



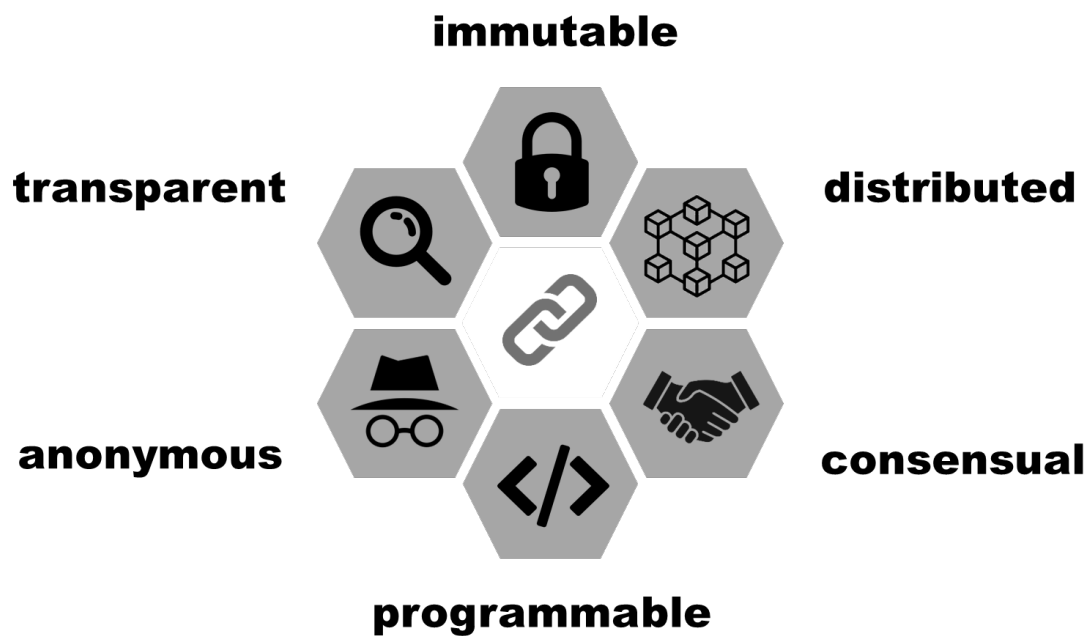
Final report dated 27 September 2021

---

# Blockchain energy consumption

## An exploratory study

---





# **ETH**

Eidgenössische Technische Hochschule Zürich  
Swiss Federal Institute of Technology Zurich

**Date:** 27 September 2021

**Location:** Bern

**Publisher:**

Swiss Federal Office of Energy SFOE  
Energy Research and Cleantech  
CH-3003 Bern  
[www.bfe.admin.ch](http://www.bfe.admin.ch)

**Co-financing:**

ETH Zürich  
Universitätsstr. 6  
8092 Zürich  
[www.ethz.ch](http://www.ethz.ch)

**Subsidy recipients:**

ETH Zürich  
Universitätsstr. 6  
8092 Zürich  
[www.ethz.ch](http://www.ethz.ch)

**Author:**

Vlad Coroamă, ETH Zürich, [vcoroama@ethz.ch](mailto:vcoroama@ethz.ch)

**SFOE project coordinators:**

Dr. Michael Moser, [michael.moser@bfe.admin.ch](mailto:michael.moser@bfe.admin.ch)  
Roland Brüniger, [roland.brueeniger@brueniger.swiss](mailto:roland.brueeniger@brueniger.swiss)

**SFOE contract number:** SI/502164-01

The author bears the entire responsibility for the content of this report and for the conclusions drawn therefrom.



## Zusammenfassung

Die Grundsätze der ersten Kryptowährung Bitcoin samt deren zugrundeliegenden *Blockchain-Technologie* wurden 2008 veröffentlicht. Blockchains bauen auf mehreren bereits bestehenden Informatikkonzepten auf (z. B. kryptografische Hashes, Hashpointer sowie Konsensmechanismen in verteilten Systemen). Sie kombinieren und ergänzen diese auf innovative Weise, um sichere und unveränderliche Transaktionen ohne eine vertrauenswürdige zentrale Autorität zu ermöglichen. Schon bald nach ihrer Einführung wurde klar, dass die Erweiterung von Blockchains um die Fähigkeit, kleine Codestücke, sogenannte *Smart Contracts*, automatisch auszuführen, ihre potenziellen Anwendungsbereiche weit über Kryptowährungen hinaus erweitert.

Während Blockchain-Technologien anfangs ein Begriff für Spezialisten waren, haben sie in jüngster Zeit aufgrund des Kursanstiegs mehrerer Kryptowährungen, aber auch aufgrund ihrer Fähigkeit, Eigentum durch eindeutige Token (Englisch: non-fungible tokens, NFTs), einer besonderen Art von Smart Contracts, nachzuweisen und zu übertragen, viel Aufmerksamkeit in der Öffentlichkeit und den Medien erhalten. Gleichzeitig wurde schnell klar, dass der vertrauenslose Konsensmechanismus, der für die Registrierung neuer Transaktionen und die Erzeugung neuer Kryptowährung erforderlich ist, das sogenannte Mining, eine große, ständig wachsende und möglicherweise nicht nachhaltige Menge an Energie benötigt.

In Anbetracht wachsender Besorgnis über diesen Energieverbrauch, aber auch angesichts der sich diversifizierenden Anwendungsbereiche der Blockchain-Technologie, wurde in dieser Studie versucht, die verschiedenen Faktoren zu analysieren, die den Energieverbrauch einer Blockchain beeinflussen, sowie die besten Hebel zu identifizieren, um diesen Energieverbrauch zu senken. Die Ergebnisse bestätigen, dass der vertrauenslose Konsensmechanismus, der auf dem so genannten *Proof-of-Work* basiert, den Energieverbrauch einer Blockchain mit Abstand dominiert: Während er für über 100 TWh jährlich verantwortlich sein kann, benötigt die Speicherung der Blockchain über 4-6 Größenordnungen weniger Energie (30 MWh – 3 GWh) und die über das Internet verschickten Koordinationsnachrichten über 7 Größenordnungen weniger (6 MWh). Energiesparmaßnahmen müssen daher auf den Proof-of-Work-Konsensmechanismus abzielen: Einzelne Blockchains können zu alternativen Konsensmechanismen wechseln, die sich nicht auf Proof-of-Work stützen, während Unternehmens- und öffentliche Maßnahmen darauf abzielen können, Proof-of-Work-basierte Blockchains zu entmutigen und die Akzeptanz von Blockchains mit alternativen Konsensmechanismen zu fördern.



## Résumé

Les principes de la première crypto-monnaie, le bitcoin, ont été publiés en 2009 et, parallèlement, la *technologie blockchain* a été développée. Les blockchains s'appuient sur plusieurs concepts informatiques préexistants (tels que les hachages cryptographiques, les pointeurs de hachage et les mécanismes de consensus dans des systèmes distribués), en les combinant et en les complétant de manière innovante pour fournir des transactions sécurisées et immuables sans avoir besoin d'une autorité centrale de confiance. Rapidement après leur création, il est devenu évident que l'extension des blockchains avec la capacité d'exécuter automatiquement de petits morceaux de code appelés *contrats intelligents*, étend leurs domaines d'application potentiels bien au-delà des crypto-monnaies.

Alors qu'il s'agissait au départ d'une notion réservée aux spécialistes, les technologies de la blockchain ont récemment fait l'objet d'une grande attention de la part du public et des médias en raison de la flambée des prix de plusieurs crypto-monnaies, mais aussi en raison de leur capacité à prouver et à transférer la propriété par des jetons non fongibles, un type particulier de contrats intelligents. Dans le même temps, il est rapidement apparu que le mécanisme de consensus sans confiance nécessaire à l'enregistrement de nouvelles transactions et à la production de nouvelles crypto-monnaies, appelé *minage*, nécessite une quantité d'énergie importante, en constante augmentation et potentiellement insoutenable.

Compte tenu des préoccupations croissantes concernant cette consommation d'énergie, mais aussi de la diversification des domaines d'application de la technologie blockchain, cette étude a pour but d'analyser les différents facteurs qui affectent la consommation d'énergie d'une blockchain et d'identifier les meilleurs leviers pour atténuer cette consommation d'énergie en conséquence. Les résultats confirment que le mécanisme de consensus sans confiance basé sur ce qu'on appelle la *preuve de travail* domine de loin la consommation d'énergie d'une blockchain : Alors qu'il peut être responsable de plus de 100 TWh par an, le stockage de la blockchain nécessite plus de 4 à 6 ordres de grandeur d'énergie en moins (30 MWh – 3 GWh) et les messages de coordination envoyés sur Internet plus de 7 ordres de grandeur en moins (6 MWh). Les mesures d'économie d'énergie doivent donc porter sur le mécanisme de consensus par preuve de travail : les blockchains individuelles peuvent passer à des mécanismes de consensus alternatifs qui ne reposent pas sur la preuve de travail, tandis que les politiques des entreprises et des pouvoirs publics peuvent viser à décourager les blockchains basées sur la preuve de travail et à encourager l'adoption de blockchains dotées de mécanismes de consensus alternatifs.



## Summary

The principles of the first cryptocurrency Bitcoin were published in 2008; along with them the enabling *blockchain technology* was developed. Blockchains build on several pre-existing computing concepts (such as cryptographic hashes, hashpointers, and consensus mechanisms in distributed systems), combining and adding to them in an innovative way to provide secure and immutable transactions without the need of a trusted central authority. Quickly after their inception, it became clear that extending blockchains with the capability to automatically execute small pieces of code called *smart contracts*, extends their potential application domains far beyond cryptocurrencies.

While in the beginning a notion for specialists, blockchain technologies have recently received much public and media attention due to price surges of several cryptocurrencies, but also due to their ability to prove and transfer ownership through non-fungible tokens (NFTs), a special type of smart contracts. At the same time, it quickly became clear that the trustless consensus mechanism needed to register new transactions and produce new cryptocurrency, called *mining*, needs a large, continuously growing, and possibly unsustainable amount of energy.

Given growing concerns about this energy consumption, but also the diversifying application domains for blockchain technology, this study set out to analyse the different factors that affect the energy consumption of one blockchain, and to identify the best levers to mitigate this energy consumption as a consequence. The results confirm that the trustless consensus mechanism based on what is called *proof-of-work* dominates the energy consumption of a blockchain by a margin: While it can be responsible for over 100 TWh per year, the storage of the blockchain requires over 4-6 orders of magnitude less energy (30 MWh – 3 GWh) and the coordination messages sent across the Internet over 7 orders of magnitude less (6 MWh). Energy conservation measures must thus address the proof-of-work consensus mechanism: Individual blockchains can switch to alternative consensus mechanisms that do not rely on proof-of-work, while company and public policies can aim at discouraging proof-of-work-based blockchains, and encourage the uptake of blockchains with alternative consensus mechanisms.

## Main findings

- Beyond cryptocurrencies and some hyped but unrealistic applications, several blockchain application domains are already established or potentially meaningful. They include: immutable proof and transfer of (digital or physical) asset ownership, autonomous governance and decentralised finance, revenue hedging and fostering investment in renewable electricity markets.
- In a blockchain deploying a proof-of-work consensus mechanism, the energy required by this mechanism dominates the blockchain's energy consumption with currently more than 100 TWh per year for the most popular blockchain. Meanwhile, the other components can be ignored, as the storage requires 33 MWh – 3 GWh yearly, and the coordination messages only 6 MWh; both many orders of magnitude less.
- The energy consumption of the proof-of-work mechanism is independent of the energy efficiency of the mining hardware, and only determined by the price of the cryptocurrency. Any efficiency gains of the mining hardware will be eaten up by more mining and resulting harder cryptographic puzzles. There is thus no theoretical limit to the energy consumption used for mining.
- Energy conservation measures must address the proof-of-work consensus mechanism: Individual blockchains can switch to alternative consensus mechanisms, while company and public policies can aim at discouraging proof-of-work-based blockchains, and encourage the uptake of blockchains with alternative consensus mechanisms.



# Contents

<b>Zusammenfassung</b> .....	<b>3</b>
<b>Résumé</b> .....	<b>4</b>
<b>Summary</b> .....	<b>5</b>
<b>Main findings</b> .....	<b>5</b>
<b>Contents</b> .....	<b>6</b>
<b>Abbreviations</b> .....	<b>8</b>
<b>1 Introduction</b> .....	<b>9</b>
1.1 Background information and current situation.....	9
1.2 Purpose of the study and structure of the report.....	9
<b>2 Blockchain technology and its cryptocurrency origins</b> .....	<b>10</b>
2.1 Peculiarities of a cryptocurrency.....	10
2.2 Distributed ledger .....	10
2.3 Proof-of-Work consensus mechanism, miners, and blockchain structure .....	11
2.4 Alternative consensus mechanisms .....	13
2.5 Permissioned and permissionless chains .....	13
<b>3 Further blockchain developments: Smart contracts and non-fungible tokens (NFTs)</b> ...	<b>14</b>
3.1 Smart contracts .....	14
3.2 Real-world connection and the Oracle problem .....	15
3.3 Non-fungible tokens (NFTs) .....	16
<b>4 Application domains beyond cryptocurrencies</b> .....	<b>17</b>
4.1 Crypto art and other digital collectibles .....	17
4.2 Traditional art and other physical assets.....	18
4.3 Company ownership.....	20
4.4 Autonomous governance and the first decentralised autonomous organisation (DAO) .....	21
4.5 More advanced DAOs and the emergence of decentralised finance (DeFi) .....	22
4.6 Disputed further domains .....	23
4.7 Electricity markets .....	24
<b>5 The energy consumption of a (permissionless) blockchain</b> .....	<b>25</b>
5.1 Energy required for storing the distributed virtual machine (distributed ledger).....	25
5.2 Energy required by the computational complexity of consensus mechanisms.....	27
5.2.1 Computing an upper bound for the PoW energy consumption .....	27
5.2.2 Further estimates for Bitcoin's PoW energy consumption .....	30
5.2.3 Energy for the computational complexity of other consensus mechanisms .....	31
5.3 Energy due to the communication of consensus mechanisms .....	31
5.4 Comparative analysis .....	32



<b>6</b>	<b>Discussion.....</b>	<b>32</b>
6.1	Main levers that influence the energy consumption of blockchains .....	32
6.2	Energy needed by the Proof-of-Work mechanisms .....	32
<b>7</b>	<b>Conclusions and future research .....</b>	<b>34</b>
7.1	Summary of the analysis .....	34
7.2	Policies and technological measures to discourage PoW.....	34
7.3	Future research .....	35
	<b>Acknowledgements.....</b>	<b>35</b>
	<b>References .....</b>	<b>35</b>



## Abbreviations

BTC	Bitcoin, one unit of the currency of the homonymous blockchain
CHF	Swiss francs
DAO	decentralised autonomous organisation
DC	data centre
DeFi	decentralised finance
DLT	distributed ledger technology
EH	exahash, $10^{18}$ hashes
ETH	Ether, one unit of the Ethereum blockchain
EVM	Ethereum virtual machine
GB	gigabyte
GH	gigahash
HDD	hard-drive disk
IS	information system
IoT	Internet of things
kWh	kilowatt-hours
MB	megabyte
NFT	non-fungible token
PKI	public key infrastructure
PoS	proof-of-stake
PoW	proof-of-work
PUE	power usage effectiveness
RAM	random access memory
SSD	solid-state disk
TA	trusted authority
USD	US dollars
VM	virtual machine





# 1 Introduction

Building on several pre-existing computing concepts (such as cryptographic hashes, hashpointers and Merkle trees, consensus mechanisms in distributed systems), combining and adding to them in an innovative way, the person or group of persons pseudonomously known as Satoshi Nakamoto introduced in 2008 the principles of Bitcoin, the first successful cryptocurrency. This weaving of existing and newly developed technological components into a new technology has come to be known as 'blockchain'.

## 1.1 Background information and current situation

As enabling technology, blockchain allowed Bitcoin to become the first digital currency to be secure and prevent multiple spending while at the same time not requiring a central authority trusted by all players. As consequence, Bitcoin was able to easily scale to a worldwide presence, to work well as tool for value preservation and transfer, and to continuously gain momentum. Despite being entirely distributed and without a central authority, its transactions are immutable, transparent, and for anyone verifiable.

