Towards an Analytical Discipline of Forkonomy[⋆]

Dr. Wassim Z. Alsindi[1][0000−0002−6701−0655]

Parallel Industries wassim@pllel.com
http://www.pllel.com

Abstract. This work introduces a novel field of cryptocurrency research that the author terms *forkonomy*, and provides a general overview of recent phenomena in this area. Attention is directed towards the first UTXO consolidation *fork-merge* combining Zclassic and Bitcoin ledger histories into the so-called Bitcoin Private network. Potential implications for ageing blockchain ecosystems, prominent minority cryptocurrency network fragments and divergent factions are discussed.

Keywords: Forkonomy · Cryptocurrency · Blockchain.

[⋆] Final revision: 10th August 2018. Funded by Parallel Industries. *Secundum Aquinas: caveat lector, veritas est adequatio rei et intellectum.*

Table of Contents

Sections

1 Introduction and Literature Review: The Hitherto Canon of Forkonomy

1.1 Forkonomy, Forks and Forkability

With respect to cryptocurrencies, *forkonomy* can be considered to constitute the study of the fragmentation of software codebases and protocol networks comprising distributed communities and/or stakeholders operating in a *permissionless* or *trust-minimised* manner. Much as astronomy utilises observation and theory to understand and predict cosmological characteristics and phenomena, here follows an analogous attempt to apply blockchain analytics and historical precedent to with a view to understanding fundamental and emergent characteristics of the forking tendencies of divergent monetary network factions.

In the open source computer science domain, the notion of project codebase forks is well established and occurs when an existing piece of software develops in diverging paths by independent developer constituencies, creating separate and distinct pieces of software. Torvalds' original *Linux* kernel from 1991 has been forked into countless descendant projects [1]. With the launch of the *Bitcoin* network in 2009, the prospect of provable digital scarcity and secure decentralised open source value transfer protocols was realised. This was implemented through the novel combination of systems networking, *UTXO* (Unspent Transaction Output) based accounting, resilient data architecture, cryptography and thermodynamic elements [2]. With a permissionless ledger system employing a *blockchain* and triple-entry accounting to reach a high degree of probabilistic transaction finality over time, there exists the prospect of both *codebase* and *ledger* forks [3]. For the purposes of this work, a blockchain is defined as a temporally sequenced, linear and append-only data structure employing cryptography to facilitate the implementation of a high assurance, tamper-evident transaction ledger.

A codebase fork of a cryptocurrency corresponds closely to the relationship between Linux kernel forks, creating an independent project typically launched with a new genesis block which may share consensus rules but with an entirely different transaction history than its progenitor. An example of this relationship type is that between Bitcoin (BTC) and Litecoin (LTC) and this method may be thought of as a *static fork* insofar as there is little time-sensitivity to the process. By contrast, a ledger fork creates a separate incompatible network, sharing its history with the progenitor network until the divergent event, commonly referred to as a *chain split.* Consensus rule changes or alteration of the network transaction history may be the cause of such a fracture, deliberate or unplanned. This occurrence may be regarded as a *dynamic fork* since the process takes place in real time. Often when networks upgrade software, consensus rules or implement new features a portion of the network participants may be left behind on a *vestigial* timeline that lacks developer, community, wallet or exchange support. Recently a fifth of nodes running Bitcoin Cash (BCH) - a *SHA-256* minority ledger fork of BTC with significantly relaxed block size limitations - were separated from the BCH network and a non-trivial number of would-be nodes remain disconnected from the canonical BCH blockchain at time of writing weeks later [4].

1.2 What Maketh a Fork?

The distinction between what constitutes a vestigial network and a viable breakaway faction is unclear and difficult to objectively parameterise. There is a significant element of adversarial strategy, political gamesmanship and public signalling of (real or synthetic) intent and support *via* social media platforms. The notions of critical mass and stakeholder buy-in are ostensibly at play since ecosystem fragmentations would be characterised as *strongly negative sum* through the invocation of Metcalfe's Law as regards network effects and hence value proposition [5]. Any blockchain secured thermodynamically by Proof-of-Work (PoW) is susceptible to attack vectors such as so-called *51 %* or *majority attacks*, leading to *re-orgs* (chain re-organisations) as multiple candidates satisfying chain selection rules emerge. These can result in the potential for *double-spending* the same funds more than once against entities such as exchanges who do not require sufficient confirmations for transaction finality to be reliable in an adversarial context. Should a network fragment into multiple disconnected populations, adversaries with control of much less significant computational resource would be in reach of majority hashrate either using permanent or rented computation from sources such as *Nicehash* or *Amazon EC3* [6].

A striking example of this was the divergence of the *Ethereum* developer and leadership cadre (ETH) from the canonical account-oriented Ethereum blockchain (ETC) due to the exploitation of a flawed smart contract project resembling a *quasi-securitised* decentralised investment fund known as *The DAO* (Decentralised Autonomous Organisation) [7]. In this case the Ethereum insiders decided to sacrifice immutability and by extension censorship-resistance in order to conduct an effective bailout of DAO participants which came to exercise *Too-Big-To-Fail* influence over the overall Ethereum network, insider asset holdings, token supply and mindshare [8]. A social media consultation process in conjunction with *on-chain voting* was employed to arrive at this conclusion though both methods are known to be flawed and gameable [9]. During the *irregular state transition* process akin to a rollback, a co-ordinated effort between miners, exchanges and developers took place on private channels, exposing the degree of centralisation inherent in the power structures of constituent network participants. The key event which transformed the canonical Ethereum blockchain (where the DAO attacker kept their spoils) from a vestigial *wiped out* chain to a viable if contentious minority fork was the decision by Bitsquare and Poloniex exchanges to list the attacker's timeline as Ethereum Classic (ETC) alongside high-profile mining participants such as Chandler Guo, well resourced financial organisations such as Grayscale Invest (a subsidiary of Digital Currency Group) and former development team members such as Charles Hoskinson to publically declare and deploy support, developers and significant hashrate to defend the original Ethereum network [10]. ETC now exists as an independent and sovereign network with diverging priorities, characteristics and goals to ETH as discussed in Section 4.

1.3 Transient Fork Dynamics in PoW Networks

At a granular level, blockchains grow in height incrementally as new valid blocks are found by miners or validators and added to the canonical chain as determined by the network's chain selection rules. In PoW consensus mechanisms this leaderless race is conducted through the combination of *nonces* (an arbitrary variable cycled through sequentially) with the proposed *block header* to generate hashes which are then compared against the *network difficulty* which is closely related to the quantity of computational resource directed at the network. Should a hash be found that is *below* the network's difficulty requirements, given that no other consensus rules have been violated in the process of constructing the candidate block then it is typically considered valid by the network. As the miner announces the proposed block it propagates across the network typically *via* a gossip protocol, whereby nodes broadcast all messages to connected peers.

