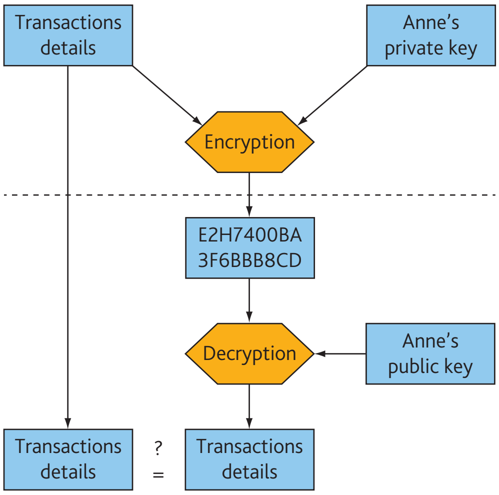

Cryptographic hash functions and Bitcoin's proof of work scheme

As discussed in the main text, Bitcoin miners must demonstrate a proof of work before their proposed block of transactions is accepted by the network. Given that typically, all users need to know all previous transactions to figure out account balances, it becomes important that all users agree on which transactions have actually happened and in which order. If two users observe different transaction histories, they will be unable to come to the same conclusion regarding balances and double spends. The block chain serves as a way for all users to come to a consensus regarding which transactions have already happened and in which order. In Bitcoin, the way in which users agree on a set history of transactions is to pick the history which users have put the most work into creating. **The 'work' must be a task that is hard for a computer to complete, but easy for other computers to verify**.

A simple example would be a requirement that people repeatedly roll three six-sided dice until they roll three ones. When somebody does this, everybody accepts their message as true and moves on to the next message. This is a time-consuming exercise in trial and error, but one where success is immediately visible to everyone. The time taken for somebody to successfully roll three ones is random, but the expected number of attempts is known. The more people that take part, or the faster that each person makes each attempt, the shorter the time until somebody succeeds. To offset this, each person might be required to roll four dice and to get four ones. With careful calibration, by making the problem harder as more people join, the average time taken for somebody to succeed can be made to stay roughly constant.

The proof of work scheme used by Bitcoin makes use of a special algorithm called a 'cryptographic hash function', which takes any amount of information as an input and creates an output of a standard length (the 'hash value'). The function is cryptographic because the hash value produced is different for any change in the input (even of a single character), and it is almost impossible to know in advance what hash value will be produced for a given input. For example, the hash function used by Bitcoin (called 'SHA-256') generates the following:

Input (case sensitive)	Output (the 'hash value')
Bank of England	6b31489400146361800f1f67cfb003f6ba5734b645c30a68d1888b4a19c9d64c
Bank of England1	38f0f960648853c9675951b10cf55acb3f5696bfb183d398782d4f32e99905fe
Bank of England2	ba9745451de288a04fcbf08fdbd43fa429e9c5f7d6ce436e5adf97f10dd22836
Bank of England3	028aa80090f374aaed153ac5b3ab199f3cd63b1e409f55be777c1189dd4b23f1

The Bitcoin protocol requires that miners combine three inputs and feed them into a SHA-256 hash function:

- A reference to the previous block.
- Details of their candidate block of transactions.
- A special number called a ‘nonce’.

If the hash value produced is below a certain threshold, the proof of work is complete. If it is not, the miner must try again with another value for the nonce. Because there is no way to tell what value of the nonce, when combined with the other two inputs, will produce a satisfactory hash value, miners are forced to simply cycle through nonce values in trial and error (**Figure B**).

Figure B Bitcoin’s proof of work scheme

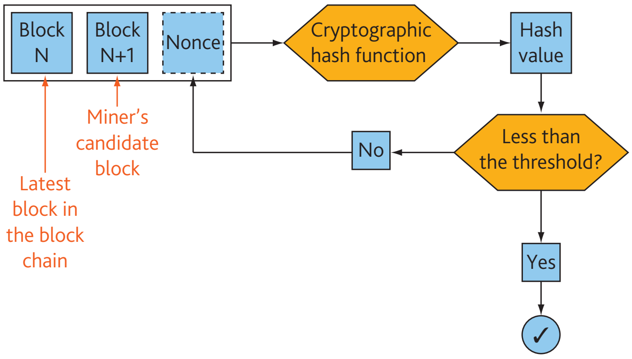

Are digital currencies fraud-proof?

The current design of digital currencies is predicated on the assumption that fraud — the creation of false transactions — can only be achieved by an agent, or coalition of agents, controlling a majority of computing resources on the mining network over a sustained period of time (a ‘50%+1 attack’). However, a number of researchers have suggested that it may be possible to defraud such schemes while possessing less than a strict majority of computing power. Potential weaknesses have been identified in two key areas: (i) the position of an attacker in the network; and (ii) the strategic timing of when an attacker chooses to release messages to the rest of the network.

To appreciate these weaknesses, it may be helpful to consider a simple example of a verification network. **Figure C** provides one such example. Individual miners are arranged in a peer-to-peer network, with each of them controlling a different share of the total computing power. Note that although **Clare** controls the smallest share of the network’s computing resources, she is quite ‘central’ to the network in that she is immediately connected to other nodes that together represent a majority.

