




BANK FOR INTERNATIONAL SETTLEMENTS



BIS Working Papers No 765

Beyond the doomsday economics of “proof-of- work” in cryptocurrencies

by Raphael Auer

Monetary and Economic Department

January 2019

JEL classification: D40, D20, E42, E51, F31, G12, G28,
G32, G38, L10, L50

Keywords: Cryptocurrencies, cryptoassets, digital
currencies, blockchain, proof-of-work, proof-of-stake,
distributed ledger technology, consensus, bitcoin,
ethereum, money, digitalisation, finance, history of
money

BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in them are those of their authors and not necessarily the views of the BIS.

This publication is available on the BIS website (www.bis.org).

© *Bank for International Settlements 2019. All rights reserved. Brief excerpts may be reproduced or translated provided the source is stated.*

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

Beyond the doomsday economics of “proof-of-work” in cryptocurrencies¹

Raphael Auer

Abstract

This paper discusses the economics of how Bitcoin achieves data immutability, and thus payment finality, via costly computations, ie “proof-of-work”. Further, it explores what the future might hold for cryptocurrencies modelled on this type of consensus algorithm. The conclusions are, first, that Bitcoin counterfeiting via “double-spending” attacks is inherently profitable, making payment finality based on proof-of-work extremely expensive. Second, the transaction market cannot generate an adequate level of “mining” income via fees as users free-ride on the fees of other transactions in a block and in the subsequent blockchain. Instead, newly minted bitcoins, known as block rewards, have made up the bulk of mining income to date. Looking ahead, these two limitations imply that liquidity is set to fall dramatically as these block rewards are phased out. Simple calculations suggest that once block rewards are zero, it could take months before a Bitcoin payment is final, unless new technologies are deployed to speed up payment finality. Second-layer solutions such as the Lightning Network might help, but the only fundamental remedy would be to depart from proof-of-work, which would probably require some form of social coordination or institutionalisation.

JEL classification: D40, D20, E42, E51, F31, G12, G28, G32, G38, L10, L50.

Keywords: cryptocurrencies, crypto-assets, digital currencies, blockchain, proof-of-work, proof-of-stake, distributed ledger technology, consensus, bitcoin, ethereum, money, digitalisation, finance, history of money.


¹ Email: Raphael.auer@bis.org. An earlier version of this paper was titled “The mechanics of decentralised trust in Bitcoin and the blockchain”. I thank David Archer, Morten Bech, Stijn Claessens, Jon Frost, Giulio Cornelli, Henry Holden, Krista Hughes and Paul Wong for comments and seminar participants at the Bank for International Settlements and the 2018 Annual Meeting of the Central Bank Research Association for comments; and Codruta Boar, Alan Villegas and Giulio Cornelli for outstanding research support. The views expressed in this paper are those of the author and not necessarily those of the Bank for International Settlements.

Introduction

Judged by internet searches, popular fascination with Bitcoin and other cryptocurrencies soared in late 2017, outstripping interest in sovereign currencies, or even gold (see Graph 1, left-hand and centre panels).² Yet few people were actually using Bitcoin to buy things (see Graph 1, right-hand panel).³

Global interest in Bitcoin, valuations, and use in retail payments

Graph 1



¹ Numbers represent search interest relative to the highest point on the chart for the given search term and time. A value of 100 is the peak popularity for the term. A value of 50 means that the term attracts half as many searches as a value of 100. A score of 0 means there were not enough data for this term. Google trends was accessed on 5 November 2018 with searches for bitcoin, euro, gold, usd ("search term" and worldwide search interest, respectively). ² Data from the CoinDesk Bitcoin Price Index. BPI value data returned as USD. ³ The displayed line shows the monthly volume of global retail payment transactions made in bitcoins and handled by the major bitcoin payment-processing firms (volumes are expressed in dollars).

Sources: CoinDesk; Google trends (site accessed on 05.11.2018); chainanalysis.com (site accessed on 05.11.2018).

Much of the allure surrounding cryptocurrencies stems from the fact that no government is needed to issue them. And they can be held and traded without a bank account. Instead, they are exchanged via simple technical protocols for communication between participants, as well as a publicly shared ledger of transactions (the "blockchain") that is updated by a decentralised network of "miners" via costly computations, ie "proof-of-work".

What is the economic potential of this new means of exchange? This paper analyses the underlying economics of how Bitcoin achieves payment finality, ie how it seeks to make a payment unalterable once included in the blockchain, so that it can be considered as irrevocable. It then discusses the future of this type of cryptocurrency in general. The focus lies on the technical elements underlying Bitcoin and its blockchain, as devised by Nakamoto (2008). But its conclusions extend to cryptocurrencies that are slightly modified clones of Bitcoin (eg Bitcoin Cash, Bitcoin SV, or Litecoin) or digital tokens that, so far at least, share the crucial reliance on proof-of-work to underpin their payment finality (eg Ethereum or Monero).

² International bodies have also turned their attention to cryptocurrencies, see, for example, Bank for International Settlements (2018), Carstens (2018a,b,c), Committee on Payments and Market Infrastructures (CPMI) (2015, 2017), Financial Action Task Force (2015), Carney (2018), Financial Stability Board (2018a,b) and G20 Finance Ministers and Central Bank Governors (2018).

³ To put the magnitudes in the right-hand panel of Graph 1 in context, the peak of USD 400 million bitcoin payments processed compares with around USD 500 billion processed on average in a month by just one conventional payment processor, VISA.

Nakamoto's⁴ key innovation is to balance the cost and reward for updating the blockchain, by creating incentives to ensure that updates are correct. The updating process deters forgeries by imposing a cost on updating the blockchain. At the same time, accurate updating of the blockchain confers a reward on the so-called miners who do the updating. Miners, or their computers, effectively compete to solve a mathematical problem. Presenting a solution proves that they have done a certain amount of computational work. Such "proof-of-work" allows a miner to add a block of newly processed transactions to the blockchain, collecting fees from the subject transactions as well as "block rewards" – newly minted bitcoins that increase the outstanding supply.

The costs and rewards of Nakamoto's updating process are the focus of our discussion here. Two questions are raised. First, how efficient is the fundamental architecture of deterring forgeries via costly proof-of-work? And second, can the market for transactions actually generate rewards that are valuable enough to ensure that payment finality is really achieved?

Analysing these two elements uncovers fundamental economic limitations that cloud the future of cryptocurrencies based on proof-of-work. In sum, with the current technology, it is not even clear whether such cryptocurrencies can keep functioning as they do at the time of writing. This statement is unrelated to well known restrictions on the scale of such payment systems or the volatility of cryptocurrencies.⁵ Rather, it concerns the fundamentals of Nakamoto's updating process, which has two limitations that interact in a fateful manner.

The first limitation is that proof-of-work axiomatically requires high transaction costs to ensure payment finality (see also the important contribution by Budish (2018) on this issue). Counterfeiters can attack bitcoin via a "double-spending" strategy, ie spending in one block and later undoing this by releasing a forged blockchain in which the transactions are erased.

This paper starts by introducing the concept of "*economic payment finality*" in the blockchain. That is, a payment can be considered final only once it is unprofitable for any potential adversary to undo it with a double-spending attack. This economic concept differs starkly from the operational considerations of finality in Nakamoto (2008), who examines a double-spending attack by a large miner controlling a significant fraction of the network's computational power. Nakamoto's definition of payment finality (although not explicitly spelled out as such) is thus operational: the deeper a payment is buried in the ledger, the less likely an adversary with given computational resources will succeed in a double-spending attack.

If the incentives of potential attackers are analysed, it is clear that the cost of economic payment finality is extreme. For example, to achieve economic payment finality within six blocks (one hour), back of the envelope calculations suggest that mining income must amount to 8.3% of the transaction volume – a multiple of transaction fees in today's mainstream payment services. The underlying intuition is simple: double-spending is very profitable. In fact, attackers stand to gain a much higher bitcoin income than does an honest miner. While honest miners simply collect block rewards and transaction fees, counterfeiters collect not only any block rewards and transaction fees in the forged chain, but also the amount that was double-spent, ie the value of the voided transactions. This "*attacker advantage*" ultimately translates into a very high required ratio for miners' income as compared with the transaction volume (the amount that can be double-spent).

⁴ Nakamoto (2008) – a pseudonym for a hitherto unknown person or group of persons – did not invent the individual technological elements of bitcoin but rather made use of a novel combination of existing technologies. Proposals for digital forms of cash date include eg Chaum (1983). The proof-of-work concept is commonly attributed to Dwork and Naor (1992), while Szabo (2005) too recognised that this principle (initially developed to deter spam) could be used in digital payment systems.

⁵ On limited scale and volatility, see in particular Bank for International Settlements (2018). For other limitations, see eg Biais et al (2017), Huberman et al (2017), Budish (2018), and Morris and Shin (2018).

The second fundamental economic limitation is that the system cannot generate transaction fees in line with the goal of guaranteeing payment security. Either, the system works below capacity and users' incentives to set transaction fees are very low, or the system becomes congested (see Huberman et al (2017) and Easley et al (2018) for analysis of the case of congestion and associated queuing). Underlying this is a key externality: the proof-of-work and hence the level of security is determined at the level of the block one's transaction is included in, with protection also being provided by the proofs-of-work for subsequent blocks. In contrast, the fee is set by each user privately, hence creating a classical free-rider problem, amounting to a veritable "*tragedy of the common chain*".⁶ While each user would benefit from high transaction fee income for the miner, the incentives to contribute with one's own fee are low.

The key takeaway of this paper⁷ concerns the interaction of these two limitations: proof-of-work can only achieve payment security if mining income is high, but the transaction market cannot generate an adequate level of income. As a result, liquidity is set to deteriorate substantially in years to come. The backdrop is that the bulk of miners' current income consists of block rewards. But block rewards are being phased out. For example, in Bitcoin and many of the clones that have "forked" from it, the next time block rewards will halve is in 2020. Whenever block rewards decrease, the security of payments decreases and transaction fees become more important to guarantee the finality of payments. However the economic design of the transaction market fails to generate high enough fees. A simple model suggests that ultimately, it could take nearly a year, or 50,000 blocks, before a payment could be considered "final".

Given these considerations, the paper concludes with a discussion of how technological progress is set to affect the efficiency of Bitcoin and related cryptocurrencies. So-called second-layer solutions such as the Lightning Network can improve the economics of payment security (in addition to mitigating scaling limits). However, they are no magic bullets, as they face their own scaling issues.

In order to prevent liquidity from ebbing away, Bitcoin and other cryptocurrencies would need to depart from using proof-of-work – a system that is not sustainable without block rewards – and embrace other methods for achieving consensus on blockchain updates. Among many proposed developments, the most prominent one is "proof-of-stake," a system in which coordination on blockchain updates is enforced by ensuring that transaction verifiers pledge their coin holdings as guarantees that their payment confirmations are accurate. Yet, because such a system lacks the solid grounding offered by proof-of-work (which proves actual offline activity), its success may rest on additional overarching coordination mechanisms, ie some degree of implicit or explicit coordination by an institution.⁸

Judging based on the current technology, the overall conclusion is that in the digital age too, good money is likely to remain a social rather than a purely technological construct (see eg Carstens (2018a) and Borio (2018)).⁹ That cryptocurrencies might in future profit from social coordination or institutions is also

⁶ The tragedy of the commons is a frequently encountered problem in economics when individuals try to reap benefits from a commonly accessible resource in fixed supply without taking into account the effect of one's consumption on the well-being of others. The name originates from over-grazing of common land (see Lloyd (1833)).

⁷ Note that Huberman et al (2017) examine congestion in the market for transaction fees while assuming that "the mining resources are sufficient to guarantee the system's reliability and security" (see p 4), a focus very similar to Easley et al (2018). In contrast, Budish (2018) examines the economics of security, ie of double-spending attacks, but not how mining income is determined. In this paper, I combine these approaches to show how the economics of security and the market for transaction fees interact, ie how the market of transactions determines payment security and what this implies for the future liquidity of bitcoin.

⁸ Other proposals, such as "delegated proof of stake" or "proof of importance", directly aim at implementing such institutionalisation via a variety of voting mechanisms.

⁹ Certainly, above all, even if cryptocurrencies should one day become an economically efficient payment means, the economic and ethical problem of whether this is actually desirable remains. Given the impossibility of dealing with this issue in a few pages, this paper sidesteps it entirely. One aspect is illicit payments. Landau and Genais (2017), Auer and Claessens (2018, 2019), Clayton (2018), Fanusie, and Robinson (2018) and Foley et al (2018) discuss cases where Bitcoin has been used for illicit payments. Another aspect regards the macroeconomic implications of privately issued currencies (see Amihud and Cukierman (2018), Fernández-Villaverde and Sanches (2016), and Schilling and Uhlig (2018) on the implications of currency competition).

highlighted by the very same algebra that shows the doomsday economics of pure proof-of-work. The point is that their payment efficiency could be greatly improved by introducing an institutional underpinning to undo double-spending attacks should they occur. In this light, one key question for future research is whether and how technology-supported distributed exchange could complement the existing monetary and financial infrastructure.


Technology basics and the economics of mining

This section starts by giving an overview of how bitcoin transactions work. It then zooms in on the basic economics of proof-of-work, mining, and equilibrium "difficulty".¹⁰

A payment transaction in Bitcoin and its blockchain: system overview

At face value, the idea underlying Bitcoin is simple: instead of a bank centrally recording transactions, a publicly shared blockchain (a chain of files) records successive transactions. Graph 2 depicts the basic concepts of a purchase with a central ledger updated by a bank (left-hand panel) and a distributed blockchain (right-hand panel).

Making a payment transaction via a bank account and via Bitcoin Graph 2



A buyer purchases a good from the seller, who initiates shipment upon perceived confirmation of the payment. If the payment takes place via bank accounts (left-hand panel) the buyer sends the payment instruction to the bank, which adjusts account balances by debiting the amount paid from the buyer's account and crediting it to the seller's account. The bank then confirms payment to the seller. In contrast, if payment takes place via Bitcoin (right-hand panel), the buyer first publicly announces a payment instruction stating that the bitcoin holdings of the buyer are reduced by one, while those of the seller are increased by one. After a delay, a so-called miner includes this payment information in a new block of transactions, which is added to the blockchain. The updated blockchain is subsequently shared with other miners and users, each verifying that the newly added payment instruction is authorised by the buyer and is not a double-spending attempt. Finally, the seller observes that the blockchain including the payment instruction emerges as the one used by the entire network of miners and users.

Source: Author's elaboration.

¹⁰ For other introductions by economists, see Andalfatto (2013), Böhme et al (2015), Athey et al (2016), Bolt and van Oordt (2016), Bech and Garrat (2017), Catlini and Gans (2017), and Chiu and Koepl (2017), Andalfatto (2017), Berentsen and Schaer (2017, 2018), Pichler et al (2018), Abadi and Brunnermeier (2018), and Lewis (2018).

In more detail, the transaction on the right-hand side of Graph 2 plays out as follows:

1. The buyer's "cryptographic digital signature" publicly announces the payment transaction, including the payee, the paid amount, and the transaction fee the payer is willing to pay to the miner.
2. Miners select the unprocessed transactions that will maximise their income from fees and engage in computations until the first miner emerges with a valid proof-of-work.
3. The successful proof-of-work allows the miner to add a block of transactions to the blockchain, collecting the fees of the included transactions and the block reward.
4. The new blockchain is shared among the network of miners and other users, who also verify the update (verify the proof-of-work, the signatures, and the absence of double-spending). If this new blockchain emerges as the consensus version, the majority of miners keep on adding to it.

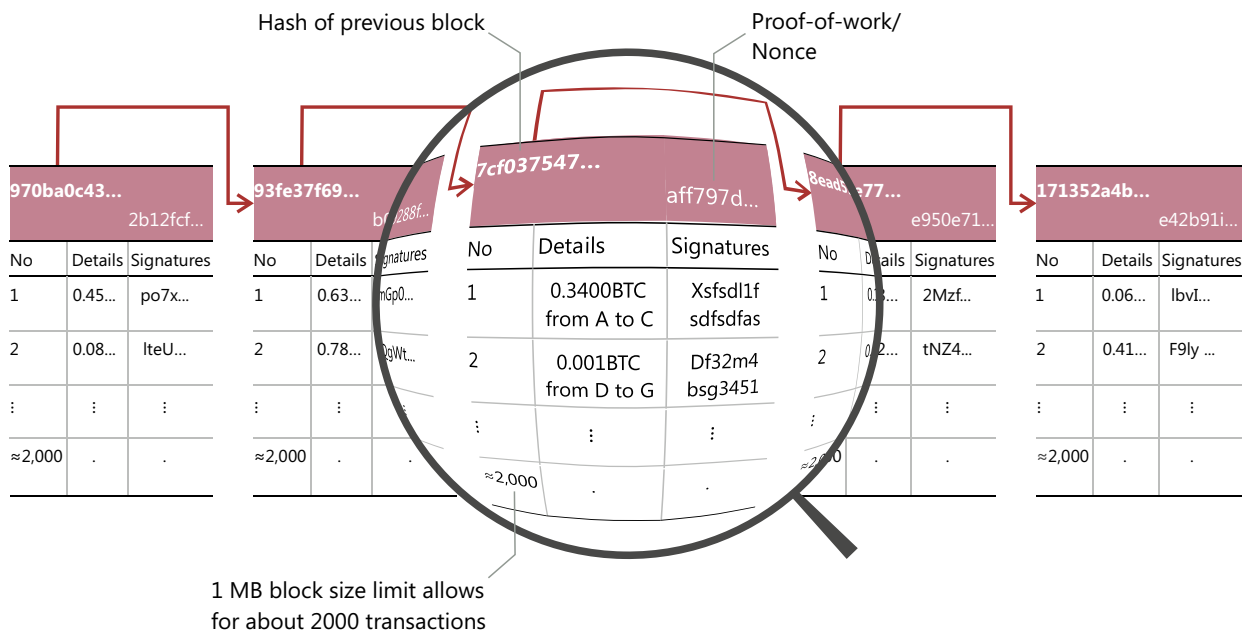
In the above steps, any bitcoin owner should be able to spend their funds, but only once so. "Cryptographic digital signatures" are used to verify payment transactions such as "C pays 1 to S." This digital signature proves that the payment has been authorised by whoever controls the bitcoins that are being spent.

That leaves the problem of double-spending, which might occur, for example, if C were to simultaneously broadcast the payment instructions "C pays 1 to S" and "C pays 1 to Y" for one and the same bitcoin. Because C uses the correct digital signature to sign both payment messages, both are valid. But the blockchain would be compromised if both transactions are entered into it, and a way needs to be found to stop this from happening. One part of the solution is that merchants check the public blockchain, verifying that their counterparties actually own the amount of funds they claim to be transferring.

The second, and crucial, part is an algorithm that incentivises miners to add only correct updates to the blockchain. In a decentralised system, there is no reputation to lose and laws cannot be enforced easily. The risk is that counterfeiters would spend bitcoins and simultaneously disseminate fake versions of the blockchain to the network, in the hope that merchants will accept these fake blockchain versions and transfer goods to the counterfeiters. Hence, updating the blockchain must be expensive enough to deter fake updating attempts. Yet, if updating the blockchain is costly, there must also be a reward in place to incentivise truthful updates. The bitcoin protocol solves this by creating a class of agents known as miners, who update the blockchain via computational work, and in return receive block rewards and transaction fees when they add batches of valid transactions ("blocks") to the blockchain.¹¹

Graph 3 gives a schematic overview of the resulting blockchain and its main elements: the publicly available blockchain is updated in blocks of transactions. Each block is a small file that includes a number of payment transactions, stating the amount, the payer and the payee. Blocks, in turn, are chained to each other sequentially, thus forming the blockchain.

¹¹ Miners face strong incentives to check the validity of the transactions that they include in their block, for if any of the included transactions turns out to be invalid (because either the signature is invalid or somebody has spent funds that they don't actually own) the entire block is invalid, thus also invalidating the transaction fees and the block reward. But the validation itself is not computationally intensive when compared with the computational effort involved in proof-of-work.



The publicly available ledger is updated in bunches of transactions, and each update is termed a “block.” Blocks, in turn, are chained to each other sequentially, thus forming the “blockchain.” The blockchain is updated much like adding individual pages with new transactions to a ledger, with page numbers determining the order of the individual pages. Each block is a small file that includes a number of payment transactions, stating the amount, the payer and the payee, and also the transaction fee. The original Bitcoin protocol restricts each block to a maximum file size of 1 MB, which in practice implies that around 2,000 transactions can be included in each block. Only transactions including the valid digital signature associated with the transferred funds are accepted into a block. A new block is added to the blockchain only about once every 10 minutes. Adding a block to the existing block chain requires a valid proof-of-work (also called a “nonce”), which involves a hash function that takes a random text input and produces from this an output according to set rules. The key property of the SHA256 hash function used in the Bitcoin protocol is that the output is unpredictable: to get a desired result, the only solution is thus to try many starting values randomly, which creates a computing cost. Cryptographic chaining of blocks is achieved by including summary information from the previous block in the proof-of-work of the current block.

Source: Author’s elaboration.

Proof-of-work: rolling a dice

Proof-of-work is a simple cryptographic tool that allows to send a credible signal to others that a certain amount of money has been wasted on electricity and equipment. An analogy is a dice with a large number of sides that each have an equal probability of coming out on top. If there are 1,000 sides numbered from 1 to 1,000, on average, one would have to roll the dice 100 times until a number between 1 and 10 comes out on top. Showing a dice with any number between 1 to 10 on top thus shows that one in all likelihood has rolled the dice about 100 times.

Proof-of-work is the mathematical equivalent of credibly rolling the dice. It relies on asymmetrical mathematical problems with solutions that are difficult to come up with, but easy to verify. This is the process of hashing. A “hash” function takes a random text input and produces from it a hash output according to set rules.¹² The hash function used in the Bitcoin protocol – known as SHA256 – satisfies the property that it is not possible to deduce the input from the specific hash output.¹³

¹² For example, a simple hash function is to take the second and fourth letter from any input. For the input “ABcDSEffdf...”, the output of this hash function is “BD”.

¹³ Note that the security of the SHA256 is not guaranteed axiomatically. Other hash functions once thought safe have, in fact, been broken.

A specific SHA256 output can thus be found only by trial and error, which proves that a miner has done a certain amount of computational work. Conceptually, the Bitcoin protocol will only add blocks to the blockchain that are accompanied by a rare hash output. This is defined as one starting with many 0s (or more precisely, the value of which expressed in binary numbers is below some “target” level). The expected number of hashes that needs to be performed to obtain a hash result below target is called “difficulty” (difficulty is thus proportional to $1/\text{target}$).

By adjusting the target level, it is possible to change the cost of adding a block to the blockchain. To translate this difficulty into the expected cost to add block b to the chain, one needs to know the cost per hash, ie the cost of performing one SHA256 computation:

$$\text{Expected cost of proof of work } b = \frac{\text{Cost per Hash}}{\text{Chance of success}_b} = \text{Cost per Hash} * \text{Difficulty}_b$$

The cost per hash is the cost of the computational equipment required to perform the hash calculations, as well as the cost of electricity and other operating costs. To get a sense of the magnitudes, in mid-2018, the Antminer S9, a frequently purchased item of mining hardware, could perform an advertised 13.5×10^{12} SHA256 hashes per second while using around 1,300 watts of power. Assuming an electricity cost of $0.05 \frac{\text{cents}}{\text{Kwh}}$, a price tag of USD 1,000 for the hardware, no other costs, and a life expectancy for the equipment of three years, the cost per hash emerges as

$$\frac{1.3\text{Kw} * 0.05 \frac{\text{USD}}{\text{Kwh}} + \frac{1000 \text{ USD}}{3 * 365 * 24 \text{ h}}}{13.5 * 10^{12} * 3600 \text{ hash/h}} \approx 2.12 * 10^{-18} \text{ USD/hash} \quad (1)$$

If a miner comes up with an input that solves to a low hash value, this signals that the person, group, or company has done a certain amount of computational work. An input text that solves for a hash result below target is called a proof-of-work (or “nonce”), and it allows a block of transactions to be added to the blockchain.

The final piece of the updating game is a coordination algorithm ensuring that a unique consensus emerges between the many actors exchanging information. Note that, although the above elements show how one blockchain is updated, it cannot be taken for granted that there exists only one version of the blockchain. In fact, due to errors, coordination issues, and attacks, there are often competing versions. Bitcoin solves this issue by adopting the rule that, if competing versions are observed, the one which is the most expensive to forge continues to be used. Since this is generally the blockchain with the most blocks, this rule can be summarised as: “follow the longest chain.”¹⁴ Bitcoin thus solves a coordination issue via an economic approach: it coordinates on the version into which the most resources have been invested.¹⁵

Mining economics and equilibrium “difficulty”

How is equilibrium on the market for blockchain updates determined? A formal analysis requires some notation to be introduced. Let blocks be represented by an integer number b that starts at 0 (the “genesis” block) and then increases in steps of one. Each transaction in a block is indexed by t . The revenue to the miner (denoted by $\text{Mining revenue}_b^{BTC}$ and expressed in bitcoin (BTC)) for coming up with the proof-of-work for block b is:

$$\text{Mining revenue}_b^{BTC} = \text{Block reward}_b + \sum_{t \text{ in } b} \text{Fee}_t \quad (2)$$

¹⁴ Calling this the rule of following the longest chain is slightly misleading: the actual rule is not that the longest chain should continue to be used but the one that is the most difficult to forge. A chain with a small number of blocks but of high difficulty can be more costly to forge than a long chain with low average difficulty.

¹⁵ Note that game-theoretic analysis, as presented in Biais et al (2018), shows that the consensus mechanism developed by Nakamoto (2008) is not unique: other equilibria in which miners coordinate on random “sunspot” events and branch off (ie “fork”) the blockchain can emerge.

The evolution of block rewards ($Block\ reward_b$), was specified at the outset in Nakamoto (2008) (see Graph 4 for their path) and is given only by the block number: block rewards were originally set at 50 bitcoins per block, and halve every 210,000 blocks (around every four years). They will be set to zero once the halving results in less than $1/100,000,000$ of bitcoin (the smallest denomination possible). By then (probably at some point in the year 2140), the total supply of bitcoin will reach its upper limit of 21,000,000.


Miners also collect the transactions fees ($\sum_{t\ in\ b} Fee_t$) in equation (2), but these are currently very meagre compared with the block rewards (see Graph 4). Fees are determined endogenously by the system (see the section below on the transaction market).

Importantly, the *difficulty* of adding a new block to the blockchain is self-calibrating. This ensures that the number of blocks added tends to be stable over time. The number of miners in Bitcoin may fluctuate over time, while technological advances are likely to reduce the cost of hashing. The Bitcoin protocol has an in-built formula that, every 2,016 blocks (about every two weeks), adjusts the difficulty of finding a rare hash result. The difficulty is increased if blocks have been added more quickly than one block every 10 minutes, and reduces the difficulty otherwise. This means that the implicit cost of finding a valid proof-of-work, and also the required break-even mining revenue, fluctuate over time and with the entry and exit of miners.


Block rewards, bitcoin circulation, and mining income

Graph 4

Bitcoins are brought into circulation via block rewards, but at a decelerating speed



Miners' income is made up of block rewards and transaction fees



All bitcoins in existence have been issued via "block rewards." Every new block added to the block chain increases the total supply, with the newly created bitcoins being credited to the miner who adds the block. Block rewards were set to 50 bitcoins per block initially and are halving every 210,000 blocks. They will be set to zero once the halving results in less than $1/100,000,000$ of bitcoin (one Satoshi), meaning that the total supply of bitcoins will be 21,000,000. Miners' income is made up of block rewards and transaction fees. Dashed pattern indicates estimated future values.

¹ Thirty-day moving average of the sum of all transactions fees in each block (in bitcoin).

Source: <https://bitinfocharts.com>; <https://coinmetrics.io>; author's calculation.

If we assume that miners are risk-neutral and that the mining process is competitive, the break-even or free entry condition is that the expected mining revenue for block b is equal to the expected mining cost:

$$P_{USD} Mining\ revenue_b^{BTC} = Expected\ Mining\ cost_b$$

where P_{USD} is the value of one bitcoin in US dollars.

Taking into account the above free entry condition together with the determinants of the mining revenue, this solves to the equilibrium difficulty of Bitcoin. In this equilibrium, the cost of updating the blockchain (the difficulty of the updating game times the cost to produce one hash in USD) is equal to the reward (the sum of all transaction fees in a block plus the block reward, and all this multiplied by the USD price of one bitcoin):

$$Difficulty_b * Cost\ per\ Hash = P_{USD} (Block\ reward_b + \sum_{t\ in\ b} Fee_t) \quad (3)$$

Equation (3) shows that, with endogenous difficulty, proof-of-work becomes a purely economic concept. The difficulty adjusts so that miners, on average, break even and the average expenses of the computational work equal the block reward plus the transaction fees times the price of bitcoin.¹⁶


Developments in late 2018 offer a window of opportunity to study how the economics underlying equation (3) play out in practice. The price of bitcoin, which had fluctuated around USD 6,000–6,500 from April to early November, collapsed first to around USD 4,500 in late November, and then to USD 3,500 and below in mid-December. This resulted in a massive exodus of miners who were no longer able to recover the power costs of running their mining equipment. As a result, proof-of-works were discovered at a speed of less than one every 10 minutes, and Bitcoin's difficulty decreased substantially until the free entry condition in (3) was restored.¹⁷

Bitcoin price developments, difficulty, and block discovery time during late 2018

Graph 5

Proof-of-work difficulty follows the USD price of bitcoin as...

... falling bitcoin prices cause miners to shut down equipment, resulting in fewer block discoveries, and thus a downward re-calibration of difficulty²



¹ Data from the CoinDesk Bitcoin Price Index. ² The bitcoin protocol adjusts the difficulty of the proof-of-work required to add a block to the blockchain such as to keep the average arrival time steady at one per 10 minutes. The adjustment of the difficulty takes place every 2,016 blocks (ie around every two weeks). If over the most recent 2,016 blocks, the average arrival time was faster than one per 10 minutes, the difficulty increases. If the arrival time was slower, it eases. During the price collapse in late 2018, many miners shut down their equipment as they could not recover their power costs. This led to a decline in the block arrival time (see right-hand panel), and thus eventually a decline in difficulty.

Sources: bitcoinwisdom.com; CoinDesk; data.bitcoinity.org (site accessed on 4 Jan 2019).

Attacker advantage: the high cost of economic finality via proof-of-work

How good is the security provided by a specific "difficulty" and what is the economic cost of achieving finality via such costly signals? This section starts by introducing the concept of "economic payment finality", by which a payment can be considered irrevocable only once it is unprofitable to reverse it. This concept takes as its starting point the double-spending attacks outlined in Nakamoto (2008). However,

¹⁶ Equation (3) also shows that technological progress, for example in the form of cheaper computing power and a lower cost per hash, is simply offset by higher difficulty: self-calibrating proof-of-work ultimately proves that a certain amount of actual resources has been spent and, on balance, this amount of money is equal to the expected reward.

¹⁷ Note that the difficulty was still increasing until August 2018, despite prices being much below the 2017 peak. This is evidence of the time it takes time to accumulate a stock of mining equipment. Prat and Benjamin (2017) provide an in-depth analysis of the dynamics of entry into the mining industry.

instead of asking what the chances of an attack succeeding are, as Nakamoto does, it raises the question of what the incentives of potential adversaries might be (see also Budish (2018) for a closely related analysis).

An economic definition of payment finality

Finality in exchange generally means that “a transfer of funds [or] a transfer of securities that have become irrevocable and unconditional” (see CPSS (2003, p 496)). The key here is that a traditional payment or other transfer is not protected by market mechanisms, but rather by the legal system: once a payment has made its way through the national payment system and into the books of the central bank, it is final by law and cannot be revoked.¹⁸

In a cryptocurrency, finality is a starkly different concept. Broadly, it signifies that once a transaction is included in the blockchain, there is certainty that it will not be undone later by the emergence of an alternative “longer” blockchain which does not include the subject transaction.

Nakamoto (2008) considers a change-of-history attack by a large miner controlling a significant fraction of computational power of the network. In this type of “operational attack”, the adversary spends bitcoins while simultaneously and covertly mining and building an alternative “longest chain” that does not include these transactions. In this scenario, merchants would wait for the payments to be included in a block, and then wait for a certain number of subsequent blocks to be added to the blockchain (each additional block is called a “confirmation”). The attacker would wait until all merchants accept the payment, and then release the secretly mined blockchain in which the original payment instructions are not included. If the attacker was successful in outmining the rest of the network, the secret chain would be accepted as the consensus upon release.

Nakamoto’s definition of payment finality (although not explicitly spelled out) is thus probabilistic: if a payment is buried deep in the ledger, it is unlikely (though not impossible) that a longer rival chain without the payment transaction exists. Thus, the probability of an operational attack succeeding depends on the adversary’s share of the total computational power available, and on how fast payments are accepted as being final. It declines exponentially with the length of time that merchants are willing to wait before releasing the payment. For example, if merchants follow a rule to release merchandise after a waiting time of n confirmations and the attacker controls a share $0 < x < 0.5$ of the mining power, the chance of this type of attack succeeding is related to $(x/x-1)^{n+1}$.¹⁹

By contrast, the economic notion of payment finality can be defined as follows: **a cryptocurrency payment can be considered as final once it is certain that from a certain moment of time onwards, it will never be profitable to undo the payment via a double-spending attack.**

Thus, to establish payment finality in this way, it is necessary to evaluate what the costs and gains to a potential attacker might be. In other words, a payment is considered safe from attack as soon as an attack would no longer be profitable. This perspective gives a radically different answer as to when a payment can be considered final compared to the operational considerations in Nakamoto (2008) (see Table 1 for an overview).

To exemplify an economic attack on bitcoin, take the following stylised example of an adversary who rents mining equipment for a short period of time to conduct a double-spending attack. Consider a simple strategy to undo all transactions in block number b by renting computational equipment and mining faster than the rest of the network until block number $b + waittime$, where *waittime* is the time expressed in blocks

¹⁸ This is not to say that errors cannot be corrected ex post. Erroneous payments can be reversed via legal challenge and associated re-payments, but this does not invalidate the original payment.

¹⁹ The chance of an attack succeeding is not exactly equal to x^{n+1} as the adversary might lose ground against the network of miners, but eventually catch up and out-mine the network later on. For example, the chance of successfully out-mining the rest of the network for six blocks is around 5% for an adversary controlling 25% of the network’s CPU power.

until merchants irreversibly release the merchandise plus one block (as the forged block chain needs to be longer than the one the rest of the network has been mining).

Approaches to payment finality: legal, operational, and economic definitions Table 1

Legal finality	<p>Concept: Legal or procedural definition</p> <p>Logic: A transaction is considered final once specific legal requirements have been met implying that a transfer is unconditional and irrevocable even if one of the involved parties becomes insolvent or enters into bankruptcy (see CPMI-IOSCO (2012)).</p> <p>Criterion: The specific criteria for settlement finality differ by jurisdiction, involved counterparties, mode of payment, and asset class (see CPMI (2012) for examples).</p>
Operational/probabilistic finality	<p>Concept: Probabilistic – no formal definition of finality, only the idea that a transaction buried “deep” in the ledger is unlikely to be reversed.</p> <p>Logic: If an adversary controls a given share of the total mining power and merchants wait for <i>Waittime</i> blocks before shipping goods, the probability can be calculated that the adversary can mine enough blocks to overtake the rest of the network of miners and undo a payment via the release of a longer blockchain.</p> <p>Formula for likelihood of finality (from Nakamoto (2008)):</p> $1 - \sum_{k=0}^{Waittime} \frac{\lambda^k e^{-\lambda}}{k!} \left(1 - \left(\frac{x}{1-x}\right)^{Waittime-k}\right)$ <p>Parameters:</p> <ul style="list-style-type: none"> - Share x ($0 < x < 50\%$) of the network’s hashing capacity that is controlled by the adversary. - Merchants wait for <i>Waittime</i> blocks until goods are shipped.
Economic finality	<p>Concept: Incentives – a transaction is final once it is no longer profitable to reverse it.</p> <p>Logic: If potential attackers can rent mining equipment on a short-term basis, how long do merchants have to wait until is unprofitable to undo the payment via a double-spending attack?</p> <p>Formula: see equation (7):</p> $\underbrace{\sum_{i=b}^{i=b+Waittime} Mining\ revenue_t^{BTC}}_{Cost\ of\ a\ forgery} > \underbrace{\sum_{t\ in\ b} Amount_t}_{Amount\ that\ is\ double-spent} > \underbrace{\left(\frac{Cost\ per\ rented\ hash}{Cost\ per\ hash} \frac{P_{USD}}{(1 - \Pi^{HF}) P_{USD}^{Attack}} - 1\right)^{-1}}_{Attacker\ disadvantage}$ <p>Parameters:</p> <ul style="list-style-type: none"> - Ratio of miner’s income compared to transaction volume. - Adversary cost disadvantage for short-term mining equipment rentals. - Price decline following successful attacks. - Probability of social coordination to undo a double-spending attack.

Sources: Nakamoto (2008); author’s calculations.

The cost of an attack would be as follows: if the attacker is able to rent equipment for a *Cost per rented hash* (which would likely exceed the previously introduced *Cost per Hash* of normal miners), the expected cost to forge a blockchain from block b onwards and until block $b+waittime$ is equal to²⁰

$$Cost\ Attack_b = Cost\ per\ rented\ hash * \sum_{i=b}^{i=b+Waittime} difficulty_i \tag{4}$$

Inserting the free entry condition (3) that relates the mining difficulty to mining income yields that the cost of an attack increases with the waiting time (ie how long any forged block chain would have to be to

²⁰ Note that, in this attack vector, it is crucially assumed that attackers can rent any equipment they want at a stated price. The attack vector is thus certain to succeed. The calculations also assume that the difficulty is determined instantaneously and constantly for the subsequent blocks.

convince counterparties to release the merchandise), with the attacker's cost disadvantage $\frac{\text{Cost per rented hash}}{\text{Cost per hash}}$, and with the revenues from mining:

$$\text{Cost Attack}_b = \frac{\text{Cost per rented hash}}{\text{Cost per hash}} * P_{USD} * \sum_{i=b}^{i=b+Waittime} \text{Mining revenue}_i^{BTC} \quad (5)$$

Equation (5) shows that the higher the miner's income per block is, the higher are the equilibrium expenses that an adversary would need to incur in order to forge the blockchain.

On the other hand, the gain from an attack is not only the double-spent coins (the sum of all transactions t in block b , $\sum_{t \text{ in } b} \text{Amount}_t$), but also the mining income (block rewards plus transaction fees) for the forged wait-time blocks. The gain from this economic attack is thus

$$\text{Gain Attack}_b = P_{USD}^{Attack} ((1 - \Pi^{HF})) (\sum_{t \text{ in } b} \text{Amount}_t + \sum_{i=b}^{i=b+Waittime} \text{Mining revenue}_i^{BTC}) \quad (6)$$

One important thing to note is that, because the attacker forging the blockchain collects not only the double-spent bitcoins, but also the block rewards and transaction fees in the forged chain, the attacker collects a higher bitcoin income than an honest miner.

However, an offsetting force is that P_{USD}^{Attack} , the price of bitcoin in USD after an attack, is potentially much lower than the pre-attack price P_{USD} . This price drop reflects the collapse of confidence in Bitcoin that would probably ensue after a successful double-spending attack.

Also overarching coordination mechanisms by the network of users provide protection for payments (see more on this in the conclusion). This is captured by the term $((1 - \Pi^{HF}))$, where Π^{HF} represents the probability that, following a successful double-spending attack, the network of users would collaborate to ignore the forged chain (even though it is the longest chain), ie by initiating a so-called hard fork.

An attack is unprofitable as long as the expected cost of an attack exceeds the expected gain:

$$\text{Cost Attack}_b > \text{Gain Attack}_b$$

We can rearrange this expression to show that bitcoin or any other proof-of-work-based cryptocurrency is safe from such an attack as long as:

$$\begin{aligned} & \underbrace{\sum_{i=b}^{i=b+Waittime} \text{Mining revenue}_i^{BTC}}_{\text{Cost of a forgery}} \\ & > \underbrace{\sum_{t \text{ in } b} \text{Amount}_t}_{\text{Amount that is double-spent}} * \underbrace{\left(\frac{\text{Cost per rented hash}}{\text{Cost per hash}} \frac{P_{USD}}{(1 - \Pi^{HF}) P_{USD}^{Attack}} - 1 \right)^{-1}}_{\text{Attacker disadvantage}} \end{aligned} \quad (7)$$

The left-hand side of equation (7) shows that Bitcoin is safe from an economic attack if the costs of forging the blockchain is high, which can either be a result of each block coming with high block reward and transaction fees or because the number of blocks that need to be forged is large (high *waittime*).

On the other hand, the first term on the right-hand side of equation (7) shows that Bitcoin is more susceptible to an attack if the total value of the transactions included in this block is large, ie if the amount that can be double-spent is large.²¹ The second term on the right-hand side of equation (7) summarises that the attacker is at a cost disadvantage: renting mining equipment at short notice is expensive, and $\frac{\text{Cost per rented hash}}{\text{Cost per hash}}$ is likely above one.

Offering further protection from an economic attack is the consideration that $\frac{P_{USD}}{P_{USD}^{Attack}} < 1$, ie that the value of bitcoin would collapse after a successful attack. This proceeds from the fact that part of the gains from

²¹ Note that many guidelines regarding the safe use of bitcoin recommend making the waiting time dependent on the amount of transactions, but equation 5 shows that the size of an individual transaction is not relevant. The maximal gain from a successful double-spending attack depends on the total value of all the transactions in a block. Thus, the waiting time depends on the total value of all transactions in a block rather than on the value of an individual transaction.

a double-spending attack come from bitcoins that have been “unspent” in the forged blockchain and which can be spent again in the future. Yet, if bitcoin were to lose its value after an attack on Bitcoin, there would be no point in attacking it in the first place. While nobody can say for sure what would happen after a successful attack, a series of attacks on “Bitcoin Gold” shows what might happen. After the attacks took place on 10–20 May 2018, the cumulative price of this smaller cryptocurrency fell by more than one fifth.

Last, Bitcoin would be harder to attack if $\Pi^{HF} > 0$, ie if users could be persuaded to ignore the rules of Nakamoto (2008) in case of a successful double-spending attack. It is noteworthy that, while nobody can say whether this would actually happen, there have been instances in the past where the Bitcoin community has ignored the rule to follow the longest chain (see below).

Equation (7) offers three key insights regarding the security of payments in the blockchain. First, assuming that users are true to Nakamoto (2008) and will not initiate a hard fork ($\Pi^{HF} = 0$), and assuming that mining revenue and the amount that is spent is constant, rewriting (7) highlights the high required transaction costs:

$$\frac{\text{Mining revenue}_b^{BTC}}{\sum_{t \text{ in } b} \text{Amount}_t} > (\text{Wait time})^{-1} \left(\frac{\text{Cost per rented hash}}{\text{Cost per hash}} * \frac{P_{USD}}{P_{USD}^{Attack}} - 1 \right)^{-1} \quad (8)$$

avg. transaction cost in %

Equation (8) documents the high cost of decentralised payment security. For example, say that Bitcoin users are on average prepared to pay transaction costs of 1%, that rented hash power is twice as expensive as the underlying price of equipment and electricity for honest miners, and that bitcoin would lose one third of its value after a successful attack. Then, the required waiting time is 50 blocks (over eight hours). But if the average transaction cost is 0.1%, the required waiting time is 500 blocks, ie around three and a half days! Equation (8) also shows that one needs to be quite wary about the recommendation seen on many websites that users should wait until their payment is included in a block, and then for at least five additional subsequent confirmations, implying a waiting time of about one hour on average. If users wait for six blocks or fewer (about one hour), the required average transaction costs (as a percentage of the transaction amount) are about 8.3%!

Note that the economic attack vector outlined above crucially assumes that any amount of hash power can be rented at short notice. Up to December 2018, this would only have been a realistic possibility for cryptocurrencies with a modest network of users (some of which were, in fact, attacked), but not for Bitcoin.²² Since then, however, as the price of bitcoin has collapsed, many miners located in countries with high energy costs can no longer recover the cost of electricity and have turned off their equipment, as is evident from the decline in the total hash power of Bitcoin’s network of miners (see Graph 6, left-hand side, and also Graph 5 above). As a result, the surplus of mining equipment that could be switched on any time for high-return double-spending attacks might even bring an attack on Bitcoin within the realms of possibility. And this issue is set to intensify during mid-2020, when the block rewards are halved, pushing further mining equipment out of the regular mining market.


Furthermore, other forms of attack on bitcoin have become substantially cheaper in recent months. As an extreme example, consider how expensive it would be to amass equipment that, in total, would yield 101% of the hash power of all current bitcoin miners. This could then be used to launch double-spending attacks and essentially hold Bitcoin hostage. Although substantial, the cost of doing this has come down dramatically, not just because the hash rate of bitcoin’s network of miners has peaked, but thanks mainly to the steep recent fall in the price of mining equipment (see Graph 6, centre and right-hand panels).

Aside from the discussion on the type of attack vector, two fundamental insights emerge from all economic considerations of an attack. The first is that attackers gain not only the double-spent coins but also, like regular miners, the transaction fees and the block rewards. This makes an attack inherently more profitable

²² On the other hand, it should be noted that the above example describes only a very basic economic attack in which the transaction is undone in a single block and legitimate blocks are mined thereafter. An even more profitable strategy would be to spend coins in several successive blocks and then undo all of these transactions.

than honest mining unless there are strong disadvantages in terms of costs for short-term rentals, a price collapse following any double-spending, or deterrence through overarching coordination. Second, from an economic point of view, the consideration in Nakamoto (2008) that waiting time adds exponentially to the security of bitcoin payments does not hold true: equation (8) shows that waiting times only add linearly to the cost of a forgery, so that the system can sustain low transaction costs only by means of extremely long waiting times. This second consideration is crucial for Bitcoin’s future, as explained in detail below.

Overall hashing rate, equipment prices, and the cost of a large-scale attack Graph 6



¹ Price history of Antminer S9 mining equipment with advertised capacity of 13.5 tera-hash per second on Geizhals.de (price in Germany). ² The cost of acquiring equipment capable of performing calculations at the same rate as bitcoin’s total hashing rate. Calculated as (network hashing rate in TH/S * price per Antminer S9) / 13.5 TH/s.

Sources: <https://bitinfocharts.com>; <https://geizhals.de/bitmain-antminer-s9-a1768361.html>; author’s calculations.

The “tragedy of the common chain” in the market for transactions

The second main economic limitation relates to the inability of bitcoin to generate non-negligible transaction fees other than via congestion. That is, although miners compete to update the blockchain, they cannot affect the maximum number of transactions that are being processed.²³

The basics of the transaction market

Each bitcoin owner wanting to transact sets a transaction fee.²⁴ Miners see all pending transactions and choose those maximising their fee income, thus generating an endogenous average transaction cost. But as long as blocks still have free space, the marginal cost to the miner to include a transaction is 0, and the miners include any transaction with a non-zero fee.

Because the supply of transaction throughput is fixed, while demand for transactions has fluctuated substantially over time, the market for transactions fluctuates between two extremes, as seen in Graph 7.

²³ Although a nuance is that, as the difficulty of the proof of work adjusts only once every two weeks, rapid entry by miners can in fact affect the number of transactions that are being processed for up to two weeks at a time.

²⁴ Any payment transaction also includes a separate transaction fee payable to whomever successfully mines the block in which the transaction is included.

On the one hand, positive and very high fees can result when the system gets congested (on the case of congestion, see in particular Easley et al (2018) and Huberman et al (2017)). When newly added blocks are already at the maximum size permitted by the protocol, the system congests and many transactions go into a queue. Users who want to have their fees transacted immediately start setting higher fees. During peak crypto-hype in late 2017, transaction fees spiked in this way at more than USD 50 per transaction (!), a situation that persisted for some time (see Graph 8).

On the other hand, whenever demand for transactions is such that, even with a fee of zero, blocks are not full, and the equilibrium fee has remained at around zero.

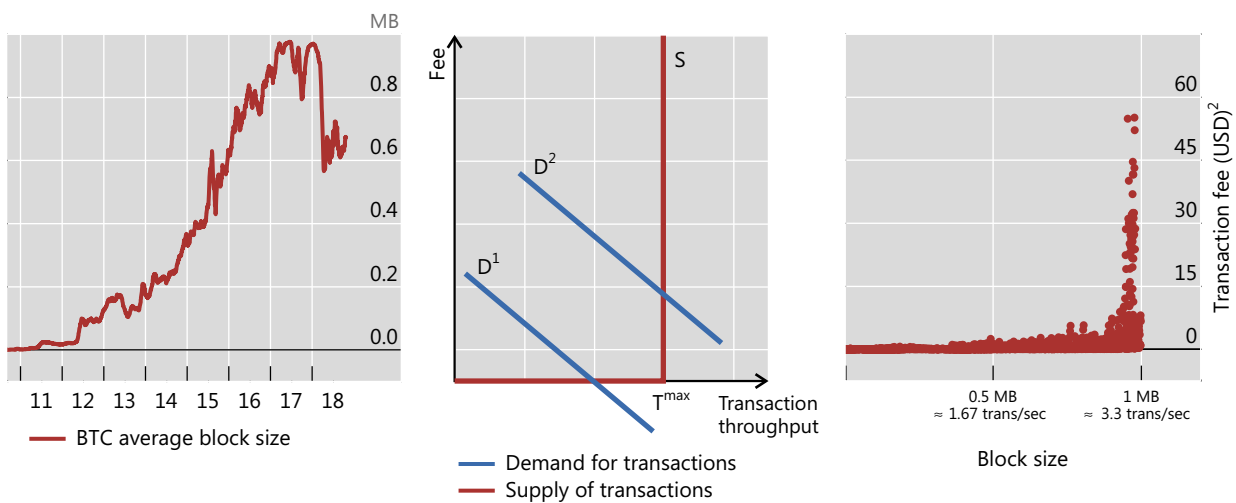
The market for Bitcoin transactions

Graph 7

Transaction demand fluctuates widely¹

With capped supply, demand fluctuations shift fees only when blocks are full...

...which explains the kinked relationship between block size and fees



¹ Thirty-day moving average. ² Transaction fee paid to miners over the period 1 Aug 2010–22 Oct 2018; daily averages.


Sources: www.bitinfocharts.com; author's elaboration.

Spiking fees and congestion of the payment process

Graph 8

As fees spike...¹

... waiting times lengthen²



¹ Average fee per transaction (in USD). ² Seven-day moving average of how long it would take for Bitcoin miners to process all pending transactions.

Source: www.bitinfocharts.com; author's calculations.

A simple model of free-riding in the market for transactions

Underlying very low fees is a key externality that might be termed the “*tragedy of the common chain*”. The problem is that, although higher transaction fees offer higher security, the benefits accrue to all transactions in the block equally (because the cost of counterfeiting is to reproduce the proof-of-work), while the fee accrues to each transaction individually. Even worse, not only does one’s security depend on the sum of fees in the block the transaction is included in (which one can affect), but also on the fees for future blocks (over which one has no control). In sum, this is a classical free-rider problem.

To put this free-rider game into a formal context, consider again equation (7) and, for simplicity, assume that block rewards are 0 (so that fees are miners’ only income), that there are N pending transactions that are all of equal size S (so that the total amount being spent is equal to SN), and that the users waiting to be processed are impatient: each additional block for which they need to wait until the payment can be considered final²⁵ has a cost of μS . Rewriting equation (7) with these assumptions yields

$$\sum_{i=b}^{i=b+Waittime} \sum_{t \text{ in } i} Fee_t > SN * \left(\frac{\text{Cost per rented hash}}{\text{Cost per hash}} \frac{P_{USD}}{P_{USD}^{Attack}} - 1 \right)^{-1} \quad (9)$$

Let us first consider a game in which the user wanting to transact can coordinate and agree on a common fee \bar{F} , resulting in a *waittime* of

$$Waittime > \frac{S}{\bar{F}} \left(\frac{\text{Cost per rented hash}}{\text{Cost per hash}} \frac{P_{USD}}{P_{USD}^{Attack}} - 1 \right)^{-1} \quad (10)$$

Users reduce to a minimum the sum of the common fee and the cost of waiting, $\mu Waittime$:

$$\min_{\bar{F}} \bar{F} + \mu S Waittime,$$

Optimisation results in an optimal fee²⁶ equal to $S \sqrt{\mu \left(\frac{\text{Cost per rented hash}}{\text{Cost per hash}} \frac{P_{USD}}{P_{USD}^{Attack}} - 1 \right)^{-1}}$, ie the optimal choice is to pay a fraction of the transaction size S , which is increasing in impatience (high μ) and increasing to the attackers’ disadvantage (if attackers are at a low disadvantage, high fees are required).

On the other hand, consider a decentralised game in which each user sets their fee privately, taking into account only the benefit to themselves. Users are symmetrical, so in this game too, the equilibrium will be one in which all users post a fee. But, when setting this fee, each user considers deviating from the common rule. Denote the fee that individual j is paying by F_j and that everyone else is paying by \tilde{F} , it holds that the *waittime* has to be such that

$$Waittime N \tilde{F} + (F_j - \tilde{F}) > SN * \left(\frac{\text{Cost per rented hash}}{\text{Cost per hash}} \frac{P_{USD}}{P_{USD}^{Attack}} - 1 \right)^{-1}$$

And the minimisation problem of each individual is

$$\min_{F_j} F_j + \mu S Waittime,$$

which solves to a peculiar first-order condition:

$$\frac{\partial (F_j + \mu S Waittime)}{\partial F_j} = 1 - \frac{\mu S}{N \tilde{F}} \quad (11)$$

²⁵ Note that this waiting game is distinct from the congestion games of Easley et al (2018) and Huberman et al (2017), where blocks may be at their maximum size, thus resulting in a waiting time until a transaction is included in the block chain, but once this has happened a transaction is considered to be final. Here, the number of transactions is assumed to be such that any transaction is included in the next block, but those receiving the funds need to wait for some time until they can consider the payment to be final.

²⁶ This is not the social optimum, but the constrained optimum satisfying the condition of deterring attacks. The calculation below ignores any integer constraints for *waittimes*.

Consider first the problem of user j taking as given the fee set by others. The first-order condition is such that, if an individual is impatient (high μ), or if the average block fee income (equal to the fee \bar{F} times the number of transactions N) the others pay is low, the user anticipates that, in the absence of a high own fee F_j , the waiting time will be extremely long. If $\mu S > N\bar{F}$, the user thus decides to set a fee such that the waiting time is the minimum (one block), essentially putting up the entire bill by themselves. Another extreme – that the users set the lowest non-zero fee possible (one Satoshi, or 1/100'000'000 bitcoin) – results whenever $\mu S < N\bar{F}$.

Of course, everybody else is making just the same calculation, which regulates the fee market. If everybody else was to set the fee just marginally above 0, future blocks would provide almost no security, and j would decide to set a high fee to protect their payment. Vice versa, if everybody else was to set a high fee, future blocks would provide ample security anyway, and, since the expected waiting time is very short, incentives are such that j will set a fee as close as possible to zero. In equilibrium, agents are symmetrical and the fee adjusts such that $F_j = \bar{F} = \frac{\mu S}{N}$.

The key result is that the fee set on a decentralised basis is much lower than the optimal fee \bar{F} , resulting in extreme waiting times. This is simply due to the presence of $\frac{1}{N}$, typically a high number (up to around 2,000, due to the block size limit), as well as the presence of μ (instead of $\sqrt{\mu}$ in the centralised game), typically a very low number as the cost to wait for one additional block – 10 minutes – should be small compared with the size of the payment. With the above example of an attacker disadvantage of $\frac{1}{2}$, 1,000 transactions, and a cost of waiting of 1% per block, the optimal fee \bar{F} is set at 7.07% of the payment with \bar{F} (resulting *waittime* around seven blocks), while it is set at 1%/1000=0.001% in the decentralised game. The resulting waiting time is 50,000 blocks, equivalent to almost a year!

Some considerations on the road ahead

Ingenious though it is, the Bitcoin protocol has severe limitations, as revealed by the above analysis of the underlying economics. What does this imply for the road ahead and the outlook for Bitcoin and related cryptocurrencies?

The doomsday economics of proof-of-work


Putting the pieces of the above analysis together shows that Bitcoin's liquidity will fall substantially in the years to come in the absence of relevant technological advances. One needs to keep in mind that, together with self-calibrating difficulty, proof-of-work becomes a deeply economic concept that ultimately proves that a certain amount of real resources has been used for computations. This amount is insensitive to general technological progress by its very design. And with block rewards – which, at present, represent the vast majority of miners' income and thus underpin the security of payments – being gradually phased out (see Graph 4), the security of payments is also set to deteriorate. Graph 9 gives an outlook regarding how waiting times could increase in the years to come, based on the above considerations of what is required to deter an attack (see, in particular, equation (7)).

Against this dire backdrop, it should be noted that the code of Bitcoin is far from being set in stone. For example, Graph 10 shows how new versions of the leading Bitcoin software client ("Bitcoin Core") have been introduced and adopted by the network of full nodes over time.²⁷ These updates are mostly true to the spirit of the original protocol developed by Nakamoto (2008) (although there have been controversies), but improve technical aspects such as how nodes of the network communicate with one another.

²⁷ A "full node" is a computer or server that maintains a full, up-to-date copy of the Bitcoin blockchain. A full node is "reachable" if it not only receives blockchain updates from other reachable full nodes, but propagates the full blockchain to other nodes.

Substantially longer waiting times result when block rewards decline

Graph 9




Dashed pattern indicates predicted values.

The lines displayed in this graph show the implied waiting time (number of block confirmations before merchants assume the payment is irreversible and release the merchandise) required to make an economic attack unprofitable: the attacker rents mining equipment on a short-term basis and executes a change-of-history attack. The waiting time depends on the attacker disadvantage, which consists in the high price of short-term rentals for hash power or in the likelihood that the price of bitcoin will collapse following an attack. Calculations of the implied waiting times are based on equation (7) and assume transaction fees of 0.18 bitcoin per block, which corresponds to average transaction fees during the period 30 Apr 2018–31 Oct 2018.

Sources: <https://bitinfocharts.com>; <https://coinmetrics.io>; author’s calculations.

The evolution of Bitcoin full-node clients (Bitcoin Core versions)

Graph 10



Distribution of reachable full nodes across leading user agents.

Source: bitnodes.earn.com (site accessed on 18.05.2018).

While there are many other technological developments, in the context of this paper it needs to be noted that many are aimed at boosting the volume of exchange rather than improving the economics of payment finality.²⁸ So-called pruning replaces older parts of the history of past transactions by the netted representation of ownership, which reduces the size of the blockchain. Other solutions are to store transactions more efficiently by shifting some of the information off the blockchain (in particular,

²⁸ In discussing these innovations, one must move beyond Bitcoin as most new technologies are being developed for other cryptocurrencies. However, should one of these technologies demonstrate substantial efficiency gains, it could be adopted by Bitcoin.

transaction capacity has been boosted by the “Segwit” digital signatures, as available in Bitcoin Core clients 16.0 and higher). Last, there are proposals to do away with the strictly linear database structure of the blockchain²⁹ or partition it via so-called sharding.

Second-layer solutions

One technology that can improve upon the economics of payment security is the development of “second-layer” solutions that essentially aim to add liquid methods of exchange on top of the Bitcoin blockchain (which then serves to evidence the underlying value). Most prominent is the “Lightning Network,” proposed by Poon and Dryja (2016). The idea is that two parties A and B jointly lock in one bitcoin on the blockchain via a joint digital signature. A and B can then run a side-contract that keeps track of how the one bitcoin is split between A and B. In this side contract, A and B can shift funds back and forth by digitally signing off changes in the balance without creating any traffic on the blockchain itself. The subcontract is netted on the blockchain only when one of the parties wants to settle.

The second element of the Lightning Network is to connect many of these prefunded bilateral payment channels, with the goal of scaling to a working micro-payment network. This would, for example, let A route a micro-payment to C via B. The Lightning Network is already working in a test environment. As of early 2019, over 500 bitcoins have been committed to bilateral payment channels (see Graph 11), and the network was able to route small payments.

The Lightning Network opens up some new options for cryptocurrencies, as it theoretically allows bilateral transactions to be final before a block is added to the blockchain. And because these transactions occur in separate bilateral contracts, they can somewhat reduce the cost of decentralised exchange: while both opening and netting bilateral payment channels creates “on-chain” transactions (requiring entries in a block and proof-of-work security), the transactions themselves occur “off-chain” (without any corresponding entry into the blockchain) and thus do not directly require proof-of-work security. For example, if every channel is used for 10 payments between opening and netting, it could offer a scaling factor of (very roughly) five, as compared with purely on-chain transactions.³⁰

It is unclear, however, whether second-layer solutions can themselves scale, ie whether they will be restricted to serving small subnetworks of users or, as envisioned by Poon and Dryja (2016), scale to a network that, by the law of six degrees of separation, connects everyone with everyone else. There are two sets of key concerns. The first is of a technical nature and beyond the scope of this paper – it relates to what is required to deter potential attacks on this specific architecture and whether all participants need to be online all the time for payments to be routable.

The second concern relates to economic network theory on the trade-off between efficiency and centralisation. If the Lightning Network remains truly distributed, it would require substantial pre-funding. For example, if routing a payment from A to B typically involves four intermediate channels, it would in total require preloaded values five times as large as the actual payment amount. And it is uncertain whether a typical user, who might upload, say, USD 200 to finance small expenses, would also be willing to foot another USD 800 just to support the network’s routing capacity.

On the other hand, a more efficient network structure is that of hub and spoke: each normal user connects to one larger, highly interconnected intermediary, thus allowing for shorter routes and more netting of payment flows. Such a hub-and-spoke Lightning Network might not look too different from the setup of today’s financial infrastructure.³¹ Recent developments indeed show that the Lightning Network is prone

²⁹ See eg Sompolinsky et al (2016) and Sompolinsky and Zohar (2018). The cryptocurrency IOTA employs such a non-linear ledger, but so far it requires a central coordinator.

³⁰ Dryja (2015) presents some back-of-the-envelope calculations arguing that, with the current block size limit of 1 MB, the Lightning Network could serve between 20,000 (with high security and 150 new channels per user and year) and 8.3 million bitcoin users (assuming six channels per year and user) while consuming half the blocksize on average.


³¹ In the taxonomy of Buterin (2017), a hub-and-spoke Lightning Network could be considered decentralised, but not distributed.

to centralisation, with as of as of 3 January 2019, 362 of a total of 544 committed bitcoins being associated with a single website. In other words, at that point in time, two thirds of the network's capacity was controlled by a single entity.


The Lightning Network: a second layer for Bitcoin

Graph 11

Representation of the Lightning Network



A beta version is in use



Source: <https://bitcoinvisuals.com/lightning> (site accessed on 3 Jan 2019).

Proof-of-stake: gambling on or institutionalising the truth?

Above all, it is the economic considerations which highlight that the vital technical development for Bitcoin and related cryptocurrencies would be to do away with proof-of-work and embrace a different consensus model. While several alternatives have been put forward, conceptually the most important one is to replace proof-of-work by "proof-of-stake", a scheme whereby coordination on blockchain updates is achieved via staking claims to particular coin holdings.

There exist multiple proposed proof-of-stake protocols, see eg the Ouroboros protocol of the cryptocurrency Cardano (Kiayias et al (2017)), or the Casper protocol for Ethereum (see Buterin and Griffith (2017) and Zamfir et al (2018)). One simple implementation (of many possible ones) is such that, for each newly added block, a number of holders of the cryptocurrency are randomly selected to verify the block. But in order to do so, they must pledge some of their cryptocurrency holdings. Cheating is deterred by the threat of losing the pledged cryptocurrency in the case that one user's verified update differs from that of others.

There is one big caveat with the idea of replacing costly computations by an essentially resource-free betting game, sometimes termed the "nothing-at-stake" problem. This caveat concerns the lack of clear criteria for distinguishing between different blockchains with alternative payment histories (see eg Poelstra (2014, 2015)). With proof-of-work, the rule of following-the-longest-chain allows the winning blockchain to be selected based on a hard and externally verifiable criterion. With proof-of-stake, the absence of an actual cost means that users can secretly bet on alternative blockchain histories at no cost. And if alternative blockchain histories ever emerge, there is no hard criterion for choosing between them, thus requiring an overarching selection mechanism. These considerations have led to ample discussion in the cryptocurrency developer community regarding the feasibility of proof-of-stake (see eg the discussion in Muneeb et al's (2018) review of Zamfir et al (2018)).

The nothing-at-stake caveat may mean that successful proof-of-stake implementation might indeed rest on some degree of institutionalisation or reliance on social conventions (as indicated in Buterin (2014 a and b)).

Other proposed consensus algorithms, such as “delegated proof of stake” or “proof of importance”, aim to guarantee finality directly via additional social coordination mechanisms, eg a range of voting mechanisms by current coin holders.


Whether or not moving beyond pure-proof-of-work will require overarching coordination, it is noteworthy that much has already been done to protect Bitcoin and other cryptocurrencies beyond just applying the rule to follow the longest chain. For example, in March 2013, an erroneous software update caused the Bitcoin blockchain to fork into two branches. The latter was undone via coordinated action by large mining pools to ignore the rule to follow the longest chain and instead coordinate on the one that would reunite all Bitcoin users. A second example was the failed introduction of the so-called Segwit2X protocol update in November 2017, when a change of the Bitcoin protocol was implemented by the majority of miners but failed to convince other stakeholders. This led to a situation in which miners ultimately abandoned the longest chain (because nobody was willing to transact on it). A third instance occurred when the cryptocurrency Ethereum split over the undoing of the “DAO hack” in mid-2016. At the time, an application based on Ethereum protocol proved faulty (but not the cryptocurrency itself), which allowed a hacker to successfully steal ether tokens worth roughly USD 70 million at the time. Following heated discussion, on online forums and elsewhere, most miners and users decided to undo the hack by creating an alternative blockchain that would start just before the hack occurred. A minority of users, however, decided to stay true to the rule to follow the longest chain, thus giving rise to the Ethereum Classic cryptocurrency.

These three episodes, and others like them, show that social coordination has to be a key element of smoothly functioning cryptocurrencies. In the future, if novel consensus concepts such as proof-of-stake gain momentum, such social coordination may take a more central role, so that effective cryptocurrencies might ultimately require institutional backing of some form.

How readily institutional backing could improve the efficiency of cryptocurrencies can also be shown in the above analysis of payment security (in particular equation (7)). If it were possible to pre-commit to undo any double-spending attack, this would negate any incentives to attack bitcoin in the first place. More generally, the higher the probability that there will be coordination to undo any attack (as captured by Π^{HF} in equation (7)), the lower transaction fees and waiting times can be set in order to deter attacks. Graph 12 shows the quantitative impact of greater social coordination (higher Π^{HF}) on the trade-off between transaction costs and *waittime*.

Institutionalisation is key to the efficiency of cryptocurrencies

Graph 12



This graph shows the impact on the required waiting times in the case that social coordination is used to undo a double-spending attack. Calculations are based on equation (7) in the main text, assuming that block rewards are 0. The horizontal axis denotes the probability that the network of bitcoin users will coordinate and undo any double-spending attack (Π^{HF} in equation (7)). The vertical axis shows the resultant required waiting times for various levels of transaction fees.

Source: Author’s calculations.

Conclusion: towards a world of semi-decentralised exchange?

The overall conclusion from this paper is that, at least judging based on current technologies, in the digital age too, good money is likely to remain a social construct rather than a purely technological one: the efficiency of decentralised exchange via proof-of-work exclusively is much lower than would appear at first sight, and alternative technologies still need to demonstrate that they can function without institutional backing.

But claiming that technology alone cannot do the trick is not to say that it is useless. It simply means that the focus could shift away from the issue of whether the technology can replace traditional sovereign money and financial institutions.

One key question for future research is whether and how technology-supported distributed exchange can complement and improve upon existing monetary and financial infrastructure. For example, in mixed systems, normal market functioning could be guaranteed by decentralised economic consensus, yet should it fail there would be overarching coordination mechanisms that are also tied to the legal system. What would be the gains regarding efficiency, transparency, and resilience from such semi-decentralised exchange compared to current market designs?

Outside the world of cryptocurrencies aficionados, answering these and related questions will require a more widely distributed understanding of the new technology and how it might be used in existing markets. On the other hand, for those already involved in distributed ledger technology, what is needed is an awareness of how institutions have sustained trust throughout mankind's history, an issue that lies at heart of central banking and financial regulation (see, for example, Lewis (1969), Giannini (2011), Graeber (2011), Schnabel and Shin (2018), Bank for International Settlements (2018), Carstens (2018a), and Borio (2018)).

All this aside, the societal value of Nakamoto (2008) and his followers is substantial for reasons that extend well beyond the technology's use as means of exchange. Bitcoin's developers have created the backbone of a first-generation decentralised infrastructure that, over the past decade, has survived many attacks. In addition, Bitcoin has inspired an entire cohort to study the underlying technology, spotlighting fields as diverse as payments, cryptography, and database management. In the long run, the value of cryptocurrencies might be to catalyse our thinking on how society can handle access to data and the right to edit it, a much-needed impulse at a time characterised by loss of privacy and the rise of technology-driven disinformation campaigns.

References

- Abadi, J and M Brunnermeier (2018): "Blockchain economics", Princeton University, mimeo, May.
- Amihud, Y and A Cukierman (2018): "The macroeconomic perils of a world with a private digital currency and how to address them", VoxEU.org, 9 October.
- Andalfatto, D (2013) "Why gold and bitcoin make lousy money", Wednesday, 24 April, 2013, MacroMania Blog.
- (2017) "My perspective on the Bitcoin Project (collected works)", 21 December, 2017, MacroMania Blog.
- Athey, S, I Parashkevov, V Sarukkai, and J Xia (2016): "Bitcoin pricing, adoption, and usage: theory and evidence", *SIEPR Working Papers*, no 17-033.
- Muneeb, A, J Nelson and A Blankstein (2018): "Peer Review: CBC Casper", Medium.com, 6 Dec.
- Auer, R and S Claessens (2018): "Regulating cryptocurrencies: Assessing market reactions", *BIS Quarterly Review*, September.
- (2019): "Cryptocurrencies: why not (to) regulate?", *VoxEU EBook*, February.
- Bank for International Settlements (2018): *Annual Economic Report*, June.
- Bech, M and R Garratt (2017): "Central bank cryptocurrencies", *BIS Quarterly Review*, September 2017, pp 55–70.
- Berentsen, A and F Schär (2017): *Bitcoin, Blockchain und Kryptoassets*, Universität Basel.
- (2018): "A short introduction to the world of cryptocurrencies", *Federal Reserve Bank of St Louis Review*, vol 100, no 1.
- Biais, B, C Bisière, M Bouvard and C Casamatta (2017): "The blockchain folk theorem", *TSE Working Papers*, no 17-817.
- Biais, B, C Bisière, M Bouvard, C Casamatta and A Menkveld (2018): "Equilibrium bitcoin pricing", *TSE Working Papers*, no 18-973, December.
- Böhme, R, N Christin, B Edelman, and T Moore (2015): "Bitcoin: economics, technology, and governance", *Journal of Economic Perspectives*, vol 29, no 2, pp 213–38.
- Bolt, W and M van Oordt (2016): "On the value of virtual currencies", *Bank of Canada Staff Working Papers*, no 42, 2016.
- Borio (2018): "On money, debt, trust and central banking", keynote speech at 36th Annual Monetary Conference, Cato Institute, 15 November, Washington DC.
- Budish, E (2018): "The economic limits of bitcoin and the blockchain", *NBER Working Papers*, no 24717, June.
- Buterin, V (2014a): "Proof of stake: how I learned to love weak subjectivity", 25 November.
- (2014b): "On stake", ethereum.org, 5 July.
- (2016): "A proof of stake design philosophy", Medium.com, 30 December.
- (2017): "The meaning of decentralization", 6 February.
- Buterin, V and V Griffith (2017): "Casper the friendly finality gadget" (submitted on 25 Oct 2017 (v1), last revised 22 November 2018 (this version, v3)), arXiv:1710.09437.
- Carney, M (2018): "FSB Chair's letter to G20 finance ministers and central bank Governors", 13 March.

- Carstens, A (2018a): "Money in the digital age: what role for central banks?", lecture at the House of Finance, Goethe University, Frankfurt, 6 February.
- (2018b): "Central banks and cryptocurrencies: guarding trust in a digital age", remarks at Brookings Institution, Washington DC, 17 April.
- (2018c): "Technology is no substitute for trust", *Börsen-Zeitung*, 23 May.
- Catalini, C and J Gans (2017): "Some simple economics of the blockchain", *MIT Sloan Research Papers*, no 5191-16.
- Chaum, D (1983): "Blind signatures for untraceable payments", *Proceedings of the Springer-Verlag Crypto'82 conference*, vol 82, no 3, pp 199–203.
- Clayton, J (2018): "Chairman's testimony on virtual currencies: the roles of the SEC and CFTC".
- Chiu, Jonathan and Koepl, Thorsten, (2017), *The Economics of Cryptocurrencies - Bitcoin and Beyond*, No 1389, Working Papers, Queen's University, Department of Economics,
- Committee on Payments and Market Infrastructures (2012): "Payment, clearing and settlement systems in the CPSS countries – Volume 2", *CPMI Papers*, no 105, November.
- (2015): *Digital currencies*, November.
- (2017): *Statistics on payment, clearing and settlement systems in the CPMI countries*, December.
- Committee on Payments and Market Infrastructures and Markets Committee (2018): *Central bank digital currencies*, March.
- Committee on Payments and Market Infrastructures and International Organization of Securities Commissions (2012): *CPMI-IOSCO Principles for financial market infrastructures*, April.
- Committee on Payment and Settlement Systems (2003): *Payment and settlement systems in selected countries*, April.
- Dryja, T (2015): "Scalability of lightning with different BIPs and some back-of-the-envelope calculations", https://scalingbitcoin.org/hongkong2015/presentations/DAY2/1_layer2_2_dryja.pdf (accessed 19.11.2018).
- Dwork, C and M Naor (1992): "Pricing via processing or combatting junk mail", *Proceedings of the Annual International Cryptology Conference*, Springer.
- Easley, D., O'Hara, M., Basu, S., (2018). From mining to markets: the evolution of bitcoin transaction fees, Forthcoming, *Journal of Financial Economics*
- Fanusie, Y and T Robinson (2018): "Bitcoin laundering: an analysis of illicit flows into digital currency services", Center on Sanctions and Illicit Finance memorandum, January.
- Fernández-Villaverde, J and D Sanches (2016): "Can currency competition work?", *CEPR Discussion Papers*, no 11095.
- Financial Action Task Force (2015): *Guidance for a risk-based approach to virtual currencies*, June.
- Financial Stability Board (2018a): *Crypto-assets: report to the G20 on the work of the FSB and standard-setting bodies*, July.
- (2018b): *Crypto-asset markets Potential channels for future financial stability implications*, October.
- Foley, S, J Karlsen and T Putniņš (2018): "Sex, drugs, and bitcoin: how much illegal activity is financed through cryptocurrencies?", *Review of Financial Studies*, forthcoming.
- Lloyd, W (1833): *Two Lectures on Population*, unpublished manuscript.
- G20 Finance Ministers and Central Bank Governors (2018): *Buenos Aires Summit communiqué*, 19–20 March.

- Giannini, C (2011): *The age of central banks*, Edward Elgar.
- Graeber, D (2011): *Debt: The First 5,000 Years*, Melville House.
- Huberman, G, J Leshno and C Moelleme (2017): "Monopoly without a monopolist: an economic analysis of the Bitcoin payment system", *Columbia Business School Research Papers*, no 17-92.
- Kiayias, A, A Russell, B David and R Oliynykov (2017): "Ouroboros: a provably secure proof-of-stake blockchain protocol", EPrint Archive.
- Landau, J-P and A Genais (2018): *Les crypto-monnaies*, rapport au Ministre de l'Économie et des Finances, 4 July.
- Lewis, D (1969): *Convention: a philosophical study*, Blackwell Publishing.
- Lewis, J (2018): "The seven deadly paradoxes of cryptocurrency", *Bank Underground*, Bank of England, 13 November.
- Morris, S and H S Shin (2018): "Distributed ledger technology and large value payments: a global game approach", mimeo, Princeton University, November.
- Nakamoto, S (2008): "Bitcoin: a peer-to-peer electronic cash system", white paper, <https://bitcoin.org/bitcoin.pdf>.
- Pichler, P, A Schierlinger-Brandmayr and M Summer (2018): "Digital money", *Monetary Policy & the Economy*, Oesterreichische Nationalbank, Q3/18, pp 23–35.
- Poelstra, A (2014) "Distributed consensus from proof of stake is impossible", 2014, <https://download.wpsoftware.net/bitcoin/old-pos.pdf>, May
- (2015) "On Stake and Consensus" <https://download.wpsoftware.net/bitcoin/pos.pdf>, March
- Poon, J and T Dryja (2016): "The Bitcoin Lightning Network: Scalable off-chain instant payments", DRAFT Version 0.5.9.2, www.lightning.network (accessed 19.11.2018).
- Prat, J and W Benjamin (2017): "An equilibrium model of the market for bitcoin mining", Working Paper, 2017-15, Center for Research in Economics and Statistics.
- Schilling, L and H Uhlig (2018): "Some Simple Bitcoin Economics", *NBER Working Papers*, no 24483, National Bureau of Economic Research.
- Schnabel, I and H S Shin (2018): "Money and trust: lessons from the 1620s for money in the digital age", *BIS Working Papers*, no 698, February.
- Sompolinsky, Y, Y Lewenberg and A Zohar (2016): "SPECTRE: Serialization of proof-of-work events: confirming transactions via recursive elections", no 1159, Cryptology ePrint Archive, IACR.
- Sompolinsky, Y and A Zohar (2018): "PHANTOM, GHOSTDAG: Two Scalable BlockDAG protocols", no 1004, Cryptology ePrint Archive, IACR.
- Sompolinsky, Y and A Zohar (2005): "Bit Gold", Unenumerated blog, <http://unenumerated.blogspot.com/2005/12/bit-gold.html>, retrieved 25 October 2018.
- Zamfir, V, N Rush, A Asgaonkar and G Piliouras (2018): "Introducing the minimal CBC Casper family of consensus protocols", DRAFT v1.0, 5 November, Ethereum Research.

Appendix: glossary

Glossary		Table A1
Attacker advantage	Advantage for a double-spending attacker (see equation (7) in the main text) deriving from the fact that the attacker profits from a higher bitcoin income (block rewards, transaction fees, and double-spent amount) than does an honest miner (collecting only block rewards and fees).	
Attacker disadvantage	Counterforce advantages over an attacker: double-spending attackers likely have higher costs than regular miners do; the price of a cryptocurrency collapses following a double-spending attack; and social coordination might undo such an attack. Attacker disadvantage (see equation (7) in the main text) summarises these three considerations.	
Block reward	Newly minted bitcoins that increase the outstanding supply whenever a new block is added to the blockchain.	
Confirmation	Each additional block added to the blockchain after the block that contains the payment in question. If a transaction is included in block b and the blockchain currently includes $b+2$ blocks, the transaction is said to have three confirmations (also see waiting time).	
Consensus on the longest chain	Economic consensus algorithm in cryptocurrencies that are based on "proof-of-work." If conflicting versions of the blockchain are ever observed, the blockchain that is the most costly to forge (often the one with the most blocks, ie the longest one) is the one that the network coordinates on.	
Crypto-asset	A type of private digital asset that depends primarily on cryptography and distributed ledger or similar technology as part of their perceived or inherent value.	
Cryptocurrency	A crypto-asset used exclusively/primarily for payments.	
Cryptographic digital signature	A public/private key digital signature technology used to verify payment transactions. The digital signature verifiably proves that the payment has been authorised by whoever controls the cryptocurrency units that are being spent.	
Difficulty	The expected number of hashes that needs to be performed to obtain a hash result to find a valid proof-of-work.	
Double-spending	Strategy that consist of spending in one block and later undoing this by releasing a forged blockchain in which the transactions are erased. Requires access to enough computational power to overwhelm the rest of a cryptocurrency's network of miners.	
Economic payment finality	Definition of payment finality in blockchain transactions developed in this paper. A cryptocurrency payment can be considered as final once it is certain that, from a certain moment of time onwards, it will never be profitable to undo the payment via a double-spending attack (see also Table 1).	
Follow the longest chain	Rule that establishes that if competing versions are observed, the one which is the most expensive to forge continues to be used. This is generally the blockchain with the most blocks.	
Hash function	Function that takes a random text input and produces from this an output according to set rules.	
Lightning Network	Second-layer solution in which two parties A and B jointly lock in one bitcoin on the blockchain via a joint digital signature.	
Miner	Class of agents, who update the blockchain via computational work, and in return receive block rewards and transaction fees when they add batches of valid transactions to the blockchain.	
Proof-of-work	Mathematical evidence that a certain amount of computational work has been done, in turn calling for costly equipment and electricity use.	
Proof-of-stake	A system in which coordination on blockchain updates is enforced by ensuring that transaction verifiers pledge their coin holdings as guarantees that their payment confirmations are accurate.	
Protocol	The coded "laws" of a cryptocurrency. Set of rules that governs what constitutes a blockchain that is accepted by the network of users.	
Second-layer	A technology that aims to improve upon the economics of payment security by adding liquid methods of exchange on top of the Bitcoin blockchain.	

Target value	A proof-of-work is valid if the hash solves for a hash output below the target value. The lower the target value, the more difficult it is to find a valid proof-of-work. (see also difficulty, which is proportional to $1/\text{target}$.)
Tragedy of the common chain	Concept developed in this paper that users free ride on the security provided by the transaction fees of other transactions in the chain (see equations 9–11 in the main text). The proof-of-work and hence the level of security is determined at the level of the block in which a transaction is included, whereas the transaction fee is set by each user privately.
Transaction fee	Fee set by the paying party of a transaction. The fee is paid to the miner who includes the transaction in a block that is added to the blockchain.
Waiting time	One plus the time (in blocks) that merchants wait before assuming a payment is final and releasing the merchandise. A double-spending attacker thus has to forge a chain that is at least equal in length to this waiting time.

Sources: Bank for International Settlements (2018); Financial Stability Board (2018a,b); author's elaboration.

Previous volumes in this series

764 January 2019	Global Banking, Financial Spillovers, and Macroprudential Policy Coordination	Pierre-Richard Agénor and Luiz A. Pereira da Silva
763 January 2019	On money, debt, trust and central banking	Claudio Borio
762 December 2018	A key currency view of global imbalances	Hiro Ito and Robert N McCauley
761 December 2018	Non-Monetary News in Central Bank Communication	Anna Cieslak and Andreas Schrimpf
760 December 2018	Gross Capital Flows by Banks, Corporates and Sovereigns	Stefan Avdjiev, Bryan Hardy, Şebnem Kalemli-Özcan, Luis Servén
759 November 2018	Assessing inflation expectations anchoring for heterogeneous agents: analysts, businesses and trade unions	Ken Miyajima and James Yetman
758 November 2018	Foreign currency borrowing, balance sheet shocks and real outcomes	Bryan Hardy
757 November 2018	Explaining Monetary Spillovers: The Matrix Reloaded	Jonathan Kearns, Andreas Schrimpf and Fan Dora Xia
756 November 2018	Financial structure and income inequality	Michael Brei, Giovanni Ferri and Leonardo Gambacorta
755 November 2018	Measuring financial cycle time	Andrew Filardo, Marco Lombardi and Marek Raczko
754 November 2018	Euro area unconventional monetary policy and bank resilience	Fernando Avalos and Emmanuel C Mamatzakis
753 October 2018	Currency depreciation and emerging market corporate distress	Valentina Bruno and Hyun Song Shin
752 October 2018	The effects of prudential regulation, financial development and financial openness on economic growth	Pierre-Richard Agénor , Leonardo Gambacorta , Enisse Kharroubi and Luiz Awazu Pereira da Silva
751 October 2018	Exchange rates and prices: evidence from the 2015 Swiss franc appreciation	Raphael Auer, Ariel Burstein and Sarah M Lein
750 October 2018	Forward guidance and heterogeneous beliefs	Philippe Andrade, Gaetano Gaballo, Eric Mengus and Benoit Mojon

All volumes are available on our website www.bis.org.