Since cryptographic hash functions are deterministic (albeit with with unpredictable outputs) and a broad subset of possible hash values satisfying the difficulty requirements exist, it is entirely plausible that more than one valid candidate block may be found by competing miners at very similar times. In such an eventuality there begins a *block propagation competition* of sorts which serves to allow the network to reach consensus on the latest state of the transaction ledger. Since there can only be one block with a particular height, should multiple candidates emerge the prospect of network partition arises if subsets of the population of validating nodes do not overwhelmingly agree on the latest block. Such partitions may be short-lived in the case of *orphans* and *uncles* which represent discarded timelines as the canonical chain built upon another candidate block. The term uncle is used primarily in Ethereum-based networks, as a partial subsidy is allocated to orphaned blocks and therefore acts as a consolation prize for producing a valid block which does not become part of the canonical chain. Ethereum currently subsidises uncles with approximately 3000 ETH per day which equates to over 1 million USD at time of writing [11]. Increasing orphan rates may also be indicative of malicious behaviour on a network such as *51 % attacks*, *selfish mining* or *distributed denial of service* vectors on reachable nodes which accept incoming peer connections.

Due to the message propagation characteristics of *partially synchronous* distributed systems such as *peer-to-peer* (P2P) cryptocurrency networks, there exists an inverse relationship between the median inter-block time (more commonly referred to as the *block time*) as set by the protocol - 600 seconds in BTC/BCH and 15 seconds in ETH/ETC - and the incidence of orphans and uncles. With shorter block times the likelihood of orphan blocks increases, with some mitigating effect possible through miners aggregating together co-operatively into so-called *mining pools*. A similar effect of increasing orphan rate would also be expected should the utilisation of block capacity also increase, as larger amounts of information must propagate around the network nodes. ETH uncle rates have been increasing since October 2017 due to mining subsidy reduction, network congestion and increasing block size, whilst ETC's has remained more consistent, due at least in part to the lower transactional volume on the canonical Ethereum chain [12].

There may be a fundamental basis rooted in natural science that provides insight into the increasing forking tendencies of blockchains. These phenomena may be a result of *entropic bias*, that is to say divergent paths are those of least resistance in accordance with Newtonian physics. The second law of thermodynamics states that the total entropy (energy unavailable to do useful work) of a closed system undergoing an irreversible process can never decrease. In other words, all that can be done is to arrest the descent of order into chaos is to continue applying effort so as not to allow the amount of available energy to decrease. In the context of network forks, a simple model may be constructed of a PoW cryptocurrency network as a closed thermodynamic system with a growing blockchain (an irreversible process) with mining participants' cryptographic hashing as the work going into the system. Taking this a step further, despite the ongoing work in the system a chain split would satisfy the second law of thermodynamics as it pertains to increasing disorder in a system. Therefore it may be the case that the *energetic dynamics* of cryptocurrency networks provides a rational basis for the eventuality of ledger forks in networks which do not strongly penalise or prevent them.

Another issue widely encountered with ledger forks are *replay attacks*. In the case where two recently partitioned network fragments share identical or very similar codebases and transaction histories, unless specific measures are taken there exists the very real prospect that a network user wanting to send cryptocurrency may inadvertently send the transaction on *both* network fragments and therefore have the transaction accidentally replayed. Replay protection may be achieved through a small codebase change which allows networks to distinguish transactions as arising from one particular fragment. A related issue which may see an increase in incidence as a result of the development of protocols facilitating the issuance of non-native assets, tokens and off-chain payment channels atop blockchains is the lack of precedence in the event of a fork and chain split in the base layer. As off-chain protocols proliferate and grow in intricacy, functionality and interoperability this issue is likely to increase in complexity.

Selfish mining - also known as *block withholding* - is a postulated attack vector most effectively employed by mining oligopolists on a PoW network with relatively long block times. It may be conducted by a miner who finds a valid block but instead of immediately broadcasting to peers, the block is withheld and kept secret. The miner then begins to search for a valid block atop the previous clandestine block, with the aim of finding a valid second block before another participant finds an alternative valid first block. It has been claimed that this strategy is more beneficial than *honest mining* for a sufficiently well-resourced adversary, with 2013 research finding that Bitcoin is vulnerable to block withholding attacks when an adversarial entity controls as little as a quarter of the total computational resource possessed by the network. [13] Naturally this is a far lower bound than the majority hashrate required for 51 % attacks. However the efficacy of this attack vector has been disputed more recently with findings that the strategy only performs well in the period immediately after a difficulty adjustment. With that in mind, a fairly minor change to the Bitcoin protocol (albeit requiring upgrade consensus) could be effected to mitigate the possibility of this attack [14].

Selfish mining is potentially relevant to forks as chain splits may be more likely in the presence of selfish mining participants. A possible heuristic for selfish mining is the issuance of empty blocks (to capture efficiency in propagation time) that Bitmain-controlled mining pool Antpool regularly mined for long periods of time despite network congestion and foregoing transaction fees, indicating a potential benefit greater than an honest miner's payoff of block reward and transaction fees [15]. There is evidence that a selfish mining attack possibly took place in May 2018 on Monacoin, a Japanese cryptocurrency network, with a succession of blocks only containing the coinbase (mining subsidy) transaction between block heights 1329837 and 1329846. However it is not straightforward to differentiate between 51 % and selfish mining attack vectors as the culprit definitively. As Monacoin's difficulty adjustment occurring every block the window of opportunity for selfish mining is somewhat limited and the attacker's spoils corresponded to less than 100000 USD at time of the attack [16]. Stubborn mining builds on this methodology to facilitate a wider range of hybrid strategies between honest and selfish mining extremes [17].

Zhang *et al.* proposed a selfish mining disincentivisation and fork-resolving policy improvement for BTC chain selection ruleset having explored *censorship-attack* vectors such as *blacklisting via feather-forking* [18] as originally characterised hypothetically by Miller in 2013 [19]. Feather-forking can be understood as a strategy available to mining participants (more likely pools than individual entities) to refuse to construct blocks atop a timeline which contains unfavourable transactions within the recent history. By doing so the feather-forking participant may also incentivise other mining participants to also join the feather-fork for a short time. However this vector is rendered ineffective provided that a majority of the computational resource remains honest. Zhang and coauthors propose a mitigating upgrade to Bitcoin named *Publish or Perish* which would slightly modify the chain selection rule to include all hashes of orphaned blocks in the block currently being worked upon. However the stringent synchronicity assumptions in the proposed initial framework do no match the characteristics of typical cryptocurrency networks and no provision is made against chain splits or intentional forks [20].

1.4 Forks and Network Governance

For a range of reasons, there is often strident resistance to *hard forks* - irreversible protocol upgrades or relaxing of the existing consensus ruleset - in trust-minimised cryptocurrency networks such as BTC. The lack of controlling entities may lead to a chain split and network partition if the delicate balance of orthogonal stakeholder incentives fails in the presence of a potential divergent event. The implementation of Segregated Witness (SegWit) by the BTC network was eventually achieved in 2017 as a backward-compatible *soft fork* following several years of intense political and strategic manoeuvring by the constituent stakeholders in the BTC network. This off-chain governance process of *emergent consensus* requiring *de facto* supermajority or unanimity measured by miner signalling has proven to be an inefficient and gameable mechanism for administering the BTC network [21].

Certain stakeholder constituencies such as the developers maintaining the reference Bitcoin Core software client implementation of BTC could not easily reach agreement with mining oligopolists and so-called *big block advocates* over the optimum technological trajectory for the BTC network. The solution combined a fix for *transaction malleability* and network capacity increase through the restructuring of block contents, principally through the addition of a second Merkle tree which includes *witness* (signature) data but excludes coinbase transactions. This was initially conceived as a hard fork, and was only found to be implementable as an opt-in soft fork due to inventive engineering. Despite this, major stakeholders of the mining constituency strongly opposed SegWit as it would render a previously clandestine proprietary efficiency advantage known as *covert ASICBoost* ineffective on the canonical BTC chain [22]. A grassroots BTC community movement campaigning for a so-called *User Activated Soft Fork* (UASF) for SegWit implementation and a face-saving Bitcoin Improvement Proposal (BIP91) from mining farm operator James Hilliard in tandem facilitated the eventual lock-in of the SegWit upgrade in the summer of 2017 [23].

A new and contentious network partition took place in August 2017 as SegWit locked in for later activation, giving rise to the Bitcoin Cash (BCH) network which rejected SegWit and instead opted for linear on-chain scaling. This was implemented in the form of block size increases which have the effect of externalising network resource burden onto node operators, chiefly in the form of increased bandwidth and storage performance requirements. BCH continues to be regarded as a hostile ledger fork of BTC owing to its constituency of high-profile personalities claiming that their network more closely resembles the initial whitepaper specification of the Bitcoin protocol [24] and therefore qualifies as the ”real Bitcoin”. By contrast, PoW - also known as *Nakamoto consensus* - selects the canonical BTC blockchain as the chain with the most accumulated difficulty that satisfies the consensus rules as laid out in the original Satoshi client codebase and Bitcoin whitepaper. By changing the block size and loosening the consensus ruleset without overwhelming agreement from all constituencies of the BTC network, it is difficult to find a basis for BCH proponents’ claims to be the canonical Bitcoin blockchain without invoking appeals to emotion, authority or other logical fallacies. The continuing presence of Craig S. Wright and his claims to be a progenitor of Bitcoin are an example of these attempts at legitimacy [25], though these claims do appear to be substantially weakening.

1.5 Forks and Networks Employing Proof-of-Stake

Alternatives to Nakamoto consensus such as Proof-of-Stake (PoS) and various approaches to *Byzantine Fault Tolerance* (BFT) are the subject of active exploration in distributed systems research and development. In foregoing the utilisation of brute thermodynamic force to secure the network, PoS consensus protocols must satisfy through alternative means the properties of *persistence* and *liveness*. Persistence pertains to the immutability of the transaction history and liveness relates to network synchrony, in that valid transactions will be included in the ledger reliably.

Algorand promises *fork-resistance* through a novel block minting process employing an accelerated BFT mechanism with constantly changing committees being tasked with block proposal privileges. This protocol has yet to be implemented in a permissionless setting and concerns persist over intellectual property protection and the architecture of stakeholder incentives within the network [26] as there is currently no provision for *validator subsidy* upon block creation. In pure Proof-of-Stake systems such as *Ouroboros* there is no thermodynamic element to assign block creation privileges and instead rights are conferred based on control of coin balances. This results in a different set of fork-based challenges to PoW-oriented networks discussed above.

The *nothing-at-stake problem* arises from the lack of significant resource cost in maintaining multiple timelines in a pure PoS network. In PoW networks resource must be committed to find valid blocks and therefore a significant penalty exists for malicious actors to maintain multiple blockchain timelines. In PoS this penalty is small or absent and therefore it is feasible to proliferate multiple timelines branching from various points in the chain with little drawback if one such fork fails and is not built upon substantially. Nothing-at-stake also raises the possibility of re-orgs should an adversary acquire enough "*old stake*" from wallets that no longer control balances in the current ledger but previously did. Once sufficient old stake is amassed, the user can then begin to build upon alternative timelines in order to outrun the honest timeline and therefore become the canonical chain should the selection rules not provide protection against this approach. The *long-range attack* employs nothing-at-stake to seed Byzantine network nodes with dishonest timelines such that a node joining the network can face significant challenges in determining which is the canonical blockchain. *Stake grinding* is an attack vector class observed in early PoS implementations employed by Blackcoin, Peercoin and NXT, where block validators take measures to game the "randomness" of validator selection and/or block creation privileges in their favour by *grinding* - or sequentially searching through parameter space - for a dishonest edge over the intended working of the block creation mechanism [27]. The *Cardano* network's proposed PoS-based consensus mechanism family *Ouroboros* claims to have addressed these attack vectors by employing sophisticated cryptographic elements such as *Verifiable Random Functions* and *Genesis Proofs* to facilitate *stake-based finality*, *provable security* and *dynamic availability* such that nodes may join the network at any time and *bootstrap from genesis*. However implementation into the public Cardano network has yet to take place, so the security model of Ouroboros is yet to be tested in the wild [28].

Given the significant downside potential of real and perceived threats to the resilience and legitimacy of a fragmenting network and loss of associated network effects, the ability of a blockchain-based protocol network to demonstrate fork resistance provides significant strength to its value proposition. Decred is an example of a hybrid PoW/PoS monetary network which is implementing an off-chain proposal and governance mechanism termed *Politeia* [29]. Since coinholders have voting rights based on stake weight, they have the ability to keep miners and developer constituencies honest through the mechanism to reach decisions by majority stakeholder consensus on matters including hard forks. These lessons were ostensibly learned through the developer team's experiences in writing a BTC client which they felt was not appraised objectively by the Bitcoin Core developer ecosystem. Decred's fork resistance is effectively achieved by the fact that most stakeholders would be non-voting on a minority chain, it would remain stalled as blocks would not be created or propagated across the upstart network.

Recently another class of fork has emerged, caused by factionalisation *before* networks launch and/or code is open sourced. These appear similar to contentious political factions in existing blockchain networks though there is little concrete information in the public sphere. Recently several distinct entities have arisen within the pre-functional *Tezos* ecosystem who do not support the decisions of Dynamic Leger Solutions (DLS) as they move towards launching their mainnet, particularly regarding the recent decision to require de-anonymising *Know-Your-Customer* (KYC) information from their 2017 token offering donations taken last year which raised the equivalent of several hundred million USD. Aside from the ostensible paradox of rather security-like donations requiring *Anti-Money Laundering* (AML) procedures for future claims on the DLS-Tezos network, at the time of writing three alternative proposed non-KYC implementations exist: *TzLibre*, *nTezos* and *OpenTezos*. Little is publically known about these groups, but the effective bifurcation of the pre-functional network into *white* KYC and *black* non-KYC populations is a phenomenon likely to repeat as blockchain forensic tools become more widely adopted by law enforcement agencies [30]. At time of writing, Tezos has an operational *betanet* and TZLibre appears to have adjusted strategy, becoming a leading *delegated staker* - or *baker* in the Tezos parlance - within the DLS-Tezos network and campaigning for a reversal of the KYC implementation decision.