Figure C An illustrative digital currency verification network

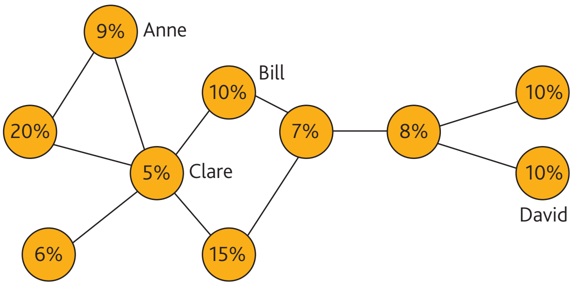

Note: Percentages indicate the share of total computing resources controlled by each node. For simplicity, links are assumed to be undirected (eg if **Bill** is connected to **Clare**, then **Clare** is also connected to **Bill**), although this may not be true in practice.

An attacker’s **position in the network** is important because the longer it takes for messages to propagate across a digital currency’s network, the greater the probability that a fork in the block chain (with two candidates for the next block being successfully verified at similar times) will emerge. A hypothetical attacker that is centrally located in the network (such as **Clare**) will be able to communicate to most of the network very quickly, and so may not strictly require a majority if other users (such as **David**) are, relatively speaking, quite distant.[1] More generally, even honest users in central positions will, for the same reason, be expected, over time, to earn shares of total payments (by successfully adding blocks to the chain) that exceed their shares of computing power on the network.

An incentive also exists for miners to **strategically choose the time** when they broadcast their success at verifying transaction blocks. For example, suppose that when **Bill** successfully verifies a candidate block N, he does not reveal his success immediately. Instead, he starts work verifying block N+1 and only discloses his success to the rest of the network after a short delay. **Bill’s** strategy will force other miners to waste extra time attempting to verify their own candidates for block N and grant **Bill** a head start in trying to verify the next block. Over time, **Bill’s** share of total payments will, on average, exceed his share of total computing power.[2]

Since mining is a zero-sum game — extra earnings for one miner must come at the expense of another — then it is sometimes argued that when one miner receives outsized returns, this creates an incentive for other miners to either drop out or to join the first in a pool, eventually leading to the pool controlling a majority of the network’s computing resources (and so expose the system to the risk of fraud). Complete analysis of these settings is not yet complete,[3] but research done to date does suffice to illustrate that the incentives surrounding fraud prevention in digital currency networks have not been fully explored.

[1] Decker and Wattenhofer (2013) examine propagation times for the Bitcoin network and conclude that a perfectly centrally located attacker would indeed require less than a strict majority of the total computing resources.

[2] Eyal and Sirer (2013) discuss a variant of this strategy in which a single ‘selfish’ miner seeks to establish and maintain an undisclosed lead of at least two in the number of blocks verified over the other, honest miners in the network. In their model, they show that even if the selfish miner is only distantly connected to the rest of the network, their share of total earnings will exceed their share of computing resources when controlling only one third of the network’s computing power.

[3] For example, the Bank is not aware at the current time of any research that has (i) derived the optimal (that is, profit maximising) strategy for each self-interested miner; (ii) established a Nash equilibrium when all agents are individually self-interested and profit-maximising; or (iii) considered the problem of non co-operative bargaining between multiple self-interested miners that seek to pool their resources.

References

Brynjolfsson, E and McAfee, A (2014), *The Second Machine Age: work, progress and prosperity in a time of brilliant technologies*, W. W. Norton & Company, New York.

Decker, C and Wattenhofer, R (2013), 'Information propagation in the Bitcoin network', 13th IEEE International Conference on peer-to-peer computing.

Duncan, G and Elliott, F (2009), 'Chancellor on brink of second bailout for banks', *The Times*, 3 January.

Eyal, I and Sirer, E (2013), 'Majority is not enough: Bitcoin mining is vulnerable', *mimeo*, available at http://arxiv.org/abs/1311.0243.

Finan, K, Lasaosa, A and Sunderland, J (2013), 'Tiering in CHAPS', *Bank of England Quarterly Bulletin*, Vol. 53, No. 4, pages 371–78, available at www.bankofengland.co.uk/publications/Documents/quarterlybulletin/2013/qb130408.pdf.

Goodhart, C (1988), *The evolution of central banks*, MIT Press, Cambridge (Mass).

Haldane, A and Qvigstad, J (2014), *The evolution of central banks: a practitioner's perspective*, available at www.norges-bank.no/pages/100044/14_2_Haldane_and_Qvigstad_28_June_2014.pdf.

Lamport, L, Shostak, R and Pease, M (1982), 'The Byzantine generals problem', *ACM Transactions on Programming Languages and Systems*, Vol. 4, No. 3, pages 382–401.

Manning, M, Nier, E and Schanz, J (2009), *The economics of large-value payments and settlement: theory and policy issues for central banks*, Oxford University Press, Oxford.

McLeay, M, Radia, A and Thomas, R (2014), 'Money creation in the modern economy', *Bank of England Quarterly Bulletin*, Vol. 54, No. 1, pages 14–27, available at www.bankofengland.co.uk/publications/Documents/quarterlybulletin/2014/qb14q102.pdf.

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

Naqvi, M and Southgate, J (2013), 'Banknotes, local currencies and central bank objectives', *Bank of England Quarterly Bulletin*, Vol. 53, No. 4, pages 317–25, available at www.bankofengland.co.uk/publications/Documents/quarterlybulletin/2013/qb130403.pdf.

Norman, B, Shaw, R and Speight, G (2011), 'The history of interbank settlement arrangements: exploring central banks' role in the payment system', *Bank of England Working Paper No. 412*, available at www.bankofengland.co.uk/research/Documents/workingpapers/2011/wp412.pdf.

Wenger, A (2013), *Bitcoin as protocol*, available at www.usv.com/posts/bitcoin-as-protocol.