It quickly became clear, however, that its characteristics make the blockchain technology suitable for many more application domains beyond cryptocurrencies. In particular, it can be deployed not only to make financial transfers and cryptocurrency ownership immutable and verifiable, but ownership of digital and physical assets more generally; any notarial action can be stored by a blockchain in a similar way to the storage of cryptocurrency. Moreover, in a similar way they can reflect changes of ownership, blockchains can encode state changes more broadly, and even encode the conditions that can automatically trigger those state changes. This feature is known as 'smart contracts' and can arguably enable much more transformational consequences of blockchains than the original cryptocurrency application domain.

At the same time, however, it also increasingly became evident that the very feature setting blockchain apart (i.e., the trust and security it enables without the need of a central trusted authority), has been bought at a large, continuously growing, and potentially unsustainable environmental price. The original mechanism that guarantees the trustworthiness and immutability of its transactions (or state changes, more generally), called proof-of-work, need a large amount of (distributed) energy consumption. What is more, the more successful a blockchain system becomes, the larger this energy footprint becomes.

## 1.2 Purpose of the study and structure of the report

In this context, the current study aims to shed light on two parallel, but interconnected, dimensions: the likely (current and future) blockchain application domains, and the technological drivers behind a blockchain's energy consumption. The combination of the two – adoption rate and the specific energy consumptions of blockchain sub-technologies – represents a qualitative analysis for the expected magnitude of the overall energy consumption of blockchains.

The report is organised as follows: Section 2 represents an introduction to the early stages of blockchain technology, Bitcoin as distributed ledger, the security-guaranteeing proof-of-work consensus mechanisms, the reasons for its existence and possible alternatives, and the distinction between permissionless and permissioned blockchains. Section 3 introduces further developments of blockchain, specifically smart contracts, the Oracle problem they entail, and non-fungible tokens. Given these technological components of blockchains and their features, Section 4 discusses several blockchain application domains beyond cryptocurrencies, and argues for those we consider realistic amidst the hype surrounding blockchain. Section 5 assesses the overall energy consumption of a blockchain, by analysing its individual sources (i.e., distributed storage of the blockchain, computation



needed by the consensus mechanism, and messages sent across the network). Building on these results, Section 6 compares their magnitudes showing that the proof-of-work consensus mechanism dominates a blockchain's energy consumption, and argues that there are no economic mechanisms that will necessarily limit its energy consumption. Finally, Section 7 discusses the areas in need of more research and perhaps of regulating policies.

## 2 Blockchain technology and its cryptocurrency origins

### 2.1 Peculiarities of a cryptocurrency

As with any digital data, the value units (i.e., 'coins') of electronic currencies can be trivially multiplied, stored, and transmitted. This issue is known as double spending or multiple spending. Simply having coins certified by a trusted authority (TA) – such as a bank – through a public-key infrastructure (PKI), does not address the multiple spending problem: Not having access to the private key of the TA, an attacker indeed cannot generate new coins; however, once owing a legit coin, it could be multiplied and re-spent any number of times.

Electronic currencies thus need a unanimously trusted ledger, which can certify the ownership structure of the electronic currency and in particular ownership changes, i.e., value transfers (Sedlmeir et al. 2020a). By registering each transaction and reflecting current ownership rights, such a ledger prevents the multiple spending problem addressed above. In early attempts at electronic currencies, this ledger was often held by the same central TA that also issued the electronic currency, even for anonymous electronic currencies such as the well-known eCash proposed by David Chaum in 1983 (Chaum 1983) and used by several banks for micropayments throughout the 1990s.

### 2.2 Distributed ledger

While preventing multiple spending, such centralised solutions also have drawbacks: They rely on the safety of the TA, which often uses proprietary algorithms. The TA can become a bottleneck, preventing scalability. In particular, as argued by the person or group of persons pseudonomously known as Satoshi Nakamoto, all transactions can be contested and must be reversible, and the cost of this mediation increases the transaction costs, yielding centralised solutions not feasible “for small casual transactions” (Nakamoto 2008) – given today's Bitcoin prices, the irony is, of course, strong here.

To enable a digital currency that does not require a central TA and yet provides a secure solution against multiple spending, the same (Nakamoto 2008) introduced the principles of the digital currency (often called 'cryptocurrency') *Bitcoin* and together with it the *blockchain technology*. Bitcoin is “a decentralized payment system in which all participating computers (“nodes”) store a copy – or, more precisely, a replica, since there is no distinguished master – of the associated ledger” (Sedlmeir et al. 2020b); this is why the technology is often called 'distributed ledger technology' (DLT). The coins of the Bitcoin system have the homonymous name Bitcoin (BTC). All BTCs that are issued are registered as *transactions* on the distributed ledger; each payment (i.e., transfer of funds) is also registered as a transaction on the distributed ledger. There is thus no electronic 'ownership' of any bits comprising the (digitally signed) coins, as in earlier digital currency systems such as eCash (Chaum 1983). Instead, BTCs (or subdivisions thereof) are publicly associated on the ledger with specific account numbers; the 'wallets' so often advertised by companies mediating access to cryptocurrency do not hold any currency per se; they only hold the private key of the account (accounts are based on a public-key-infrastructure, PKI), which proves that the key holder is the rightful account owner.



As there is no central authority, the distributed solution needs a *consensus mechanism* by which individual nodes can agree which are the valid transactions of the last period. Reaching a consensus among distributed nodes with no central authority, some of which might fail, others of which might disagree, is not a new problem in Computer Science – it is a known topic for fault-tolerant distributed systems. It has been theorised in the early 1980s by later Turing Award winner Leslie Lamport as the ‘byzantine generals’ problem (Lamport 1983). Several solutions for the Byzantine generals problem exist, such as one inspired by the antique ‘part-time parliament’ from the Greek island of Paxos, where a consistent state of the parliamentary resolutions was kept, despite the fact that the honorary members of parliament were mostly going about their trading businesses and often missing parliamentary sessions (Lamport 1998).

The Byzantine generals problem, however, assumes distributed nodes that might fail or disagree, but which belong to the same closed system and are thus benevolent and intend to reach the ‘correct’ consensus. In a distributed system with free access and no TAs whatsoever, which aims at establishing a new type of electronic currency, such as the one envisioned by (Nakamoto 2008), the potential for malevolence is obvious. The problem is hardened by the fact that the number of nodes is not restricted (see discussion on such *permissionless blockchains* versus *permissioned blockchains* below). If there are no – or only marginal – costs for taking part in the consensus mechanisms, a so-called Sybil attack (Douceur 2002) becomes possible: The attacker manipulates the outcome by creating a large number of nodes under pseudonyms, which all put forward the same result desired by the attacker, thus manipulating the outcome of the consensus mechanism.

## 2.3 Proof-of-Work consensus mechanism, miners, and blockchain structure

To address the danger of Sybil attacks, a system without access restrictions has to bind the relative weight of a vote to a limited resource (Sedlmeir et al. 2020a). For Bitcoin, as well as other later cryptocurrencies, this is achieved via a mechanism known as Proof-of-Work (PoW), which functions as follows: Account holders who want to transfer BTCs to another account, submit their desired transaction to a pool of waiting transactions. Specialised nodes within the blockchain network (which will be discussed in more detail shortly) bundle such transactions submitted to the pool into a *block* of transactions. They then add an aleatory number (called *nonce*) to the block, and perform a one-way cryptographic function on the result. One-way functions, also known as hash functions or simply *hashes*, have been used in cryptography for some time: They are relatively straightforward to compute (although still requiring computing effort), but their inverse is much more difficult to compute (Diffie and Hellman 1976), usually unfeasible with current computing technology within reasonable time.

The challenge built into the system is that the result of the hash function (itself called *hash* of the original value) must be of a certain form – i.e., to have a defined number of leading zeroes. To find such a number, the specialised nodes mentioned above repeat this process of adding a nonce to the block, performing the hash function, and checking the result, while choosing a new nonce for each cycle, until they find a hash value that has the required leading number of zeroes. As the one-way function does not allow any inference back from the desired result to the possible inputs that would satisfy the condition, it is a brute force, trial-and-error effort to find a suited nonce. As the results are computed in the hexadecimal system, where a zero is one of 16 possible outcomes, and assuming an equal distribution of the possible outcomes of the one-way function, the likelihood of finding a number with  $N$  leading zeroes is  $1/16^N$  or  $1/2^{4N}$ .

The trial-and-error process involved in finding a suitable nonce is what constitutes the PoW. The number of leading zeroes defines the likelihood of finding a fitting hash, and thus the overall effort needed to do so. And it is no trivial effort: For 10 leading zeroes, this implies already a likelihood of less than one in one trillion hashes (or ‘terahashes’, in the jargon of cryptocurrency); currently (June 2021), hash values need 20 leading zeroes, yielding a probability of finding a suitable nonce slightly smaller than one in a septillion. The PoW places an entry hurdle into the system: only those who can prove they have performed the task (hence ‘proof of work’) and found a suitable nonce (and thus



hash) for a block of transactions are able to publish them on the distributed ledger. Its principle has been engraved into Bitcoin since its inception: “The proof-of-work involves scanning for a value that when hashed, such as with SHA-256, the hash begins with a number of zero bit” (Nakamoto 2008).

Due to the PoW, participating in the consensus mechanism is associated with effort and costs; the specialised nodes performing this Herculean task need to be incentivised to do so (Sedlmeir et al. 2020a). The Bitcoin protocol foresees two types of such financial incentives: For one, the node finding the correct hash receives from each of the transactions that were bundled together in the block *transaction fees*. Additionally, the same node receives a *block reward*, in form of a predefined (and over time diminishing) number of new BTCs per new block. Citing again the original Bitcoin paper: “This adds an incentive for nodes to support the network, and provides a way to initially distribute coins into circulation, since there is no central authority to issue them. The steady addition of a constant amount of new coins is analogous to gold miners expending resources to add gold to circulation. In our case, it is CPU time and electricity that is expended. The incentive can also be funded with transaction fees. If the output value of a transaction is less than its input value, the difference is a transaction fee that is added to the incentive value of the block containing the transaction” (Nakamoto 2008). Due to this analogy, the ‘specialised nodes’ we kept mentioning so far are called *miners* in the Bitcoin jargon.

Once a miner finds a suitable nonce and thus a hash of the required form, they add it to their local copy of the ledger and make it public on the Bitcoin network. All other nodes of the Bitcoin network (currently around 12,000) can trivially verify the claim of the miner, by performing once the one-way function on the block with the corresponding nonce. Being accepted by a majority of nodes in the network is the implicit consensus mechanism that makes the new block (and all transactions contained in it) accepted; and also what keeps the distributed copies of the same ledger consistent. Miners, of course, are in competition to each other for finding a suitable nonce and gain both block reward and transaction fees. Once the new block was found and added, the race restarts for the finding a nonce for the new block, and so on.

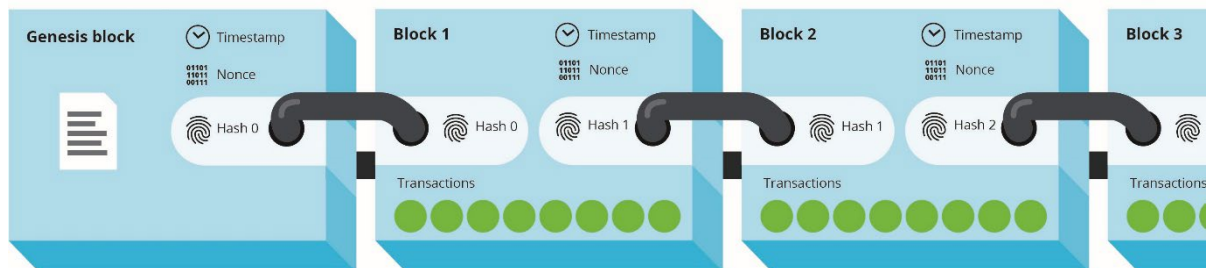


Figure 1 Schematic structure of a blockchain, showing how each block contains not only its own hash value (with the appropriate number of leading zeroes, for which a suitable nonce had to be found first), but also the hash value of the previous block. Source: Swiss Foundation for Technology Assessment (TA-Swiss), illustration: Hannes Saxer. Reprinted with permission.

The distributed ledgers behind Bitcoin and other cryptocurrencies are thus sequences of blocks added by various miners in their eternal race for the next nonce, and accepted by the other nodes of the mechanism. It remains to be seen why the technology is a chain of blocks, and thus called ‘blockchain’. This is in fact due to the structure of the individual blocks: they contain not only the bundled transactions, the nonce, the own hash, a timestamp and some metadata, but also the hash value of the previous block (Braun-Dubler et al. 2020), which acts as a pointer and is thus called ‘hashpointer’. This structure, schematically shown in Figure 1, adds to the security of the system, as it prevents later changes to the ledger: If a block somewhere in the ledger were to be changed, the next block would no longer entail its correct hash, and it would need to be modified as well (including the search for a new nonce satisfying the leading zeroes condition of the current block), and so on; all subsequent blocks would need to be recomputed. Given how difficult it is to find the correct nonce for



a single block, this is obviously an insurmountable task. This pointer structure is also responsible for the name 'blockchain', although strictly speaking the blockchain is not a sequential list, but has – for the efficiency of searches – the structure of a Merkle tree (Merkle 1988).

## 2.4 Alternative consensus mechanisms

As will be discussed in detail in Section 5 below, PoW is currently the main culprit for the high energy consumption of blockchains (de Vries 2018), and has come under increasingly critical scrutiny by environmental (Palmer 2021), governmental (Guerra 2021) and international (McKenz 2021) organisations, the press (Martin and Nauman 2021) as well as businesses (Cellan-Jones 2021).

Due to these environmental concerns, but also other reasons such as scalability, the blockchain community is investigating alternative consensus mechanisms. The search is not trivial as any alternative consensus mechanisms still needs to be both tied to a scarce resource and yet attractive for a large pool of nodes (to thereby guarantee the safety of the overall system and prevent Sybil attacks), and at the same time much more energy efficient than PoW.

The most promising such alternative is called 'proof-of-stake' (PoS), to which the Ethereum blockchain plans to transition to (Ethereum 2021d). In PoS, the scarce resource is no longer computing power (proven by the brute force search for desired hashes) as in PoW, but capital (or *stake*) as proven by the ownership of cryptocurrency native to the corresponding blockchain (e.g., ETHs for Ethereum). PoS gets rid of the concept of miners entirely. Instead, users interested to validate new blocks on the blockchain stake a minimum amount of native cryptocurrency (for Ethereum, 32 ETHs) and become validators (Finematics 2020). An algorithm making use of pseudo-random numbers chooses who the next validator will be; the likelihood of being chosen is proportional to the amount of currency staked (Sedlmeir et al. 2020b). As with PoW, validators receive both block reward and transaction fees; depending on how much currency will ultimately be stashed, the estimated return on 'investment' is estimated to lie yearly between 1.8 and 18% of the stashed sum (Finematics 2020).

The stashed sum (which is blocked in a smart contract on the traditional, PoW-based blockchain), is also intended to serve as a safety against dishonest participants; if the (transparent) actions of the validator are not accepted by a majority of other randomly chosen members of a committee of validators, a mechanism called 'slashing' makes the validator lose a part of the stake and he can be removed from the mechanism altogether. In theory, PoS should also provide a similarly high level of security as PoW, as a successful Sybil attack would require more than 50% of the stashed sum – a very unlikely scenario. On a sidenote, PoS could also help making Ethereum more scalable through a process called *sharding*, which essentially splits a database in several pieces ('shards') while maintaining the consistency of the overall database (Luu et al. 2016).

Some cryptocurrencies – such as EOS, Tezos, and TRON, which are all among the top 20 cryptocurrencies in terms of market capitalisation – already successfully deploy PoS (Sedlmeir et al. 2020b). As briefly discussed in Section 5.2.3 further down, by eliminating the energy intensive cryptographic puzzle competition, PoS is several orders of magnitude more energy efficient than PoW.

Next to PoW and PoS, there are several other consensus paradigms (Eklund and Beck 2019), but most can only be deployed in 'permissioned' blockchains, not in 'permissionless' ones – these notions are discussed in Section 2.5 below. Although not further discussed, these mechanisms also do not involve competitively solving puzzles, and are thus far more energy efficient than PoW.

## 2.5 Permissioned and permissionless chains

The two blockchains mainly discussed so far, Bitcoin and Ethereum, are both entirely open: The nodes taking part in the consensus process (miners or validators) can join or leave at any time, without the oversight of any central TA, and their identities are not mutually known (Wüst and Gervais 2018). Due to these characteristics, such blockchains, which have been the first ones to emerge, are called





*permissionless*, blockchains (Sedlmeir et al. 2020b). By contrast, in *permissioned* blockchains, only a restricted group of participants can take part in the consensus mechanism (Sedlmeir et al. 2020b), and a central entity decides who they are and attributes them the right to write on the blockchain (Wüst and Gervais 2018). The two types are sometimes also called *public* and *private* blockchain, respectively (Braun-Dubler et al. 2020).

Several application domains have been put forward for permissioned blockchains. They mostly revolve around certifying the origin (Everledger 2021) of products, their entire lifetime in form of an electronic product passport (Smart Energy International 2021), the ways they have been handled during transport (Wüst and Gervais 2018), or the genuineness of data provided by Internet of Things (IoT) sensors (Joos and Schmitz 2021).

The meaningfulness of deploying permissioned blockchains in such applications, however, has been questioned. For one, while a blockchain can provide a very high amount of security within its virtual world, it has inherently no direct control over the correct representation of real-world events in the virtual one. For the supply chain examples above, for example, while data on the origin or the condition of the goods can be securely and immutably written on the blockchain, the system's real-world backend could be manipulated: sensors could have been tampered with or simply broken, the good beautifully represented through a product passport long ago exchanged with another, inferior one. The challenges brought by this uncertain connection between the real and the virtual world (also typical for earlier computing paradigms, such as Ubiquitous Computing, Ambient Intelligence, or IoT), is known as 'Oracle problem'; it will be discussed in Section 3.2 below.

Furthermore, in many such systems a centralised database could also be deployed instead of a blockchain (Wüst and Gervais 2018). In fact, only a relatively narrow set of conditions seems to justify a permissioned blockchain: when multiple nodes, all of which are known (otherwise a permissionless blockchain would be required) but do not trust each other (otherwise a centralised database would suffice) need to store information about some status (e.g., of goods along a supply chain) and additionally cannot rely on the continuous online availability of the central TA which accredited them in the first place (otherwise a centralised database at the TA would be sufficient) (Wüst and Gervais 2018). Additionally, the consensus mechanisms specific for permissioned blockchains are similar to those of distributed databases (Eklund and Beck 2019) and do not require any PoW (Sedlmeir et al. 2020b). We will thus not consider them within the scope of this exploratory study, but will nevertheless shortly discuss them in Section 6.

## 3 Further blockchain developments: Smart contracts and non-fungible tokens (NFTs)

In Bitcoin, the blockchain is used to store BTC transactions, and thus implicitly the state of the individual wallets, which can be accessed via the account holders' private key. Instead of the usual analogy of a distributed ledger, on a more technical level Bitcoin can thus be conceptualised as a finite state machine, in which at any moment the blockchain reflects the state of ownership of BTCs in use, and each new block appended to the chain (actually, each transaction included in that specific block, but we ignore this rather academic difference) defines a state change reflecting a partially changed ownership structure. Albeit the Bitcoin state machine does reflect ownership over financially considerable assets, its functionality is obviously quite restricted.

### 3.1 Smart contracts

Envisioning blockchain technology deployed for far more diverse applications than cryptocurrencies, Vitalik Buterin published 2013 the whitepaper for the Ethereum blockchain (Buterin 2013). Although there are several notable differences between the two, the main novelty of Ethereum over Bitcoin is



that the data stored in its blockchain does not only represent currency transactions, but can also be executable pieces of code, which are called *smart contracts*. Transactions submitted to the blockchain can thus be either financial transactions of Ethereum's own currency, Ether (ETH), submitting a smart contract (which after compilation becomes part of the blockchain as any other transaction) or calling the execution of a specific smart contract with individual parameters.

As with Bitcoin, each such transaction (or each bundle of transactions contained in a new block) changes the machine from one state to another (due to its active smart contract components, however, the analogy of a distributed ledger no longer works). Ethereum's condition as a state machine is stressed through the concept of the Ethereum virtual machine (EVM), which is the state machine defining the blockchain's current status and rules, reflected in all participating nodes. The specification of the EVM are defined in the Ethereum 'yellowpaper' (Wood 2014), to which all EVM implementations (including all Ethereum clients) have to adhere to (Ethereum 2021b).

Smart contracts are represented within Ethereum as a particular type of *accounts*. There are two types of accounts: those held by physical persons (and named within Ethereum *externally owned*), and *contracts*, which can also receive, hold, and send ETH currency. Smart contracts can even initiate transactions; however, they can only do so in response to receiving a transaction – any such chain of transactions has to be initially triggered by an externally owned account (i.e., physical person). There are, however, no limits to the new transaction that a contract can initiate as response to the triggering incoming transaction: it can transfer currency, call another contract, and even create a new contract (Ethereum 2021a).

Several languages that can be deployed for coding smart contracts, most popular among them being the full-fledged, object-oriented 'Solidity' and the Python-inspired, deliberately simply kept 'Vyper', which does not support several features such as inheritance, function or operator overloading or recursion and whose contracts are thus more secure and easier to audit (Ethereum 2021e). Before they can be deployed to the network, smart contracts need to be compiled so that the EVM can interpret and store them.

### 3.2 Real-world connection and the Oracle problem

The reach of smart contracts can be fully contained within the blockchain universe, whether this happens to be Ethereum or another blockchain technology supporting smart contracts. They can, for example, hold a certain amount of ETHs and release them to a predefined account after a specific amount of time, or when they are called with certain parameters.

Smart contracts can be more universal and arguably more powerful, however, when reacting to changing conditions in the real world and when, conversely, they can trigger actions in the real world. Their algorithmic rules running on the EVM (or another similar virtual machine, VM) would then respond to external inputs, which could be fully autonomous sensors such as a digital thermometer, online inputs such as a change in stock price, or decisions by human agents. They could also initiate (immutable, since on the blockchain) actions in response, such as performing a computation, making a cryptocurrency payment, or triggering real-world actions in return through corresponding actuators (DuPont 2017).

This potential ability, however, raises legitimate questions about the sources of information on the real world. As intricate a structure as smart contracts might be, they are fully contained in the virtual world of the blockchain they reside on. They thus need to rely on trusted sources for information about the outside world (and, conversely, trust the actuators they might trigger outside the blockchain). Such sources can, of course, be manipulated. This fundamental limitation and challenge of smart contracts is known as the *Oracle problem*. It is perhaps best described in the vivid words of Paul Cuffe, who uses a logistics example from (Wüst and Gervais 2018): "How does that smart contract, that bit of code that's running on a blockchain, ..., what are its eyes? How does it get meaningful information from the wider world? You could have a lovely blockchain system for supply chain management to



give farm-to-fork provenance for the food that you buy. And they may be able to tell you: 'In that truck the temperature that it was transported at was -3.2 degrees.' The blockchain can make that data immutable and it can't be interfered with, and it will be perfectly time stamped. The blockchain can do those things, which is great. The blockchain **cannot** stop the driver getting a bucket of ice water and sitting it next to that thermometer. This is the Oracle problem" (Cuffe 2021). Of course, next to the possible dishonesty of the truck driver, the temperature sensor itself could also have been tampered with, or the publish-subscribe mechanism that reads out the sensor and publishes this data as transaction on the blockchain (Wüst and Gervais 2018).

To address this issue, Ethereum uses homonymous (and hopefully trustworthy) *Oracles*. Oracles represent the interface to the outside world. Technically, they are themselves contracts, their state can be updated via transactions, and other contracts query them about the state of the outside world they reflect (Bartoletti and Pompianu 2017). Obviously, *Oracles* do not (and cannot) represent a conceptual solution to the *Oracle problem*; they merely shift the problem and concentrate the need for trust: In Paul Cuffe's example above, we need to trust the entity behind the temperature Oracle that it employs both trustworthy temperature sensors and drivers. They do, however, make the execution of smart contracts deterministic and homogeneous: If the single Oracle was not determined from the outset, different nodes could receive different results from distinct sources for the same query, which would contradict the necessity of determinism for the EVM (Bartoletti and Pompianu 2017).

### 3.3 Non-fungible tokens (NFTs)

Although Bitcoin might have been originally developed with other purposes in mind (Nakamoto 2008), given its limited supply of BTCs (there is an exponential decrease hardwired into Bitcoin, leading to a limit of 21 million BTCs that will ever be mined), the system has quickly become a tool for investment, speculation, but also a hedge against the inflation inherent to fiat currencies (Harper 2013; Brière, Oosterlinck, and Szafarz 2015).

Expanding this (for the owners financially desirable) scarcity beyond cryptocurrencies, smart contracts can be used to both create scarcity for various types of digital assets, and to attest the ownership over scarce assets, whether digital or physical (Yasar 2021). These *non-fungible tokens* (NFTs) represent a particular type of smart contracts, which establish and attest ownership of a digital or physical asset, typically located outside the blockchain itself (even when it is a digital asset). How the ownership stipulated by NFTs on the blockchain translates into legally accepted and guaranteed ownership is a specific part of the Oracle problem and will be addressed in Sections 4.2 and 4.3 below. An NFT is issued by submitting a transaction to the blockchain that claims ownership of the corresponding asset – a process known (in analogy to the issuance of physical coins) as *minting* (Ethereum 2021c). Once issued, the token (and with it the ownership it represents) can be traded via the blockchain's embedded trading mechanism.

NFTs are opposed to fungible – or exchangeable – tokens, and have been from the outset at the very core of Ethereum's vision. The original white paper already mentions them in its Introduction: "Commonly cited applications include using on-blockchain digital assets to represent custom currencies and financial instruments ('colored coins'), the ownership of an underlying physical device ('smart property'), non-fungible assets such as domain names ('Namecoin') as well as more advanced applications such as decentralized exchange, financial derivatives, peer-to-peer gambling and on-blockchain identity and reputation systems" (Buterin 2013).

The concept of NFTs predated Ethereum, and several attempts at establishing non-fungible tokens on the Bitcoin chain existed since 2012; however, they ultimately failed due to Bitcoin's limited capabilities to support the concept (Yasar 2021). The NFT market really gained momentum in wake of the Ethereum blockchain, and in particular two standards defining two types of Ethereum tokens: ERC-20 (Vogelsteller and Buterin 2015) and ERC-721 (Entriken et al. 2018).





Next to the scarcity of the assets they represent, tokens share another attribute with cryptocurrencies: Their transactions (and implicitly the ownership structure of the assets they represent) are immutable. NFTs allow many new application domains for blockchains, some of which are presented in Section 4 below.

## 4 Application domains beyond cryptocurrencies

Smart contracts and NFTs bring entirely new application possibilities for blockchains, which go beyond the cryptocurrencies the technology was originally developed for. A comprehensive collection is shown in Figure 2. Some of the better-known blockchain application domains or those with future potential are discussed in this section. The selection of the latter is of course to some extent subjective; as an old Danish saying, often attributed to Niels Bohr, goes, “it is difficult to make predictions, particularly about the future” (Quote Investigator 2013). We try, however, to bring arguments on why we consider each individual application domain relevant and worth including in this section.



Figure 2 A detailed graph of possible blockchain application domains, some of which are discussed in this section as well. Source: Prof. Dr. Tim Weingärtner, Lucerne University of Applied Sciences and Arts. Reprinted with the author's permission.

### 4.1 Crypto art and other digital collectibles

NFTs can be used to claim, transfer, and attest ownership over almost any type of digital content. The ‘Star Trek’ actor William Shatner tokenised around 125,000 personal digital memorabilia into 10,000 packs, which all found buyers in the summer of 2020 (Wright 2020). The US basketball league, National Basketball Association (NBA), launched in October 2020 the platform ‘NBA Top Shot’, where it sells minted short video clips of NBA highlights, such as memorable slam dunks, reaching over 300m of sales during its first 5 months (Crockett 2021). A viral youtube video from 2007, in which a toddler bits his brother’s finger, also sold for an impressive amount (Sherman 2021). The founder of Twitter, Jack Dorsey, minted his first-ever tweet from 21 May 2006 and put it up for sale in March 2021 – it was sold for over 2.5 million US dollars (Dale and Reynolds 2021). In an interesting recursive spin,



the inventor and entrepreneur Elon Musk minted and put up for sale a video about NFTs – but later changed his mind and stopped the sale (Musk 2021).

CryptoKitties, a game of virtual breedable cats, each of which possessing unique features (encoded in a 256-bit ‘genome’) that are passed on to the next generation of virtual cats, gained notoriety in mid-2017 when individual cats were sold for hundreds of ETHs, equivalent to over 100,000 US dollars at the time. The wide attention received as a consequence has been claimed to have represented the pivotal moment when NFTs became mainstream, paving the way for the subsequent developments (Yasar 2021).

So far the greatest success story for NFTs, however, has been in the field of digital art. When minted, digital art is often referred to as *crypto art* (Romeo 2021; Campbell and Whitaker 2021). The first crypto art (i.e., minted digital art) achieving some notoriety was a collection of 10,000 so-called CryptoPunks (Cornish 2018), simple pixel faces algorithmically generated in such a way that each is a unique combination of characteristics such as gender, skin colour, hair shape and colour, worn hats, glasses, or jewellery, pipes or cigarettes some of them smoke, etc. The project started just as an experiment by two enthusiasts, who gave away the CryptoPunks for free to anyone with an Ethereum account in June 2017. However, they started being traded within days, reaching within an year (by mid-2018) prices equivalent to a few thousands of dollars (Cornish 2018). By mid-2021, however, they saw dramatical price increases, the two most expensive CryptoPunks so far (July 2021) being sold on 11 March 2021 for 4,200 ETHs each, equivalent at the time to approximately 7.5 million US dollars (“CryptoPunks” 2021).

In March 2021, Canadian singer and artist Grimes sold crypto art worth over 6 million US dollars at an auction. While the highest selling piece was a unique video, the bulk of the sum was realised from hundreds of individually minted (comparable to a limited-edition physical artwork such as a print) copies of two short videos entitled ‘Earth’ and ‘Mars’ (Kastrenakes 2021a). The first NFT ever to be minted (not on Ethereum but another blockchain, and reminted in 2021 according to ERC-721), artist’s Kevin McCoy ‘Quantum’, was recently sold at an auction (Cascone 2021), reaching almost 1.5 million US dollars. The highest prices so far, however, were achieved by graphic designer and digital artist Mike Winkelmann, better known under his artist’s name of Beeple: After starting selling minted versions of his ‘Everyday’ artwork series in late 2020 (for tens of thousands of dollars each), he could sell a short video for over 6 million in 2021, and then a minted collage of 5000 of his ‘Everyday’ series (comprising over 13 years of daily digital artwork) for a whopping 69 million US dollars (Kastrenakes 2021b).

This development may seem quite odd and counterintuitive. By contrast to physical artforms and collectibles, digital art and other collectibles can by their very essence be copied, distributed, and enjoyed any number of times without loss of quality. Nonetheless, NFTs provide something unique that cannot be copied: proven ownership of the respective artwork (Clark 2021). The scarcity of ownership is what drives prices for crypto art in such heights, quite possibly combined with either or all of the following factors: i) the expansionary fiscal policy in wake of the 2008 financial crisis first and of the later economic downturn triggered by the Covid pandemic, which brought inflationary fears, ii) the hype that often comes with new technologies and financial possibilities, and iii) the risk friendliness and the ease of money spending of the tech-awy recent cryptocurrency millionaires, as noted also by e.g. (Kochkodin and Kharif 2021).

## 4.2 Traditional art and other physical assets

Cryptocurrencies, and in particular Bitcoin, have been shown to act as a hedge against market uncertainties (Brière, Oosterlinck, and Szafarz 2015; Bouri et al. 2017). NFTs have only been around for a couple of years and the spectacular developments of crypto art and other digital collectibles are at the time of writing (June 2021) less than one year old. Hence, academic research on the economic properties of NFTs is just emerging. As argued in the end of Section 4.1, however, it is nevertheless



reasonable to assume that the scarcity of ownership that NFTs enable is what essentially contributed to their recent popularity.

Nevertheless, crypto art (and digital assets more generally) also have important drawbacks. Perhaps most importantly, in the recent hype surrounding investments in crypto art, its quality might not have been sufficiently considered. The artistic significance of the 5,000 artworks comprising Beeple's collage 'Everydays: The First 5,000 Days', for example, has been questioned; in particular whether any of the artist's four main themes comprised by the collage would stand the test of time (Davis 2021). Art experts also notice that the interest in crypto art comes mainly "from people who, as far as I can tell, aren't that interested in art" but mainly in investments; as a consequence "art's cultural value tends to get lost. When you buy an artwork, you're not clicking to buy the piece, you're falling in love with it. This is a different kind of transaction, it's based on relationships" (Botz 2018).

In this context, and given that Beeple's artwork sold for a higher sum than Monet's waterlilies painting 'Nymphéas' did in 2014 (Clark 2021), it seems quite possible that current prices reflect only partially the market value of the digital art and to some extent the momentary hype around them. A market consolidation might follow, reducing the relative importance of digital collectibles and crypto art among the blockchain application domains.

By contrast, traditional art (such as paintings or sculptures) has a long-established marketplace. It has been traditionally used as hedge against inflation; between 1977-1982, for example, when consumer prices rose by 80%, an art index art rose even more, by 130% (Segal and Goldfarb 2009). Investment in traditional art, however, is also confronted with a couple of issues: the inherent unpredictability of the art market, the illiquidity of the very expensive pieces for which the pool of potential buyers is very small, and the high transactions costs that can reach 25% or even more (Segal and Goldfarb 2009). Additionally issues are the authenticity of artwork – estimates speak about 50% or more of the art being forgeries or misattributions (Artnet News 2014) – and its ownership, and thus rightful sale.

For most of these issues – with the exception of the unpredictability – NFTs could improve the current situation. A Swiss bank specialised in digital assets (but otherwise fully regulated) has recently announced the purchase and tokenisation of a Picasso painting (Grundlehner 2021b). While the physical painting will be securely stored in a specialised safe (with the possibility to be loaned to Unesco-recognised museums for temporary exhibitions, which would also increase the notoriety of the painting and ultimately its price), the bank will issue 4000 tokens with a starting value of 1,000 Swiss francs (CHF), which will amount to the total currently estimated value of the painting of 4 million CHF. The tokens (which are all equal, and do not represent particular portions of the painting) will then be freely traded on a blockchain-based platform (Grundlehner 2021b). Such shared ownership of otherwise expensive artwork has the potential to vastly expand the marketplace for art.

Additionally, and as other digital technologies as well, NFTs have the potential to offer (much) more efficient ways to conducting traditional businesses such as art trade (Sherman 2021), thereby reducing the costs: The aforementioned Swiss bank, for example, plans to ask transactions costs of 0.30-0.50% (Grundlehner 2021b) – two orders of magnitude less than the 25% transaction costs of traditional auction houses (Segal and Goldfarb 2009). Finally, as pointed out by (Sherman 2021), "the blockchain also keeps a record of all the transactions connected to the NFT and the property it represents. In art sales, for example, that could represent the provenance of something going back to the creator." The immutability and verifiability of the blockchain has proof of ownership embedded at its very core. The chain of verifiable ownership can also reduce the authenticity issue to trusting only i) the first evaluation provided by experts (such as for the Picasso now being tokenised), together with ii) the safe keeping of the artwork. When a certified bank provides these two services, the trust level can arguably be much higher than with anonymous account holders: While the blockchain provides very good traceability of the tokens themselves, the connection to the real assets these NFTs can represent (i.e., the Oracle problem) needs to be managed outside the blockchain.

Given these considerations, the tokenisation of scarce and valuable physical assets for shared ownership (as means of investment, speculation, hedge against inflation, etc) seems a promising



future application field for blockchain. A trusted authority (TA) is still needed for the initial minting of the tokens and the safekeeping of the physical assets (including their insurance); beyond that the blockchain can offer a secure, fine-granular, efficient, and affordable trading mechanism. These physical assets can (and probably will) go beyond artwork: The same bank has already tokenised for shared ownership a collection of rare wines (Grundlehner 2021b), and other assets such as classic cars or any other physical collectibles are conceivable.

### 4.3 Company ownership

In September 2020, the Swiss parliament adopted a new law setting the legal framework around blockchains, named “Federal Act on the Adaptation of Federal Law to Developments in Distributed Ledger Technology” (DLT), or “DLT bill” in short (Swiss Confederation 2020). The bill essentially equates asset ownership via tokens on a blockchain to the traditional proofs of ownership (such as shares, bonds, etc) and the blockchain-based trade of tokens to the traditional marketplaces such as the stock market, and thus brought amendments to the Code of Obligations, the Federal Intermediated Securities Act, the Federal Act on International Private Law, and further Swiss federal laws (Wettstein 2020). The parts of the bill regulating token-based ownership came into force on 1 February 2021, while the parts regulating token-based trading will come into force on 1 August 2021.

The main contribution of the Swiss DLT bill is to have addressed, in a legally binding manner, the Oracle problem of NFTs issued for physical assets. According to the bill, an NFT minted according to the legal requirements stipulated by the bill, is equivalent to any other form of ownership. In particular, tokens on physical assets – for example art or a rare wine collection, as presented in Section 4.2 – guarantee token holders the legally recognised ownership over those assets, with all rights and duties involved by this status for all involved stakeholders, in particular for the keeper of the tokenised physical assets.

Similarly, the bill states the legal equivalence of tokenised company shares to traditional shares. Such tokenised shares predated the bill; as for any digital or physical item, a token could of course be minted to represent ownership of that particular share. While such ownership relation is all that is needed within the closed world of the blockchain, however, before the bill there was no certainty about the legal status of the blockchain state. A legally binding contract was thus needed to link the token with the share, a cumbersome and uncertain detour, as contracts could be challenged or not honoured (Grundlehner 2021a). The DLT bill brings legal clarity by equating the token to the share.

While there was a general optimism in the market, the development of the market for tokenised shares has been rather sobering. To mark the day, during the first minutes of 1 February 2021, the first Swiss company issued tokenised versions of its own shares (on the same day, a Swiss bank issued tokenised ownership over a collection of rare wines, as described in Section 4.2 above). However, not many others have followed. As one market participant puts it, “currently, it would almost seem that there are more companies out there offering tokenisation than those demanding the service” (Grundlehner 2021a).

As analysed in Section 4.2 above, the tokenisation of physical assets such wine collections, rare art, or classic cars, makes these otherwise prohibitively expensive assets accessible to a substantially larger pool of investors, their trade much more efficient and thus cheaper, and helps in the traceability of their authenticity and ownership. The stock market, however, is already virtual, safe and efficient, shares typically already have a relatively low value which is accessible to most potential investors, and transaction costs are also rather low. The relative lack of success of tokenised company shares thus comes perhaps to little surprise.



#### 4.4 Autonomous governance and the first decentralised autonomous organisation (DAO)

From the very beginning of the Ethereum blockchain, one of its founding visions was that of organisations autonomously acting according to predefined rules programmed into the blockchain: “The logical extension of this is decentralized autonomous organizations (DAOs) – long-term smart contracts that contain the assets and encode the bylaws of an entire organization” (Buterin 2013). The algorithmic rules running on the EVM would respond to external inputs, which could be fully autonomous sensors such as a digital thermometer, online inputs such as a change in stock price, or decisions by human agents. They could then initiate (immutable, since on the blockchain) actions in response, such as performing a computation, making a cryptocurrency payment, or triggering real-world actions in return through corresponding actuators (DuPont 2017).

Based on this vision, the development of the first prominent DAO framework started in June 2015, and by March 2016 its whitepaper had been published (Jentzsch 2016). The project, given the placeholder name “The DAO” which was in fact to remain its definitive name, was an autonomous enterprise aiming at enabling cryptocurrency-based funding of business ideas and other enterprises – a “democratic hedge fund”, in the words of Paul Cuffe (Cuffe 2021). Gathering starting capital through crowdfunding itself, ‘The DAO’ allowed investors-to-be the exchange of ETHs to DAO tokens during a “creation phase” in May 2016 – and did so quite successfully, raising almost 12 million ETHs, corresponding to 14% of all available ETHs at the time and to around 250 million USD (DuPont 2017). Once this initial funding period ended, ‘The DAO’ shareholders could submit project ideas, while the other shareholders could decide whether to back these projects or not; if accepted, the investors would be also involved through voting in the management of the newly backed enterprises – all transparent and permanent on the blockchain, according to the rules encoded in “The DAO”’s smart contract: “Tokens would be used to directly fund and control ‘proposals’ on The DAO platform. Anyone with a (refundable) minimum token deposit could create a proposal to be voted on by token holders. Investors voted by allocating DAO tokens for specific proposals. Since tokens would be valuable (comprised of exchange-convertible ETH cryptocurrency), ‘voting’ for a proposal was conceptually the same as funding it, in much the same way that projects are funded on Kickstarter. Unlike Kickstarter, however, DAO voting members would have significant control over projects” (DuPont 2017).

While the funding phase of ‘The DAO’ was approaching its end, three prominent figures of the blockchain community revealed some fundamental issues with its voting and financing mechanisms, inducing several biases and making a couple of attacks conceptually possible. The three thus called for a temporary moratorium for ‘The DAO’, allowing it to address these issues in its smart contract before going live (Mark, Zamfir, and Sier 2016). Moreover, with only a couple of days left until the activities of ‘The DAO’ were to start, a more urgent problem was found: a technical bug allowing a so-called “race-to-empty” attack, an exploit which would allow an attacker already owing some DAO tokens to repeatedly call a withdrawal function while pretending to be leaving the system and before the DAO tokens’ account could be updated (DuPont 2017). Despite these warnings, a moratorium was not put in place, ‘The DAO’ started its activities as planned, and only 2 days later a race-to-empty attack took indeed place, draining the enterprise of more than 3.5 million ETHs or about 30% of its entire funds (DuPont 2017).

In the hours after this exploit, many prominent figures of the Ethereum community, led by Vitalik Buterin, asked all Ethereum exchanges to stop the exchange of DAO tokens back to ETH. Over the following weeks, a so-called “hard fork” has been agreed upon in the community, which conceptually (if not technically) reverted all the conversions from ETH to DAO tokens, thus implicitly dissolving ‘The DAO’ as if it had never existed (DuPont 2017), while also making the DAO tokens held by the attacker worthless. While this way self-regulation worked, it was not undisputed: “a vocal minority in the community argued that ‘code is law,’ echoing Lawrence Lessig’s (1999) influential slogan. Therefore, they argued, any effort to block the ‘attacker’ would be morally wrong and against the very spirit of decentralized autonomous organizations [...] ‘Moderates’ saw the hard fork as evidence of the flexibility and practicality of Ethereum and its leaders, while the more ideological saw the hard fork as





ensorship by a powerful cabal, or proof that block chain technology was unable to live up to its idealistic promises” (DuPont 2017). This event and the related premature disappearance of ‘The DAO’ revealed pragmatical, but perhaps also ontological limitations of the DAO paradigm, as discussed in Section 4.5 below.

#### 4.5 More advanced DAOs and the emergence of decentralised finance (DeFi)

After this initial setback, the popularity of DAOs has again risen since around 2018, with the advent of *MakerDAO*. Its principles are quite similar to ‘The DAO’ two years earlier, but its financial aims and means are more diversified and sophisticated. The main idea behind *MakerDAO* is to enable holders of cryptocurrency (mainly ETHs, since it runs on the Ethereum blockchain, but a few other cryptocurrencies are accepted as well) to not only store the currency or use it for buying goods, but to lend it, invest it, or otherwise economically put it to economic use. Similarly to ‘The DAO’, prospective investors need to block an amount of ETHs or other cryptocurrency in a smart contract, which they can use for economic activities – which are, however, much more diverse than only crowdfunding. By contrast to ‘The DAO’, however, *MakerDAO* has its own currency, called DAI, which are being lent to investors in return for their ETH deposit. The DAI is pegged against the US dollar at the ratio 1:1, and the ratio is being kept through automatic mechanisms programmed in *MakerDAO*’s smart contract, which also makes use of Oracles (Manrique 2019).

For economic activities, the benefit of a stable currency (in terms of fiat currencies) over the very volatile cryptocurrencies, is that it allows more planning security for investors, creditors, and debtors alike: Very few would be probably willing to lend a cryptocurrency directly, exchange it to a fiat currency such as USD, invest those USDs into a business they intend to grow, and have to buy back cryptocurrency 6 months later to pay the debt, when that cryptocurrency could have valorised by hundreds of percent in the meanwhile. By introducing DAI and decoupling the economic activity from the fluctuations of the cryptocurrency, *MakerDAO* offers an elegant way for investors and potential creditors to keep their ETHs (or other cryptocurrencies) as value storage and yet use them for economic activities; potential debtors, on the other hand, can borrow money with the normal caveats but not the additional currency risk.

These varied possibilities together with the planning security (if the system does not lose the trust of participants and no other exploit is successful, that is), let an entire ecosystem of economic activities emerge around *MakerDAO*, and also led to the quick appearance of similarly-intended and built DAOs, such as *Aave*, *Curve Finance*, or *Compound*. Altogether, at the time of writing (July 2021), they managed a total locked cryptocurrency value equivalent to around 60 billion USD (DeFi Pulse 2021).

This plethora of possibilities also led to the usage of the more generic term ‘decentralised finance’ (DeFi) for the description of such blockchain-based enterprises. DeFi has been defined as “an umbrella term for a variety of financial applications in cryptocurrency or blockchain geared toward disrupting financial intermediaries” by its proponents (Hertig 2020), or, more soberly, “(DeFi) is neither a legal nor a technical term. It is nonetheless increasingly used in the context of discussions about the future evolution of finance and its regulation. Common usage incorporates one or more elements of: (i) decentralization; (ii) distributed ledger technology and blockchain; (iii) smart contracts; (iv) disintermediation; and (v) open banking” by established economists (Zetsche, Arner, and Buckley 2020). More technically, “DeFi instruments are essentially a second layer application developed on a – blockchain – an immutable and decentralized ledger that facilitates transparent and secure transactions between multiple, disparate and anonymous parties” (Alao and Cuffe 2020b).

As many other digital systems, DeFi makes the process of lending more efficient and thus ultimately more affordable. It achieves this mainly by eliminating the need for intermediaries (i.e., traditionally banks and other financial institutions) in the lending process and thus reducing transaction and other costs (Chen and Bellavitis 2020). Other potential advantages of DeFi are: decentralisation (reducing the risks for financial monopolies), innovativeness (as DeFi platforms often publicly share their core technologies allowing other to innovate on top of them), transparency (which centralised finance can



by definition not fully achieved due to the need to secure the centralised ledgers), and borderlessness (Chen and Bellavitis 2020).

This last point seems of particular relevance: In traditional financial markets, financing is usually geographically limited, as small and medium enterprises cannot participate in the global marketplace. For actors from less developed or stable countries, financing is difficult and costly: “Clients from countries with weaker institutional environments need to rely on costly workarounds” (Zetzsche, Arner, and Buckley 2020). By contrast, DeFi allows a quick, unbureaucratic, and affordable access to global financing. Such access could prove crucial for those economies that need it most – and not only for small and local projects, but also for larger infrastructure projects. In particular for electricity markets, DeFi is generally well-suited as the Oracle problem is manageable, electricity being a fungible and easily measured value, and the trust in the few actors traditionally high (Alao and Cuffe 2020a); the usage for DeFi in promoting independent renewable electricity projects in sub-Saharan Africa has also been explored (Alao and Cuffe 2020b).

While proponents are enthusiastic about the perspectives, there are also reasons for scepticism. It has been suggested by economists that “decentralization has the potential to undermine traditional forms of accountability and erode the effectiveness of traditional financial regulation and enforcement” (Zetzsche, Arner, and Buckley 2020). When MakerDAO, for example, starts accepting real-world assets as security for DAI loans (Allison 2020; Dale 2021), and these assets are to be mainly US real estate, largely evading any regulation, painful memories of the subprime crisis emerge. And while DAOs have been mainly financially-oriented, they could one day cover broader governance activities, with further-reaching implications into the society – while continuing being purely virtual organisations, driven by immutable algorithms and the input of anonymous account holders; companies that have no headquarters, CEO, and that are beyond the reach of regulation and governments (Kühl 2016). Whether such organisations are socially desirable, whether self-regulation such as for ‘The DAO’s exploit can be expected, or whether and to which extent they should (and can) be regulated, remain open topics. For the near future, given both the economic interests and their potential societal benefits, DAOs seem likely to continue expanding.

## 4.6 Disputed further domains

The application examples so far have all been more or less related to financial markets, whether the blockchain technology was deployed to mediate ownership, value preservation, or governance. Numerous further-reaching application domains for blockchain have been proposed, from using blockchain to store national registers such as the public land registry (Braun-Dubler et al. 2020) all the way to handling the payments for food distribution in refugee camps (Wong 2017).

At closer look, however, such applications often seem to make an (artificial) effort to use blockchain. From the two examples above, the traditional public land registry not only works flawlessly, but might actually display important advantages over a blockchain (Braun-Dubler et al. 2020). And while using an own electronic registry for cash transfers to refugees and for their subsequent food purchases, has indeed slashed the clearing costs by over 98% as compared to using regular banks accounts, and this electronic registry is loosely connected to blockchain technologies, it actually has one single copy kept by a centralised trusted authority (Braun-Dubler et al. 2020), and is thus essentially a centralised database and not a distributed virtual machine in need of consensus mechanisms.

For applications outside the financial sector (understood in a wide sense, comprising cryptocurrencies as well as the domains from Sections 4.1 – 4.5), the blockchain technology has too often been a “solution without a problem” (Braun-Dubler et al. 2020). Either traditional, proven technologies, already do a good job and do not need replacement, or the complexity of a blockchain setup is unnecessary and even detrimental, or what has been claimed as being blockchain was actually not.



## 4.7 Electricity markets

While this might of course change in the future, there is at least one non-financial application domain where the deployment of blockchain appears quite meaningful already: electricity markets, and in particular dynamic markets comprising both volatile renewable production and a quickly changing ecosystem of producers and consumers. A good review of possible applications of blockchains in electricity markets is provided by (de Villiers and Cuffe 2020). They provide a three-tier taxonomy, categorising blockchain applications in the electricity domain according to their disruption potential to the industry, from a mere sustaining technology to a potentially deep disruption of the sector:

- On tier 1 ('sustaining technology'), blockchain is merely in a supporting role. Here, cryptocurrencies might be used for payments to the utility or among peers, for the digital notarisation of assets, or for managing the supply-chain management.
- Tier 2 ('evolutionary technology') revolves mainly around smart contracts for the electricity market among its actors: large-scale utilities, small-scale producers, consumers, and those colloquially called 'prosumers' (i.e, producers and simultaneously consumers).
- Tier 3 ('disruptive technology') emphasises the radical new ownership and governance aspects brought by blockchain technology for the electricity domain. This would not only imply decentralised governance (which, to some extent, can already be included in the smart contracts on tier 2), but even blockchain-based decentralised regulation.

According to this taxonomy, tier 1 is not only not transformational, but also not specific to the electricity market. Tier 3, on the other hand, is highly speculative, would imply considerable and partly difficult to assess societal impacts (de Villiers and Cuffe 2020), and is thus not necessarily desirable and quite uncertain. The most interesting blockchain applications in electricity markets might thus lie on the middle tier of 'evolutionary technology'.

Electricity markets have a couple of characteristics that make the adoption of blockchain technology both more meaningful and more likely to be accepted than in other domains: For one, electricity is a fungible as well as easily and consistently measurable unit. Each kWh of electricity is identical to another, so there is no product differentiation (as for, say, apples or shoes), which otherwise makes the autonomous assessment of a product's quality quite challenging for differentiated products. Second, electricity cannot be stored (not directly, at least, but only transformed in another form of energy such as the potential energy of water in pumped hydro storage or chemical potential in batteries) and is thus not subject to potential loss of quality due to unsuitable storage (as would be the case for apples, fish, or shoes), easing thus the Oracle problem even further. Finally, there is already a ground level of trust in electricity markets: We tend to trust our utility company, the (traditional or smart) metres it installs on our premises, and the subsequent billing. Taken together, all these features make the Oracle problem in electricity markets much more limited (and thus manageable) as compared to other domains: Trust must basically only be placed into a (couple of) smart metre(s) acting as oracles.

There are also solid reasons to deploy blockchain in electricity markets. One of the main problems of renewable generation is its volatility, which exposes renewable electricity producers to both price and volume risks (Cuffe 2021). Combined with the gradual downscaling of subsidies for renewable electricity, this negatively affects the 'bankability' of (i.e., the keenness of banks to finance) renewable energy projects (Alao and Cuffe 2020a). To hedge these revenue risks in the short term, renewable electricity producers typically use contract-for-difference (CfD) financial instruments in day-ahead markets. While these help mitigate the revenue risks, they are susceptible to other downsides such as counterparty credit, margining and third-party risks, slow settlement time, high overhead, bilateral friction, and human errors. Smart contracts could avoid or reduce most of these issues (Alao and Cuffe 2020a). A well-known localised energy management system (EMS) that successfully deployed smart contracts and other blockchain-enabled features for the trade of locally produced electricity is the 'Brooklyn Microgrid' (Mengelkamp et al. 2018). Moreover, beyond the hedging of production, smart





contracts could be deployed for financing the development of renewable electricity projects in the developing world, e.g., in Sub-Saharan Africa (Alao and Cuffe 2020b).

## 5 The energy consumption of a (permissionless) blockchain

On the most general level, there are three main sources within a blockchain that determine its energy consumption:

- the *storage* of the distributed VM (or distributed ledger),
- the *computation* triggered by the consensus mechanism, in particular PoW (if applicable), and
- the *communication* among nodes, which can be triggered by some or all of the following events: i) the initial download of the entire blockchain by a new node entering the system, ii) the transactions submitted by individual nodes, and iii) the messages of the consensus mechanism.

Sections 5.1, 5.2, and 5.3 address these three components individually; Section 5.4 then qualitatively discusses and compares the three.

### 5.1 Energy required for storing the distributed virtual machine (distributed ledger)

As discussed in Sections 2 and 3, blockchains can be conceptualised as distributed ledgers or, more technically, as distributed VMs. So far, for all mainstream blockchains, this distribution also implies replication of the entire blockchain since its inception. Bitcoin, for example, is at the time of writing (June 2021), 350 gigabyte (GB) large (Blockchain.com 2021b) and grows with about 53,000 blocks per year (as there is one block every 10 minutes on average and an year has 525,960 minutes on average, accounting for leap years). As the average block size as of 2021 is about 1.25 megabytes (MB) (Blockchain.com 2021a), this means that the Bitcoin blockchain currently grows at a rate of about 66 GB / year.

At the same time, there are currently about 11,000 (Bitnodes 2021) – 13,000 Bitcoin (Avan-Nomayo 2021b) nodes; all of them have the entire blockchain replicated. These nodes typically run 24/7. Estimating their energy consumption is not trivial, as they could run on a variety of hardware platforms, from enterprise servers to – due to the manageable size of the blockchain – personal computers running Linux, Mac OS, or Windows. Estimates are more difficult as around half of the nodes run behind the Tor anonymisation network (Avan-Nomayo 2021b).

The annual energy consumption triggered by the storage of a fully replicated blockchain can be estimated with Equation (1):

$$E_{St} \left[ \frac{kWh}{year} \right] = \#Repl_{Av} * BC_{St} [GB] * EI_{St} \left[ \frac{kWh}{year * GB} \right] \quad (1),$$

where  $E_{St}$  is the yearly energy for storing the blockchain,  $BC_{St}$  is the size of the stored blockchain (in GB),  $\#Repl_{Av}$  the average number of replicas (as weighted over the year), and  $EI_{St}$  the average energy intensity of storing a unit of data (1 GB) for one year.

For Equation (1), the size of the (permissionless) blockchain is typically public and well known, while the number of nodes (and thus blockchain replicas) can be well estimated (Bitnodes 2021; Avan-Nomayo 2021b). The largest uncertainty lies in the energy intensity of the nodes. For a first, conservative estimate, we suggest the following assumptions:



- Nodes are continuously on. Although the minimum requirement is that they are running at least 6h/day, it is recommended they are on 24/7 (Bitcoin.org 2021). Moreover, statistics of running full nodes such as (Bitnodes 2021; Avan-Nomayo 2021b) are snapshots that abstract from the not reachable nodes and are thus a good estimate of the number of nodes running on average.
- Nodes run mainly on laptop and desktop computers. Given the relatively modest requirements of 2 GB of RAM memory and a few hundred GBs of disk space, and the existence of Bitcoin full node software for the three main PC operating systems (Linux, Mac OS, and Windows) (Bitcoin.org 2021), nodes can certainly run on PCs. Nodes can obviously also run on more efficient servers and even on small hardware such as Raspberry Pis (Jones 2021), so this is a conservative assumption yielding most likely an overestimate; however, not by a large margin.
- The use of the hardware is dedicated to the blockchain node. This is also a conservative assumption, as some of the PCs storing the blockchain might well be used for other tasks as well. As above, however, it does probably not overestimate the result by a huge margin, as from the 24h assumption, the PCs will typically only be used for a couple of hours per days for other tasks.

Using these assumptions in Equation (1), and an average power consumption of laptops and desktops of 30 W (between the 10W of a typical laptop and the 50W of an efficient desktop without screen), yields 263 kWh/year for a node over the 8,760 hours of one year. This number represents the product of the last two factors on the right side of Equation (1). We could, of course, divide this energy consumption by the size of the blockchain  $BC_{St}$  to arrive at the storage intensity per GB  $EI_{St}$ , and then remultiply with  $EI_{St}$  as part of Equation (1). When blockchains are stored entirely on dedicated nodes, however, we can deploy the simpler Equation (2) instead, in which the last two factors of Equation (1) have already been multiplied:

$$E_{St} \left[ \frac{kWh}{year} \right] = \#Repl_{Av} * E_N \left[ \frac{kWh}{year} \right] \quad (2),$$

where the yearly energy for storing the blockchain  $E_{St}$  is directly computed from the average number of blockchain replicas  $\#Repl_{Av}$  and the energy required on average for a blockchain node,  $E_N$ .

For Bitcoin, given a current estimated average of 12,000 full nodes (Bitnodes 2021; Avan-Nomayo 2021b) and the 263 kWh per node and year computed above, this would yield a yearly storage energy  **$E_{st} = 3.15$  GWh/year**; as discussed above, probably an overestimate.

A lower bound for the energy required for storing a blockchain can be computed by modifying the second assumption above to the opposite, that all nodes are stored in large data centres (DCs). In hyperscale DCs, the power consumption of hard-disk drives (HDDs) and solid-state disks (SSDs) are converging in 2020 to about 6.5 W for 10 TB HDDs and 6 W for 5 TB SSDs, respectively – see Fig. 15 on page 16 in (Shehabi et al. 2016). This yields an yearly energy consumption of 57 kWh for 10 TB HDDs, and 52.5 kWh for 5 TB SSDs in hyperscale DC, respectively

DCs being optimised for data storage, there is obviously no exclusive usage of their disk storage by the blockchain data; on the contrary, only a share proportional to its size has to be attributed to the blockchain. The yearly energy computed above has to be distributed over the 10 TB of an HDD and the 5 TB of an SSD, leading to 5.7 kWh/TB and 10.5 kWh/TB annually, respectively. Choosing a number between the two to reflect the variety of storage mediums within a DC, results in an energy intensity of storage  $EI_{St}$  of 8 kWh/TB and year in DCs (based on the same data from (Shehabi et al. 2016); (Masanet et al. 2020) compute a DC energy intensity of storage  $EI_{St}$  of 0.11 kWh/TB and year, a factor of 72 lower than our result (but it could stem from a mistaken computation).

Using  $EI_{St} = 0.008$  kWh/GB in Equation (1) for Bitcoin results in a lower boundary of only **33.6 MWh per year** for storing the storage of the entire Bitcoin blockchain – a factor of almost 100 lower than the upper bound of 3.15 GWh yearly, computed with the different storage assumption.



So far, the implicit assumption was that the state of the distributed virtual machine is entirely replicated at all participating nodes. For reasons of scalability and efficiency, however, blockchains can be truly distributed databases, being split into several ‘shards’ (Luu et al. 2016). For the Ethereum blockchain, for example, 63 new shards are soon to be introduced next to the current blockchain, for a total of 64 shards (Finematics 2020). In such a sharded blockchain, each node will no longer store the entire blockchain but just one of the shards. Equation (1) can be rewritten to cover this more general case as well; the resulting equation is:

$$E_{St} \left[ \frac{kWh}{year} \right] = \sum_{i=1}^{\#shards} (\#Repl_i * S_i [GB]) * EI_{St} \left[ \frac{kWh}{year * GB} \right] \quad (3),$$

where  $\#shards$  is the number of shards of the blockchain,  $\#Repl_i$  is the number of nodes replicating shard  $i$  of the blockchain, and  $S_i$  is the size of shard  $i$ . The assumption implicit to Equation (3) is that the energy intensity of storage is independent of the shard; hence  $EI_{St}$  can be outside the sum. This should be a reasonable assumption for most use cases (i.e., sharded blockchains); otherwise the energy intensity needs to be included within the parentheses and below the sum, with a corresponding index  $i$ .

The introduction of shards is not likely to massively influence the total energy needed for the storage of a blockchain. As one of the main purposes behind their introduction is to allow more transactions and general scalability, it does seem probable that a sharded blockchain will grow faster in size than an unsharded one. Faster growth would probably mean both a larger overall size of the (cumulative) size of all shards taken together, and more overall nodes (each storing one particular shard). Not both these values, however, can be relevant for the overall storage energy: if DCs are mainly used for data storage, then the growth in size is relevant, if (to a large extent exclusively dedicated) PCs are used, then the growth in nodes is relevant. If the Bitcoin blockchain were to have 64 shards, for example, and by that i) quickly grow to a total of 64 times more nodes than today, ii) each of which would be exclusively dedicated to storing the shard, its worst-case energy consumption would also grow 64-fold, reaching 200 GWh yearly. Both assumptions i) and ii) above are unrealistically conservative; for new shards, the number of nodes storing each of them would grow slowly, and it is quite possible that some of the physical machines storing nowadays the single shard would be used as virtual machines to store one or more additional shards as well, at an essentially unchanged energy consumption. New nodes storing on a PC one (in the beginning rather small) shard are unlikely to be exclusively dedicated to that one shard only, but most likely used for further activities as well.

## 5.2 Energy required by the computational complexity of consensus mechanisms

As discussed in Section 2.2, the PoW consensus mechanism is computation (and thus energy) intensive by design; the aim of PoW being to put a burden on the participation in the consensus, and thus prevent Sybil attacks. For the system to work, however, participants need conversely to be incentivised to perform this work. This typically happens by being awarded new, ‘mined’ currency to the participant succeeding first to solve the mathematical puzzle (i.e., finding the correct nonce that results in a hash value with the required number of leading zeroes).

When the price of a cryptocurrency relying on such PoW-based mining grows, such as during the first Bitcoin rush in 2017 (Higgins 2017), this mechanism can quickly grow into an environmental problem (de Vries 2018). In this section, we first show how to compute an *upper bound* for the energy consumption of the PoW mechanism, and then embed our result for Bitcoin mining in the literature.

### 5.2.1 Computing an upper bound for the PoW energy consumption

As long as the expected gain is higher than the costs, it is a reasonable for participants to keep mining, or to freshly join the mining community. For each period and individual miner, the variable



costs (which abstract from the fixed costs such as buying the hardware, rent of the premises, etc) comprise the price of the electricity needed for computing, which can be expressed as

$$C = \#H[hash] * EI_h \left[ \frac{J}{hash} \right] * PUE * P_E \left[ \frac{USD}{J} \right] \quad (4)$$

where  $C$  are the miner's costs for a period,  $\#H$  the number of hashes the miner performs during that period (which is just a number and thus unitless; for clarity, we write nevertheless the unit 'hash' next to it),  $EI_h$  the energy intensity of hashes (expressed in Joule per hash operation),  $PUE$  the 'power usage effectiveness' (a measure that relates the total power consumption of an IT facility to the power used strictly by the IT equipment and thus accounts for energy overhead, in particular for the cooling of DCs), and  $P_E$  the price of electricity (usually expressed in a currency such as USD per kWh; for unit consistency we express it here as USD per Joule).

To reflect, as often done in the literature, the state-of-the-art of the mining hardware (shown by the speed of hash computations), the number of hashes over a mining period from Equation (4) can be further expanded as

$$C = t[s] * HR \left[ \frac{hash}{s} \right] * EI_h \left[ \frac{J}{hash} \right] * PUE * P_E \left[ \frac{USD}{J} \right] \quad (5)$$

where  $t$  is the length of the mining period in seconds, and  $HR$  the hash rate expressed in hashes per second.

The expected revenue of a miner, on the other hand, depends on its share of hash operations among all hash operations of all the miners needed to find the correct nonce. The expected value for the latter, as discussed in Section 2.2, is  $16^N$  (or the equivalent  $2^{4N}$ ),  $N$  being the number of leading zeroes aimed for. The revenue a miner can expect is thus

$$E(R) = \frac{t[s] * HR \left[ \frac{hash}{s} \right]}{2^{4N} [hash]} * \#CC * P_{CC} [USD] \quad (6)$$

where  $E(R)$  is the expected value for the revenue,  $N$  is the number of needed leading zeroes needed,  $\#CC$  the number of awarded cryptocurrency units awarded to the miner first solving the puzzle, and  $P_{CC}$  the price of one cryptocurrency unit awarded.

The fraction on the right side of Equation (6) is the share of hashes one particular miner can perform before the expected number of total network hashes is reached, and thus the probability that this specific miner will solve the puzzle first. This probability multiplied by the total reward (terms 2 and 3 on right side of Equation (6)) represents the expected revenue of this particular miner.

According to microeconomic theory, the market equilibrium in the short-run is when marginal costs equal the expected returns (Samuelson and Nordhaus 2009). In theory then, while  $E(R) > C$ , mining is a worthy pursuit and rational actors will keep or start mining. This is obviously an abstract model, however; in particular low hash rates would result in very low probabilities – despite the same expected revenue, hardly anyone would mine with little investment for an expected win in several millennia or much later: “Back in January 2011, a miner with an up-to-date GPU (2 GH/s) could expect to find more than two blocks a day. In November 2018, because of the increasing difficulty of the search puzzle, the same miner could expect to find a block every 472,339 years” (Stoll, Klaaßen, and Gallersdörfer 2019).

Although it will not always apply at the individual level, this rule is a good approximation for the general, systemic behaviour. Adam Smith's invisible hand (Smith 1776) makes miners indeed mine while the expected revenue is greater than the costs – and increasingly larger mining companies as well as pools of smaller miners (which distribute computation load and revenues) keep the times to revenue manageable (Stoll, Klaaßen, and Gallersdörfer 2019).



Given this need for competitiveness, the hardware used for mining also evolved over the last decade, becoming both vastly faster and much more efficient: “First-generation miners used central processing units (CPUs) in conventional personal computers with computing power of less than 0.01 gigahashes per second (GH/s) and an efficiency of 9,000 joule per gigahash (J/GH). Over time, miners switched to graphic processing units (GPUs), with 0.2–2 GH/s and 1,500–400 J/GH in 2010 and, starting in 2011, moved to field-programmable gate arrays (FPGA) with 0.1–25 GH/s and 100–45 J/GH. Since 2013, application-specific integrated circuit (ASIC)-based mining systems, with up to 44,000 GH/s and less than 0.05 J/GH have prevailed” (Stoll, Klaaßen, and Gallersdörfer 2019).

Despite becoming both faster and more efficient, for a cryptocurrency whose value increases, the latter will not be able to keep pace with the former, and the overall energy consumption of that cryptocurrency will increase. Why this necessarily needs to be so becomes clear if writing out the system equilibrium  $C = E(R)$ , according to Equations (5) and (6):

$$t * HR * EI_h * PUE * P_E = \frac{t * HR}{2^{4N}} * \#CC * P_{CC} \quad (7)$$

The term  $t * HR$  appears on both sides of this identity and cancels out. From the other factors, which in the short run can be regarded as constant, the equilibrium energy intensity for a hashing operation mining can be derived:

$$EI_h = \frac{1}{2^{4N} * PUE} * \frac{\#CC * P_{CC}}{P_E} \quad (8)$$

This equilibrium energy intensity is also a threshold: If a miner owns hardware that is more energy efficient than this threshold (i.e., with a lower energy intensity of hashes), it is worth joining the pool of miners; otherwise not. Equation (8) also shows the dependencies of the threshold energy intensity:

- unsurprisingly, it is directly correlated to the reward expected,  $\#CC * P_{CC}$ , and thus the price of the currency. It has been often shown in the past that the (estimated) energy consumption of Bitcoin mining is directly related to the price of BTC.
- it is inversely related to everything that makes mining more expensive: the (average) price of electricity  $P_E$  and the  $PUE$  (which, the larger it is, the more electricity consumption and thus costs it implies). Above all, however, it is of course inversely correlated to the (exponent of) the complexity of the cryptographic puzzle that needs to be solved. This also underlines how in the early days, when only a few leading zeroes needed to be found, inefficient GPUs or even CPUs could be used for mining, and how today this is no longer an option.

Using current (July 2021) data for Bitcoin would mean  $N = 20$ ,  $\#CC = 6.25$ , and approximately  $PUE = 1.3$ ,  $P_{CC} = 40k USD$ , and  $P_E = 0.05 USD/kWh$ , as also deployed by (CBECI 2021) and representative for the countries where most mining takes place. Dividing the last value by 3.6 million (to transform it from USD/kWh into USD/J) yields a resulting energy intensity  $EI_h = 1.14 * 10^{-11}$  J/hash, or  **$EI_h = 0.011$  J/GH**.

We can further use this estimate of the threshold energy intensity of hashes to compute an upper bound estimate for the total power used in the PoW consensus mechanism, and thus also for the overall energy consumption over a certain period of time. Analogue to Equation (4), we can write the energy used (either by a miner or the entire network) during one period:

$$E = \#H * EI_h * PUE \quad (9)$$

This is in fact almost identical to Equation (4); only the last term  $P_E$  that was transforming the energy into costs has been left out. As the energy during one period is less relevant, we can easily transform this into an equation computing the average power, by dividing both sides of the equation by the time. Accounting also for the fact that the expected number of hashes during one period is  $\#H = 2^{4N}$ , we can now rewrite Equation (9) as



$$P_{PoW} = \frac{2^{4N} * EI_h * PUE}{t} \quad (10)$$

Using  $EI_h$  as computed in Equation (8) yields:

$$P_{PoW} = \frac{2^{4N} * EI_h * PUE}{t} = \frac{2^{4N} * PUE}{t} * \frac{1}{2^{4N} * PUE} * \frac{\#CC * P_{CC}}{P_E} = \frac{\#CC * P_{CC}}{t * P_E} \quad (11)$$

Current Bitcoin data in Equation (11) are  $\#CC = 6.25$ ,  $P_{CC} = 40k \text{ USD}$ ,  $t = 600s$  (as a period is about 10 minutes long), and  $P_E = 0.05 \text{ USD/kWh}$  (transformed into USD/J by multiplying with 3.6 million). Using this data yields a current power consumption for Bitcoin's PoW mechanism of **30 GW** – or an energy  $E_{PoW}$  of about **263 TWh annually**.

### 5.2.2 Further estimates for Bitcoin's PoW energy consumption

Over the last years, the literature has provided three types of estimates for Bitcoin's PoW power or energy consumption: a lower bound, an upper bound, or an estimated value somewhere between these two. Some of these estimates are presented in Table 1.

Table 1 Some well-known estimates for the energy consumption of Bitcoin's PoW consensus mechanism: (de Vries 2018), (Krause and Tolaymat 2018), (Stoll, Klaaßen, and Gallersdörfer 2019), (de Vries 2020), (Bendiksen and Gibbons 2019), (Sedlmeir et al. 2020b), (CBECI 2021), (Digiconomist 2021), and the current study, respectively. All numbers are expressed in [TWh/year]. Numbers have been rounded to the nearest integer. Average power values (typically indicated in GW) were transformed to TWh/year by being multiplied with 8,760 (the number of hours in an year).

Period	Publication	Date	Lower bound	Expected	Upper bound
Earlier studies	(de Vries 2018)	May-18	22	67	78
	(Krause and Tolaymat 2018)	Nov-18		30	
	(Stoll et al. 2019)	Jul-19		46	
	(de Vries 2020)	Sep-19		87	
	(Bendiksen and Gibbons 2019)	Dec-19		61	
	(Sedlmeir et al. 2020)	Feb-20	60		125
Current studies	(Cambridge BECI 2021)	Jul-21	28	85	312
	(Digiconomist 2021)	Aug-21	38	152	
	this study	Jul-21			263

The lower bound is typically computed via methods similar to our Equation (10), which asserts that the average overall mining power equals the expected number of hashes until a solution is found ( $2^{4N}$ ) times the energy intensity of a hash operation ( $EI_h$ ) and the PUE, divided by the average amount of time of a round. Assuming in this equation that all miners use the currently most efficiently available mining hardware, and that  $EI_h$  is thus not an average value but the lowest possible value according to current technology, leads to the lower bound (de Vries 2018; Sedlmeir et al. 2020b; CBECI 2021). The expected value can be computed in several ways; one common method is to adjust the lower bound to more realistic assumptions (de Vries 2020).

By contrast, the upper bound is independent of HW performances and computed purely economically. As we have done while deriving Equation (11), and as initially put forward by (Sedlmeir et al. 2020b), the upper bound depends solely on the expected reward, the average duration of the mining period, and the price of the electricity used in mining. As argued in Section 5.2.1 above, miners do not need to





have the most efficient hardware; as long as the efficiency of devices is above the threshold, it is rational to use them for mining. The upper bound thus has the opposite (implicit) assumption to the lower bound: that all devices are the least efficient that can still be economically deployed.

The true number will be somewhere between these two boundaries: with a large factor of more than 10 between lower and upper bound, (CBECI 2021) places its expected value a factor of 3 more than the lower bound and 3.67 less than the upper one. As can be further seen in Table 1, (Digiconomist 2021) places the expected value higher, a factor of 4 above its lower bound (which in itself is slightly larger). Overall, it is probably fair to say that currently, Bitcoin mining accounts for a consumption of  $E_{PoW} = 100 - 150 \text{ TWh annually}$ , and that this value is largely independent of the energy efficiency of the hardware – any HW efficiency gains will be quickly consumed through digital rebound (Coroamă and Mattern 2019).

### 5.2.3 Energy for the computational complexity of other consensus mechanisms

By contrast, the other consensus mechanisms – and PoS in particular – do not require any computationally intensive steps equivalent to solving cryptographic puzzles in PoW. Moreover, the computation complexity of the PoS consensus mechanisms is independent of the number of nodes in the network and thus very efficiency for large-scale systems (Sedlmeir et al. 2020b). Overall, the energy consumption of the PoS consensus mechanism is not only several order of magnitude lower than the consensus mechanism of a PoW consensus mechanism (Sedlmeir et al. 2020b), but does also not grow substantially with the network size or – crucially – the value of its underlying cryptocurrency.

## 5.3 Energy due to the communication of consensus mechanisms

Next to the computation, a PoW-based consensus mechanism also generates traffic: Once a miner has found a fitting nonce, it will broadcast the respective block to the entire network of nodes. The power consumption induced globally by these coordination messages can be approximated with:

$$P_C = \frac{Bl [GB] * \#N * (EI_{WAN} + EI_{FAN}) \left[ \frac{kWh}{GB} \right]}{t [s]} \quad (12)$$

where  $Bl$  is the size of one block,  $\#N$  the number of nodes on the blockchain participating in the consensus mechanism,  $EI_{FAN}$  and  $EI_{WAN}$  the energy intensities (typically expressed in kWh/GB) of the fixed access network and the wide-area network, respectively (Coroamă 2021), and  $t$  represents the average time between two blocks.

It is worth noting that this number correlates linearly with the number of nodes  $N$ , as the message complexity of the PoW consensus mechanism itself correlates linearly to it – in the language of algorithmic complexity theory, it is  $O(N)$ .

Current data for Bitcoin are:  $Bl = 1.25 \text{ MB}$  (Blockchain.com 2021a),  $\#N = 12,000$  full nodes (Bitnodes 2021; Avan-Nomayo 2021b),  $EI_{WAN} = 0.02 \text{ kWh/GB}$  and  $EI_{FAN} = 0.07 \text{ kWh/GB}$  (Coroamă 2021),  $t = 600s$  (Sedlmeir et al. 2020a). Using these values in the fraction counter of Equation (12) yields an estimated **global communication energy usage of just 0.11 kWh / block**, and dividing this average power consumption by time yields the average power consumption due to Bitcoin's communication complexity of  $P_C = 0.675 \text{ kW}$ , less than a boiling kettle turned on somewhere in the world.

As one year comprises 8760 hours, the global yearly energy consumption of Blockchain's communications  $E_C$  at just under **6 MWh**.



## 5.4 Comparative analysis

The three values for Bitcoin's  $E_{St}$ ,  $E_{PoW}$ , and  $E_C$  estimated in Sections 5.1, 5.2, and 5.3, respectively, could hardly be more different:

- The communication energy  $E_C$  of yearly about 6 MWh for Bitcoin does not need to concern us further: Scaling linearly with the number of nodes of the blockchain, and requiring on average less power than a kettle or an electric oven, it is entirely negligible.
- Depending on the chosen assumptions, the yearly energy needed for the storage of all of Bitcoin's replicas can be between 33 MWh – 3 GWh, one to three orders of magnitude more than the energy induced by the coordination messages. As argued in Section 5.1, with today's technology, the value is likely closer to the upper margin of this wide interval. Nonetheless, even this upper value is negligible when compared to the energy needed for the PoW computations, which – as shown for Bitcoin in Section 5.2 – lies 5 orders of magnitude above this worst-case storage energy.
- As has often been argued before, the truly worrisome component of blockchain is the PoW consensus mechanism. Measured hundreds of Terawatt-hours for Bitcoin only, it is 5-7 orders of magnitude higher than the energy required for storage. While two years ago, in 2019, it was estimated to amount to 0.1 – 0.3% of worldwide electricity consumption (Kamiya 2019), today it seems to be closer to 0.5% of the world's yearly 25,000 TWh electricity consumption – and thus close to the energy consumption of all data centres in the world, estimated at between 200 TWh (Masanet et al. 2020) and 400 TWh (Hintemann and Hinterholzer 2019).

# 6 Discussion

## 6.1 Main levers that influence the energy consumption of blockchains

Given this comparative analysis, not only the energy needed for coordination messages, but also the storage energy can be ignored when it comes to global blockchain energy consumption. As analysed in Section 5.1, under the most conservative assumption (i.e., all copies are stored on PCs dedicated exclusively to the one blockchain), the currently largest-sized blockchain with most nodes, Bitcoin, would require some 3 GWh of storage energy yearly. This represents a factor of about 0.000012% of the world's yearly 25,000 TWh electricity consumption.

Even the storage of thousands, or tens of thousands of blockchains that would coexist, would hardly represent a worrisome environmental issue, particularly as most would probably be substantially smaller than the current size of Bitcoin and much less replicated. Additionally, for such a massive-scale deployment of blockchain, most of them would probably be stored on energy-efficient servers, pulling the average per-blockchain storage energy towards the lower end of our estimated range. Both from an environmental and also from a policy perspective, the only technological blockchain component that deserves indeed a closer look is the PoW consensus mechanism.

## 6.2 Energy needed by the Proof-of-Work mechanisms

In Section 5.2, we argued that the right-hand side values from Equations (8) and (11) are constant in the short run (other than the very volatile price of some cryptocurrencies, that is). While this is true, most of them do change over time. Moreover, they are not historically independent, their values at any given time being often the consequence of the evolution over time of some of the other values.

The current threshold energy intensity of hashing  $EI_h$ , for example, is in a causal loop with both the price of the cryptocurrency  $P_{CC}$  and the complexity of the cryptographic puzzle defined by the parameter  $N$ : When the price goes up, it attracts more miners and the competition increases. In this





context, there is a strong incentive to deploy more energy efficient (i.e., less energy intense) hardware, as more of it can then be used at the same variable electricity costs, which constitutes a clear competitive advantage. Both more miners and more hardware per miner, in turn, bring about a shortening of the time elapsed until a suited nonce is found; to rebalance the algorithm, the cryptographic challenge needs to be made more difficult by increasing the parameter  $N$ ; then, the same game starts all over again.

The efficiency gains of the hardware used for hashing thus seem to be a textbook example for a 'digital rebound' (Coroamă and Mattern 2019); a digitally induced rebound effect, in which initial efficiency gains backfire and ultimately result in increased overall energy consumption.

It is also noteworthy that in Equation (11), the power needed for the PoW consensus mechanism rather surprisingly does not seem to depend on the complexity of the cryptographic puzzle, but only on the ratio between potential reward and the cost of electricity. Due to the causal loop described above, however, the current value of the parameter  $P_{CC}$  very much depends on the historic development of the cryptographic puzzle; if the puzzle was still almost trivially solvable for anyone with a usable CPU, as it was in 2010, a BTC could hardly be worth tens of thousands of US dollars.

As discussed in Section 5.2.2, the literature often puts forward two numbers, a lower bound and an upper bound for the mining energy of blockchains in general, and of Bitcoin in particular. The lower bound is usually computed from the current overall hashrate  $HR$  (which can be derived as  $HR = 2^{4N}/t$ , the expected value of hashes divided by the current length of one period) multiplied by the known hashing energy intensity of the most efficient hardware on the market today. This is necessarily a lower bound as some miners use less efficient equipment. The upper bound, on the other hand, often relies on market data on the sold equipment combined with assumptions about its lifetime and a (conservatively chosen) mix of older, less efficient, and newer, more efficient technologies (CBECI 2021).

In line with (Sedlmeir et al. 2020b), our estimate from Section 5.2.1 also puts forward an upper bound, as in reality not all equipment will work at the threshold, but some will be more efficient. By tying the threshold to one parameter only, the energy intensity of hashing, the approach is quite similar to the typical computation of lower bounds, and crispier than many of the upper bound approaches based on market data and dynamics. It is probably also more dynamic than other approaches and might thus generally result in lower upper bounds (i.e., closer to reality), as the one parameter it depends on follows price changes of the currency immediately, and might thus model the decisions of individual miners more accurately. Given its simplicity, it is also trivially computable – the number of coins received per successful finding of a suitable nonce, the average time of one period, and the current price of the cryptocurrency are all well known; only for the average electricity price used by the miners some assumptions need to be made.

As the difference between the lower and upper bound is quite large – often a factor of 10 (CBECI 2021) – the question arises where in this range the exact value resides. While the literature indicates this to be somewhat close to the geometric mean, a factor of 3-4 away from both the lower and the upper bound (CBECI 2021; Digiconomist 2021), it also seems reasonable to expect the true value closer to the lower bound in a period of falling cryptocurrency prices (as for uncompetitive HW, the expected revenue stops matching the costs) and closer to the upper bound in times of increasing cryptocurrency prices (as it would become again economic to use less efficient HW). If this reasonable assumption is true, it implies that the energy consumption of the PoW mechanism exacerbates the swings of the price of the blockchain's underlying cryptocurrency.

Crucially, what our analysis shows is that there is no economic mechanism that could with certainty limit the energy consumption of PoW cryptocurrency mining. As long as the currency price will increase, mining activities and their energy consumption will follow.



## 7 Conclusions and future research

### 7.1 Summary of the analysis

Given the growing concerns about the energy consumption of cryptocurrency mining, but also the diversification of application domains for its underlying blockchain technology, this study set out to analyse the factors that affect the energy consumption of a blockchain, irrespective of its deployment domain. The results show unambiguously that from the three energy-consumption sources of a blockchain (storage, communication, computation), the computations of the PoW-based consensus mechanism immanent to numerous blockchain instantiations is by far the dominating factor. While

- all the coordination messages of Bitcoin, for example, require on average under 1 kW of power (equivalent to the production of a small rooftop photovoltaic system) and amount to a few MWh of electricity per year, and
- the storage of the entirety of more than 10,000 Bitcoin replicas may be responsible for an average power of 4-400 kW (but a fraction of the peak production of one modern wind turbine) inducing an energy consumption of tens of MWh to a few GWh per year,
- the PoW mechanism is responsible for an average power of around 10 GW (equivalent to the power produced by about 10 nuclear power plants) and yielding an yearly energy consumption of more than 100 TWh.

While the fact that the PoW mechanism dominates the overall energy consumption comes to little surprise, the result is more unambiguous than that: Given the numerous orders of magnitude of difference, for a blockchain deploying PoW, the other factors can be confidently neglected.

Furthermore, we have argued that the energy consumption of PoW depends only on the price of the cryptocurrency reward (in relation to the price of electricity), and not on the efficiency of the mining hardware. Without a theoretical price limit for the underlying cryptocurrency, there is thus also no theoretical upper bound for the energy consumption of the PoW mechanism.

### 7.2 Policies and technological measures to discourage PoW

As a consequence, for individual permissionless blockchains, the single crucial energy efficiency intervention is the substitution of another consensus mechanisms for PoW. At organisational or societal level, the policies can also target fostering such blockchains that do not deploy PoW, while discouraging or hindering those that do. Without the PoW consensus mechanism, a blockchain is just a replicated database, which does consume more energy than a traditional DB due to the copies of the virtual machine, but is for the moment no particular source of concern amongst the diversity of human activities, nor within the digital technologies.

Some important cryptocurrencies, such as EOS, Tezos, and TRON, already successfully deploy a PoS consensus mechanism (Sedlmeir et al. 2020b). The second-most important cryptocurrency and the foremost platform for smart contracts and NFTs, Ethereum, is transitioning towards a split into 64 shards and a PoS consensus instead of the current PoW (Finematics 2020). Together with the smart contracts and NFTs it enables, the perspective of a switch to PoS might be a reason behind ETH's apparent price decoupling from BTC over the last months.

Through the decision of accepting or not PoW-based cryptocurrencies, companies can have a large influence on the value of the respective cryptocurrencies, and thus of the energy consumption used in their mining. When Tesla announced in May 2021 it would no longer accept BTCs over climate concerns (Cellan-Jones 2021), the currency lost substantial value as a consequence, and the mining energy decreased correspondingly.



Governments can also exert substantial influence: If they decide to impose clear regulations on blockchain technology (Avan-Nomayo 2021a), perhaps posing substantial burdens on PoW-based blockchains, this would certainly make them less attractive in relation to blockchains using other consensus mechanisms. Clear regulations and legal clarity surrounding the deployment of blockchain, such as in Switzerland (Wettstein 2020) are also likely to contribute to more inherent trust in the blockchain technology, and thus less pressure in favour of the (already well tested) utmost level of security provided by PoW.

### 7.3 Future research

As ground-breaking as it indisputably was, the Swiss DLT bill did not make any technological recommendations regarding individual blockchain components, and in particular for the deployed consensus mechanism. As a consequence, even Swiss government bodies, such as the tax authority in the Canton of Zug, started accepting the two most popular cryptocurrencies, which are BTC and ETH (SRF 2020), thus unwillingly contributing to the growing energy demand of PoW-based cryptocurrencies. One of the most important areas of research is thus how and to which extent laws and regulations can and should discourage PoW-based consensus mechanisms, perhaps in a similar manner as energy efficiency mandates regulate the energy consumption of household appliances or the emissions of new vehicles.

Another domain worth exploring is the projected energy consumption of blockchain technology beyond PoW. Assuming the wide-spread adoption of PoW-free permissioned or permissionless blockchains in one application domain (such as NFT trading, DeFi, or electricity markets) with hundreds of millions or billions of transactions annually, are there any other reasons for concern? How likely is such wide-spread adoption among the application domains listed in Section 4 and possibly others? As blockchains might scale up across these domains, are there further foreseeable reasons for concern (e.g., for the message complexity of alternative consensus mechanisms or the storage of an increasing number of blockchain copies)? Or can the benefits of blockchain be carefree taken advantage of, once PoW is out of the way?

## Acknowledgements

The author wishes to thank Michael Moser and Roland Brüniger for the opportunity to perform this study, Tim Weingärtner for the insightful personal communications on the application domains of blockchain, Quinn DuPont for organising the inspiring “Uncommon Economies” lecture series at the University College Dublin and kindly providing access to it, especially to Paul Cuffe’s remarkable presentation, Andrei Bogdan Sterescu for a helpful discussion on DeFi from an economist’s perspective, Tim Weingärtner, Fabian Schlupe and Hannes Saxer for generously granting permission to use their authored or commissioned figures, and in particular Roland Brüniger for the unusually thoughtful comments and far-reaching feedback to earlier versions of this manuscript.

## References

- Alao, Olakunle, and Paul Cuffe. 2020a. “Towards a Blockchain Contract-for-Difference Financial Instrument for Hedging Renewable Electricity Transactions.” In *2020 6th IEEE International Energy Conference (ENERGYCon)*, 858–63. Gammarth, Tunis, Tunisia: IEEE. <https://doi.org/10.1109/ENERGYCon48941.2020.9236436>.
- . 2020b. “Towards a Blockchain Special Purpose Vehicle for Financing Independent Renewable Electricity Projects in Sub-Saharan Africa.” In *6th IEEE International Energy Conference*



- (ENERGYCon), 1041–46. Gammarth, Tunis, Tunisia: IEEE.  
<https://doi.org/10.1109/ENERGYCon48941.2020.9236599>.
- Allison, Ian. 2020. "MakerDAO Weighs Accepting Real-World Assets as Crypto Loan Collateral." *CoinDesk* (blog). June 4, 2020. <https://www.coindesk.com/makerdao-weighs-accepting-real-world-assets-as-crypto-loan-collateral>.
- Artnet News. 2014. "Over 50 Percent of Art Is Fake." *Artnet News* (blog). October 13, 2014. <https://news.artnet.com/market/over-50-percent-of-art-is-fake-130821>.
- Avan-Nomayo, Osato. 2021a. "Irish MEP Calls for Stringent Crypto Regulations in Europe." *Cointelegraph* (blog). June 2, 2021. <https://cointelegraph.com/news/irish-mep-calls-for-stringent-crypto-regulations-in-europe>.
- . 2021b. "Bitcoin Network Node Count Sets New All-Time High." *Cointelegraph* (blog). July 15, 2021. <https://cointelegraph.com/news/bitcoin-network-node-count-sets-new-all-time-high>.
- Bartoletti, Massimo, and Livio Pompianu. 2017. "An Empirical Analysis of Smart Contracts: Platforms, Applications, and Design Patterns." In *Financial Cryptography and Data Security*, edited by Michael Brenner, Kurt Rohloff, Joseph Bonneau, Andrew Miller, Peter Y.A. Ryan, Vanessa Teague, Andrea Bracciali, Massimiliano Sala, Federico Pintore, and Markus Jakobsson, 10323:494–509. Lecture Notes in Computer Science. Cham: Springer International Publishing. [https://doi.org/10.1007/978-3-319-70278-0\\_31](https://doi.org/10.1007/978-3-319-70278-0_31).
- Bendiksen, Christopher, and Samuel Gibbons. 2019. "The Bitcoin Mining Network - Trends, Composition, Average Creation Cost, Electricity Consumption & Sources." *CoinShares*. December 3, 2019. <https://coinshares.com/research/bitcoin-mining-network-december-2019>.
- Bitcoin.org. 2021. "Running A Full Node - Bitcoin." 2021. <https://bitcoin.org/en/full-node>.
- Bitnodes. 2021. "Global Bitcoin Nodes Distribution." 2021. <https://bitnodes.io/>.
- Blockchain.com. 2021a. "Average Block Size (MB)." *Blockchain.Com*. 2021. <https://www.blockchain.com/charts/avg-block-size>.
- . 2021b. "Blockchain Size (MB)." *Blockchain.Com*. 2021. <https://www.blockchain.com/charts/blocks-size>.
- Botz, Anneli. 2018. "Is Blockchain the Future of Art? Four Experts Weigh In." *Art Basel* (blog). 2018. <https://www.artbasel.com/stories/blockchain-artworld-cryptocurrency-cryptokitties>.
- Bouri, Elie, Rangan Gupta, Aviral Kumar Tiwari, and David Roubaud. 2017. "Does Bitcoin Hedge Global Uncertainty? Evidence from Wavelet-Based Quantile-in-Quantile Regressions." *Finance Research Letters* 23 (November): 87–95. <https://doi.org/10.1016/j.frl.2017.02.009>.
- Braun-Dubler, Nils, Hans-Peter Gier, Tetiana Bulatnikova, Manuel Langhart, Manuela Merki, Florian Roth, Antoine Burret, and Simon Perdrisat. 2020. "Blockchain: Capabilities, Economic Viability, and the Socio-Technical Environment." *TA Swiss*. <https://www.ta-swiss.ch/en/blockchain>.
- Brière, Marie, Kim Oosterlinck, and Ariane Szafarz. 2015. "Virtual Currency, Tangible Return: Portfolio Diversification with Bitcoin." *Journal of Asset Management* 16 (6): 365–73. <https://doi.org/10.1057/jam.2015.5>.
- Buterin, Vitalik. 2013. "Ethereum Whitepaper: A Next-Generation Smart Contract and Decentralized Application Platform." <https://whitepaper.io/document/5/ethereum-whitepaper>.
- Campbell, Sophie, and Amy Whitaker. 2021. "Crypto-Art Goes Mainstream: A Guide to Blockchain, NFTs & Co." *LGT Group* (blog). April 28, 2021. <https://www.lgt.com/en/magnet/lifestyle/crypto-art-goes-mainstream-a-guide-to-blockchain-nfts-amp-co/>.
- Cascone, Sarah. 2021. "Sotheby's Is Selling the First NFT Ever Minted — and Bidding Starts at \$100." *Artnet News* (blog). May 7, 2021. <https://news.artnet.com/market/sothebys-is-hosting-its-first-curated-nft-sale-featuring-the-very-first-nft-ever-minted-1966003>.
- CBECI. 2021. "Cambridge Bitcoin Electricity Consumption Index (CBECI)." 2021. <https://cbeci.org/>.
- Cellan-Jones, Rory. 2021. "Tesla Will No Longer Accept Bitcoin over Climate Concerns, Says Musk." *BBC News*, May 13, 2021, sec. Business. <https://www.bbc.com/news/business-57096305>.
- Chaum, David. 1983. "Blind Signatures for Untraceable Payments." In *Advances in Cryptology*, edited by David Chaum, Ronald L. Rivest, and Alan T. Sherman, 199–203. Boston, MA: Springer US. [https://doi.org/10.1007/978-1-4757-0602-4\\_18](https://doi.org/10.1007/978-1-4757-0602-4_18).



- Chen, Yan, and Cristiano Bellavitis. 2020. "Blockchain Disruption and Decentralized Finance: The Rise of Decentralized Business Models." *Journal of Business Venturing Insights* 13 (June): e00151. <https://doi.org/10.1016/j.jbvi.2019.e00151>.
- Clark, Mitchell. 2021. "NFTs, Explained." *The Verge*, March 11, 2021. <https://www.theverge.com/22310188/nft-explainer-what-is-blockchain-crypto-art-faq>.
- Cornish, Chloe. 2018. "CryptoKitties, CryptoPunks and the Birth of a Cottage Industry." *Financial Times*, June 6, 2018. <https://www.ft.com/content/f9c1422a-47c9-11e8-8c77-ff51caedcde6>.
- Coroamă, Vlad C. 2021. "Investigating the Inconsistencies among Energy and Energy Intensity Estimates of the Internet – Metrics and Harmonising Values." 67656. Swiss Federal Office of Energy SFOE. <https://www.aramis.admin.ch/Default?DocumentID=67656>.
- Coroamă, Vlad C, and Friedemann Mattern. 2019. "Digital Rebound – Why Digitalization Will Not Redeem Us Our Environmental Sins." In *Proc. of the 6th Int. Conf. on ICT for Sustainability (ICT4S)*. Lappeenranta, Finland. [http://ceur-ws.org/Vol-2382/ICT4S2019\\_paper\\_31.pdf](http://ceur-ws.org/Vol-2382/ICT4S2019_paper_31.pdf).
- Crockett, Zachary. 2021. "Why One Guy Paid \$208k for a Video Clip of LeBron James Dunking." *The Hustle*, March 7, 2021. <https://thehustle.co/why-one-guy-paid-208k-for-a-video-clip-of-lebron-james-dunking/>.
- "CryptoPunks." 2021. <https://www.larvalabs.com/cryptopunks>.
- Cuffe, Paul. 2021. "The Electricity Industry: It's Boringly Reliable and We Mostly Trust the Incumbents. So, Is There a Role for Blockchain at All?" UC Dublin, January 12. <https://www.smurfitschool.ie/facultyresearch/cito/activities/>.
- Dale, Brady. 2021. "MakerDAO Moves to Full Decentralization; Maker Foundation to Close in 'Months.'" *CoinDesk* (blog). July 20, 2021. <https://www.coindesk.com/makerdao-moves-to-full-decentralization-maker-foundation-to-close-in-months>.
- Dale, Brady, and Kevin Reynolds. 2021. "Bidding Reaches \$2.5M as Twitter's Dorsey Highlights NFT Version of First-Ever Tweet." *CoinDesk*, March 6, 2021. <https://www.coindesk.com/twitter-ceo-jack-dorsey-is-offering-to-sell-the-first-ever-tweet>.
- Davis, Ben. 2021. "I Looked Through All 5,000 Images in Beeple's \$69 Million Magnum Opus. What I Found Isn't So Pretty." *Artnet News* (blog). March 17, 2021. <https://news.artnet.com/opinion/beeple-everydays-review-1951656>.
- DeFi Pulse. 2021. "DeFi Pulse | The DeFi Leaderboard | Stats, Charts and Guides." 2021. <https://defipulse.com>.
- Diffie, W., and M. Hellman. 1976. "New Directions in Cryptography." *IEEE Transactions on Information Theory* 22 (6): 644–54. <https://doi.org/10.1109/TIT.1976.1055638>.
- Digiconomist. 2021. "Bitcoin Energy Consumption Index." Digiconomist. 2021. <https://digiconomist.net/bitcoin-energy-consumption/>.
- Douceur, John R. 2002. "The Sybil Attack." In *Peer-to-Peer Systems*, edited by Peter Druschel, Frans Kaashoek, and Antony Rowstron, 2429:251–60. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer Berlin Heidelberg. [https://doi.org/10.1007/3-540-45748-8\\_24](https://doi.org/10.1007/3-540-45748-8_24).
- DuPont, Quinn. 2017. "Experiments in Algorithmic Governance: A History and Ethnography of 'The DAO,' a Failed Decentralized Autonomous Organization." In *Bitcoin and Beyond: Cryptocurrencies, Blockchains and Global Governance*, 157–77. Routledge.
- Eklund, Peter W., and Roman Beck. 2019. "Factors That Impact Blockchain Scalability." In *Proceedings of the 11th International Conference on Management of Digital EcoSystems*, 126–33. Limassol Cyprus: ACM. <https://doi.org/10.1145/3297662.3365818>.
- Entriiken, William, Dieter Shirley, Jacob Evans, and Nastassia Sachs. 2018. "ECR-721 Non-Fungible Token Standard." EIP-721. <https://eips.ethereum.org/EIPS/eip-721>.
- Ethereum, Developer Resources. 2021a. "Ethereum Accounts." 2021. <https://ethereum.org/en/developers/docs/accounts/>.
- . 2021b. "Ethereum Virtual Machine (EVM)." 2021. <https://ethereum.org/en/developers/docs/evm/>.
- . 2021c. "Non-Fungible Tokens (NFT)." 2021. <https://ethereum.org/en/nft/>.
- . 2021d. "Proof-of-Stake (PoS)." 2021. <https://ethereum.org/en/developers/docs/consensus-mechanisms/pos/>.





- . 2021e. “Smart Contract Languages.” 2021. <https://ethereum.org/en/developers/docs/smart-contracts/languages/>.
- Everledger. 2021. “Everledger Wins Major Australian Government Critical Minerals Blockchain Pilot Project.” Everledger. July 12, 2021. <https://everledger.io/everledger-wins-major-australian-government-critical-minerals-blockchain-pilot-project/>.
- Finematics. 2020. *Ethereum 2.0 -A Game Changer? Proof Of Stake, The Beacon Chain, Sharding, Docking Explained*. [https://www.youtube.com/watch?v=ctzGr58\\_jel](https://www.youtube.com/watch?v=ctzGr58_jel).
- Grundlehner, Werner. 2021a. “Die digitale Aktie ist da – und es läuft nicht wie erwartet.” *Neue Zürcher Zeitung*, March 16, 2021. <https://www.nzz.ch/finanzen/etwas-picasso-ins-portfolio-legen-id.1635718>.
- . 2021b. “Schweizer Bank digitalisiert Meisterwerk: Ein Stück Picasso ins Portfolio legen.” *Neue Zürcher Zeitung*, July 15, 2021. <https://www.nzz.ch/finanzen/etwas-picasso-ins-portfolio-legen-id.1635718>.
- Guerra, Raquel. 2021. “MEPs Push for Cryptocurrencies Generated via Sustainable Technology.” June 4, 2021. <https://www.endseurope.com/article/1718188/meps-push-cryptocurrencies-generated-via-sustainable-technology>.
- Harper, Jim. 2013. “What Is the Value of Bitcoin?” *Cato at Liberty* (blog). April 5, 2013. <https://www.cato.org/blog/what-value-bitcoin>.
- Hertig, Alyssa. 2020. “What Is DeFi?” *CoinDesk* (blog). September 18, 2020. <https://www.coindesk.com/what-is-defi>.
- Higgins, Stan. 2017. “From \$900 to \$20,000: Bitcoin’s Historic 2017 Price Run Revisited.” *CoinDesk* (blog). December 29, 2017. <https://www.coindesk.com/900-20000-bitcoins-historic-2017-price-run-revisited>.
- Hintemann, Ralph, and Simon Hinterholzer. 2019. “Energy Consumption of Data Centers Worldwide - How Will the Internet Become Green?” In *Proc. of the 6th Int. Conf. on ICT for Sustainability (ICT4S)*. Lappeenranta, Finland. [http://ceur-ws.org/Vol-2382/ICT4S2019\\_paper\\_16.pdf](http://ceur-ws.org/Vol-2382/ICT4S2019_paper_16.pdf).
- Jentzsch, Christoph. 2016. “Decentralized Autonomous Organization to Automate Governance.” *slock.it*. <https://lawofthelevel.lexblogplatformthree.com/wp-content/uploads/sites/187/2017/07/WhitePaper-1.pdf>.
- Jones, Tyler. 2021. “How to Run a Bitcoin Full Node on a Raspberry Pi.” *Howchoo* (blog). March 30, 2021. <https://howchoo.com/bitcoin/run-bitcoin-full-node-raspberry-pi>.
- Joos, Thomas, and Peter Schmitz. 2021. “Use Cases für den Einsatz der Blockchain.” *IT-Business*, March 10, 2021. <https://www.it-business.de/use-cases-fuer-den-einsatz-der-blockchain-a-1005933/>.
- Kamiya, George. 2019. “Bitcoin Energy Use - Mined the Gap.” IEA. July 5, 2019. <https://www.iea.org/commentaries/bitcoin-energy-use-mined-the-gap>.
- Kastrenakes, Jacob. 2021a. “Grimes Sold \$6 Million Worth of Digital Art as NFTs.” *The Verge*, March 1, 2021. <https://www.theverge.com/2021/3/1/22308075/grimes-nft-6-million-sales-nifty-gateway-warnymph>.
- . 2021b. “Beeple Sold an NFT for \$69 Million.” *The Verge*, March 11, 2021. <https://www.theverge.com/2021/3/11/22325054/beeple-christies-nft-sale-cost-everydays-69-million>.
- Kochkodin, Brandon, and Olga Kharif. 2021. “Cryptocurrency Millionaires Fuel a Boom in Digital Art Market.” *Bloomberg.Com*, February 21, 2021. <https://www.bloomberg.com/news/articles/2021-02-26/cryptocurrency-millionaires-fuel-a-boom-in-digital-art-market>.
- Krause, Max J., and Thabet Tolaymat. 2018. “Quantification of Energy and Carbon Costs for Mining Cryptocurrencies.” *Nature Sustainability* 1 (11): 711–18. <https://doi.org/10.1038/s41893-018-0152-7>.
- Kühl, Eike. 2016. “Blockchain: Und plötzlich fehlen 50 Millionen Dollar.” *Die Zeit*, June 20, 2016, sec. Digital. [https://www.zeit.de/digital/internet/2016-06/the-dao-blockchain-ether-hack?utm\\_referrer=https%3A%2F%2Fwww.google.com%2F](https://www.zeit.de/digital/internet/2016-06/the-dao-blockchain-ether-hack?utm_referrer=https%3A%2F%2Fwww.google.com%2F).
- Lamport, Leslie. 1983. “The Weak Byzantine Generals Problem.” *Journal of the ACM* 30 (3): 668–76. <https://doi.org/10.1145/2402.322398>.



- . 1998. “The Part-Time Parliament.” *ACM Transactions on Computer Systems* 16 (2): 133–69. <https://doi.org/10.1145/279227.279229>.
- Luu, Loi, Viswesh Narayanan, Chaodong Zheng, Kunal Baweja, Seth Gilbert, and Prateek Saxena. 2016. “A Secure Sharding Protocol For Open Blockchains.” In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 17–30. Vienna Austria: ACM. <https://doi.org/10.1145/2976749.2978389>.
- Manrique, Sharon. 2019. “What Is Dai, and How Does It Work?” *MyCrypto* (blog). March 6, 2019. <https://medium.com/mycrypto/what-is-dai-and-how-does-it-work-742d09ba25d6>.
- Mark, Dino, Vlad Zamfir, and Emin Gün Sirer. 2016. “A Call for a Temporary Moratorium on The DAO.” *Hacking, Distributed* (blog). May 27, 2016. <https://hackingdistributed.com/2016/05/27/dao-call-for-moratorium/>.
- Martin, Katie, and Billy Nauman. 2021. “Bitcoin’s Growing Energy Problem: ‘It’s a Dirty Currency.’” *Financial Times*, May 20, 2021. <https://www.ft.com/content/1aecb2db-8f61-427c-a413-3b929291c8ac>.
- Masanet, Eric, Arman Shehabi, Nuo Lei, Sarah Smith, and Jonathan Koomey. 2020. “Recalibrating Global Data Center Energy-Use Estimates.” *Science* 367 (6481): 984–86. <https://doi.org/10.1126/science.aba3758>.
- McKenz, André François. 2021. “Sustainability Solution or Climate Calamity? The Dangers and Promise of Cryptocurrency Technology.” *UN News* (blog). June 20, 2021. <https://news.un.org/en/story/2021/06/1094362>.
- Mengelkamp, Esther, Johannes Gärtner, Kerstin Rock, Scott Kessler, Lawrence Orsini, and Christof Weinhardt. 2018. “Designing Microgrid Energy Markets: A Case Study: The Brooklyn Microgrid.” *Applied Energy* 210 (January): 870–80. <https://doi.org/10.1016/j.apenergy.2017.06.054>.
- Merkle, Ralph C. 1988. “A Digital Signature Based on a Conventional Encryption Function.” In *Advances in Cryptology — CRYPTO ’87*, edited by Carl Pomerance, 293:369–78. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer Berlin Heidelberg. [https://doi.org/10.1007/3-540-48184-2\\_32](https://doi.org/10.1007/3-540-48184-2_32).
- Musk, Elon. 2021. “I’m Selling This Song about NFTs as an NFT.” Tweet. *@elonmusk* (blog). March 15, 2021. <https://twitter.com/elonmusk/status/1371549960030842893>.
- Nakamoto, Satoshi. 2008. “Bitcoin: A Peer-to-Peer Electronic Cash System.” <https://bitcoin.org/bitcoin.pdf>.
- Palmer, Daniel. 2021. “Greenpeace Stops Accepting Bitcoin Donations, Cites High Energy Use.” *CoinDesk* (blog). May 21, 2021. <https://www.coindesk.com/greenpeace-says-it-will-stop-accepting-bitcoin-donations-cites-energy-use>.
- Quote Investigator. 2013. “It’s Difficult to Make Predictions, Especially About the Future.” 2013. <https://quoteinvestigator.com/2013/10/20/no-predict/>.
- Romeo, Jess. 2021. “What’s the Deal with Crypto Art?” *JSTOR Daily* (blog). April 9, 2021. <https://daily.jstor.org/whats-the-deal-with-crypto-art/>.
- Samuelson, Paul, and William Nordhaus. 2009. *Economics*. 19th edition. Boston: McGraw-Hill Education.
- Sedlmeir, Johannes, Hans Ulrich Buhl, Gilbert Fridgen, and Robert Keller. 2020a. “Ein Blick auf aktuelle Entwicklungen bei Blockchains und deren Auswirkungen auf den Energieverbrauch.” *Informatik Spektrum* 43 (6): 391–404. <https://doi.org/10.1007/s00287-020-01321-z>.
- . 2020b. “The Energy Consumption of Blockchain Technology: Beyond Myth.” *Business & Information Systems Engineering* 62 (6): 599–608. <https://doi.org/10.1007/s12599-020-00656-x>.
- Segal, Jeff, and Jeffrey Goldfarb. 2009. “Art Appreciated as Inflation Hedge.” *The New York Times*, May 11, 2009, sec. Business. <https://www.nytimes.com/2009/05/12/business/12views.html>.
- Shehabi, Arman, Sarah Josephine Smith, Dale A. Sartor, Richard E. Brown, Magnus Herrlin, Jonathan G. Koomey, Eric R. Masanet, Nathaniel Horner, Inês Lima Azevedo, and William Lintner. 2016. “United States Data Center Energy Usage Report.” LBNL-1005775. Lawrence Berkeley National Laboratory. [https://eta-publications.lbl.gov/sites/default/files/lbnl-1005775\\_v2.pdf](https://eta-publications.lbl.gov/sites/default/files/lbnl-1005775_v2.pdf).



- Sherman, Erik. 2021. "After The Hype: The Future For NFTs." *Forbes*, March 30, 2021. <https://www.forbes.com/sites/zengernews/2021/05/30/after-the-hype-the-future-for-nfts/>.
- Smart Energy International. 2021. "Blockchain Battery Lifecycle Managements Olution Launched," March 18, 2021. <https://www.enlit-europe.com/news-data/blockchain-battery-lifecycle-management-solution-launched>.
- Smith, Adam. 1776. *An Inquiry into the Nature and Causes of the Wealth of Nations*. London, UK: Strahan and Cadell.
- SRF. 2020. "Steuern zahlen mit Bitcoin - Kanton Zug akzeptiert Kryptowährungen bei Steuern." Schweizer Radio und Fernsehen (SRF). September 3, 2020. <https://www.srf.ch/news/regional/zentralschweiz/steuern-zahlen-mit-bitcoin-kanton-zug-akzeptiert-kryptowaehrungen-bei-steuern>.
- Stoll, Christian, Lena Klaaßen, and Ulrich Gallersdörfer. 2019. "The Carbon Footprint of Bitcoin." *Joule* 3 (7): 1647–61. <https://doi.org/10.1016/j.joule.2019.05.012>.
- Swiss Confederation. 2020. *Bundesgesetz zur Anpassung des Bundesrechts an Entwicklungen der Technik verteilter elektronischer Register*. Vol. 19.074. <https://www.fedlex.admin.ch/eli/fga/2020/2007/de>.
- Villiers, Almero de, and Paul Cuffe. 2020. "A Three-Tier Framework for Understanding Disruption Trajectories for Blockchain in the Electricity Industry." *IEEE Access* 8: 65670–82. <https://doi.org/10.1109/ACCESS.2020.2983558>.
- Vogelsteller, Fabian, and Vitalik Buterin. 2015. "ERC-20 Token Standard." EIP-20. <https://eips.ethereum.org/EIPS/eip-20>.
- Vries, Alex de. 2018. "Bitcoin's Growing Energy Problem." *Joule* 2 (5): 801–5. <https://doi.org/10.1016/j.joule.2018.04.016>.
- . 2020. "Bitcoin's Energy Consumption Is Underestimated: A Market Dynamics Approach." *Energy Research & Social Science* 70 (December): 101721. <https://doi.org/10.1016/j.erss.2020.101721>.
- Wettstein, Frank. 2020. "Federal Council Brings Part of DLT Bill into Force." *The Federal Council* (blog). December 11, 2020. <https://www.sif.admin.ch/sif/en/home/dokumentation/medienmitteilungen/medienmitteilungen.msg-id-81563.html>.
- Wong, Joon Ian. 2017. "The UN Is Using Ethereum's Technology to Fund Food for Thousands of Refugees." *Quartz*, November 3, 2017. <https://qz.com/1118743/world-food-programmes-ethereum-based-blockchain-for-syrian-refugees-in-jordan/>.
- Wood, Gavin. 2014. "Ethereum: A Secure Decentralised Generalised Transaction Ledger." <https://ethereum.github.io/yellowpaper/paper.pdf>.
- Wright, Turner. 2020. "William Shatner's NFT Collectibles Sell Out at Warp Speed." *Cointelegraph*, July 30, 2020. <https://cointelegraph.com/news/william-shatners-nft-collectibles-sell-out-at-warp-speed>.
- Wüst, Karl, and Arthur Gervais. 2018. "Do You Need a Blockchain?" In *2018 Crypto Valley Conference on Blockchain Technology (CVCBT)*, 45–54. Zug: IEEE. <https://doi.org/10.1109/CVCBT.2018.00011>.
- Yasar, Brad Bulent. 2021. "A Brief History of NFTs and a Look into the Future." May 13, 2021. <https://www.linkedin.com/pulse/brief-history-nfts-look-future-brad-bulent-yasar/>.
- Zetsche, Dirk A, Douglas W Arner, and Ross P Buckley. 2020. "Decentralized Finance." *Journal of Financial Regulation* 6 (2): 172–203. <https://doi.org/10.1093/jfr/fjaa010>.