1.6 Forks in Favour of ASIC-Resistance

Since SHA-256 *Application Specific Integrated Circuits* (ASICs) were first developed in 2012 for the Bitcoin network, there has been a trend among upstart networks to choose alternative hashing algorithms so as to avoid the problems associated with being a minority network in relation to a particular type of computational resource. A series of existing and new algorithms such as *Scrypt*, *CryptoNight*, *Blake 2b*, *Ethash* and *Equihash* with greatly increased memory requirements relative to SHA-256 were implemented into networks such as Litecoin, *Monero*, *Siacoin*, Ethereum and *Zcash* respectively, under the supposition that *memory-hardness* would prevent the development of ASICs for these algorithms as the ability to parallelise processes would be greatly reduced *via* the system memory bottleneck. Such algorithms were commonly referred to as *ASIC-resistant*, however this does not appear to have remained the case as there now exist ASICs for all of the above hash functions.

The failure to prevent specialised hardware development was unavoidable in retrospect. As cryptocurrency network valuations increased the incentives for equipment manufacturers to allocate the substantial capital to develop specialised integrated circuits outweighed the downside risks. Other contributing factors were optimisations in mining hardware engineering, steps forward in semiconductor manufacture and margin compression in the more mature SHA-256 ASIC marketplace encouraging hardware manufacturers to diversify. As the mining hardware business is extremely competitive, development of ASICs for new algorithms was conducted with utmost secrecy so participants would not lose their early-mover advantage. Indeed it is commonly accepted (if not conclusively proven) that many mining manufacturers will mine in secret prior to announcing their equipment and offering units for sale. Light testing of electronic equipment prior to despatch is uncontroversial as part of a quality assurance process, however there have been widespread accusations that ASIC manufacturers - or partners for the purposes of plausible deniability - deploy ASICs to networks clandestinely and gradually with hashrate spread over several pools to avoid detection [31]. Further, there have been a number of instances whereby a new ASIC type would be announced (by a manufacturer such as Baikal, Innosilicon or Bitmain) and an impression of limited run scarcity would be implied, to maintain a value proposition for the profitability of the device. There would then follow what may be regarded as supply dumping where the manufacturer sells so many ASICs that the possibility of a purchaser achieving a return on investment would be *nil.* There is also a question mark over the network security of cryptocurrencies with clandestine ASICs online, as an equipment manufacturer ”testing” large batches of their equipment would have an asymmetric edge over existing participants employing *Central Processing Unit* (CPU), *Graphics Processing Unit* (GPU) or *Field-Programmable Gate Array* (FPGA) and may easily garner a majority of network hashrate making 51 % attacks trivial, with grave impact on network value proposition.

Some networks that have adopted the philosophy of ASIC-resistance - with the goal of maximising decentralisation at the mining level - reacted to the suspicion or discovery of ASICs on their network by proposing a fork (hard or soft depending on the circumstances) to change the hashing algorithm to an alternative candidate sufficiently distinct from the original so as to render the ASICs ineffective. As in all cases with forks to irreversibly change mining parameters on PoW networks, should sufficient computational resource remain on the original chain then it has a prospect of avoiding wipeout and surviving as a sovereign network. In this case where large quantities of ASICs were produced and then threatened with being rendered incompatible through hashing algorithm adjustment, these machines would most likely be obliged to remain on the original chain, or to switch to mining on a smaller network which did not undergo such a fork. It has been postulated that new CPU architectures such as *Vector Processors* may be present in current or forthcoming generations of ASICs which would allow for a greater ability to remain on their intended network after hard forks to change hashing algorithms. By analysing the limited efficiency gains in ASICs developed for memory-hard algorithms such as Ethash compared to those observed previously realised for SHA-256, an alternative technical configuration with greater computational flexibility than traditional ASICs is a plausible though unconfirmed hypothesis [32].

Providing a counterpoint to the above motivations, Daian asserts that ASICs are inevitable for algorithms which are employed on sufficiently valuable networks. Therefore they should be accepted as emergent phenomena arising from the success of networks adopting those particular hashing algorithms. As ASICs realise large efficiency gains over general-purpose hardware in terms of operational costs (energy efficiency as measured in hashes *per* Watt) and capital outlay (hashes *per* dollar cost of ASIC) therefore lending themselves to industrial mining facilities and the economies of scale they can access. Therefore the reaction of forking to change hashing algorithm only provides a temporary respite from the development of specialised hardware, and indeed regularly scheduled tweaks may become less effective as more versatile hardware is designed. Indeed such protocol changes may favour well-resourced hardware manufacturers as they will be more able to deploy capital and resources to produce new hardware. The decision making process involved in enacting such protocol changes may also be subject to corruption or sub-optimal outcome, as with Ethereum's chain split following the failure of The DAO as discussed in Section 1.3 [33].

Two recent networks which took different approaches to the manifestation of ASICs were Monero and Siacoin. Monero (XMR) is a privacy-focused cryptocurrency with a healthy community, active developer ecosystem and strong philosophy of maintaining decentralisation at the mining level through the promotion of ASIC-resistance in favour of GPU mining. As XMR nethash began to climb steeply in January and February 2018, ASIC mining was suspected to be taking place surreptitiously, followed by announcements by manufacturers Bitmain and Baikal that ASICs for XMR were available for imminent shipping [34]. In April 2018, Monero underwent its twice-annual scheduled hard fork which facilitates regular protocol upgrade and included an adjustment to the CryptoNight hashing algorithm to render the ASICs ineffective. Around the time of the hard fork, XMR experienced a sudden 80 % decline in nethash with stabilisation at around 40-50 % decline. Prior to the fork, over 90 % of hashrate was of unknown/anonymous origin, whereas post-fork the proportion of hashrate with unknown provenance had stabilised around 30-40 %. Therefore the level of transparency as to distribution and provenance of computational resource increased as much as coarse heuristics as pool activity allow inference. Some questions remain over the methods employed to achieve consensus on the algorithm change, with some appeals for patience or to maintain the *status quo*. There was also a rather surreal incidence of extreme price volatility of the mining equipment with fire sales as Monero's hard fork was implemented. Baikal was advertising a "buy one, get four free" offer on the ASICs which would have exacerbated dumping of commodity nethash on ASIC-friendly CryptoNight networks. A number of putative breakaway Monero factions announcing support for the original chain also announced themselves but do appear to have largely waned into irrelevance [35].

Siacoin (SC) is a network providing secure and censorship-resistant data storage via a decentralised P2P architecture. A hardware manufacturer named Obelisk with strong ties to the Siacoin founders had a Blake 2b ASIC under development and had taken a significant amount of pre-orders for the SC1. Bitmain appears to have intercepted information relating to this device and leveraged their economies of scale and expedience to front-run the Obelisk miners by delivering the Antminer A3 before them and furthermore offering aggressive discounts to Obelisk pre-order customers. This may have been through the utilisation of faster but sub-optimal integrated circuit development processes such as *place-and-route* rather than *fully-custom routing* as Obelisk employed. Unbeknownst to outsiders, Obelisk had engineering a second *fallback* algorithm into their equipment so that a soft fork adjustment to the Siacoin protocol would be sufficient to render the Bitmain ASICs ineffective. However this was not exercised and instead an uncontentious hard fork was conducted to recalibrate the difficult adjustment algorithm and block time in anticipation of large increase in network hashrate [36].

2 Research Aims and Methodology

2.1 What does Forkonomy Aim To Achieve?

As a putative analytical discipline in the early stages of development, forkonomy is as much a perspective as a coherent set of tools and methods at present. The notion of performing comparative analysis on ledger forks is not new, however this somewhat high-level combination of quantitative observation and qualitative inference is not commonly applied to characterise the emergent phenomena exhibited in cryptocurrencies. By taking a wider view than the present and recent past, forkonomy aims to provide insight into the possible fates of blockchain-oriented P2P monetary networks. A future aim is to build sufficiently sophisticated models such that even-handed forecasts of the probabilities of future scenarios may be elucidated from network observation and simulation. Many of the concepts employed are borrowed from the disciplines of astronomy, cosmology and physics, which the author previously researched.

2.2 Research Methods and Resources

This work has relied on numerous primary and secondary data sources as cited in the text. Blockchain analytics of BTC, BCH, ETH, ETC, XMR, MONA, ZCL and BTCP was achieved through the use of block explorers Blockchair.com, Blockchain.info, Etherscan.io, Etherhub.io, Bchain.info, Monerohash.com and Bitinfocharts.com with data exported in CSV or JSON formats. This was imported into the statistical computing suite RStudio (built upon R) for cleaning, treatment, analysis and visualisations. Network-wide observation and inference was conducted using publically available sources Coin.dance for node count and implementation versions for BTC and BCH, Crypto51.app for ZCL and BTCP network hashrates, Doublespend.cash for malleated transactions on BCH, Coinmetrics.io for high-level network heuristics and Onchainfx.com for networks' token price, supply issuance and monetary policy.

3 Case Study: Advent of the Fork-Merge

3.1 Introduction

In a 2017 presentation at *Breaking Bitcoin* conference, Eric Lombrozo postulated the theoretical possibility of a managed process of convergence of chains sharing the same provenance and similar codebase which may be thought of as a *chainmerger*. The idea was developed further by Eric Wall ostensibly as potential a mechanism for BTC and BCH to reunite post-chain split, but no prominent examples exist in the wild. This may be subject to *entropic bias*, that is to say divergent paths are those of least resistance in accordance with thermodynamics as discussed in Section 1.3 [37].

3.2 Fork-Merge through UTXO Cross-Chain Consolidation

Building on the chainmerger concept outlined above, the notion of a *fork-merge* was introduced earlier this year as the mechanism by which a ledger fork of BTC entitled Bitcoin Private (BTCP) could be artificially synthesised from an *Equihash* PoW network named Zclassic (ZCL), itself a codebase fork of Zcash (ZEC) which in turn was originally derived from the BTC codebase [38]. It is somewhat similar to the "Fork + Merge" operation in Git-based repository protocols. Since the BTC and ZCL networks possess different histories as evinced by their unique UTXO sets and the codebase had additionally diverged further, this was not a trivial process [39] and may be further hindered by entropic bias.

The UTXO model of ledger accounting introduced by Bitcoin is managed by tracking the outputs of transactions as either spent or unspent. Unspent transaction outputs contribute to coin-holders' balances whereas spent outputs do not. In order to maintain such a ledger, each transaction may be comprised of one or more inputs (UTXOs with non-zero balances) and two or more outputs. This is because UTXOs may not be partially spent, and thus any value remaining in an UTXO after transaction is completed must be returned as a new "*change*" UTXO in an analogous manner to spending a paper fiat currency banknote and being returned different notes and coins.

The quantitative parameters underlying this cross-chain UTXO consolidation warrant further examination. Both BTC and ZCL networks possess equivalent relationships controlling mining subsidy emission over time. ZCL has a target block time of 150 seconds, block reward of 12.5 ZCL, 840000 block reward halving period (not yet reached) and 21 million ZCL maximum supply. BTC has a 600 second target block time with initial reward of 50 BTC *per* block, though this has experienced two subsidy halvings to the present value of 12.5 BTC *per* block - with a current approximate BTC block height of 540000 and halving period of 210000 blocks. Figure 1 displays the characteristics of BTC mining subsidy and monetary issuance over time.

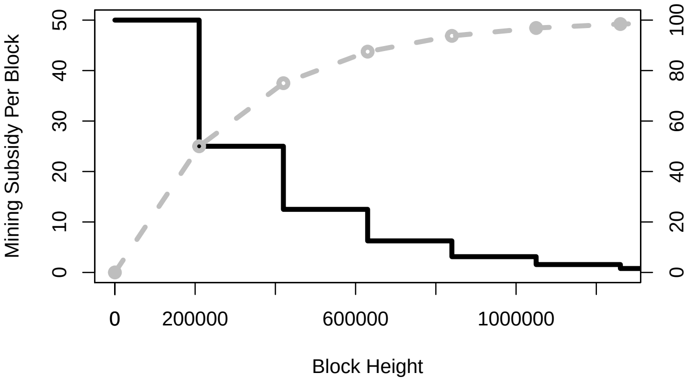

Fig. 1. The relationship between BTC block height, mining subsidy and supply issuance.

As the time-*per*-halving is broadly equal on both networks the number of halvings may be used as an approximate heuristic for the maturity of the network. ZCL having experienced no halving to date can be considered a *young network*, characterised by a high mining subsidy which incentivises miners to secure the chain at the expense of a high effective annual supply inflation rate of approximately 100 %, with approximately 4.5 of 21 million total ZCL coins issued. BTC is halfway between its second and third halvings and as such can be thought of as a *mature network*. The subsidy has already declined 75 % since network launch with approximately 17 of 21 million total BTC mined and an effective annual supply inflation of around 4 %. During periods of elevated demand for block space, a transaction fee market has emerged which at peak times has provided miners with *greater income than the block reward* [40]. This occurrence is crucial to the long-term viability of all blockchain-based monetary networks that employ PoW for security and have a fixed asymptotic supply curve, as the network must continue to incentivise miners to deliver hashpower [41]. Most

UTXO-based cryptocurrencies have also adopted BTC's monetary issuance policy to claim analogous value propositions centred around supply limitations.

By merging these UTXO sets, BTCP has synthetically created an Equihash blockchain network with approximately 500000 of 21 million coins yet to be issued, negligible annualised supply inflation and therefore a meagre mining subsidy of 1.5625 BTCP, corresponding to approximately 0.0035 BTC at time of writing. Unlike BTC however, BTCP has not been able to bootstrap a transaction fee market, and in order to properly incentivise miners to protect the network the transaction fees would have to be greater than the transaction value itself. Additional idiosyncratic risks to BTCP mining profitability arise from possible supply shocks from involuntary coin holders who would be more likely to commence liquidation in the event of sudden BTCP coin price rises, and the ongoing emergence of specialised Equihash ASIC mining hardware from multiple hardware suppliers deploying more plentiful *commodity* hashrate [42].

3.3 Forkonomics: The Impact of Fork-Merging on Monetary Networks

The fork-merge process has effectively created an *elderly* BTCP blockchain between third and fourth halvings (as seen in Figure 2), with little incentive for miners to protect and therefore minimal value proposition as a PoW monetary network. Much of the BTCP UTXOs involuntarily assigned to BTC UTXO owners have gone uncollected, undoubtedly due to the low value of the 1:1 airdrop for the BTC side or prevention of private key compromise risk. In many respects BTCP is now experiencing an *eternal post-fork hangover* caused by the lopsided incentive structures engineered into the fork-merge. The event asymmetrically benefited ZCL holders which had a much lower *per* coin price than BTC but also entitled holders to a 1:1 airdrop. This was particularly the case for those who held ZCL balances prior to the announcement of the fork-merge, as the market price of ZCL experienced an approximate *hundredfold* increase in USD terms within a 30 day period prior to the fork-merge [43].

Due to the disparity in mining subsidy value and network age (not "*effective maturity*" as discussed above) between ZCL and BTCP, ZCL appears to retain a reasonably cohesive constituency of stakeholders - miners, exchanges, users and so on - despite many developers abandoning the project at time of fork. In contrast, BTCP seems to have lost most of its pre-fork proponents and has failed to acquire listing on major exchanges to access liquidity in order to improve its value proposition as a speculative asset. BTCP vs ZCL may be considered an extreme case of *fork-induced emission curve fatigue.* That is to say that the fork-merge process has resulted in a cryptocurrency network simultaneously vulnerable to majority attacks and unable to bootstrap itself into a secure and reliable state as the block subsidy available in an elderly network does not sufficiently incentivise computational resource in the absence of an on-chain transaction fee market. The lack of evidence of such attacks on BTCP may be due to the lack of on-chain transaction volume and associated fiat equivalent value making even a low-cost attack a waste of resource. Furthermore trading platforms do appear to anticipate the likelihood of such an attack as typically 25-50 confirmations are required to consider a BTCP deposit confirmed and spendable at an exchange.

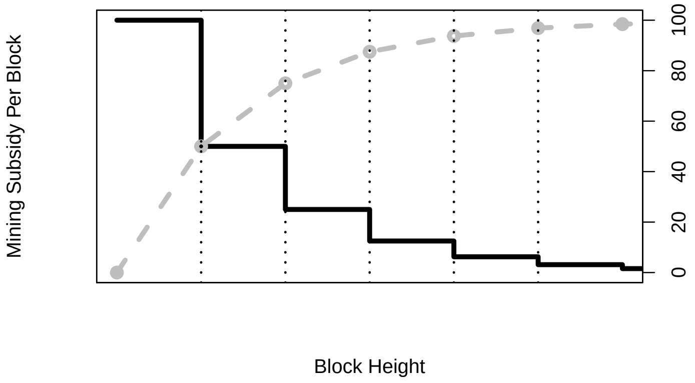

Fig. 2. Generalised emission curve and supply schema for cryptocurrency networks deriving their accounting and monetary chacateristics from Bitcoin. Each "step down" represents a halving of block subsidy, halving in effective supply inflation rate and an advancement in the lifecycle phase of a blockchain network.

In 2018 there has been an emerging trend of ledger forks of BTC possessing greatly inflated market capitalisations in comparison to codebase forks with *virgin* genesis blocks and ledgers. This is at least in part due to the effective sequestration of large proportions of the supply, essentially *attention-locked* since BTC UTXO owners have neither financial nor ideological motivation to participate at the potential expense and inconvenience of accessing private keys. Observable on-chain transaction volume (not including shielded transactions which typically constitute a tiny minority of usage) is minimal on both BTCP and ZCL networks with significantly under one million USD average daily volume, whilst BTC moves approximately several billion USD equivalent *per* day. In terms of hashrate ZCL has approximately 25 times more network hashrate than BTCP with a nominal market capitalisation of 3 times less [44]. The consequence of this is that the BTCP chain is rendered extremely vulnerable to 51 % attacks with a trivial vector employing rented hashrate - using figures at time of writing the 1 hour cost of a majority attack was approximately 200 USD. For a network with a nominal value (using market capitalisation as a coarse heuristic) of approximately one hundred million USD, the prospect for *transaction disruption* seems sufficiently high to preclude any realistic proposition of BTCP as a monetary network. If majority takeovers become trivial in a cryptocurrency network, exchanges will be reticent to list it as they would be the primary victims of double-spending attacks when not requiring sufficient confirmations for transaction finality to be beyond doubt [45].

4 Discussion: Implications for Ageing Blockchains and Prominent Minority Forks

The emission curve fatigue that BTCP is experiencing, combined with lack of transaction fee market results in an insecure network with absent value proposition. Indeed this is one of the possible futures for any elderly PoW blockchain. By analogy with stellar lifecycles, the moniker *white dwarf chain* may be applied to BTCP. In common with the celestial remnant, high maturity and low economic gravity prevent the network from attracting substantive accretion, eventually no longer possessing the critical mass to function. There is a prospect that BTCP will attempt a transition to PoS or dPoW in order to seek refuge from thermodynamic attacks. Recently the prospect of *confiscation* of "inactive" UTXOs in order to liberate coin supply from attention-locked holders of BTCP in order to provide further miner subsidy in order to attract greater hashrate has emerged [39]. The disingenuous trope of "Satoshi's Vision" was invoked by BTCP proponents in the pre-fork marketing, though it is difficult to see how Satoshi Nakamoto's *cypherpunk* principles were respected and honoured through the mechanism of confiscating UTXOs under his control.

An alternative outcome termed a *chain death spiral* is also a possibility for BTCP. Should Equihash resource be sufficiently incentivised to be directed elsewhere, the network may stop issuing blocks altogether. This was a particular concern for BTC at the time of the BCH chain split, though ironically it was BCH that produced severely tardy blocks with block intervals reaching many hours for some time. This was due to the BCH network inheriting the BTC network's difficulty whilst only possessing a fraction of the former BTC hashrate. A customised difficulty adjustment algorithm was invoked to rapidly adjust the BCH network difficulty downwards to reflect the much lower nethash of the minority SHA-256 BCH network fragment. The lack of such a difficulty adjustment mechanism in BTC beyond the original specification's 2016 block window came to be perceived as a potential attack vector from a hostile ledger fork [46].

The significance of implications arising from the BTCP case study are due to the lack of *organically elderly blockchain networks* in existence today. Emergent behaviours that are observed in these distributed environments may vary from hypothetical studies utilising cryptoeconomic, distributed systems or game theoretical perspectives. Due in part to the BCH difficulty adjustment process - and successor algorithms performing analogous functions - BTC and BCH have already diverged by approximately seven thousand blocks chain length (Figure 3) which corresponds to around 50 days greater effective age of BCH in the year since chain split. The consequence is that, *ceteris paribus*, the BCH blockchain will reach its next block subsidy halving sooner than BTC. Coupled with the fact that BCH shares the SHA-256 mining algorithm with BTC but now has approximately ten times less hashrate (Figure 4), there is declining economic incentive for miners to secure the minority BCH network [47]. With no fix currently implemented for transaction malleability due to BCH's rejection of SegWit and no alternative ready to deploy, 51 % attacks have become trivial to conduct by several BTC mining pools and double spent transactions are growing in frequency, calling any notion of monetary soundness or payment utility proposition into serious question [48].

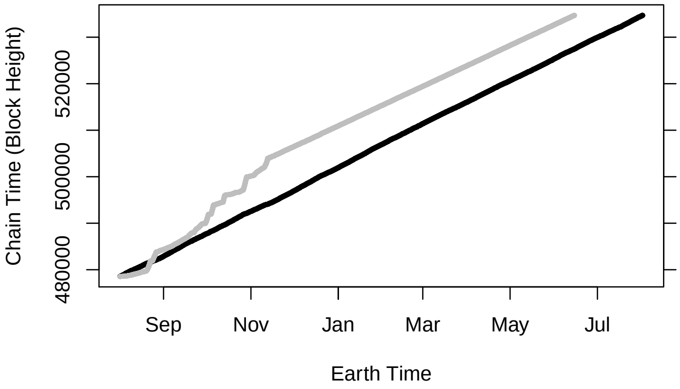

Fig. 3. Chain dynamics of BTC (*black*) and BCH (*grey*) networks August 2017-18, as visualised through the benchmarking of "chain time" versus Earth time. Data from Blockchair.com.

Through the observation of networks which in the past competed for ASIC hashrate such as Litecoin and Dogecoin, it has been observed that once the security of a PoW network sharing a mining algorithm with a dominant competitor is believed to be compromised, two main categories of remedial action may be utilised. To preserve decentralisation and network sovereignty, the adoption of an alternative and unique PoW algorithm is an option but would be unpalatable for an ASIC-oriented network such as BCH. An alternative is to implement *merge-mining* whereby PoW on the dominant network for a particular algorithm counts towards PoW on the merge-mined network [49], or periodic *checkpoint notarisation* - also known as *delayed PoW* - of latest block hash into the most secure blockchain as utilised by minority Equihash network *Komodo* [50]. Confiscation of "inactive" UTXOs or account balances has also been proposed by minority forks such as United Bitcoin and Bitcoin Private as discussed above.

The canonical Ethereum network ETC may have a different future to the typical minority branch, as development paths between forks have diverged and ETH intends to attempt transition to PoS with the *Casper* family of consensus

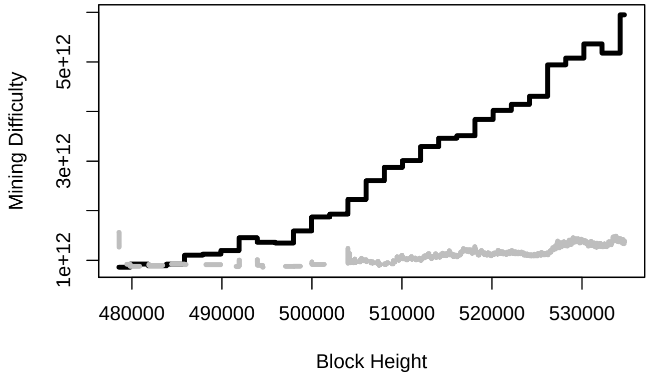

Fig. 4. Difficulty (as proxy heuristic for hashrate) comparision of BTC (*black*) and BCH (*grey*) networks August 2017-18. Data from Blockchair.com.

protocols [51], accompanied by a significant reduction in block issuance subsidy to 0.6 ETH *per* block [52]. Should this occur as multiple competing Ethash ASICs and high performance FPGA *bitstreams* are distributed more widely, ETC may retain a strong value proposition as the canonical, decentralised and immutable Ethereum network with a sound monetary policy and thermodynamically assured network security. As Figure 5 shows, ETH has an annual equivalent supply inflation of approximately 7.5 % and no maximum limit on token supply, whereas ETC's inflation is around 5.75% and projected to decrease much more rapidly due to a fixed supply limit. ETC has also removed the so-called *difficulty bomb* which is intended to disincentivise mining by making it increasingly unprofitable.

5 Future Perspectives on Forks

As with any novel field of study many open questions remain as to how new technologies, emergent phenomena and threats caused by internal factions within open source protocol networks or external entities such as rival blockchains, lawmakers and silicon foundries may influence the forking tendencies of cryptocurrency networks. Sztorc's notion of *fork futures* has merit insofar as competing visions may be assessed and priced in real time by the marketplace prior to implementation. This facilitates the assessment of support for the various options proposed by competing factions, potentially preventing quite a substantial proportion of chain splits by using the market to assess the value of competing ideas. [54].

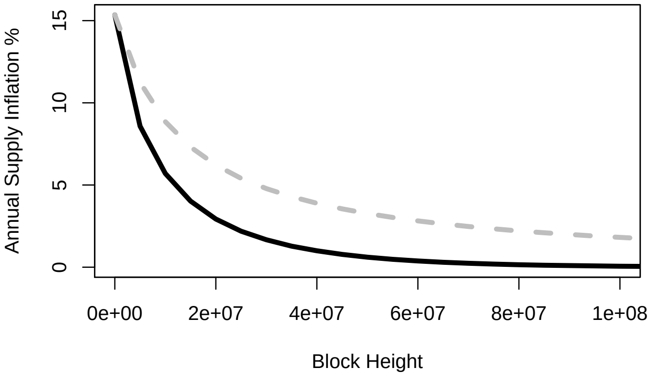

Fig. 5. All-time supply inflation comparision of ETC (*black*) and ETH (*grey*). Data from ECIP1017 [53].

Velvet forks as proposed by Kiayias *et al.* could help mitigate potential network consensus failures by increasing inclusiveness and compatibility of protocol upgrades, by being minimally invasive with respect to network participants not running the velvet fork upgrade [28]. An example of successful implementation of a velvet fork has been found in decentralised mining pool *P2Pool*'s *sharechain*, which keeps track of *mining shares* which correspond to block hashes close to but not below the network difficulty limit. In order to reduce reward variance for individual participants in the mining pool, shares are kept track of by the sharechain [55].

The ongoing litigation against the cryptocurrency exchange Bitgrail involves an attempt to legally enforce a rollback of the *Nano* (formerly *Raiblocks*) *block-lattice* network to reclaim tokens which were lost due to software vulnerabilities. It is hard to envisage an outcome whereby a legal pronouncement is made which carries sufficiently global or borderless jurisdiction to coerce large constituencies of a network to behave *contra* to their incentives. Most likely this would trigger a *factional network disintegration* event [56].

Hypothesising more broadly, as the canon of forkonomy expands to include new and emergent phenomena there may develop further aesthetic disciplines with which to codify, classify and characterise trust-minimised network partitions in all their forms. As with celestial outcomes, the interplay of enthalpy and entropy could provide a generalised basis for modelling the fate of cryptocurrency networks and further work is underway in this area. Moving from the ontological and observational basis presented here as forkonomy (by analogy with astronomy) and *forkonomics* (by analogy with economics), epistemological treatises may be considered *forkology* [57] and philosophical approaches *forkosophy.*

References

1. List of Linux Distributions, Wikipedia. http://en.wikipedia.org/wiki/List_of_Linux_distributions. Last accessed 10 August 2018.
2. Nakamoto, S. (2008) Bitcoin: A Peer-to-Peer Electronic Cash System. http://bitcoin.org/bitcoin.pdf.
3. Delgado-Segura, S., Prez-Sol, C., Navarro-Arribas, G., and Herrera-Joancomart, J. (2017) Analysis of the Bitcoin UTXO set. IACR Cryptology ePrint Archive, 1095. https://eprint.iacr.org/2017/1095.pdf
4. BCH Node Status, Coin Dance. https://cash.coin.dance/nodes#nodeVersions. Last accessed 10 August 2018.
5. Metcalfe, B. (2013) Metcalfe's Law after 40 Years of Ethernet. In: Computer, vol. 46, no. 12, 26-31. https://doi.org/:10.1109/MC.2013.374.
6. Bonneau, J. (2018) Hostile Blockchain Takeovers. In Bitcoin '18: Proceedings of the 5th Workshop on Bitcoin and Blockchain Research.
7. Mark, D., Zamfir. V., and Sirer, E. G. (2016) A Call for a Temporary Moratorium on The DAO, http://hackingdistributed.com/2016/05/27/dao-call-for-moratorium.
8. BitMEX Research (2018) Revisiting The DAO, https://blog.bitmex.com/revisiting-the-dao.
9. Buterin, V. (2017) Notes on Blockchain Governance, https://vitalik.ca/general/2017/12/17/voting.html.
10. van Wirdum, A. (2016) Ethereum Classic Community Navigates a Distinct Path to the Future, https://bitcoinmagazine.com/articles/ethereum-classic-community-navigates-a-distinct-path-to-the-future-1471620464.
11. Conner, E. (2018) EIP 1234: A Case For Ethereum Block Reward Reduction in Constantinople, https://medium.com/@eric.conner/a-case-for-ethereum-block-reward-reduction-in-constantinople-eip-1234-25732431fc77.
12. Ethereum Uncle Rates, Etherscan. https://etherscan.io/chart/uncles. Last accessed 10 August 2018.
13. Eyal, I., and Sirer, E. G. (2014) Majority is Not Enough: Bitcoin Mining is Vulnerable. In International Conference on Financial Cryptography and Data Security, 436-454. Springer, Berlin, Heidelberg.
14. Grunspan, C., and Perez-Marco, R. (2018) On Profitability of Selfish Mining. IACR Cryptology ePrint Archive, 1805.
15. Bitcoin Empty Blocks Blockchair. https://blockchair.com/bitcoin/blocks?s=size(asc). Last accessed 10 August 2018.
16. Monacoin Block Explorer. https://bchain.info/MONA/. Last accessed 10 August 2018.
17. Nayak K., Kumar, S., Miller, A., and Shi, E. (2016) Stubborn Mining: Generalizing Selfish Mining and Combining with an Eclipse Attack. IEEE European Symposium on Security and Privacym 305-320. EuroS&P, Saarbrucken. https://doi.org/:10.1109/EuroSP.2016.32.
18. Zhang, R., and Preneel, B. (2017) Publish or Perish: A Backward-compatible Defense against Selfish Mining in Bitcoin. In Cryptographers' Track at the RSA Conference, 277-292. Springer, Cham. https://doi.org/:10.1007/978-3-319-52153-4_16.
19. Miller, A. (2013) Feather-forks: enforcing a blacklist with sub-50 % hash power. https://bitcointalk.org/index.php?topic=312668.0. Last accessed 10 August 2018.
20. Hacken (2017) The Rush for Hashpower. https://hacken.io/wp-content/uploads/The-Rush-for-Hashpower.pdf. Last accessed 10 August 2018.

21. Sclavounis, O. (2017) Understanding Public Blockchain Governance, https://www.oii.ox.ac.uk/blog/understanding-public-blockchain-governance.
22. BitMEX Research (2018) Covert versus overt AsicBoost, https://blog.bitmex.com/graphical-illustration-of-a-bitcoin-block.
23. Hilliard, J. (2017) BIP91, https://github.com/bitcoin/bips/blob/master/bip-0091.mediawiki. Last accessed 10 August 2018.
24. Ver, R. (2018) Why I Think Bitcoin Cash is Bitcoin, https://www.yours.org/content/why-i-think-bitcoin-cash-is-bitcoin-6cb2dda7ca08
25. O'Hagan, A. (2016), The Satoshi Affair, https://www.lrb.co.uk/v38/n13/andrew-ohagan/the-satoshi-affair
26. Castor, A. (2018) No Incentive, https://www.coindesk.com/no-incentive-algorand-blockchain-sparks-debate-cryptography-event
27. Poelstra, A. (2016) A Treatise on Altcoins, https://download.wpsoftware.net/bitcoin/alts.pdf
28. Kiayias, A., Miller, A., and Zindros, D. (2018) Non-interactive Proofs of Proof-of-work. IACR Cryptology ePrint Archive, https://eprint.iacr.org/2017/963.pdf
29. Decred Politeia Github, https://github.com/decred/politeia. Last accessed 10 August 2018.
30. Crystal Blockchain Analytics, https://crystalblockchain.com. Last accessed 10 August 2018.
31. Sayres, N. (2018) Bitmain Faces New Accusations of Secret Mining, https://hacked.com/bitmain-faces-new-accusations-of-secret-mining/. Last accessed 10 August 2018.
32. BitMEX Research (2018) New Ethereum Miner Could be a Game Changer, https://blog.bitmex.com/nextstageinmining. Last accessed 10 August 2018.
33. Daian, P. (2018) Anti-ASIC Forks Considered Harmful, https://pdaian.com/blog/anti-asic-forks-considered-harmful. Last accessed 10 August 2018.
34. Wilmoth, J. (2018) Manufacturer Holds CryptoNight ASIC Firesale after Monero Hard Forks, https://www.ccn.com/manufacturer-holds-cryptonight-asic-firesale-after-monero-hard-forks. Last accessed 10 August 2018.
35. Source: Bitinfocharts, https://bitinfocharts.com/comparison/monero-hashrate.html. Last accessed 10 August 2018.
36. Vorick, D. (2018) The State of Cryptocurrency Mining, https://blog.sia.tech/the-state-of-cryptocurrency-mining-538004a37f9b. Last accessed 10 August 2018.
37. Lombrozo, E. (2017) Speech at Breaking Bitcoin Conference, https://www.youtube.com/watch?v=0WCaoGiAOHE
38. