

*The Microeconomics of Cryptocurrencies**

Hanna Halaburda[†]
NYU-Stern

Guillaume Haeringer[‡]
Baruch College

Joshua Gans[§]
University of Toronto

Neil Gandai[¶]

Tel Aviv University

First version: May 14, 2018

This version: December, 2020

Forthcoming Journal of Economic Literature

Abstract

Since its launch in 2009 much has been written about Bitcoin, cryptocurrencies and blockchains. While the discussions initially took place mostly on blogs and other popular media, we now are witnessing the emergence of a growing body of rigorous academic research on these topics. By the nature of the phenomenon analyzed, this research spans many academic disciplines including macroeconomics, law and economics and computer science. This survey focuses on the microeconomics of cryptocurrencies themselves. What drives their supply, demand, trading price and competition amongst them. This literature has been emerging over the past decade and the purpose of this paper is to summarize its main findings so as to establish a base upon which future research can be conducted.

*We thank the editor Steven Durlauf and six anonymous referees for suggestions and comments that greatly improved the paper. We are also grateful to Bruno Biais, Eric Budish, Michael Dickstein, Ben Fung, Rod Garratt, Jorge Cruz Lopez, Cyrus Minwalla, and Vijay Mohan for their comments and suggestions. Earlier version of this survey circulated under the title *Bitcoin and Blockchain: What we know and what questions are still open*.

[†]Email: hhalaburda@gmail.com; Web: www.halaburda.ca

[‡]Email: guillaume.haeringer@baruch.cuny.edu; Web: <https://sites.google.com/site/guillaumehaeringer/>

[§]Email: joshua.gans@utoronto.ca; Web: <https://joshuagans.com>

[¶]Email: gandal@tauex.tau.ac.il; Web: <https://neilgandal.com>

Contents

1	Introduction	4
2	Overview of the Bitcoin ecosystem	5
2.1	Motivation	5
2.2	Basic description	6
2.3	Processing transactions	7
2.4	Historical perspective	9
2.5	The meaning of the word “blockchain”	11
3	Consensus mechanism for a blockchain	12
3.1	The longest chain rule	13
3.2	The proof-of-work contest	17
3.3	Consensus as an equilibrium	20
3.4	Longest-chain attack: a formal model	23
3.4.1	Free entry equilibrium	23
3.4.2	Equilibrium with (potentially) dishonest miners	25
3.4.3	Flow vs. stock attack	29
3.5	Efficient design	30
3.6	Proof-of-Stake as an alternative consensus mechanism	31
3.6.1	The Nothing-at-Stake problem	31
3.6.2	Sustainability under proof-of-stake	32
3.7	Transaction fees	35
4	Bitcoin from the users’ side	38
4.1	Early adoption and usage	38
4.1.1	Early adoption in the U.S. and Canada	39
4.1.2	Early adoption by merchants	40
4.2	Different activities using Bitcoin	41
4.3	Recent data on use activity: where are we today	43
5	Price of Bitcoin	44
5.1	Theoretical analysis	45
5.2	Empirical papers on determinants of Bitcoin price	47

5.3	The Bitcoin bubbles	52
5.4	Price manipulation	53
5.5	Bitcoin as “digital” gold	54
6	Competition between cryptocurrencies	55
6.1	Bitcoin dominance of the market	55
6.2	Pump and dump schemes in cryptocurrencies	57
6.3	Other aspects of the ecosystem: cryptocurrency exchange markets	58
7	Conclusion	59

1 Introduction

Since the proposal for Bitcoin—a digital currency—was released in 2008 ([Nakamoto \(2008\)](#)), its robustness has been questioned. However, as of the time of this survey (end of December 2020) one bitcoin was worth around \$27,000. This is despite the fact that Bitcoin is not backed by any real asset nor any governmental claims (such as the ability to use it to settle tax debts). This persistence has generated increasing interest from economists. While macroeconomics look to it as a potential study of monetary theory, microeconomists have been interested in Bitcoin due to the seemingly robust nature of an otherwise highly decentralized network without any clear owner. The purpose of this survey is to examine this developing microeconomics literature.

We analyze Bitcoin (and cryptocurrencies at large) using the standard divisions of economics. For an asset, such as the Bitcoin digital tokens, to have economic value (i.e., to trade at a positive price), we need to examine both the supply and demand sides of a market for a cryptocurrency. We begin, however, in Section 2, with an overview of the Bitcoin ecosystem that underpins the market for the cryptocurrency.

We then turn to consider the supply and demand sides of the market in turn. Section 3 looks at how Bitcoin is supplied including the mining and transaction processing activities that make exchange possible. Our focus is on the decentralized nature of the blockchain that lays the foundation for Bitcoin and how participants in the ecosystem achieve consensus regarding what tokens are available and who the owners of those tokens are in any given period. Our focus is on the incentives of those participants to honestly support the network rather than disrupt it for some private gain.

Following this, in Section 4, we consider the demand-side of the market. What have economists learned about the reasons people use Bitcoin. In particular, is demand driven by pure speculation and, if it is used for real world transacting, what classes of transactions arise?

Section 5 then looks at the outcome of the interactions between supply and demand and how these impact the price of Bitcoin. We note that these markets have a highly volatile price that is likely driven by speculation. However, several studies show that because cryptocurrencies exchanges are largely unregulated, various forms of market manipulation have been attempted and these potentially account for some of the volatility beyond ‘normal’ fluctuations in demand and supply.

Finally, in Section 6, we examine competition between cryptocurrencies which is an evolving set of economic phenomena.

Since the introduction of Bitcoin, the literature on the blockchain and cryptocurrencies has branched out in many directions including initial coin offerings, smart contracts, governance, macroeconomic impacts, stable coins substitution, central bank digital currencies and, of course, an entire branch of computer science. Our focus is on none of those things but rather on the pure microeconomics of cryptocurrencies; namely, supply (Section 3), demand (Section 4), price (Section 5), and competition (Section 6).

2 Overview of the Bitcoin ecosystem

We offer in this section a short presentation of the Bitcoin ecosystem. Readers already familiar with it may want to skip it, and those who are interested in a more detailed description can refer to, for instance, [Velde \(2013\)](#), [Badev and Chen \(2015\)](#), [Böhme et al. \(2015\)](#), [Narayanan et al. \(2016\)](#), [Halaburda and Sarvary \(2016\)](#), [Andolfatto \(2018\)](#), or [Haeringer and Halaburda \(2018\)](#).

2.1 Motivation

Bitcoin was initially proposed in a white paper by [Nakamoto \(2008\)](#) and went into existence on January 3rd, 2009.¹ Nakamoto’s contribution was an answer to a long-standing question in the cryptography community (and to a lesser extent among libertarians): is it possible to design a fully decentralized digital currency? Having a decentralized cash system means that individuals may engage in “monetary transactions” without any third party involved (like cash, that can be given by a buyer directly to a seller) and without any authority that would for instance conduct a monetary policy. In the technological jargon, monetary transactions in such a system are called peer-to-peer.

Until the advent of Bitcoin, the problem did not have any obvious solution. By being electronic, the “coins” can in principle be easily be copied and thus used several times, that is, one faces the risk of the *double spending problem*. This problem could be theoretically avoided if, at any time, there is a *consensus* among all participants about which coin has

¹“Satoshi Nakamoto” is a pseudonym, and the true identity of the author —or authors— remains unknown.

been spent by whom. However, real-time consensus in peer-to-peer systems is known to be impossible according to Fischer et al.'s (1985) "FLP theorem," one of the most important theorems in computer science.² In a practical sense, this means that one needs what is called a "Byzantine fault tolerant system," i.e., a system that accounts for false messages and allows for temporary disagreements. The challenge is that Byzantine fault tolerant systems open the possibility of double spending attacks, for it is not possible to distinguish genuine Byzantine faults from double spending. Nakamoto's contribution is thus for many a breakthrough, explaining why it quickly generated a large enthusiasm and further development.

2.2 Basic description

Bitcoin is a digital cash ecosystem, with the *bitcoin* as unit account (with a lowercase "b"). Like many currencies, bitcoins can be used with fractional values, where the smallest fraction of a bitcoin is called the *satoshi*, which corresponds to one hundred millionth of a bitcoin. The three main components of the Bitcoin systems are the users, the miners and the blockchain.

In the Bitcoin terminology a user (i.e., any person holding and making use of bitcoins) does not have an account (or accounts) but a *wallet*, which is the combination of a *Bitcoin address* (i.e., "an account number", also called a *public key*) and a *private key* (i.e., a password). Wallets are relatively easy to create and the way the Bitcoin system defines an address implies that there is a potential for a nearly infinite number of wallets.³ This implies for instance that a user can create a new wallet for each transaction.

To send bitcoins to someone, a user needs his own private key to authenticate the transaction, that is, to prove that he is the owner of the Bitcoin address used to send the bitcoins.⁴ Of course, the user also needs the Bitcoin address of the recipient. When a user creates a transaction she also sets a voluntary transaction fee that will be paid once the transaction is processed and stored in the blockchain. Transactions are processed by the miners who

²The FLP theorem states in any asynchronous network where messages may be delayed (but not lost) there is no consensus algorithm if at least one node in the network may fail. A similar result is Brewer's (2000) "CAP" theorem for the case when messages can be lost.

³In the Bitcoin system there is room for 2^{160} different addresses. In comparison, the estimated number of grains of sand on Earth is "only" 2^{63} .

⁴To be precise, the authentication only serves to prove that the sender has the private key associated with the address; it does not prove the identity of the sender. The encryption tool used by Bitcoin is such that, at this day, it is impossible to deduce from a Bitcoin address the private key associated with it. But a security breach can occur when someone manages to steal one's address and its associated private key.

will record them into the blockchain. How this is done is the main part of [Nakamoto's](#) contribution.

2.3 Processing transactions

Transactions created by users are broadcast to the Bitcoin network via *Bitcoin nodes* and processed in *blocks* (i.e., batches) by *miners*. There are many miners in the network, each with a copy of the blockchain. A common misperception is that Bitcoin's holdings are stored in a (large) file called the *blockchain*. Bitcoin's blockchain does not contain bitcoins nor does it store the balance of each user, it is simply a ledger that records all *transactions* that have ever been made with bitcoins, and takes the form of a concatenation of blocks of validated transactions —hence its name, a *chain of blocks* (of transactions). Note that in Nakamoto's original paper, the word *blockchain* never appears, it was created later.⁵

Under Bitcoin no special authorization is needed to become a miner. In other words, there is “free entry” into Bitcoin mining. In the terminology of the literature, the Bitcoin blockchain is a *permissionless* database. To process a transaction a miner checks whether the transaction has been signed with the private key associated to the Bitcoin address of the sender. A miner also checks whether the sender has sufficient funds. This is done by parsing the blockchain and looking at all incoming and outgoing transactions corresponding to the senders' address. All those operations are relatively easy to handle because Bitcoin's blockchain is *public*, i.e., any person with access to the blockchain can parse it and observe, analyze and scrutinize transactions and, technologically speaking, Bitcoin's blockchain is a rather primitive type of database that is easy to read.

Each miner selects which transactions to process and packs them together to create a block. The block will also contain one additional transaction, called the *coinbase transaction*, that consists of the sum of all the transaction fees (associated to the selected transactions) and a *block reward*, which are newly created bitcoins. That is, processing transactions also tackles the problem of the creation of bitcoins. The amount of new bitcoins created in a block is set by the Bitcoin protocol. When Bitcoin started in 2009 the block reward was 50 bitcoins and, according to the protocol, the reward is halved approximately every 4 years

⁵Nakamoto simply referred to *chains*. The earliest known source using the phrase “block chain” (two words) is Hal Finney in a mailing list (<https://www.metzdowd.com/pipermail/cryptography/2008-November/014827.html>).

(more precisely, every 210,000 blocks).⁶ Around year 2140 the reward will drop to 0. After this date no more bitcoins will be created and the only source of income for the miners will be the transaction fees. So the total number of bitcoins that will be ever created will be slightly less than 21 million (should Bitcoin still exist in 2140).

Obviously, the perspective of a hefty reward from the creation of a block triggers a competition between miners: all miners seek to be the one adding the next block. Consensus will thus be facilitated if with very high probability, for any sufficiently large time window, only one miner wins this competition.⁷ We will see in Section 3 how consensus can be restored if several miners win “at the same time.” Mining is thus akin to a contest, or an “all-pay-auction.”⁸ The prize of the competition consists of the block reward, and the participation cost comes from a standard cryptographic tool called *hashing*.⁹ A secure cryptographic hashing function, a standard tool in computing, roughly consists of creating, for any input (e.g., a file or some data) a string of characters and digits that can uniquely identify the file (or the data).¹⁰ A hashing function is a *one way function*: it is easy to compute the hash of some data or file but it is impossible to reconstruct the input from the hash.¹¹ Also, there is no way to predict how the hash will look like before calculating it. That is, a hash looks like a random sequence of letters and digits.¹² So for a miners’ perspective the probability that a hash of a block starts with, say, the letter “a”, in hexadecimal system is 1/16.

Any block created by a miner must contain, besides the transactions (including the coinbase transaction), the hash of the last block in the blockchain and a number, called the *nonce*. Including the hash of the last block implies that blocks in the blockchain interlock with each other like pieces in jigsaw, linking them into the blockchain. Also, in case there are

⁶The last halving (from 12.5 to 6.25 bitcoins) occurred on May 11, 2020 (see <https://en.wikipedia.org/wiki/Bitcoin>).

⁷The time window just needs to be sufficient to account for possible network latencies and let the information about a new block spread throughout the Bitcoin network.

⁸Later we will see that this is the only mechanism that achieves the three key properties of a permissionless network.

⁹In addition to block awards, miners receive transaction fees. We discuss these in Section 3.

¹⁰There exist several hashing functions. The one used by Bitcoin is called SHA-256; it produces strings of 64 characters/digits.

¹¹In simplified terms, hashing is to computer files what the fingerprint is to humans: it is very easy to check if a fingerprint is someone’s fingerprint (we just have to compare the two fingerprints), but it is impossible to “recreate” a person from just observing her fingerprints.

¹²Hashing is a deterministic operator. But the complexity of the hashing function is such that it has the appearance of a random sequence.

several versions of the blockchain it helps miners identify to which version a newly broadcast block refers. The nonce is key in the Bitcoin protocol, and although it is chosen by the miner, not all nonces are acceptable. The Bitcoin protocol specifies that the hash of the block (that contains the nonce) must start with a certain number of zeros, called the *difficulty* in the Bitcoin system.¹³ Because of the unpredictability of the hash function the only possibility for a miner is to find by trial and error a nonce that will produce a hash that starts with the right number of zeros. In other words, hashing, which is a necessary step in the Bitcoin protocol, is a computing intensive task. Processing transactions with intensive computations is called *proof-of-work* (PoW). The number of zeros required for a hash is calculated every 2016 blocks so that on average it takes ten minutes before a miner finds the solution of the blocks she is processing (hence the difficulty is adjusted about every two weeks). This periodic adjustment implies that the blockchain grows at a steady pace, independently of the number of miners and their computing power (also called the *hash power*). Once a miner has found a nonce that gives an acceptable hash output, she will broadcast the new block to the network, expecting the other miners to add it to the blockchain.

2.4 Historical perspective

Bitcoin and blockchain are typically associated with the innovations of Nakamoto (2008). The history, however, begins much earlier.

The earliest source related to the notion of blockchain is [Haber and Stornetta \(1990\)](#), who addressed the question of how to timestamp a digital document and have an historical record of those timestamps. More precisely, they propose a protocol such that it would be easy to spot if the record has been tampered, even if the timestamping authority was doing the tampering. Note that this is very much in line with what Bitcoin does in the sense that transaction records can be altered, but any such alternation would be immediately noticed.¹⁴ What makes [Haber and Stornetta \(1990\)](#)'s "blockchain" tamper-evident is that hashes of the dataset are regularly publicly posted. Each new hash is the hash of the previously published hash and the hashes of all the documents that have just been timestamped. Linking the

¹³To be precise, the hashing function of Bitcoin is double SHA-256, which consists of calculating the hash of the hash. Note that the "correct" nonce is not unique, there might be several nonces that produce the desired hash.

¹⁴Many people conclude that the record is thus *immutable*. This is misleading. Electronic records can always be modified. What is important is that it is difficult to execute such modification undetected.

published hashes together allows to preserve the chronological order of the timestamps, which makes it impossible to backdate the issued timestamps. [Haber and Stornetta \(1990\)](#) then went on to demonstrate the practicality of their solution by publishing a hash of their ledger each week in the New York Times.¹⁵

While the concept of linking data in a chain using hashes ensures that the integrity of the data can be easily checked, it does not eliminate the incentives to make changes. This is where proof-of-work is helpful. This concept was initially proposed by [Dwork and Naor \(1992\)](#) as a way to prevent mail spamming. The driving force here is that while calculating one hash is trivial, calculating a large number of them requires substantial resources like computational power and electricity. In Bitcoin this is applied by setting a target value (the difficulty) and requiring that the block's hash is below this target (the hash starts with a number of zero at least as high as specified by the difficulty). Miners can find a hash fulfilling this requirement only if they try a lot of hashes (in expectation), which is costly. It is easy to see whether the block hash is below the target, as the target starts with a certain number of zeros. Changing an entry in a block changes the block hash, which makes the tempering evident. Therefore, to make an undetected change, the attacker needs to redo the work required to find a block hash below the target.

While these two concepts, [Haber and Stornetta \(1990\)](#)'s chain and [Dwork and Naor \(1992\)](#)'s PoW are important for the Bitcoin system, they are not sufficient. The reason is that both [Haber and Stornetta \(1990\)](#), and [Dwork and Naor \(1992\)](#), are not considering distributed databases (or ledgers), where there are several agents, or nodes, maintaining the data, as it is the case with Bitcoin. That is, a *consensus mechanism* is needed. Indeed, neither the hash-linked dataset nor the PoW concepts do say how to proceed when there are two or more conflicting versions of the blockchain, neither of them showing evidence of tampering. Using the “longest chain rule” as a consensus mechanism (which we will discuss at length in the next section) is perhaps the core of Nakamoto's contribution. In fact, in the computer science literature the use of PoW together with the longest chain rule is often called “Nakamoto's consensus” ([Bonneau et al., 2015](#)).

The other elements, like encryption with the concept of public and private keys were already well known, too. In the end, with a bit of scrutiny one realizes that Nakamoto's con-

¹⁵The [Haber and Stornetta \(1990\)](#) “blockchain” has been operating for almost three decades, and still operates to this day! The hash is published in the printed Sunday edition of the New York Times in the classifieds section.

tribution essentially consisted of putting together different concepts that turned out to complement each other (PoW, time-stamping, encryption). The second key insight of Nakamoto is in adding incentives for the miners. Potential rewards attract miners, which in turn increases the difficulty of the hashing problem (and may also give additional incentives to have a unique version of the blockchain), which entails in increasing the security of the system. These are key in Bitcoin’s system, and it is sometimes ironic to see that recent developments around blockchain technology pay little attention to these aspects.

Since the inception of Bitcoin, the development of blockchains has moved beyond time stamping of documents and recording of transactions and now includes broader applications of verification, including the provision of automated contracts and other decision-mechanisms. In addition, alternative designs for achieving consensus on digital ledgers, most notably, proof-of-stake (PoS), have been developed. These developments have been aimed at improving the speed, scale, resource use and complexity that can be achieved by blockchain technologies.

2.5 The meaning of the word “blockchain”

There is no one agreed-upon definition of “blockchain.” Some of the confusion regarding blockchain technologies can be traced to the origin of the term itself. As described earlier, the term “blockchain” originally literally referred to the “chain of blocks” of transactions in Bitcoin. However, usage evolved and a blockchain was used broadly for distributed ledgers regardless of whether they were used for cryptocurrencies, digital assets or something else. Thus, today, the word blockchain is often used to refer to distributed ledger technologies.¹⁶ Importantly, the use of encryption, smart contracts and even digital money itself do not require a distributed ledger (i.e., a blockchain) at all (Halaburda, 2018).

The distinctive feature of the Bitcoin system is that the distributed ledger operates and exists without any trusted parties, let alone some party who is not directly involved in transactions. It is surprising that despite the lack of trusted third parties, Bitcoin system has been relatively secure. Put simply, no one has been able to successfully hijack or rewrite entries on the Bitcoin blockchain. In this sense, it has been practically immutable. Bitcoin’s blockchain is also public (with the ledger visible to all) and permissionless (with any computer

¹⁶Computer Scientists might define a blockchain as “a database organized as a Merkle tree, updated via a strategy-proof communal consensus protocol/game.” We thank one of the referees for this definition.

permitted to validate transactions and update the ledger).

Thus, Bitcoin’s blockchain, which has demonstrated its robustness and security, has the features of being distributed, public and permissionless. But this does not mean that it must operate without a trusted third party. Also, those properties alone do not ensure that a blockchain is robust and secure. That is, it is not simply these technological choices that have allowed Bitcoin’s blockchain to persist, but the incentives that are embedded in the protocol that, in many respects, account for potential variation in human behavior. Thus, while Bitcoin uses a combination of technologies from cryptography, hashing and other developed parts in computer science, it has been too costly for participants to, despite the opportunity to do so, alter the Bitcoin blockchain and repurpose it for their own ends. The next section examines why this has been so.

3 Consensus mechanism for a blockchain

At its heart, Bitcoin and other cryptocurrencies are a register of digital assets and their assignment to an owner.¹⁷ A key feature of the blockchain approach is that changes in the register —either the addition of new digital assets/tokens and/or their re-assignment to different owners— are communicated and then, via a process, recorded so that they can be reliably set as the ‘state’ of the system. But whether its ‘state’ is used relies on consensus. That is, participants in the ecosystem ‘agree’ on the current state of the system and changes are communicated and made with respect to that state.

This has several implications of economic interest.¹⁸ First, there are questions of participation and incentives. If the state of the system is achieved by consensus, how is that consensus arrived at? Which set of agents must achieve consensus and what are their incentives? If there is disagreement (usually termed a ‘fork’) regarding the state of the system, how is that disagreement resolved? Second, there are questions of robustness in terms of ensuring that the digital assets that should perform like currency are not able to be replicated allowing agents to increase their currency holdings unilaterally.

This section examines the underlying way that a blockchain-based protocol can achieve consensus. While early interest focused specifically on Bitcoin, our discussion applies to

¹⁷The register aspect of Bitcoin has been likened to the monetary construct from [Kocherlakota \(1998\)](#).

¹⁸We focus here on questions related to economic forces. For an excellent overview of Bitcoin stability problems from computer scientific perspective, see [Bonneau et al. \(2015\)](#).

cryptocurrencies in general. Note that, as we mentioned in the previous section, consensus does not imply immutability of the record. Consensus, together with proof-of-work only ensure that it is more difficult to tamper the data in an undetected way.¹⁹

3.1 The longest chain rule

An important element in achieving consensus on a blockchain relates to what happens if there is disagreement. In Bitcoin, new blocks arise when miners collect messages of new transactions into a block and then compete to propose the appending of their block to the blockchain. Crucially, any miner who wins the competition to propose a new block can, subject to constraints of cryptography, freely propose transactions whether they received a message of them or not. As we will describe in more detail below, this opens up the possibility that the transactions proposed may not be legitimate in some respect. Thus, it is entirely possible that, in the absence of consensus, two or more alternative blockchains may be proposed and built upon. In other words, while those blockchains may ‘agree’ up to a certain point in time, there may be disagreement thereafter. Because this would appear as multiple branches from the last point of agreement, this situation is given the term **forks**. Consensus will only be achieved if one of the forks is adopted by the ecosystem and the others abandoned.

Forks are an expected occurrence in Bitcoin’s blockchain, due to its peer-to-peer nature and network latency. As shown by the FLP theorem, real-time consensus cannot be guaranteed. Because of the inevitable network latency, it is possible that two different miners find and broadcast, at approximately the same time, a correct nonce for the block they are processing, and that these two blocks do not reach all miners in the same order. Miners would, thus, accept whichever first block they receive and reject the other block. As a result we obtain a fork, i.e., two competing versions, called *branches*, of the blockchain. Such accidental forks happen frequently on Bitcoin due to the distributed nature of its mining.

How is consensus, that is, having all miners ultimately working on the same branch, achieved? Nakamoto’s solution for this problem is to have miners follow the **longest chain**

¹⁹One of the broader features of blockchain technologies is that they allow for verification of digital and potentially other transactions. Such verification is a key part of how cryptocurrencies are enabled by this technology by verifying that tokens are assignment to a single owner. However, we focus here on the cryptocurrency-enabling properties of blockchain only. For more on other potential uses see [Catalini and Gans \(2020\)](#).

rule (LCR).²⁰ It is important to note, however, that this is merely a recommendation and, although it is the selection rule that is made by default in the mining software, miners are free to override it and choose an alternate strategy. If miners (or a sufficiently high number thereof) follow the LCR, consensus can be restored relatively quickly because, since the pace at which blocks are added to the blockchain is stochastic, sooner or later one of the versions will be longer than the other versions. Once the consensus is restored the other branches become *orphaned*).²¹

Forks can arise by accident but also potentially based on the actions of one or more of the miners. In that case, the fork is termed an “attack.” There is, often, no way to distinguish an accidental from a deliberate fork which makes it hard to coordinate and mitigate such attacks. The most natural attack is the double spending attack. In this case, the attacker, after receiving the good or service (that is “outside of the blockchain”), forks the blockchain by proposing another branch that does not contain his payment (in bitcoins on the blockchain). He also then empties his wallet so that the original transaction can no longer be processed. Since the attacker does not pay in this scenario, double spending would thus occur if the community adopted this alternative branch. Suffice it to say, this undermines the efficacy of Bitcoin as a currency as one person literally has scope to print money.

Given that attacks cannot be detected and prevented directly, the robustness of the

²⁰To be more precise, the explicit purpose of the longest chain rule in Nakamoto (2008)’s white paper is to select the chain with the most computational power spent. It is possible that the two branches of the fork may have a different difficulty. Therefore, it is more precise to refer to the *heaviest* chain. Since typically two forks have the same difficulty level, this boils down to looking at the longest chain.

²¹The idea that miners might engage in a coordinated fork, outside of protocol rules has been noted by Arruñada and Garicano (2018). They argue that “within protocol” governance is unlikely and, at the very least, soft power and relationships will guide the evolution of blockchain networks. In the end, there are costs of achieving consensus and coordinating on a chain, especially one whose ledger is “correct.” Abadi and Brunnermeier (2018) model these costs and show that the costs of operating a blockchain that is both decentralized and correct are necessarily higher in order for nodes to be appropriately incentivized. Attempts to reduce those costs necessarily make the blockchain vulnerable to misreporting while decentralization is not compatible with reducing costs and ensuring correct reporting. In other words, there is a “trilemma” in achieving a blockchain that is correct, decentralized and low cost. Only two out of three of these goals can be achieved. Suffice it to say, these costs have already emerged in decentralized blockchains and it remains an open question whether a blockchain that is more efficient and operates at the scale of networks such as Visa and Mastercard are possible.

system depends on them not arising in equilibrium. Nakamoto argued that under the LCR such double spending attacks are prohibitive, and only possible if a miner is in control of at least 50% of the total hash power of the network. In such an attack the attacker privately operates their own blockchain for a period of time as the sole miner. That private blockchain will grow more slowly (in expectation) if that miner has less computing power than is being applied to the main blockchain (i.e., less than 50% of total computing power). As we will see later, researchers moving beyond Nakamoto’s reasoning discovered that such an attack can be successful with a smaller fraction of computing power controlled by the attacker. Controlling more than 50% guarantees the success but attacks are possible with a lower share of computing power. It is important to make the distinction between the so-called *majority attack* (or 51% attack) and the *longest-chain attack*. This latter simply consists of attempting to create a longer chain which will replace the ”honest” chain. With a majority attack a miner can conduct a longest-chain attack but also, for instance, include invalid blocks in the blockchain.²²

In economics terms, the possibility of an attack relates to market structure amongst the miners. The Bitcoin protocol is ‘incentive compatible’ (in the sense, that proposed blocks will represent true and legitimate transactions) if miners are not concentrated—that is, if there is sufficient competition. A more competitive market structure is assisted both by, as we will see, free entry into mining as well as the potential lure of the size of the market opportunity as represented by the price of Bitcoin itself. Below we will review the literature that demonstrates when this intuition holds for the underlying game as described by the Bitcoin protocol (including the embedded LCR).

Before exploring the game theoretic aspect of the consensus mechanism it is worth mentioning that the “coordination problem” is broader than just coordinating on a specific branch. If all miners follow the recommendations described in Nakamoto’s white paper, then the protocol will deliver what its designer claims: a consistent, fully decentralized ledger. The inevitable presence of strategic miners may thus jeopardize Bitcoin’s ecosystem

²²Citing the Bitcoin Wiki (https://en.bitcoin.it/wiki/Weaknesses#Attacker_has_a_lot_of_computing_power), Budish (2018) lists what an attacker can do (besides double spending): prevent some or all transactions to be confirmed (i.e., processed and stored in the blockchain) and prevent some or all other miners from mining. However, an attacker cannot change other people’s transactions without their cooperation (i.e., without their private key), prevent transactions from being broadcast to the network, or change the coinbase formula.

if miners or agents do not use Bitcoin as prescribed.

But Bitcoin may also fail for another reason. Users need to agree on the rules governing Bitcoin. For instance, bitcoins awarded to the miner who added a block to the blockchain must be recognized as legitimate by all participants. Disagreements about the rules can cause disruption, as it occurred with the hard fork of August 1, 2017.²³

Kroll et al. (2013) note that Bitcoin's success relies on three types of consensus: (1) consensus about rules, (2) consensus about the state (i.e., there is a unique ledger), and (3) consensus that bitcoins are valuable. These consensus elements are related to each other. The miners' source of income are the rewards and fees they obtain when adding a block, which are included in that block. If the blockchain forks, the rewards that are included in a branch are not recognized in the other branches. For those other branches, such rewards do not exist. The value attached to a bitcoin in a block, insofar as its owner plans to eventually spend it, then crucially depends on whether that block is recognized by other users. Consensus about bitcoins' "value" thus depends on the consensus about the state of the blockchain.

Here we will focus on the consensus mechanisms, their economic operation, their costs and their sustainability. The remainder of the section proceeds as follows. We first introduce and discuss the main elements of a blockchain design and then examine equilibrium with honest actors in a setting where there is free entry into the network. After that, we examine equilibrium with potentially dishonest actors. This section presents a condition that ensures that actors will be honest in equilibrium, that is, the condition insures incentive compatibility. Then, we derive the efficient design that satisfies both the free entry and incentive compatible constraints. The analysis first employs the proof-of-work (PoW) consensus mechanism. We then expand the analysis to a proof-of-stake (PoS) consensus mechanism, an alternative to PoW. Finally, we review important theoretical models that examine various aspects of the blockchain.

²³A group of miners proposed to update Bitcoin's design so as to allow for a higher throughput (i.e., being able to process more transactions per block). Not all miners followed the proposal, entailing in the creation of Bitcoin Cash. Its blockchain is the same as Bitcoin's for all transactions until August 1, 2017. After that the two blockchains diverge.

3.2 The proof-of-work contest

In the Bitcoin network, whenever two parties complete a transaction (via the matching of cryptographic public and private keys), that event is broadcast to the network. Anybody can receive those messages but to provide an incentive for at least someone to collect, record, check (that is, verify that sender of a token indeed owns that token), store and provide public access to transactions, a reward is provided in the form of newly minted tokens (as well a transaction fee offered by users). Which miner receives that reward and whose block ends up being appended to the blockchain is the result of a computational contest the basis for which we described in Section 2.3. We focus here on the reward itself and will return to specific issues associated with transaction fees below (Section 3.7).

The main challenge in the design of a blockchain network is to deal with the absence of trusted participants; in particular, any miner can tamper with the blockchain and it is not possible to identify such miners. Thus, in order for the blockchain to prevent such manipulation, the rewards offered to miners must be such that they prefer to propose legitimate rather than illegitimate blocks.²⁴ Thus, the reward to miners, which we will denote by R , has to ensure that the incentives of “bad” actors are muted.

The mechanism that is used is one of “random selection.”²⁵ In this mechanism, each participating miner is given a chance of becoming the miner that proposes a block and receives a reward for so doing. While, in principle, the reward might be anything, typically it is in the form of tokens from the network (e.g., for Bitcoin, it is a specified number of bitcoins); that is, if θ is the number of tokens awarded and e is the dollar to token exchange rate, then $R = e\theta$.²⁶ If p_i is the probability of miner i receiving the reward and c_i is miner i 's cost, then the net expected payoff to i is $p_i e\theta - c_i$. If this is positive, miner i will operate, otherwise it will not. Thus, $p_i e\theta - c_i \geq 0$ is i 's *participation constraint*.

To complete the specification of the game, we need to state who can be a miner. Bitcoin being *permissionless*, anybody can become a miner and apply computational (or hashing) power to participate in on-going contests. The same is true for other major networks such as Ethereum. At the other end of the spectrum, there are *permissioned* networks, such as

²⁴The consensus that arises is also a function of other factors including errors in the ways that messages are sent and received across the network.

²⁵Random selection is the outcome of the race to be the first to complete the assigned computational task. See Ma et al. (2018) for details.

²⁶We do not transaction fees in the formal model. It is straightforward to include them.

the one proposed by the Libra Association, whereby miners will be appointed and vetted by the association itself and assumed to be trustworthy.²⁷

What drives the outcome of the game is the selection rule that determines the probability p_i that a miner i wins the contest. In permissionless networks the identity of a miner (or the number of tokens in possession for each miner) does not determine whether a miner is selected to process a block. [Leshno and Strack \(2020\)](#) take a mechanism design inspired approach and identify several properties (or axioms) that the selection probability must satisfy that respects the notion that the network is permissionless and also the capabilities of potential coalitions that miners might be able to form.

- **Anonymity:** if any two miners can change their identities they inherit the selection probability of one another. Importantly, it does not allow the protocol to condition on the history of the miner’s behavior.
- **Robustness to Sybil Attacks:** a miner cannot split its performance into two or more entities and pose as a new entrant to increase his selection probability.²⁸ Incidentally, this condition ensures that free entry is possible and insiders cannot undertake certain actions that prevent others from entering.
- **Robust to Merging:** miners cannot increase their selection probability by merging. In other words, a permissionless network must forestall any incumbent advantages to effectively ensure that anyone can participate in the network on equal terms.

Denote by x_i the amount of work performed by miner i . [Leshno and Strack \(2020\)](#) show that the only selection mechanism that satisfies the three above properties is the proportional rule,

$$p_i = \frac{x_i}{\sum_{j=1}^N x_j}, \quad \text{for all } i = 1, \dots, N. \quad (1)$$

Hence, PoW is akin to an all-pay auction, or a Tullock contest ([Buchanan et al., 1980](#)). This contest formulation is common to most (if not all) papers analyzing Bitcoin’s mining game. What [Leshno and Strack](#) do is to provide a foundation for these.

²⁷See <https://libra.org/en-US/>.

²⁸In a Sybil attack, an attacker creates a large number of different identities and uses this to gain a disproportionately large influence. It is named after the subject of the book entitled “Sybil.” The book is a case study of a woman diagnosed with multiple personality disorder.

If the cost of computation is a function of the amount of work, $c_i = c_i(x_i)$, then in equilibrium a miner i will choose x_i^* that satisfies

$$e\theta \sum_{j \neq i} x_j = c'_i(x_i^*) \left(\sum_{j=1}^N x_j \right)^2. \quad (2)$$

Note that the above model is deliberately general and does not aim at modelling a specific cryptocurrency. The results still hold when considering richer approaches. For instance, [Ma et al. \(2018\)](#) go deeper by considering the specific case of Bitcoin (or other cryptocurrencies) that base miners' work on the hashing function. Doing this allows them to view the problem from a dynamic perspective. Recall that the only way for miners to find a valid nonce is by trial and error. Since the time it takes for this task follows a Poisson process the mining game becomes a memoryless stochastic game, which is the basis of many patent race models. Denoting by K the difficulty (the minimal number of zeros the hash should have), the expected time it will take for miner i to find a valid nonce to be distributed $\Gamma(K, x_i)$. [Ma et al. \(2018\)](#) show that we find the same equilibrium properties as in the patent race literature. While the resemblance with patent races may seem unsurprising, there is a key difference (and complication) here. For the case of Bitcoin the difficulty, K , is endogenous to the protocol. As miners increase the level of effort (e.g., investing in more powerful mining hardware), or as more miners enter the game, the hashing power of the network increases, which entails in a higher difficulty level so as to maintain a constant pace of a new block every 10 minutes. Like for [Leshno and Strack \(2020\)](#), [Ma et al. \(2018\)](#) show that the mining contest involves proportional selection. In other words, both [Leshno and Strack \(2020\)](#) and [Ma et al. \(2018\)](#) vindicate one of Nakamoto's original goal of "one CPU, one vote."²⁹

In summary, there are several design decisions that will determine miners' payoffs in a proof-of-work network. The first design decision involves setting permission rights to become a miner. The second design decision is the selection mechanism (which in turn may affect

²⁹Note that in the analysis we discussed thus far, the miners' reward, R , is assumed to be constant. This is not the case in reality. First, the exchange rate e is determined by broad market forces. The ensuing volatility is likely to affect miners' decisions, especially if Bitcoin's exchange rate is hard to predict. Second, recall that for Bitcoin (and for most cryptocurrencies) the reward is made of newly created coins and transaction fees. While the number of new coins is deterministic, transaction fees are not. As of today, the new coins usually account for between 97% and 99% of the reward in Bitcoin. But this percentage is relatively volatile. In December 2017, transaction fees reached (a peak of) 30.40% of the total award. Data on transaction fees for each block can easily be obtained on <https://btc.com/stats/fee>.

the amount of work exerted by the miners). However, as we have just seen, the form of the selection mechanism is constrained in permissionless networks. A third, and perhaps a bit less intuitive, design aspect is determining miners' cost structure. This can be done by changing the nature of the computational problem miners face when creating blocks. For example, Bitcoin miners now need access to specialized ASIC chips.³⁰ In contrast, the Ethereum's hashing algorithm was specially designed to be not transposable to an ASIC chip.³¹ ASIC mining chips can have substantial effects on the economics of a cryptocurrency. By bringing economies of scale they spur the growth of large miners, thereby increasing the centralization of the network. The last piece of the design is the block reward, θ . We will next see how these choices interplay with each other when determining the overall efficiency of blockchain networks.

3.3 Consensus as an equilibrium

The LCR is embedded in the Bitcoin protocol but there is nothing to prevent that protocol being changed as a result of a fork on the blockchain. Thus, the natural initial research question was whether abiding by the LCR was likely to be an equilibrium outcome. Recall that Nakamoto's intuition was that as long as miners followed the longest chain rule double spending attacks were prevented provided that no miner has a majority of the computing power. The early literature focused on exploring that intuition using formal game theoretic models. The answer from the literature offered mixed support. While Nakamoto's outcome could be confirmed as an equilibrium, there were other possible equilibria suggesting some fragility to the Bitcoin protocol.

The first paper to examine whether LCR was an equilibrium was [Kroll et al. \(2013\)](#). They started from the case where a fork had already occurred and examined what they called *monotonic* strategies whereby miners built on the most recent block of multiple branches, one of which was the longest one. They found that the best response for an individual miner was to choose the branch that others chose. This included the longest branch and so that could be a Nash equilibrium outcome. (See also [Barrera and Hurder \(2018\)](#)).

³⁰Application-Specific Integrated Circuit. Such chips are highly specialized for a very specific type of task. Computers built with such chips do not perform very well when used for other tasks than the one the chip was designed for.

³¹In the recent years manufacturers have been able to design ASIC chips for Ethereum. As a response the Ethereum Foundation is planning to move to Proof-of-Stake consensus mechanism.

As this was only a partial answer, there was room for a more general treatment. This was provided by [Biais et al. \(2019\)](#) with an infinite horizon dynamic game of miner behavior. The coordination game properties found by [Kroll et al. \(2013\)](#) emerged in that more general treatment and they found that the LCR was a more robust Markov perfect equilibrium. Multiple equilibria remain, however, including equilibria where forks arise as the result of deliberate action. Such forks can also be sustained indefinitely. This provides a foundation for the analysis of forks of the Bitcoin blockchain such as the one that occurred on August 1, 2017 and lead to the creation of a cryptocurrency competing with Bitcoin, Bitcoin Cash. [Biais et al. \(2019\)](#)'s result is particularly compelling because they not only rationalize the coexistence of multiple forks but also rationalize the split, i.e., the *creation and sustained existence* of the competing branches. In their model, forks can also delay the achievement of consensus and reduce the flexibility in upgrading the system and resolving problems.

While this literature demonstrated that, in principle, blockchains could fork easily if miners were all coordinated in their action (by design or happenstance), a related question was how much computing power did a miner need to unilaterally cause a deliberate fork? This issue was examined by [Kiayias et al. \(2016\)](#). They found that so long as no miner had more than 36% of the computing power, the LCR is a Nash equilibrium. However, for any miner with more than 46% of the computing power, forking is a profitable deviation. In other words, such a miner will always ignore the blocks that have been just mined by the other miners.

The intuition behind this result is the following. In Nakamoto's original white paper the threshold is obviously 50% because one only looks at whether a miner can create a branch that is longer than the main, "legitimate" branch. In contrast, [Kiayias et al. \(2016\)](#) look at the *incentives* to fork. Forking introduces a trade-off: a miner creating a new branch is potentially the only one mining on it, and thus will win the contest game with probability one for all the blocks in the alternative branch, as long as it is not the longest. In other words, the trade-off when forking has a low probability of success (the branch becomes the longest) with a high expected reward (as long as the new branch is not the longest all the block rewards accrue to the miner) versus a certain outcome (there is no fork) but with a low expected reward (i.e., block reward shared with other miners). Interestingly, they note that a miner forking the blockchain has two options. The first is to release their blocks as soon as they solved the hashing puzzle, and the second is to mine secretly. In the first case the probability of success is higher. This is because as soon as the branch is as long as the main

branch some miners may start mining on the new branch, thereby increasing the probability of success. However, this will be at the cost of sharing the block reward. In contrast, with secret mining, the miner is certain to reap all the block rewards. They show that with secret mining strategies the threshold drops from 36% to 30.8%.³²

The idea of secret mining actually dates back to [Eyal and Sirer \(2014\)](#). Their model is similar to that [Kiayias et al. \(2016\)](#) but is motivated slightly differently. They consider the case of a large miner (e.g., a pool of miners) who has just solved the hashing puzzle, and faces the decision of whether releasing the new block to the network or to mine secretly on top of it. [Eyal and Sirer \(2014\)](#) give a precise description of optimal, secret mining strategies. The pool mines secretly as long as its branch is longer than the main branch, and releases it otherwise. They find that such a strategy may pay off as soon as the pool has 10% of the hashing power.³³

Another potential fragility in Nakamoto's design is the problem of broadcasting transactions, identified by [Babaioff et al. \(2012\)](#). Based on the protocol, a miner's activity consists of selecting a set of transactions to be included in a block but also relaying transactions that it receives to other miners. If a miner broadcasts a transaction, the probability that some other miner will include it in their block and collect the fee increases. In the extreme case, if a miner is the first and only node hearing about a transaction, they may have incentives not to broadcast it at all and hold on to it until they are the one adding the block to the blockchain. It may be especially tempting if the transaction fee is large. Such hold up would cause the validation of this particular transaction to be delayed. This issue, while theoretically interesting, turned out not to be a problem in practice. This is most likely due to the presence of so-called lightweight nodes in Bitcoin network. These nodes hold all or part of the blockchain and participate in the broadcasting of new transactions. They, however, are not mining nodes, and thus they do not have incentives to withhold information.

³²[Sapirshstein et al. \(2016\)](#) found similar results.

³³Bitcoin mining, of course, does not just consist of independent entities or processors. Many have formed mining pools who, while operating independently, pool the rewards from their collective efforts so as to give miners a more reliable return. Economic analysis of mining pools is still very scarce. [Cong et al. \(2019\)](#) argue that miners form mining pools due to risk aversion, but that there is also a limit to the size of the pool, because pool coordinators tend to abuse their power. And the analysis in [Prat and Walter \(2018\)](#) indicates that pools have potentially increased the energy costs associated with Bitcoin mining.

3.4 Longest-chain attack: a formal model

Section 3.3 highlights a number of attacks, but the most dangerous one is perhaps the “longest chain rule attack” because it aims to override the history recorded on the blockchain, which, in turn, opens the possibility of an attacker motivated to double spend. Traditionally, Bitcoin is presented as a system involving on one side users (who proceed to transactions), and on the other side miners (who maintain and update the blockchain). The papers we have reviewed in the previous section fail to capture the fact that a miner can also be a user and *vice-versa*. One key assumption that we will introduce here is that the attacker’s payoff does not only consist of the block rewards obtained on the branch created by the attacker (like in Section 3.3), but also on the private benefit of the attacker. In particular, a miner may benefit by altering some of their past transactions made as a user. In a double spend attack, this latter component is the value of some transactions, but motives may be distinct from that.³⁴

3.4.1 Free entry equilibrium

While the previous contributions have the merit of providing a formal analysis of the “mining game” they do not fully capture miners’ incentives. Budish (2018), by contrast, focuses explicitly on their potential incentives. His model, which we present below, allows to formally examine when Bitcoin and other cryptocurrencies using PoW would be vulnerable to being hijacked. The model is relatively simple but rich enough to shed light on several aspects of Bitcoin (or the blockchain game in general) that have been overlooked by the literature so far. Also, the simplicity of the model makes it relatively easy to be extended.

Budish (2018) considers two limiting factors on a simple majority attack. First, some activities from dishonest miners may require more than a simple majority to implement. Second, for some activities that involve interaction outside the blockchain (such as a double spending attack), control of the blockchain cannot be confined to just the block in question but may require a time period to elapse. Thus, the dishonest miners may have to control the network for some time, which translates into adding a certain amount of blocks. He includes these elements in his model.

Under a system like Bitcoin the possibility of an attack depends directly on the difficulty

³⁴For instance, a Goldfinger attack is akin to sabotaging the system by undermining and/or destabilizing the consensus protocol. See Kroll et al. (2013) for an analysis of such attacks.

of the hashing contest, which itself depends on the number of miners (through their hashing power). One thus need to first relate the market structure with rewards and the mining costs, which we do now. To this end we will consider the case of "honest" miners who are only concerned about their expected payoff in a simple sense, i.e., within the system. For simplicity, suppose that all nodes are symmetric in their size and costs, i.e., $c_i(x_i^*) = c$, where c is a constant. Note that the cost c can serve as a proxy for the difficulty in the hashing puzzle. Since all miners are symmetric each one has has a probability $\frac{1}{N}$ of winning the block reward, $e\theta$. The participation constraint for miners is therefore

$$\frac{1}{N}e\theta \geq c. \quad (3)$$

If this constraint were not satisfied, no one would choose to become a miner.

While, as we will see, this model allows us to derive a number of interesting insights, the symmetry assumption is not without consequences. In particular, it obviates the issues related to strategic forking as the one studied by [Kiayias et al. \(2016\)](#) and others, where it was found that the incentives of a miner to enter depends on the hashing power distribution among the other miners. That is, a high concentration will increase the probability that small miners' blocks will be ignored.

One issue that is often discussed is that proof-of-work schemes require significant resource usage (particularly in terms of energy), approximately equal to Nc . Thus, one goal of a network designer would be to reduce those costs. [Ma et al. \(2018\)](#) show that, in the Bitcoin protocol, if the number of miners was regulated to be $N = 1$, then that miner would act in a manner that ensures that the computational difficulty of the puzzle would be as low as possible. In this case, Nc could be arbitrarily small. This illustrates one of the benefits of having a non-decentralized network, it can operate more efficiently. However, this would be at potential the cost of trust in the system.

By contrast, in a permissionless network, free entry will ensure that the equilibrium number of nodes, denoted N^* , will be such that (up to integer constraints):

$$\frac{1}{N^*}e\theta = c. \quad (4)$$

In this case, total resource use would be $N^*c = e\theta$. Note, however, that by changing the computational difficulty c there would no impact on total resource use as the number of miners would simply adjust accordingly. Thus, it is only by changing the block reward that total resource use will adjust. Of course, this might have an impact on the exchange rate e as well.

3.4.2 Equilibrium with (potentially) dishonest miners

The Free Entry condition (4) dictates what drives miners to enter when they are “honest” in the sense of being interested in processing transactions and validating blocks. However, miners could also be dishonest in the sense of having other goals that cause them to want to append blocks with information they know to be false (e.g., as might arise in a double spending attack or in an attempt to sabotage the network for other reasons). A sustainable blockchain has to be robust against such agents and deter their attacks.

Following Budish (2018), assume that there are N honest miners. Conducting activities that are dishonest requires effective control of the network. Setting aside the results of Kiayias et al. (2016) and others, let us assume that a dishonest miner needs to control a majority of computing capacity. That miners thus need to invest in additional computing power, equivalent to $N + \varepsilon$ miners. So the cost of conducting dishonest activities on the network is at least Nc , per block.

This type of majority attack allows one miner to gain control of the network for a period of time. For instance, to enable double spending, the attacker conducts a transaction at some point which is settled both on the blockchain and in the real world. Then, the attacker forks the blockchain and works (privately) on the new branch. Importantly, since the attacker mines alone on the fraudulent branch they bear the total cost, Nc (per block), of the attack. Provided the attacker has the sufficient computing power, the attack will eventually succeed; that is, the attacker’s branch is the longest and all the other miners switch to it.³⁵

Budish (2018) considers two limiting factors on a simple majority attack. First, some activities from dishonest miners may require more than a simple majority to implement. For instance, control to achieve a fork may require control of $\frac{A}{A+1}$ percent of the miners, where $A > 1$, thus increasing the cost per block to ANc . Second, for a double spending attack, control of the blockchain cannot be confined to just the block in question but may require a certain period of time to elapse. The reason for this is that, since accidental forks are relatively common, it is customary for transacting parties to wait several “blocks of confirmation” before settling the transaction. In the case of Bitcoin, this escrow period lasts usually six blocks, that is, the recipient of a transactions (e.g., a seller) will wait that the

³⁵One may argue that a double spending attack is at the risk of damaging the credibility of the cryptocurrency, thereby increasing its cost. Surprisingly, this may not be the case. For instance, the cryptocurrency Bitcoin Gold was hit by a double spending attack in May 2018 and in January 2020 without affecting its exchange rate (it actually increased after the second attack!).

blockchain grew by at least six blocks since their transaction appeared in the blockchain to consider the sender did indeed sent the money. Hence, a dishonest miner may have to control the network for some time, which translates into adding t blocks. Offsetting these limiting factors is the fact noted earlier that, while controlling a network, a block reward $e\theta$ will be earned for each block added. That reward accrues to the dishonest miner. Putting these together, the net cost to the dishonest miner is $(ANc - e\theta)t$.

The “forking” decision of a dishonest miner will be driven by the benefits they receive from being able to alter the blockchain as they see fit; that is, from their dishonest activities. To this end, denote by $V(e)$ the private benefit for an attacker is $V(e)$. In the case of double spending $V(e)$ would be the value of the tokens in the transaction the attacker seeks to cancel. That value depends on the exchange rate because the attackers’ gain precisely consists of tokens. The value of the attack is an increasing function of the exchange rate “e”.³⁶ Therefore, an attack is not profitable if the following incentive compatibility condition holds:

$$AtNc - te\theta \geq V(e) \quad \Rightarrow \quad Nc \geq \frac{V(e) + te\theta}{At} \quad (5)$$

The left-hand side, $AtNc - te\theta$, is the net cost of controlling NA miners for t periods, taking into account the block rewards earned during the control period. The right-hand side, $V(e)$, is the benefit of exercising that control for personal benefit. In the analysis that follows, we suppose that the private benefit for an attacker is $V(e)$, a non-decreasing function. That is, the more valuable is the token, the greater is the private benefit from dishonest activities.

A blockchain will be sustainable if both the participation constraint (3) and incentive compatibility condition (5) are met. The former says that the total network costs, Nc , must not be too high while the later says that they should not be too low. Putting these conditions together, we have:

$$e\theta \geq Nc \geq \frac{V(e) + te\theta}{At}. \quad (6)$$

Thus, a necessary condition for a blockchain to be sustained in equilibrium where there are potentially dishonest miners is that:

$$e\theta \geq \frac{V(e) + te\theta}{At} \quad \Rightarrow \quad e\theta(A - 1)t \geq V(e) \quad (7)$$

³⁶ V might also likely be increasing with some notion of the “usefulness” of the cryptocurrency, which is possibly different from the exchange rate.

Thus, the cost of an attack, $e\theta(A-1)t$, must exceed the private benefit, $V(e)$. This condition was derived by Budish (2018) and we call it the *Budish condition*.

If this condition holds then, even if the participation constraint binds (as it would in a permissionless network), the incentive compatibility condition will be satisfied. Budish derives several interesting economic insights from this condition. First, since from Eq. (3) we have $e\theta = Nc$, the security of the blockchain is *linear* in Nc , the amount of expenditure on computing power. As Budish notes, this is at odds with “standard” investments in computer security that yield convex returns. Second, the cost of maintaining the blockchain is a flow cost, similar to the trade-off from a deviation one sees in repeated games.

But perhaps the most interesting insight that Budish (2018) offers is what he calls the *economic limit* of the blockchain, which he derives from his equilibrium condition (7). Rearranging this condition we obtain

$$e\theta \geq \frac{V(e)}{(A-1)t}, \quad (8)$$

thereby relating the payoff to the miners, $e\theta$ (which includes both the new tokens and the transactions fees), with the private benefit of an attack, $V(e)$. This condition reads as follows. On the one hand, reducing the incentives to attack requires increasing the hash power of the network (so that obtaining a majority of the hash power becomes too expensive). But this can only be increasing miners’ block reward (so that either miners invest in more powerful machines or there is entry of additional miners). Crucially, Budish concludes that increasing the block reward is not enough, one also needs to increase the level of the transaction fees. When the attack consists of a double spending attack $V(e)$ essentially consists of transactions (that the attacker wants to erase). Therefore condition (8) always holds if it holds the transaction with the highest value. In other words, the presence of high value transactions may require small transactions to be associated with prohibitive transaction fees. However, note that (8) also highlights a limit of Budish’s model when considering for instance the results of Kiayias et al. (2016). Recall that Kiayias et al. (2016) found that attacks are possible with less than 50% of the hashing power. Budish’s model does not account for this possibility. Condition (8) is derived assuming $A > 1$, and according to (7) no safety condition can be derived for $A < 1$ as long as miners’ mining reward is positive.

Chiu and Koepl (2017) found a condition similar to the Budish condition (7), which they coin the *no double spending constraint*. Using our notation, Chiu and Koepl (2017)’s

condition is

$$e\theta \geq \frac{V(e)}{t(t+1)}. \quad (9)$$

The key difference from Budish is the effect that the escrow period has on the reward. It is linear with Budish (2018) and follows a power law with Chiu and Koepl (2017). The difference arises because, unlike Budish, Chiu and Koepl (2017) explicitly introduce the probability of success (when forking) in their calculations and they show that this probability decreases exponentially with the duration of the escrow period. Chiu and Koepl (2017) derive an interesting observation from their condition. Since their no double spending condition depends on the escrow period (as in the Budish condition), immediate settlements are incompatible with the absence of risk of attacks.

Moroz et al. (2020) extend Budish (2018)'s model by analyzing equilibrium in a model where the victim of the longest-chain double spending attack can counterattack in a similar way. The motivation is that if an attacker is willing to spend up to some amount $V(e)$ to attack, then the victim of the double spending, who lost $V(e)$, should also be willing to spend up to that amount to counter the attacker. In this case the problem becomes a war of attrition: the winner will be the one who will be able to sustain their (counter) attack the longest. They show that the subgame perfect equilibrium is such that the initial attack never occurs. Two assumptions are needed for their result. First, the cost of an attack increases over time and second that the victim's stake is actually higher than $V(e)$. This latter assumption is motivated by the observation that generally double spending attack victims are exchanges who incur also a reputation cost from the attack.³⁷

Interestingly, Moroz et al. (2020) highlight another strategy that may make longest chain attacks even easier than what authors like Kiayias et al. (2016) have suggested. One crucial issue for an attacker is to have enough hash power. Moroz et al. (2020) argue that an attacker can lure other (honest) miner to work on his branch by creating transactions with large fees that are only compatible with his branch. The prospect of large fees may be sufficient to lure additional miners, thus increasing the likelihood of success.

³⁷Moroz et al. (2020) provide a detailed summary of the major attacks that several cryptocurrencies have suffered.

3.4.3 Flow vs. stock attack

An implicit assumption in the analysis we conducted in the previous section was that the attacks we considered consisted of, as [Budish \(2018\)](#) calls them, “flow” attacks and not “stock” attacks. Specifically, while the private chain is in operation, the attacker’s cost only consist of the energy spent during the attack. If, instead, the cost is made of the energy cost is c_E and the cost of computer equipment, rC , then the cost function becomes $c = rC + c_E$. The idea is that once the attack stops, the attacker can stop using energy and can repurpose (or resell) the computer hardware. This latter option was, in fact, a feature of the Bitcoin system as proposed by [Nakamoto \(2008\)](#) in that it would democratize participation in the network by lowering the cost of entering to operate a node; the so-called “one CPU-one-vote” ideal.

However, as Bitcoin has evolved, miners started to have access to technological solutions that would give them an advantage in solving the cryptographic puzzle. One of these was the use of ASIC chips. Such chips are manufactured specifically for the computations needed to mine (i.e., hashing) but turn out not to be repurposable. With this technology an attacker would not be able to recoup the hardware cost following the attack and that cost would be sunk. This would make the attack a “stock” attack and the cost of the attack would be $N(C + c_E)$, modifying the incentive compatibility constraint as follows:

$$AN(1 - r)C + AtN(rC + c_E) - te\theta \geq V(e)$$

$$\Leftrightarrow Nc \geq \frac{V(e) - AN(1 - r)C + te\theta}{At}$$

Clearly, this would allow for sustainability against a higher $V(e)$ than the in the flow attack. Nonetheless, the analysis here would simply proceed by substituting $V(e) - AN(1 - r)C$ or $V(e)$.

That said, there are three reasons why this “stock” analysis may not be the appropriate one. First, for some blockchains (including, until recently, Ethereum), ASIC chips do not confer an advantage. Second, in some situations, it may be more efficient for the attacker to expand the number of nodes in the attack and use repurposable chips rather than ASIC chips in which case the “flow” analysis applies. Finally, it may be that an attack would not make the ASIC chips unusable because the network is not significantly damaged by the attack. However, this last option likely depends on the purpose of the attack, the possibility of future attacks and other real-world specifics of the blockchain in question.

3.5 Efficient design

Having identified the participation and incentive compatibility constraints that will allow for a sustainable blockchain, we can now consider the choice of design parameters that will minimize resource use subject to those constraints. That is, what are the choices of recruitment N , performance c , and reward θ that solve

$$\begin{aligned} & \min_{c, N, \theta} && cN \\ & \text{subject to} && \begin{cases} ANtc - t\theta \geq V(e) \\ Nc \leq e\theta \end{cases} \end{aligned}$$

[Gans and Gandal \(2019\)](#) note that a binding incentive constraint implies that $Nc = \frac{V(e) + t\theta}{At}$. Observe that by decreasing the reward θ , it is possible to maintain incentives at the expense of fewer resources. What prevents this from being reduced to zero is the participation constraint (3) condition that requires that there is a sufficient reward that the nodes are not making losses. The minimum possible θ is therefore $\frac{1}{e}Nc$, which results in $Nc = e\theta$. Given this, the θ that minimizes resource cost satisfies:

$$e\theta = \frac{V(e) + t\theta}{At} \quad \Rightarrow \quad \theta = \frac{V(e)}{e(A-1)t}. \quad (10)$$

By setting θ at this level, the participation and incentive constraints both bind and the cost Nc is minimized at

$$Nc = \frac{V(e)}{(A-1)t}.$$

There are several things to note about this outcome. First, if θ is chosen optimally, then the choices of N and c are determined. In other words, neither the difficulty of performance c nor whether the number of miners is fixed or left to free entry (as in a permissionless network) matters. Thus, for a sustainable blockchain, resource use is not directly impacted on by the number of miners, nor their cost under Proof-of-Work. Instead, it is the reward function that matters. Second, the reward function optimally varies with the exchange rate e . This is not something that any blockchain protocol (including Bitcoin) does at present. However, to create a sustainable blockchain, adjustment to market conditions is a useful property for the block reward. But in which direction should the reward move with e ? Note that, as e increases, the impact on sustainability depends on the relationship between e and V . If

$V(\cdot)$ is non-decreasing in e , the effect of an increase in e depends on the sign of $\frac{eV'(e)-V(e)}{e^2}$, or whether the elasticity of V with respect to e is greater than 1. If this is the case (i.e., the elasticity is greater than 1), a higher value for the cryptocurrency increases the likelihood that the blockchain will be vulnerable to attack, implying that θ should be higher.

To put this result in perspective, consider the case of Bitcoin which is designed to have θ decreasing over time and, according to Bitcoin’s aficionados, will have a higher future exchange rate e . The only way to maintain Bitcoin’s sustainability in this case is if the elasticity of V with respect to e is less than one. In other words, the value of transactions should not grow as much as Bitcoin’s exchange rate.

3.6 Proof-of-Stake as an alternative consensus mechanism

The largest blockchain networks, Bitcoin and Ethereum, rely on a Proof-of-Work consensus. The ‘work,’ however, that takes the form of a unnecessarily difficult computational contest consumes resources, particularly, energy (see [De Vries \(2018\)](#)). For this reason, there has been active work in computer science to find more cost effective consensus mechanisms. The main contender in this regard is proof-of-stake (PoS). While PoW, relies on the provision of real resources as a ‘ticket’ for entry as a participant in validating the blockchain, PoS looks instead at a financial cost in the form of a stake —usually, in the form of tokens from the network itself— as a means of signaling a miner’s intent to operate honestly (see [Buterin \(2013\)](#)). In this regard, it shares properties with financial market instruments such as escrow, collateral and other forms of bonding. Notably, the Ethereum Foundation is looking towards implementing a PoS mechanism.³⁸

3.6.1 The Nothing-at-Stake problem

PoS achieves consensus on blocks by requiring nodes to stake (or freeze) a certain quantity of tokens in order to be considered as a validator, i.e., to be the one adding the next block and grab the block reward and the fees associated to that block.³⁹ There are, however, different

³⁸A related notion is “proof-of-burn” where nodes publicly burn tokens (by sending them to invalid public addresses), which may may impact upon the monetary policy of a network; see [Saleh \(2019\)](#). In closed and small networks, computer scientists have also applied Byzantine Fault Tolerance as a means of achieving consensus. For a review see [Amoussou-Guenou et al. \(2019\)](#).

³⁹The term “validator” is used to distinguish from a PoW miner, since most of the work performed by a node consists of validating transactions.

ways in which validator nodes are selected but traditionally the selection process is similar to PoW where each node has a proportionate probability of being selected to propose a block and we will focus on that here. This probability may depend on the amount of tokens staked by miners. Also, under PoS a miner can at any moment regain access to the frozen tokens (although by doing this it would lose the opportunity to become a validator).

While proof-of-stake is motivated to economize on resource usage, it also changes other network properties. For instance, at first glance, one may think that it does not need a fork resolution mechanism like the LCR because at any moment only one miner is authorized to add a new block. That is, there is no risk of accidental fork. However, that does not mean that PoS can prevent hard forks or double spending attacks because the cost of working on two or more branches in parallel is negligible for nodes. In other words, PoS does not guarantee eventual consensus. This issue is known as the ‘Nothing-at-Stake Problem.’

[Saleh \(2019\)](#) argues, however, that this issue may not be as problematic as one may think. [Saleh \(2019\)](#)’s key insight is that disagreement on the blockchain, that is, the existence of several branches, is likely to have a significant impact on the dollar value of the token, and that validators take into account that it may affect their revenue. [Saleh \(2019\)](#) derives sufficient conditions that guarantee that consensus is an equilibrium, once we take into account the depreciation of the token in case of a fork. Saleh then derives two additional results. First, restricting the ability to large stakeholders facilitates and speeds up consensus in case of a fork. The intuition is that such stakeholders have the most to lose from a disagreement, i.e., from the persistence of two or more branches. Second, Saleh finds that the lower the miners’ reward the better. The reason behind this counter-intuitive result is that low rewards enable the accumulation of vested interest in the blockchain (i.e., miners have less incentives to cash out their tokens). Given this, preserving one’s vested interest in the blockchain (the tokens) increase the incentives to favor consensus.

3.6.2 Sustainability under proof-of-stake

How does PoS impact on the economic limits of the blockchain we considered above? This question was analyzed by [Gans and Gandal \(2019\)](#) and we outline their approach here.

Suppose that S tokens are required for a stake⁴⁰ and that the dollar interest rate is r ,

⁴⁰We make the assumption that the number of tokens is fixed in our setting. There are other PoS protocols with different assumptions. Our goal here is simply to employ the simplest PoS framework and show that it can be related to PoW.

then (in terms of our PoW notation) $c = reS$, that is, the opportunity cost per period of resources a node must hold to be staked (e.g., the lost earnings on fiat currency that is held in tokens). The stake S can be chosen in a permissionless PoS protocol, which means that, like c , it is a design decision.

This allows us to write a new free entry condition for a PoS network. Like in a PoW network, any node has a probability $\frac{1}{N}$ of earning the block reward. Thus, the expected payoff per block for a validator is $\frac{1}{N}e\theta - reS$. The free entry condition thus becomes (ignoring integer constraints)

$$S = \frac{\theta}{rN}. \quad (11)$$

Note that, unlike PoW, this condition does not depend on the exchange value e of a token. However, like PoW, the total network cost, SrN , is fixed. An increase in S causes a proportiona reduction in the number of nodes, N . A simple asset pricing equilibrium condition (without the possibility of capital gains or a risk premium) yields the same free-entry condition we have in equation (11).

What about protection against attacks by dishonest nodes? Both PoW and PoS methods are vulnerable to attack consisting of establishing a private chain with altered transactions before releasing to publicly. Under PoW, this entails a cost as a dishonest node is required to perform the PoW of the entire network in order to obtain the longest chain upon publication. With PoS, there is no such cost. The main challenge is having the other validators accepting it. Validators that were online while the alternative chain was being written will be able to identify the alternative chain. However, new validators (or validators that were offline) will not be able distinguish the legitimate chain from the alternative one. Thus, for an attack to be successful, the dishonest validators need to take actions that would shift the share of online validators. As in the case of PoW, we assume that this takes time t periods.

Such attacks rely on the attacker building on both the main chain and their alternative at the same time. This is something that is possible with PoS. However, PoS networks have implemented various methods to guard against this. One such method is called slashing, which involves the stake of a node being reduced or destroyed if it is found that they have worked on multiple chains. This is something that can be algorithmically detected. Importantly, when slashing is effective, it turns a flow attack into a stock attack as the stake itself can be taken from the attacker should they be discovered.

That said, while slashing can prevent “low scale” attacks on the network, PoS networks are still vulnerable to a “majority-attack” as we examined for PoW. Such an attack requires

the attacker to stake a supermajority of validators for t periods. If the value of an attack is $V(e)$ as before, then the cost of the attack is $ANtreS$ less the block reward $te\theta$ earned on the alternative chain. Note that block reward accrues to the attacker precisely because slashing or other mechanisms penalizes others if they work on the alternative chain leaving all of the block rewards to the attacker. Given this, the incentive constraint becomes

$$ANteSr - te\theta \geq V(e) \quad \Rightarrow \quad S \geq \frac{V(e) + te\theta}{eANtr}. \quad (12)$$

Thus, so long as the stake S is sufficiently high, an attack can be prevented.

We can now perform a comparable exercise to that for PoW to examine what will determine the sustainability of a PoS blockchain. In particular, using equations (11) and (12), if the incentive constraint is to be satisfied while the equilibrium number of nodes is determined by the free entry condition we obtain

$$e\theta(A - 1)t \geq V(e). \quad (13)$$

Two comments are in order. First, this condition is identical to the PoW Budish condition (7). Second, it is independent of the level of the stake S . In other words, despite the ability to control S , there are no design choices under PoS that will lead to greater sustainability than under proof of work. In fact, in designing a permissionless blockchain (even though S can be chosen), the minimum block reward that will create a sustainable outcome is $\theta = \frac{V(e)}{e(A-1)t}$. This also means that the same elasticity condition on $V(e)$ drives whether, for a fixed block reward, the network will be more sustainable as e grows.

It is useful to note, however, that the mechanism for sustainability is different. In particular, under PoS, the free entry condition is independent of the exchange rate e . In other words, the number of nodes will not change as the exchange rate changes and will be pinned down by the level of the stake. Thus, the size of the permissionless network can be controlled by changing the stake. This is not the case in PoW, since given θ and e , the complexity c , a design variable, determines the network size N in the permissionless PoW mechanism.

We have examined the equilibrium properties of consensus mechanisms for the blockchain under both proof of work and proof of stake when the blockchain is permissionless, i.e., when there is free entry into becoming a node. It is straightforward to extend the analysis to permissioned blockchains, in which the blockchain itself can determine who can become a node.⁴¹

⁴¹See [Gans and Gandal \(2019\)](#) for details.

3.7 Transaction fees

Up until now, we have not explicitly addressed transaction fees beyond their role as a block reward. Research that has focused on these fees has show that fees can also have a substantial impact on the robustness of the design.

Before going further note that a system where miners' revenue essentially depends on transaction fees is different from the current system where the newly generated coins constitute the bulk of the block reward. In a fee-only system the maintenance is paid for by the users. This implies then that, unless the level of the fees increases drastically and/or energy costs decrease substantially, the number of miners is likely to drop to a level that may jeopardize the security of the system (see (3)). But there are other potential issues, which we now comment on.

Carlsten et al. (2016) identify another potential source of forking by focusing on the difference between the block reward and the fees that are collected by the miner who added the last block. As of today the block reward is worth much more than the fees, so blockchain models usually only consider the block reward when modelling a miner's payoff. However, under Bitcoin's protocol there will be a time when fees will constitute most (if not all) of miners' payoffs. Carlsten et al. (2016) ask whether the incentives to follow the LCR (and thus not to fork) will still be present under a fee only regime. There is a fundamental difference between the block reward and the fees. The former is a deterministic variable, while the fees are probabilistic. But the variability of fees over time introduces new strategic opportunities that are irrelevant when the block reward largely dominates the miners' revenue. The intuition is relatively straightforward and consists of considering a miner's choice when a block has just been added to the blockchain. Before starting to hash a block a miner must choose which transactions to include in the block (among the transactions that have not been processed yet). If the set of available transactions does not include transactions with significant fees, a miner may have an incentive to fork "in the past" (i.e., create a fork before the last added block) and *reprocess* only a subset of the transactions that were already processed. By considering only a subset of past transactions the deviating miner leaves a larger set of transactions (with their associated fees) to be considered by the miners who will work on the fork.

A simple numerical example can illustrate this. Suppose that the total value of the fees included in the last block is 100 and that the total amount of the fees in the set of

unprocessed transaction is only 5. A miner working on the longest chain would then get at most 5. Carlsten *et al.*'s argument is that a miner can be better off forking just before the last block and consider transactions that bring a total fee of, say, 55. (If the fork occurs, the missing transactions will then be considered as not processed). The other miners then have the choice of mining on the “official” block and get at most 5 or mine on the fork with a potential gain of 50 ($5 + 45$). In other words, a blockchain without block rewards carries serious risks of instability.⁴² One may argue that this is a very remote possibility (block rewards are supposed to vanish circa 2140), but the problem may arise sooner: it just suffices that the fees become comparable to the block reward.

Huberman *et al.* (2019) also focuses on transaction fees. Recall that while the amount of the block reward is set by the Bitcoin protocol (and thus cannot depend on any strategic decision by the miners or the users), the fees are set by the users (more precisely, by the senders of the transactions). Therefore, fees turn out to matter for a very simple reason: when transactions are sent to the network to be processed, the Bitcoin protocol does not specify which transactions should be processed first. It is up to the miners to select the transactions to be included in the block they will work on. We have already mentioned the problems that can arise when miners have some discretion over which transaction they will process.

Huberman *et al.* (2019) consider the problem faced by the users when sending transactions and modeling as a queuing problem, where the queue is given by the transactions that are not processed yet and the order in that queue is given by the fees attached to each transaction. Such a queue is formed because transactions with the highest fees will be processed faster than those with low or zero transaction fees. Indeed, miners are not only engaged into a hashing race, but they also strategically select transactions to process in order to grab the highest fees.

The tradeoff faced by the users is thus very simple: the higher the fee, the faster the transaction will be processed. Note that due to the risk of forking (whether intentional or accidental), a user can generally not spend newly acquired bitcoins from a particular transaction until some time has elapsed, i.e., before a number of blocks have been added to the blockchain after the block containing that transaction. A receiver of a transaction who

⁴²Additionally, Carlsten and his coauthors show that in a fee-only environment the selfish mining strategy identified by Eyal and Sirer (2014), which is described later in this section, performs even better. This adds to the instability concerns.

wants to spend his bitcoins as soon as possible may thus arrange with the sender to include in the transaction a high fee. Choosing fees is, thus, equivalent to tournament because a user increasing their fee may negatively impact the other users. [Huberman et al. \(2019\)](#) analyze the equilibrium of the game between users (setting fees) and miners (who chose transactions). They show that in that unique equilibrium each user's transaction fee is equal to the externality the transaction imposes. Transactions differ in fees and equilibrium processing times, as users have incur different costs of delay.

The analysis draws attention to the implications of users and miners' strategic behavior for the future of the system. First, users will pay non-zero fees only if delays are costly to them but also only if Bitcoin is sufficiently congested. When there are only few transactions and blocks are almost empty, miners are indifferent between including no-fee transactions and not. As some miners will include them, the transactions will get written into the blockchain eventually. Alternatively, the user may resolve the miners' indifference, and offer a near-zero fee (e.g., 1 satoshi). Then all "available" transactions will get included in the next block. When, however, the blocks are filled to their capacity, only transactions with high enough fees get processed. In other words, Bitcoin needs to be sufficiently congested (miners' capacity far exceeds the needs) to generate delays, a necessary condition for users to pay substantial fees.⁴³

[Easley et al. \(2019\)](#) analyze a similar game between Bitcoin users and miners. Users decide whether to pay a voluntary fee to speed up their transaction, and miners decide which transactions to process first. [Easley et al. \(2019\)](#) find, however, that the Bitcoin fee game takes a structure of a coordination game, with two equilibria; in one all users pay the fee, and in the other no one pays.⁴⁴ The result difference lies in their different games. In [Easley et al. \(2019\)](#), a number of homogeneous users face a single positive fee, which they either decide to pay or not. The miners process no-fee transactions only after all paying transactions have been recorded. Thus, the larger the number of users paying the fee, the longer is the waiting time for the no-fee transactions. This creates consumption complementary, where the incentive to pay the fee increases with the number of other users paying the fee.

⁴³[Huberman et al.](#) also note that once the block reward disappears, the fees are necessary for the survival of the system. When there is not enough congestion, users' will pay nearly zero transaction fees, which will reduce miners' revenue. This loss of revenue will, in turn, reduce the number of miners, which may result in turn in a weakening of the system's safeguards against double-spending.

⁴⁴There is also a third, unstable equilibrium where the users are just indifferent.

Note that when all users pay the fee, all transactions have the same priority, so the processing times are the same as in the equilibrium where all users pay no fee. However, if all users pay the fee, the deviation to no fee is very costly, because it automatically puts the no fee transaction at the very end of the queue. This cost may be higher than the fee itself. Similarly, if all users do not pay the fee, the gain of jumping the queue may not justify the whole fee.⁴⁵

4 Bitcoin from the users' side

When analyzing the incentives at stake in the Bitcoin system an important part is devoted to the behavior of miners, i.e., the supply side. As we saw in the previous section, it is important to see whether Bitcoin behaves and evolves as it is expected to do. After all, the issues of miners' incentives in Bitcoin system would be of little relevance if it wasn't for the users.

It is almost tautological to say that without any users Bitcoin has no reason to exist. It is thus important to see how the public is receiving and adopting Bitcoin. Analyzing how Bitcoin is used and by whom can also shed light on the future of Bitcoin. For instance, if Bitcoin is essentially used for illegal purposes (drugs, tax evasion, etc.) it is likely that governments will act so as to minimize or hamper the use of Bitcoin. The extent of Bitcoin used for speculation is also related to current and future government policies, which may, in turn, influence how Bitcoin can be used. In this section, we will first examine early adoption by users and merchants —and then examine the types of activity that bitcoins are being used for. We then examine adoption by users and merchants in the most recent period for which there are data.

4.1 Early adoption and usage

At its inception, Bitcoin was only used by a small, expert community. One obvious reason is that participating in the Bitcoin network needed sufficient prowess in computing. Another reason, perhaps equally obvious, is that Bitcoin was not attractive due to the low value of

⁴⁵Easley et al complement their analysis with empirical results suggesting that as the waiting times increased, the no-fee equilibrium disappeared. They also show that the fee levels increased with the waiting times.

bitcoins and the high uncertainty about the success of the venture. This is not the case anymore. Bitcoin is now sufficiently established and frequently mentioned in the media, and anecdotally there is now a large number of people using or holding bitcoins. The first questions when considering usage are: Who is adopting Bitcoin? Why do they adopt it? How do they use it? Surveys give us some insight into these topics.

4.1.1 Early adoption in the U.S. and Canada

Henry et al. (2017) analyzed data obtained from a pilot study known as the *Bitcoin Omnibus Survey* (BTCOS) conducted in Canada. This survey is an online survey, done in two waves (November 9–13, 2016 and December 14–18, 2016), comprising of about 2,000 Canadians aged 18 or older. The data was then post-stratified to be representative of the Canadian population, taking into account age, gender and region. Participants in the survey were asked if they knew about Bitcoin, and whether they had or used to hold bitcoins.

The first observation coming out of the 2016 survey is that a relatively high level of awareness in the population: 64% of Canadians stated they had heard of Bitcoin. Knowledge of Bitcoin is mostly reported by the 25–34 years old group, but without significant difference between age groups. Male, income above \$100,000, or university degrees are additional factors correlated with higher Bitcoin awareness.

However, when considering Bitcoin holdings the numbers drop sharply. Only around 3% has or had bitcoins. The amounts reported to be held are, for most people, relatively small: 60% have at most 1 bitcoin, 32% between 1 and 10 bitcoins. Only 8% have more than 10 bitcoins. When asked for the main reason to own Bitcoin, almost one third reported “interest in new technology.” People who were aware of Bitcoin but did not own any, were also asked for their reasons. Most of them (60%) indicated that their current payment methods met their needs, or that they did not understand the technology. To sum up, although a significant part of the Canadian population seemed to be aware of Bitcoin in 2016, ownership and usage was not widespread at that time.

Schuh and Shy (2016) address similar questions, but for the US population. This paper nicely complements Henry *et al.*'s paper because it analyzes data from 2014–2015, which corresponds to a slightly earlier period when Bitcoin adoption was even lower. Schuh and Shy utilize data from the Survey of Consumer Payment Choice (SCPC), a dataset produced by the Federal Reserve Bank of Boston. One of the main observation made by Schuh and Shy is that, in 2015, more than half of the survey respondents have never heard of Bitcoin.

The percentage of individuals using or holding cryptocurrencies was also very low, about 1%–1.5% (with half of them holding no more than \$400, evaluated at 2015 prices). Survey respondents mostly cite first the adoption of a payment system, and then the interest in new technology and investment as the main reasons for adoption. Lack of trust in banks or government does not appear to be a primary driver of cryptocurrencies adoption. This seems a bit at odds with the initial libertarian motive of Satoshi Nakamoto and the cryptography community to build a system like Bitcoin. Even though the papers by Henry *et al.* and by Schuh and Shy analyze data from just two countries they permit us to have a broader perspective on the dynamics of Bitcoin adoption during the 2014-2016 period.

4.1.2 Early adoption by merchants

Polasik *et al.* (2015) provide another empirical study about Bitcoin adoption but, in contrast to Henry *et al.* (2017) or Schuh and Shy (2016) they focus on adoption by merchants. Analyzing adoption by vendors is important because it can help better understand users' adoption. Indeed, insofar as we consider Bitcoin as a form of *payment*, user and merchant adoption depend on network effects. That is, vendors' acceptance of Bitcoin depend on the number of users adopting it and, at the same time, users' incentives to use Bitcoin as payment system also depends on the number of merchants accepting it.⁴⁶ Polasik *et al.* conducted the survey in April 2013, selecting primarily small vendors known to accept Bitcoin as a form of payment for legal purchases. Hence, unlike Schuh and Shy or Henry *et al.* Polasik *et al.* (2015) focus exclusively on the existing Bitcoin ecosystem.

In line with the results of Henry *et al.* or Schuh and Shy, they find that the existence of alternative methods of payments (e.g., credit cards, Paypal, etc.) has a negative effect on Bitcoin usage. Further, Polasik *et al.* also observe that Bitcoin usage increases in countries with weak banking systems, where obtaining credit cards or even a bank account is difficult. Also the larger the shadow economy (or the lower the GDP), the higher the usage, which would confirm the popular belief that the shadow economy is an important playing field for Bitcoin.

⁴⁶Note that since Bitcoin is a *peer-to-peer* payment system, the presence of established merchants accepting payments in Bitcoin is not needed to foster users' adoption. Instead, Bitcoin can be used for private transactions between users.

4.2 Different activities using Bitcoin

[Tasca et al. \(2018\)](#) and [Athey et al. \(2016\)](#) identify different types of activity by analyzing transactions on the Bitcoin blockchain. *Tasca et al.* looks at the patterns of usage over time, and suggests that Bitcoin went through three phases of use: (1) early stage, dominated by miners, (2) “illegal” stage, dominated by black market and gambling, and (3) “business” or “legitimate” stage, dominated by cryptocurrency exchanges. We will discuss cryptocurrency exchanges in Section 6. Most of this third stage “legitimate” activity involves financial speculation, as we will see when we examine drivers of Bitcoin’s price in Section 5.⁴⁷

[Tasca et al. \(2018\)](#) extracted all the transactions stored in the blockchain between January 3, 2009 (the birth of Bitcoin) and May 8, 2015, and then identified how many times there is a transaction between two distinct addresses. At that time Bitcoin was a bit more than 6 years old but already consisted of more than 68 million transactions. This data can be mapped into a network that represents the relationship between any Bitcoin addresses. Their approach is relatively intuitive. They first identify “known” addresses, that is, addresses whose owners are public knowledge. This is the case, for instance, when a business or an organization publicly announces that it admits payments in Bitcoins and publishes an address. From observing which addresses have a transaction with those known addresses, they infer which type of activity the transaction belongs to (purchase of goods or services, financial transfer, etc.). Using graph theoretic techniques they identify clusters, i.e., groups of users with a high level of transactions. They focus their attention on what they call *super clusters*, i.e., groups that contain at least 100 unique addresses and that received at least 1,000 Bitcoins before May 8, 2015. They identify in total 2,850 such super clusters.

[Tasca et al. \(2018\)](#) links 209 super-clusters (out of 2,850 identified clusters) to real world entities and specific business categories, such as exchanges, mining pools, online gambling and black markets. With these 209 identified clusters they then look at the other, unidentified clusters that have transacted with them. Observing the type of activity of the identified cluster (e.g., gambling, exchange) they infer how users in these unidentified clusters utilize Bitcoin. For instance, if an unidentified cluster has most of its interactions with a cluster known to be specialized in gambling they deduce that the former is also likely to be involving in gambling.

⁴⁷See also [Christin \(2013\)](#) for an extensive analysis of early use of Bitcoin for drugs and other contraband.

Considering the date stamps of the transactions a pattern emerges, suggesting that Bitcoin went through three successive phases: an early stage, an “illegal” stage, and a “business, legitimate” stage. Tasca and his coauthors identify a first period of approximately 3 years (until around March, 2012) where Bitcoin was mostly used by miners and with virtually no commercial activity. Early commercial adopters started to appear between April, 2012 and October, 2013. With the existence of exchanges, black markets and gambling start to be the main drivers of Bitcoin usage. The picture in [Tasca et al. \(2018\)](#)’s analysis changed substantially early 2014, with gambling and black markets becoming marginal and cryptocurrency exchanges taking more than 80% of the income generated by Bitcoin, suggesting that financial speculation is now main driver of Bitcoin’s use.

The analysis of [Athey et al. \(2016\)](#) focuses rather on regional use differences, and on the relation between user adoption and the fluctuation of Bitcoin’s exchange rate. Analyzing data they find that until mid-2015, transaction traffic on Bitcoin’s blockchain was growing slowly. They note that adoption can be severely affected by the uncertainty about the system’s reliability or Bitcoin’s volatility. An interesting finding relates to the usage frequency. The proportion of frequent users does not grow over time; but the proportion of short term or one-time users grew from less than 20% on July, 2012, to more than 40% in May, 2015. According to Athey *et al.* this could be worrisome because it suggests that Bitcoin velocity may decline over time.⁴⁸ Another related observation is that the time that users’ bitcoins remain in their wallet is increasing. That is, bitcoins’ owners tend not to spend them.

Similarly to [Tasca et al. \(2018\)](#), [Athey et al. \(2016\)](#) make use of publicly identified Bitcoin addresses. Then they use machine learning techniques to identify clusters and classify them into different types of industries and regions, based on the connection with those publicly identified addresses. They find that over the whole sample, North America and Europe count for nearly 70% of the Bitcoin activity, followed by Asia-Pacific with a bit less than 10%.⁴⁹ Overall about 20% of transactions in the network are classified as gambling and contraband; but it accounts for only 4% of the value. Further 13% of transactions are exchanges. All the other classified categories are very small. However, over half of all transactions — making up

⁴⁸This is in contrast with the theoretical result if [Bolt and Van Oordt \(2020\)](#) (described in Section 5), where velocity is expected to increase with usage.

⁴⁹Those percentages are obtained from Table 5 in [Athey et al. \(2016\)](#). They manage to assign a region for a each address from a sample of 3407 addresses. Out of those, 1312 are from North America, 1091 from Europe and 332 from Asia-Pacific. [Athey et al. \(2016\)](#) note however, their sample is not representative ”as it has many more long-term users and miners than the full dataset.”

75% of value transferred — are not classified. Recent papers, we examine next, shed some light on this issue.

4.3 Recent data on use activity: where are we today

As the papers above reveal, illegal activity was a substantial part of Bitcoin usage. [Tasca et al. \(2018\)](#) find that the second stage of Bitcoin adoption was driven by black market and gambling, and [Athey et al. \(2016\)](#) observe that at least 20% of Bitcoin activity in their study can be attributed to gambling and contraband. If we fast forward to today, we do not see many changes. That is illegal activity continues to be significant.

[Foley et al. \(2019\)](#) find that approximately one-quarter of bitcoin users are involved in illegal activity, which they estimate to represent 46% of bitcoin transactions. Based on their estimates, the illegal use of bitcoin generates approximately \$76 billion of illegal activity per year. In terms of comparison, they note that the scale of the US and European markets for illegal drugs is only slightly larger! They do find that since 2016 the proportion of bitcoin activity associated with illegal trade has declined, but the absolute amount of activity (in USD) has continued to increase.⁵⁰

Their first approach exploits the trade networks of users known to be involved in illegal activity (“illegal users”). They use the bitcoin blockchain to reconstruct the complete network of transactions between market participants. They then apply a type of network cluster analysis to identify two distinct communities in the data—the legal and illegal communities. Their second approach exploits certain characteristics that distinguish between legal and illegal bitcoin users. For example, they measure the extent to which individual bitcoin users take actions to conceal their identity and trading records, which predicts involvement in illegal activity.

One example of illegal activity that currently flourishes with Bitcoin is “ransomware” attacks in which criminals exploit vulnerabilities in computer networks to “lock” files so that the user cannot access them. As documented in an article in the New York Times by Nathaniel Popper,⁵¹ in 2019, more than 200,000 organizations submitted files that had

⁵⁰One possible explanation is that the huge rise in the Bitcoin price, especially in 2017 when Bitcoin increased from \$1,000 in value to \$19,000 led to large investments and speculative trade. We will discuss this in the following section.

⁵¹“Ransomware Attacks Grow, Crippling Cities and Businesses,” by Nathaniel Popper, New York Times, February 10, 2020, available at <https://www.nytimes.com/2020/02/09/technology/ransomware-attacks.html>

been hacked in a ransomware attack. This was a 40 percent increase from the year before, according to information provided to The New York Times by Emsisoft, a security firm that helps companies hit by ransomware. The average “ransom” payment to release files spiked reached more than \$80,000 in the last quarter of 2019, according to data from Coveware, another security firm. In the last month of 2019, several organizations faced ransom demands in the millions of dollars. The criminals carrying out the ransomware attacks are difficult to locate (and hence prosecute) because they typically demand payments in Bitcoin.

Their results are consistent with what we know about adoption by large merchants. According to the Economist magazine using data from Morgan Stanley, in 2018, only three of the largest 500 online retailers accept Bitcoin for payments.⁵² In 2017, five such retailers accepted Bitcoin. The conventional wisdom for the lack of adoption of Bitcoin as a payment system is that very few “legal” goods are purchased using Bitcoin because its value is not stable and the system is very slow in processing transactions. This is likely a disappointment for those who envisioned as a payment system competing with Visa and MasterCard. We are nowhere near that point. The rapid growth of ever more user-friendly payment systems does not seem to bode well for Bitcoin’s adoption.⁵³ Given little use of cryptocurrencies to date, the majority attack possibilities might not be exciting enough to tempt an attacker for most coins.⁵⁴ And a majority attack on Bitcoin would likely be prohibitively costly.

5 Price of Bitcoin

There are many reasons for Bitcoin to be a subject of interest in the media, social networks or just daily conversations: It is a new currency, solely digital, not backed by any government

⁵²Cryptocurrencies: Riding the Rollercoaster, December 1, 2018, p.69, Economist.

⁵³Although we do not review the literature on privacy in this survey, one of the reasons for such extensive illegal activity in the Bitcoin ecosystem is the perception that Bitcoin transactions are anonymous. There are other cryptocurrencies (e.g., Namecoin, Monero, Dash) that specifically aim at improving anonymity of transactions. But because the transacting parties are identified only by their addresses, and not by a personal identifiers (Bitcoin is said to be “pseudonymous”), Bitcoin seems to offer substantial anonymity—and it is very costly to trace activity to “real” people. Reid and Harrigan (2013) identify a bit more than 1.2 million different public keys, but after regrouping (e.g., two public keys used in the same transaction must belong to the same user) they end up with only 86,641 users. Using various external databases they are also able to link users and IP addresses and explain that using such information could allow us to also analyze the activity of identified users. See also Androulaki et al. (2013).

⁵⁴Moroz et al. (2020) document majority attacks on some minor coins.

or central bank, etc. The price of a bitcoin is perhaps the main reason for Bitcoin's media coverage. While the market took off slowly, a massive spike in the price of a bitcoin in late 2013 led to wider interest in what had been until then a niche industry. The value of Bitcoin (BTC) increased from around \$150 in mid 2013 to over \$1,000 in late 2013. The fall was dramatic as well and bitcoin fell to \$400 in a very short period of time.

In 2017, things again changed dramatically. Bitcoin began rising again and by early 2017, the value of bitcoin was again more than \$1,000. It had taken more than three years for the value of bitcoin to return to the 2013 peak level, but that was only the beginning. Eventually, in December 2017, Bitcoin reached a peak of more than \$19,000 before plummeting over the next few months to \$6,000. Despite the dramatic fall, cryptocurrencies were on the map and massive entry (as well as non-trivial exit) has occurred in the industry in the last five years. We will have more to say about other cryptocurrencies in the next section.

While very volatile, the Bitcoin price has increased rapidly from 2009 to the present. In other words, people buy bitcoins. What are then users' intentions when buying bitcoins? Is it for transaction use or for investment? Even though many papers (see below), offer arguments that Bitcoin would not make a good currency, some people may still want to use it for transactions, e.g. black market transactions. The rapid evolution of Bitcoin's price is not the only staggering feature: it is also extremely volatile. The price of a bitcoin can easily drop by 10% in one day and jump again 10% or 15% a few days later (and sometimes this yo-yo like movement happens in a single day). The determinant of Bitcoin's price and the source of its volatility are the focus of many papers.

5.1 Theoretical analysis

[Yermack \(2013\)](#) examined the determinants of Bitcoin's price; offering a general discussion about Bitcoin with a number of interesting insights. Between May and September 2013 the price of a bitcoin was around \$100-120, with a peak at \$140 and a low at \$67. Even in that early period, Bitcoin was very volatile. According to [Yermack \(2013\)](#), this is one of the reasons why Bitcoin does not qualify as currency. For economists a currency serves three purposes: (1) it serves as a medium of exchange; (2) it serves as a unit of account; and (3) it serves as a store value.⁵⁵ The high volatility of Bitcoin makes it impractical to serve as

⁵⁵More precisely, the currency's main role is to serve as a medium of exchange. To serve that purpose well, it also needs to be a reliable unit of account, for otherwise it introduces additional frictions to price

a store of value. [Yermack \(2013\)](#) notes that Bitcoin faces hacking and theft risks and lacks access to a banking system with deposit insurance. He concludes that Bitcoin appears to behave more like a speculative investment than a currency. This early conclusion has been reinforced by other later papers.

[Bolt and Van Oordt \(2020\)](#) develop a theoretical model to separate how much of Bitcoin's price is due to speculation and how much due to transaction use. The subject of their analysis is a currency in the early stage of adoption. They assume the following properties in their model: (1) the prices are set in "established" currency units (like the USD) and instantly adjusted with the current exchange rate to give a price in the currency of interest; (2) growth of supply of the currency is predetermined, and (3) Bitcoin exhibits two-sided network effects. These properties seem to fit Bitcoin well.

They take an innovative step of combining the monetary economics framework with the literature on network effects. Their argument starts with the famous quantity theory of money, which states that the volume of transactions (prices multiplied by the quantity) is equal to the monetary mass multiplied by its velocity. [Bolt and Van Oordt \(2020\)](#) pay a special attention to the monetary mass, explaining that it should not be considered as the total amount of the cryptocurrency in existence. They differentiate between dormant bitcoins and used bitcoins. Only used bitcoins should be considered as the currency's monetary mass.

[Bolt and Van Oordt \(2020\)](#) interpret any bitcoins not transferred in a given period as speculation, and any bitcoins transferred are assumed to be payments for goods and services. There are, of course, some problems with this interpretation. When speculators are buying or selling the currency, this is measured as a "transaction," rather than speculation. The second problem is that people who hold bitcoins for transaction purposes in the future are counted as speculators. There is no good reason to believe that these two opposing effects will balance each other out. Additionally "transactions" could be for illegal purposes as well.

[Bolt and Van Oordt \(2020\)](#) show (theoretically) that the current exchange rate depends on (1) the current adoption (i.e., number of merchants accepting the currency, and the number of users paying with it), (2) future adoption, and also (3) speculators' decision to hold the currency. They conclude that, as the adoption of the cryptocurrency increases, the exchange rate increases as well, but the holdings of speculators decrease. When the network adoption reaches its final size, the price will not increase any more. Thus, the speculators

settlement between the trading parties. Also, it needs to allow for store of value between the trades (see [Halaburda and Sarvary \(2016\)](#)).

have no reason to hold bitcoins. In the early stage, however, they hold large proportion of the currency, and shocks in their expectations can have a large impact on the exchange rate. However, as the use of the cryptocurrency for transactions increases, the impact of the speculators diminishes, which translates into a more stable exchange rate. (Of course, we are nowhere near a stable price, even though there is less volatility currently than in earlier periods.)

In our opinion, the primary appeal of this paper is the formal methodology and the theoretical model they develop, not the precise conclusions they draw. Indeed, most papers below find that speculation and illegal use generate the prime demand for Bitcoins. As the authors note, their study is only a first step in trying to understand the underlying economics of virtual currencies —and this is how we think the paper should be viewed.

[Biais et al. \(2018\)](#) also examines the theoretical determinants of Bitcoin’s price. Unlike Bolt and van Oordt they address the interplay between investors and hackers. Investors find value in Bitcoin (either for transactional purposes or as an investment) and hackers simply steal bitcoins. In an overlapping generation model with rational expectations they show that small noise in beliefs about futures prices can have a large impact on bitcoin’s volatility that is unrelated to fundamentals. They then empirically estimate their model using data on prices (from exchanges), transactions fees and volumes (from the blockchain), as well as hacks, losses and theft (manually collected from the web). They find that fundamentals only explain a relatively small share of return variations on bitcoin. In the context of their model, this suggests that observe that most of Bitcoin’s volatility is due to extrinsic noise. The papers we discuss below suggest two potential sources for that noise: the impact of the media and price manipulations.

5.2 Empirical papers on determinants of Bitcoin price

[Urquhart \(2018\)](#) looks at informativeness of Bitcoin’s price. Specifically, he asks whether Bitcoin market is informationally efficient, meaning that the price incorporates all the information that is relevant to determine the value of an asset. Urquhart checks this using daily data on Bitcoin’s price between August 2010 and July 2016, and finds that the data generating process exhibits predictability that he concludes is inconsistent with market efficiency.⁵⁶

⁵⁶Technically, some predictability can be explained by time variation in risk premia, but this is perhaps less likely over relatively short horizons considered in [Urquhart \(2018\)](#) than over longer horizons considered

However, for the second part of the sample (between August 2013 and July 2016,) the hypothesis that price fluctuations are random cannot be rejected. This may be interpreted as market becoming more efficient, i.e., the Bitcoin market maturing.

A general note about this approach is that if the market was immature in the early years of Bitcoin’s existence, (2009–2013) a lot of volatility may come from newness and lack of information. In such settings one may expect that when any new information is released, it might on average have a larger impact on prices than new information released in more established settings. This would be the case for example if market participants update in Bayesian sense, and arguably in Bitcoin’s early years their priors were probably quite diffuse. Another possibility is that the dearth of information makes it easier for market participants to “agree to disagree,” which may lead to excess volatility in the [Shiller \(1981\)](#) sense.

The initial volatility and the possibility that the Bitcoin market was evolving in its early years means that we must look at the results based on data from this early period with caution. Many of the standard tools may be not adequate for the analysis of such inefficient market. And whatever conclusions we draw, it is not appropriate to extrapolate them into the later period. In any case, the bubble of 2017 (discussed below) suggests that any such conclusion may have been too preliminary.⁵⁷

Many papers examine the effect of the media on Bitcoin’s price. [Glaser et al. \(2014\)](#) aim at separating the two motives for buying bitcoins —as an investment asset, or as means of payment— and they propose to do so by analyzing how volumes and prices are affected by news. They conclude that especially uninformed users approaching digital currencies are not primarily interested in an alternative transaction system but seek to participate in an alternative *investment* vehicle.

To get a sense of the interest about Bitcoin by non-expert users [Glaser et al. \(2014\)](#) look at the number of visits on Bitcoin’s page on Wikipedia. This is a reasonable approach as Wikipedia looks like an obvious entry point for most people. The authors show that the traffic on the Wikipedia Bitcoin page increases from 2011 to 2013. At the same time, transactions on Bitcoin’s blockchain and exchange volumes on Mt.Gox, the dominant exchange at the time, increase. However, they do not determine whether this is correlation or causation. The problem is that there is no way to check whether the people visiting Bitcoin’s page on

for example in [Campbell and Shiller \(1988\)](#).

⁵⁷We do not address trading and arbitrage issues in this survey. For coverage of this issue, see [Makorov and Schoar \(2020\)](#) and the references cited within.

Wikipedia are those buying bitcoins on an exchange (and vice-versa).

Glaser *et al.* find that the exchange volume is much larger than the transaction volume recorded on the blockchain. In their view, this is evidence that Bitcoin is mostly used as a speculative asset. This is because transactions performed on an exchange like Mt.Gox are not recorded on the blockchain.⁵⁸ Unless a user withdraws his bitcoins from an exchange, holdings are stored in the users' account at the exchange. They conclude that the new users are largely non-expert users, consistent with the increased Wikipedia traffic, and suggests that Bitcoin is a speculative asset, which accounts for the disproportionate increase in exchange volume as compared to the volume of transactions on the blockchain.

While Glaser *et al.*'s study is certainly illuminating in many aspects, it is important to note that their attempt to separate speculative and purchasing transactions cannot yield precise estimates. A thorough analysis of the blockchain would be needed to disentangle these two activities. The reason is that we cannot attribute volume on the blockchain only to purchasing transactions. There are two channels through which a speculative activity would leave a trace on the blockchain. First, users buying bitcoins for speculative purposes may transfer their bitcoins from one exchange to another exchange. Second, users may simply store their newly acquired bitcoins on their personal wallets. Hacks and thefts on exchanges constitute a reasonable motive for a user to do so.

The study of [Glaser et al. \(2014\)](#) establish a positive correlation between users' interest (captured through Wikipedia visits) and the price of Bitcoin. The question that remains, though, is whether we could identify the causal direction. In other words, is it an increase of interest that boosts the price or is it the price (and its changes) that leads people to inquire about Bitcoin? This question is addressed by [Kristoufek \(2015\)](#) who uses wavelet analysis to analyze the fluctuations of Bitcoin's price. He finds that until the first half of 2012, Bitcoin's price jumps boosted Google searches, while after 2013 the relation is reversed, i.e., the number of Google searches seems to influence positively the price.

[Garcia et al. \(2014\)](#) examine the role of the circulation of information about Bitcoin. They argue that Google searches (or Wikipedia usage) mostly capture what they call the "information gathering stage." In contrast, social networks such as Facebook or Twitter allow to capture word-of-mouth communication by looking at the number of posts that are shared.

⁵⁸The exchange as a whole may show up as one address on the Bitcoin's blockchain. Then all the transactions on the exchange are so called off-chain transactions that are recorded internally by the exchange, but not on the blockchain.

Combining search data and social network data allows them to distinguish between different dynamics of the “buzz” around Bitcoin, and of its price. Using standard econometrics methods for time-series [Garcia et al. \(2014\)](#) draw a number of interesting insights. First, they observe that search activity is positively correlated with the price of Bitcoin. Second, they identify two positive feedback loops that generate bubbles. The first loop is driven by word-of-mouth communication. An increase in search volume triggers an increase in the “social debate,” which eventually leads to price increase. The second loop is driven by new Bitcoin users. Garcia and his coauthors argue that adoption can be observed first by looking at the blockchain (with the appearance of new wallets) and also the number of times Bitcoin client software is downloaded. In this loop, an increase of search volume attracts new users who increase the demand for bitcoins, and thus its price.

[Garcia et al. \(2014\)](#) show that mentions of Bitcoin in social networks are related to Bitcoin’s price. In their analysis, they only include the number of shares or re-tweets, not their content. We could expect that positive and negative mentions may have different effects. [Kaminski \(2014\)](#) focuses on the relation between posts’ sentiments and Bitcoin price. Using a semantic analysis of Twitter posts containing the word “bitcoin” he sorts tweets into three categories depending on their emotional content: positive, negative and uncertain. Positive messages contain words like “happy” or “good”, negative messages contain words like “bad” or “sad” and uncertainty is captured by words like “hope” or “fear.” A simple correlation analysis shows that there is a positive correlation between Bitcoin price and the number of “emotional” tweets, i.e., positive or negative. But correlation does not mean causation. Analyzing dynamic Granger causality shows that emotional messages are in fact more the consequence of Bitcoin market activity, not the other way around. Moreover, [Kaminski \(2014\)](#) finds that a higher trading volume induces more tweets reflecting uncertainty. In other words, social network activity should not be considered as a predictor of market activity or the price of a bitcoin.

[Mai et al. \(2015\)](#) also analyze how the sentiment of posts on social networks influence Bitcoin’s price. They conduct a semantic analysis of posts on Twitter, and also on a Bitcoin discussion forum, [Bitcointalk.org](#). They also find a positive correlation between posts’ sentiments and Bitcoin’s price. In contrast to [Kaminski \(2014\)](#), they do not find Granger causality for Twitter posts. It may be a result of a somewhat different sentiment analysis procedure. [Mai et al. \(2015\)](#), however, find the causal relation where the sentiment estimated from the internet discussion forum posts affect the Bitcoin price. This may be attributed to

the differences in how the information is presented and how users interact in the two type of media. Longer posts on the internet forum allow for more substantial arguments, and also discussion. Thus, may have more impact on the behavior of the market.

Even within the internet forum, [Mai et al. \(2015\)](#) show that the impact of the post may depend on who is posting it. They identify two types of users, active users who contribute most content (which they call the “vocal minority”) and the relatively inactive users who contribute less often (“silent majority”). They find that the posts of silent majority have stronger impact on Bitcoin price that the posts of vocal minority. This result is surprising, because as they note, in financial markets one expects that the vocal minority plays a crucial role in information cascades, which can lead to herding behavior, because leaders’ opinions are usually widely observed and assumed to convey significant information. The explaining factor could be that silent majority users are not particularly interested in generating buzz. And therefore, the sentiments of the silent majority might tend to be more concise, relevant, and less noisy.

[Wang and Vergne \(2017\)](#) offer another study of the impact of media on the price of Bitcoin. Unlike [Mai et al. \(2015\)](#) or [Kaminski \(2014\)](#), they use traditional media mentions. Interestingly, [Wang and Vergne \(2017\)](#) aim at separating the effect of media “buzz” and the effect of technological advancement on the price of Bitcoin, and other major cryptocurrencies.⁵⁹ For the technological aspect of the cryptocurrencies, they use “technological development” measure developed by CoinGecko.com. The measure is a weighted average of eight indicators, such as the number of individuals contributing code to the project, or the number of versions (i.e., updates) of the software.⁶⁰

To capture the “buzz factor,” they use two indicators, “public interest” and “negative publicity.” Public interest is measured by the search volume on Bing for given cryptocurrency term and the traffic on the official cryptocurrency website. Negative publicity looks at the number of media articles in Factiva database (which covers traditional media, not social networks), that mention a given cryptocurrency together with negative words such as “fraud,” “hack,” or “Ponzi.”

They conclude that the innovation and technological potential —not the buzz— is the key aspect that drives Bitcoin’s (and other cryptocurrencies’) price. The technological development measure is positively and significantly associated with the returns. There was no

⁵⁹[Wang and Vergne \(2017\)](#) also look at Litecoin, Peercoin, Ripple and Stellar.

⁶⁰Though, the weights are not revealed to the reader.

such positive association with the “buzz” indicators. Wang and Vergne (2017) even find negative association between public interest indicator and cryptocurrencies returns. Lastly, the media coverage measured by their negative publicity measure is not significantly associated with the returns.

This last result indicates that social media has a stronger impact on Bitcoin price than traditional media. Going back to Henry et al. (2017) survey, while general population may be widely exposed to the news on Bitcoin, only a small fraction actually decides to actively participate in the market. Common sense and anecdotal evidence tells us that these would be (on average) more technologically inclined people. They would be more likely to pay attention to aspects that come into the “technological involvement” measure in Wang and Vergne (2017). And they would also be more likely to actively participate in an internet forum discussing Bitcoin. Twitter, in contrast, is more similar to traditional media, in the sense that it is focused on statements rather than on discussion.

An additional possibility for the volatility of Bitcoin’s price focuses on various types of price manipulation, and we now turn to that aspect.

5.3 The Bitcoin bubbles

For many people the extreme volatility and price increase of Bitcoin is a telltale sign of a bubble: the price of a bitcoin does not reflect the “fundamentals” and is essentially driven by optimistic beliefs about future demand for bitcoins.⁶¹ Cheung et al. (2015) analyze to this end data from July, 2010 to February, 2014. They identify three large bubbles that lasted between 66 and 106 days between 2011 and 2013, and a number of short lived bubbles (that only lasted a few days). In fact, we could almost interpret their results as saying that there is a bubble every day. The collapses of the three major bubbles they identify coincide with major events around Bitcoin: a theft in June 2011, trading suspended in April 2013, and Mt. Gox’s shutdown in February 2014.⁶²

To obtain these results they use a technique that Phillips et al. (2015) developed for the S&P500, which roughly amounts to identifying stochastic explosive behaviors in time series. It should be noted, however, that while Cheung et al. (2015)’s results are in line with the public opinion —that Bitcoin is a bubble— their analysis should be taken with caution.

⁶¹For an accessible discussion on the impact of beliefs on price, see Andolfatto and Spewak (2019).

⁶²At that time Mt.Gox was the main Bitcoin exchange.

The technique they use has been developed for mature markets and, as Urquhart suggests, Bitcoin has not been a mature market during [Cheung et al. \(2015\)](#)'s sample period.

During the period they examined, the most dramatic bubble was in 2013-2014 the price of a bitcoin rose from \$150 to \$1,000 before falling quickly to \$400 after the collapse of Mt. Gox. [Gandal et al. \(2018\)](#) find that this bubble was due to price manipulation. See section 5.4 below.

In 2017, Bitcoin rose from approximately \$1000 in January to \$19,000 in December of that year. Like 2014, there was a spectacular fall in the price of Bitcoin and in March 2018 Bitcoin's price was "just" slightly above \$5,000.

In March 2020, during the early part of the Covid-19 pandemic, Bitcoin's price was again around \$5,000. Afterwards, there was another meteoric rise: by the end of December 2020, Bitcoin's price reached highest price ever: a bit above \$28,000.

To the best of our knowledge, there has been no rigorous analysis of these recent bubbles. The fact that Bitcoin is a new type of "asset" is indeed a source of problem in the literature. Because it does not behave like other, more classical assets (e.g., commodities, equities, bonds, currency, etc) it is not a surprise to see that no consensus has emerged yet in the literature about the methodology one should employ when analyzing Bitcoin.⁶³

5.4 Price manipulation

[Gandal et al. \(2018\)](#) examined the period in which price of a bitcoin rose from \$150 to \$1,000 before falling quickly to \$400. (2013-2014.) They identified and analyzed the impact of suspicious trading activity on the Mt.Gox Bitcoin currency exchange, which was the largest cryptocurrency exchange and accounted for more than half of all Bitcoin trades for several years.

In early 2014, in the midst of theft allegations, the Mt.Gox transaction history was leaked. The Mt.Gox data dump gave access to approximately 18 million matching buy and sell transactions which span April 2011 to November 2013. These data are much more finely grained than data one could obtain from the blockchain or public APIs for two reasons. First, a majority of the trading activity is recorded only by the exchange. Second, the exchange links transactions by the user account. Data from the dump include fields such

⁶³In section 5.5, we make a calculation showing that if all the money in "investment gold" was converted to Bitcoin, the Bitcoin price would be greater than \$24,000.

as transaction ID, amount, time, currency, and user country and state codes. Also included is the user ID, which is the internal number associated with Mt.Gox users. The user ID is crucial as it enabled the authors to link transactions by the same actor.

Gandal et al. (2018) find that the USD-BTC exchange rate rose by an average of four percent on days when suspicious trades took place, compared to a slight decline on days without suspicious activity. They conclude that the suspicious trading activity by the Mt.Gox exchange itself likely caused the unprecedented spike in the USD-BTC exchange rate in late 2013, when the rate jumped from around \$150 to more than \$1,000 in two months.⁶⁴

5.5 Bitcoin as “digital” gold

Most of the empirical research we discussed in this section suggest that currently, bitcoin demand is driven by speculation alongside likely illegal intent. A broader claim about bitcoin demand is that it is used as a hedge against inflation. Many people invest in gold for that reason. Enthusiastic supporters of cryptocurrencies often argue that Bitcoin will replace gold as the hedge against inflation. Let’s run some numbers to see how that squares.⁶⁵

Simulation 1: How much of “investment gold” would have to be transferred to cryptocurrencies to maintain the total market capitalization (282 Billion USD in July 2019) of cryptocurrencies?

There are currently 190,000 tons of gold in the world. According to Wikipedia, 19% of all gold is held for investment. At the current price 1,444 USD per ounce of gold, the total investment in gold is currently 1.668 Trillion USD. (Central banks hold another 17% of the world’s gold. Roughly 49% of the hold is held in the form of jewelry, and an additional 12% of the gold is used in industry. See https://en.wikipedia.org/wiki/Gold_holdings.)

The value of all cryptocurrencies 282 Billion USD. If 17% of gold investments were moved to cryptocurrencies, this investment demand could sustain July 2019 current prices without speculation or criminal activity. ($1.668 \text{ Trillion} * 0.17 = 282 \text{ Billion}$.)

One problem, however, with investing in cryptocurrencies is the huge swings in valuation. Total valuation in the cryptocurrency market has ranged from 100 to 800 Billion USD between mid 2019 and the end of 2020. Bitcoin, the cryptocurrency market leader with a 66%

⁶⁴In a “subsequent” trial in Japan, the former Mt.Gox, CEO Mark Karpeles, confirmed that the exchange itself operated the suspicious accounts, but claimed that the trading by these accounts was “legal.” Griffin and Shams (2020) examined price manipulation in the Tether ecosystem.

⁶⁵These calculations were made in July 2019.

market share has seen swings in valuation between 3,000 and 19,000 USD during the period from December 2017 to July 2019. Gold, on the other hand, has traded in a much narrower range in the same period: between 1,180 and 1,444 USD. The wild swings in valuation in cryptocurrency prices probably make cryptocurrencies less attractive than gold (at least at this stage) for those who invest in gold as a hedge against inflation.

Simulation 2: Suppose that all money in investment gold was converted into Bitcoins. What would be the Bitcoin price? Given that there were 17.3 million Bitcoins in circulation in July 2019, under such a scenario, each Bitcoin would be worth 96,422 USD (1.668 Trillion USD/17.3 million Bitcoins.) So, when people mention \$100,000 per Bitcoin, they are probably making calculations like this one. However, that also means that if the Bitcoin price is around \$10,000, then holders of Bitcoin may believe there is a one in ten chance of Bitcoin replacing gold in the near future.

6 Competition between cryptocurrencies

While in the previous sections, we primarily focused on Bitcoin, many other cryptocurrencies were created since 2010. The market capitalization of cryptocurrency grew stunningly in the past few years. In February 2014, following the bursting of the Bitcoin “bubble” associated with the collapse of Mt. Gox, the market capitalization of all cryptocurrencies was approximately \$14 Billion. In January 2018, near Bitcoin’s peak price of \$19,000, the total market capitalization of all cryptocurrencies reached \$825 Billion.

In March 2020, total market capitalization had fallen to approximately \$150 billion and Bitcoin’s price was slightly above \$5,000. Thus despite the spectacular decline in Bitcoin’s price from its peak (of \$19,000), total market capitalization of all cryptocurrencies was still 10 times as large as it was in 2014.

6.1 Bitcoin dominance of the market

While Bitcoin dominated the market through most of the 2009–2016 period, in 2013, a few other cryptocurrencies “competed” with Bitcoin. These coins began appreciating much more quickly than Bitcoin during the price rise at the end of 2014. The prices of some of these cryptocurrencies were increasing faster than Bitcoin between during the rise in Bitcoin’s price at the end of 2013, but the dynamics of relative price movements when the price of

Bitcoin began to decline in 2014.

[Gandal and Halaburda \(2016\)](#) analyzed how network effects affected competition in the cryptocurrency market during the price spike and subsequent fall in the price of Bitcoin. They look at the competition between cryptocurrencies through the lens of competition of products with network effects, so called platform competition. They recognize that cryptocurrencies —much more than traditional currencies— may differ technologically. By employing better technology, some cryptocurrencies may be more secure, less costly to operate, or represent “higher quality” on some other dimension.

At the same time, cryptocurrencies exhibit network effects. Literature on competition with network effects points to the possibility that inferior product can dominate the market if it enjoys the advantage of network effects. Such advantage can be achieved by being earlier to the market; it is the so called first-mover advantage.

Their analysis suggests that there were strong network effects and winner-take-all dynamics following the fall in the price of Bitcoin in early 2014. From July 2014 to February 2016, Bitcoin’s value was essentially constant against the USD, while the other currencies depreciated dramatically against the USD. Litecoin, the number two coin in the market at the time, declined by 70% in value, while other “main” coins declined by more than 90% in value. In early 2016, Bitcoin accounted for 94% of the total market capitalization, while Litecoin (the number two cryptocurrency) accounted for 2%. Despite its shortcomings, Bitcoin had emerged at that point as the clear winner and beneficiary of network effects.

But Bitcoin faced increased competition over time and its dominance significantly declined from 2016 to early 2018. At the end of 2017, At that point, Bitcoin had “only” 37 percent of the total cryptocurrency market capitalization, while Ethereum had 20 percent and Ripple had another 10 percent. Ethereum has been able to challenge Bitcoin based on its extensibility — 19 of the top 20 tokens are built on top of Ethereum. Ripple, has been able to attract over 100 banks as well as Western Union to its platform. This is a stark comparison to the earlier days of Bitcoin: from its inception through 2016, Bitcoin had more than 90 percent of the market. This “low” point in Bitcoin dominance occurred at the point when Bitcoin’s value reached its peak of \$19,000. With the dramatic fall in the price of Bitcoin (and other cryptocurrencies as well, by the end of 2018, Bitcoin’s share had risen to 55 percent.

The general pattern over time has been roughly as follows: when Bitcoin’s price rises (falls,) the price of other cryptocurrencies rise (fall) more and Bitcoin’s dominance declines

(increases.) Thus, in some sense, “Bitcoin” is the safe asset the the cryptocurrency ecosystem. Bitcoin’s current share (March 2020) is 64 percent, while Ethereum and Ripple have (respectively) ten and four percent of the market. Despite the dominance of Bitcoin, twelve coins had a market capitalization of more than one Billion USD.⁶⁶

6.2 Pump and dump schemes in cryptocurrencies

As of March 2020, there were more than 700 cryptocurrencies with market capitalization between \$1 million and \$100 million. In January 2014, there were less than 30 coins with market capitalization between \$1 million and \$100 million.

This sharp rise in cryptocurrencies with “moderate” market capitalization raises concerns of an increased potential for price manipulation via “Pump-and-Dump” schemes.⁶⁷ Such schemes consist of buying suddenly large quantities so as to create a price increase momentum, and then sell it back once it has attracted sufficient traction. This strategy is more likely to succeed if a sufficiently large number of people coordinate their buying decisions in order to *pump* the price of a particular asset or currency. It has been suspected for a while that many cryptocurrencies (and especially the lesser known ones) are victims of such price manipulations.

[Hamrick et al. \(2018\)](#) present compelling evidence of pervasive pump-and-dump schemes resulting from a systematic analysis of multiple datasets. They manually collected an impressive amount of data consisting of announcements of forthcoming attacks that were broadcast on Telegram and Discord, two chat and messaging platforms. Combining this original data with cryptocurrency prices from various exchanges they identify more than 3,000 pump-and-dump schemes over a just 6 month period in 2018.

[Kamps and Kleinberg \(2018\)](#) use market data to identify suspected pump and dumps based on sudden price and volume spikes (and the following sharp decreases). They evaluate the accuracy of their predictions using a small sample of manually identified pump signals. Employing a similar approach with a different dataset, [Mirtaheeri et al. \(2019\)](#) use data

⁶⁶There is a growing empirical literature that examines the performance of cryptocurrencies as financial assets. [Corbet et al. \(2019\)](#) offers a thorough review of more than ninety, mostly empirical, papers on this topic.

⁶⁷Like Pump and Dump Schemes, price manipulation is a form of market manipulation as well. [Gandal et al. \(2018\)](#) , which we discussed in Section 5, showed that there was price manipulation during the 2013 bubble.

collected from Twitter on cryptocurrencies cross-referenced with pump signal data from Telegram and market data. They note that a lot of the tweets are automated and attempt to predict pumps using only the Twitter traffic. [Xu and Livshits \(2018\)](#) use data on just over 200 pump signals to build a model to predict which coins will be pumped. Their model distinguishes between highly successful pumps and all other trading activity on the exchange. [Li et al. \(2018\)](#) use a difference-in-difference model to show that pump and dumps lower the trading price of affected coins.⁶⁸

6.3 Other aspects of the ecosystem: cryptocurrency exchange markets

The exchange markets play an important role in the cryptocurrency ecosystem, and the industry is very dynamic. Ironically, Bitcoin was designed to be decentralized/trustless/etc., but now a huge percentage of the transaction volume occurs through financial intermediaries!

Market leaders rise and fall at a very fast pace. Although reliable data on trading volume are hard to obtain (see below,) the trend (in trading volume) has been a movement from a very concentrated industry to a very competitive one over time.

From 2011 until the end of 2013, Mt Gox was the dominant exchange. It was essentially a monopoly in 2011, and held more than 80% of the market in that year. In 2012, it held more than 60% of the market. After the Mt Gox collapse in late 2013, other market leaders emerged.

In February 2014, the three major exchanges were BTC-e (30 percent market share), Bitstamp (28 percent market share), and Bitfinex (26 percent market share). Together these three exchanges had 84 percent of the market. (This excludes the volatile Chinese exchanges, where verifying trading volume was difficult.)

Fast forward to 2018. The exchange market was more competitive with the leader Binance holding 18 percent of the market and the next two exchanges (OKEx and Huobi) holding 14 and 10 percent respectively. Overall, the top ten exchanges held slightly less than 70 percent of the market. (These data come from Coinmarketcap, a website that provides daily data on cryptocurrency prices, market capitalization, and trading volumes by exchange.) The three

⁶⁸There have been media articles about the pump and dump phenomenon as well. [Mac \(2018\)](#) reported on pump and dump schemes in a BuzzFeed article published in January 2018. This was followed by work by [Shifflett and Vigna \(2018\)](#) in a Wall Street Journal article published in August 2018.

leaders were the same in August 2019, but their combined market share was only 13 percent (vs. 42 percent in 2018.) Overall the top 10 exchanges held 37 percent of the market (vs. nearly 70 percent in 2018.)

This discussion is subject to a caveat, since recently (March 2019,) two analyses suggested that much of the cryptocurrency exchanges' reported trade volumes may be grossly exaggerated. Bitwise Asset Management, wrote in a submission to the US SEC that 95 percent of the cryptocurrency exchanges' reported volumes is "suspect."⁶⁹ Nevertheless, the trend seems to be towards a more competitive exchange market. It will be interesting to see how the dynamics proceed in the cryptocurrency exchange markets in light of greater scrutiny and the introduction of regulatory policies.

7 Conclusion

Cryptocurrencies are new to the world and have generated considerable trading volume along with sizeable costs—in terms of energy resources—to support their availability. As with any new good or service, the microeconomics research task is to understand its supply (i.e., what technological properties allow it to operate), its demand (i.e., to what uses are agents putting it), its value (i.e., what determines its trading price in the market) and the nature of competition (i.e., how strong is substitution between different varieties of the new good or service and others with similar functionality). This paper has examined research into cryptocurrencies from this perspective. While broad trends and understanding has emerged, it is also clear that the market continues to evolve and its precise place within the broader economy is yet to be established.

What is interesting about cryptocurrencies like Bitcoin is that they have emerged in ways that many believed was not possible. While it relied on a few important advances in cryptography, even though it uses modern microprocessors, there were few components that required large advances in technologies or substantial reductions in costs of supply. Instead, the protocols that govern Bitcoin are built on the back of longstanding contributions in computer science and could be said to scale with existing technology rather than rely upon it. For instance, had computing power been more expensive, then fewer computations would have been done in support of the Bitcoin network but otherwise the same broad functionality would have been possible. In our analysis here, most research has been to explain how Bitcoin

⁶⁹<https://www.sec.gov/comments/sr-nysearca-2019-01/srnysearca201901-5164833-183434.pdf>

works but the answer to that question has not, to our knowledge, really been predicated in the historic context for the emergence of a new technology. We speculate that it is this disconnect that may govern future research into this topic over the next decade.

References

- Abadi, J. and Brunnermeier, M. (2018). Blockchain economics. Technical report, National Bureau of Economic Research.
- Amoussou-Guenou, Y., Biais, B., Potop-Butucaru, M., and Tucci-Piergiovanni, S. (2019). Rationals vs byzantines in consensus-based blockchains. *arXiv preprint arXiv:1902.07895*.
- Andolfatto, D. (2018). Blockchain: What it is, what it does, and why you probably don't need one. *Federal Reserve Bank of St. Louis Review*, 100(2).
- Andolfatto, D. and Spewak, A. (2019). Whither the price of bitcoin? *Federal Reserve Bank of St. Louis Economic Synopsis*, No. 1.
- Androulaki, E., Karame, G. O., Roeschlin, M., Scherer, T., and Capkun, S. (2013). Evaluating user privacy in bitcoin. In *International Conference on Financial Cryptography and Data Security*, pages 34–51. Springer.
- Arruñada, B. and Garicano, L. (2018). Blockchain: The birth of decentralized governance. *Pompeu Fabra University, Economics and Business Working Paper Series*, 1608.
- Athey, S., Parashkevov, I., Sarukkai, V., and Xia, J. (2016). Bitcoin pricing, adoption, and usage: Theory and evidence. *mimeo*.
- Babaioff, M., Dobzinski, S., Oren, S., and Zohar, A. (2012). On bitcoin and red balloons. In *Proceedings of the 13th ACM conference on electronic commerce*, pages 56–73. ACM.
- Badev, A. and Chen, M. (2015). *Bitcoin: technical background and data analysis*. Lulu. com.
- Barrera, C. and Hurder, S. (2018). Blockchain upgrade as a coordination game. *Available at SSRN 3192208*.
- Biais, B., Bisiere, C., Bouvard, M., and Casamatta, C. (2019). The blockchain folk theorem. *The Review of Financial Studies*, 32(5):1662–1715.
- Biais, B., Bisière, C., Bouvard, M., Casamatta, C., and Menkveld, A. J. (2018). Equilibrium bitcoin pricing. *Available at SSRN*.
- Böhme, R., Christin, N., Edelman, B., and Moore, T. (2015). Bitcoin: Economics, technology, and governance. *Journal of Economic Perspectives*, 29(2):213–38.

- Bolt, W. and Van Oordt, M. R. (2020). On the value of virtual currencies. *Journal of Money, Credit and Banking*, 52(4):835–862.
- Bonneau, J., Miller, A., Clark, J., Narayanan, A., Kroll, J. A., and Felten, E. W. (2015). Sok: Research perspectives and challenges for bitcoin and cryptocurrencies. In *2015 IEEE symposium on security and privacy*, pages 104–121. IEEE.
- Brewer, E. A. (2000). Towards robust distributed systems. In *PODC*, volume 7. Portland, OR.
- Buchanan, J. M., Tollison, R. D., and Tullock, G. (1980). *Toward a theory of the rent-seeking society*. Number 4. Texas A & M Univ Pr.
- Budish, E. (2018). The economic limits of bitcoin and the blockchain. *mimeo, University of Chicago Booth School of Business*.
- Buterin, V. (2013). What proof of stake is and why it matters. *Bitcoin Magazine*, 26.
- Campbell, J. Y. and Shiller, R. J. (1988). The dividend-price ratio and expectations of future dividends and discount factors. *The Review of Financial Studies*, 1(3):195–228.
- Carlsten, M., Kalodner, H. A., Weinberg, S. M., and Narayanan, A. (2016). On the instability of bitcoin without the block reward. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*, pages 154–167.
- Catalini, C. and Gans, J. S. (2020). Some simple economics of the blockchain. *Proceedings of the ACM*.
- Cheung, A., Roca, E., and Su, J.-J. (2015). Crypto-currency bubbles: an application of the phillips–shi–yu (2013) methodology on mt. gox bitcoin prices. *Applied Economics*, 47(23):2348–2358.
- Chiu, J. and Koepl, T. V. (2017). The economics of cryptocurrencies–bitcoin and beyond. *Available at SSRN 3048124*.
- Christin, N. (2013). Traveling the silk road: A measurement analysis of a large anonymous online marketplace. In *Proceedings of the 22nd international conference on World Wide Web*, pages 213–224. ACM.

- Cong, L. W., He, Z., and Li, J. (2019). Decentralized mining in centralized pools. *The Review of Financial Studies*.
- Corbet, S., Lucey, B., Urquhart, A., and Yarovaya, L. (2019). Cryptocurrencies as a financial asset: A systematic analysis. *International Review of Financial Analysis*, 62:182–199.
- De Vries, A. (2018). Bitcoin’s growing energy problem. *Joule*, 2(5):801–805.
- Dwork, C. and Naor, M. (1992). Pricing via processing or combatting junk mail. In *Annual International Cryptology Conference*, pages 139–147. Springer.
- Easley, D., O’Hara, M., and Basu, S. (2019). From mining to markets: The evolution of bitcoin transaction fees. *Journal of Financial Economics*, 134(1):91–109.
- Eyal, I. and Sirer, E. G. (2014). Majority is not enough: Bitcoin mining is vulnerable. In *International conference on financial cryptography and data security*, pages 436–454. Springer.
- Fischer, M. J., Lynch, N. A., and Paterson, M. S. (1985). Impossibility of distributed consensus with one faulty process. *Journal of the ACM (JACM)*, 32(2):374–382.
- Foley, S., Karlsen, J. R., and Putniņš, T. J. (2019). Sex, drugs, and bitcoin: How much illegal activity is financed through cryptocurrencies? *The Review of Financial Studies*, 32(5):1798–1853.
- Gandal, N. and Halaburda, H. (2016). Can we predict the winner in a market with network effects? Competition in cryptocurrency market. *Games*, 7(3):16.
- Gandal, N., Hamrick, J., Moore, T., and Obermann, T. (2018). Price manipulation in the Bitcoin ecosystem. *Journal of Monetary Economics*, 95:86–96.
- Gans, J. S. and Gandal, N. (2019). More (or less) economic limits of the blockchain. Technical report, National Bureau of Economic Research.
- Garcia, D., Tessone, C. J., Mavrodiev, P., and Perony, N. (2014). The digital traces of bubbles: feedback cycles between socio-economic signals in the bitcoin economy. *Journal of the Royal Society Interface*, 11(99):20140623.

- Glaser, F., Zimmermann, K., Haferkorn, M., Weber, M. C., and Siering, M. (2014). Bitcoin-asset or currency? revealing users' hidden intentions. In *ECIS 2014 Tel Aviv*.
- Griffin, J. M. and Shams, A. (2020). Is bitcoin really un-tethered? *Journal of Finance*, 55:1913–64.
- Haber, S. and Stornetta, W. S. (1990). How to time-stamp a digital document. In *Conference on the Theory and Application of Cryptography*, pages 437–455. Springer.
- Haeringer, G. and Halaburda, H. (2018). Bitcoin: A revolution? In *Economic analysis of the digital revolution. J. Ganuza and G. Llobert, Eds.*, pages 397–421. FUNCAS.
- Halaburda, H. (2018). Blockchain revolution without the blockchain? *Communications of the ACM*, 61(7):27–29.
- Halaburda, H. and Sarvary, M. (2016). Beyond bitcoin. *The Economics of Digital Currencies*.
- Hamrick, J., Rouhi, F., Mukherjee, A., Feder, A., Gandal, N., Moore, T., and Vasek, M. (2018). An examination of the cryptocurrency pump and dump ecosystem. *Available at SSRN 3303365*.
- Henry, C. S., Huynh, K. P., Nicholls, G., et al. (2017). Bitcoin awareness and usage in canada. *Bank of Canada, staff Working Paper 2017-56*.
- Huberman, G., Leshno, J., and Moallemi, C. C. (2019). An economic analysis of the bitcoin payment system. *Columbia Business School Research Paper*, 17-92.
- Kaminski, J. (2014). Nowcasting the bitcoin market with twitter signals. *arXiv preprint arXiv:1406.7577*.
- Kamps, J. and Kleinberg, B. (2018). To the moon: defining and detecting cryptocurrency pump-and-dumps. *Crime Science*, 7(1):18.
- Kiayias, A., Koutsoupias, E., Kyropoulou, M., and Tselekounis, Y. (2016). Blockchain mining games. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 365–382. ACM.
- Kocherlakota, N. R. (1998). Money is memory. *Journal of Economic Theory*, 2(81):232–251.

- Kristoufek, L. (2015). What are the main drivers of the bitcoin price? evidence from wavelet coherence analysis. *PloS one*, 10(4):e0123923.
- Kroll, J. A., Davey, I. C., and Felten, E. W. (2013). The economics of bitcoin mining, or bitcoin in the presence of adversaries. In *Proceedings of WEIS*, volume 2013.
- Leshno, J. D. and Strack, P. (2020). Bitcoin: An axiomatic approach and an impossibility theorem. *American Economic Review: Insights*, forthcoming.
- Li, T., Shin, D., and Wang, B. (2018). Cryptocurrency pump-and-dump schemes. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3267041.
- Ma, J., Gans, J. S., and Tourky, R. (2018). Market structure in bitcoin mining. Technical report, National Bureau of Economic Research.
- Mac, R. (2018). Bitcoin scammers are using this app to fleece people. *Buzzfeed News*. Available at: <https://www.buzzfeednews.com/article/ryanmac/cryptocurrency-scammers-are-running-wild-on-telegram>.
- Mai, F., Bai, Q., Shan, Z., Wang, X. S., and Chiang, R. H. (2015). The impacts of social media on bitcoin performance. In *Proceedings of the 2015 International Conference on Information Systems*.
- Makorov, I. and Schoar, A. (2020). Trading and arbitrage in cryptocurrency markets. *The Journal of Financial Economics*, 135(2):293–319.
- Mirtaheri, M., Abu-El-Haija, S., Morstatter, F., Steeg, G. V., and Galstyan, A. (2019). Identifying and analyzing cryptocurrency manipulations in social media. Available at <https://arxiv.org/abs/1902.03110>.
- Moroz, D. J., Aronoff, D. J., Narula, N., and Parkes, D. C. (2020). Double-spend counter-attacks: Threat of retaliation in proof-of-work systems. *CoRR abs/2002.10736*.
- Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system.
- Narayanan, A., Bonneau, J., Felten, E., Miller, A., and Goldfeder, S. (2016). *Bitcoin and cryptocurrency technologies: a comprehensive introduction*. Princeton University Press.

- Phillips, P. C., Shi, S., and Yu, J. (2015). Testing for multiple bubbles: Historical episodes of exuberance and collapse in the S&P 500. *International Economic Review*, 56(4):1043–1078.
- Polasik, M., Piotrowska, A. I., Wisniewski, T. P., Kotkowski, R., and Lightfoot, G. (2015). Price fluctuations and the use of bitcoin: An empirical inquiry. *International Journal of Electronic Commerce*, 20(1):9–49.
- Prat, J. and Walter, B. (2018). An equilibrium model of the market for bitcoin mining.
- Reid, F. and Harrigan, M. (2013). An analysis of anonymity in the bitcoin system. In *Security and privacy in social networks*, pages 197–223. Springer.
- Saleh, F. (2019). Blockchain without waste: Proof-of-stake. *mimeo, McGill University*.
- Sapirshstein, A., Sompolinsky, Y., and Zohar, A. (2016). Optimal selfish mining strategies in bitcoin. In *International Conference on Financial Cryptography and Data Security*, pages 515–532. Springer.
- Schuh, S. and Shy, O. (2016). US consumers’ adoption and use of bitcoin and other virtual currencies. In *DeNederlandsche bank, Conference entitled “Retail payments: mapping out the road ahead*.
- Shifflett, S. and Vigna, P. (2018). Some traders are talking up cryptocurrencies, then dumping them, costing others millions. *The Wall Street Journal*. Available at: <https://www.wsj.com/graphics/cryptocurrency-schemes-generate-big-coin/>.
- Shiller, R. J. (1981). The use of volatility measures in assessing market efficiency. *The Journal of Finance*, 36(2):291–304.
- Tasca, P., Hayes, A., and Liu, S. (2018). The evolution of the bitcoin economy. *The Journal of Risk Finance*.
- Urquhart, A. (2018). What causes the attention of bitcoin? *Economics Letters*, 166:40–44.
- Velde, F. R. (2013). Bitcoin: A primer, the federal reserve bank of chicago. *Chicago Fed Letter (Dec. 2013)*.
- Wang, S. and Vergne, J.-P. (2017). Buzz factor or innovation potential: What explains cryptocurrencies’ returns? *PloS one*, 12(1):e0169556.

Xu, J. and Livshits, B. (2018). The anatomy of a cryptocurrency pump-and-dump scheme.
Available at: <https://arxiv.org/abs/1811.10109>.

Yermack, D. (2013). Is bitcoin a real currency? an economic appraisal. In *Handbook of Digital Currency*, pages 31–43. Springer.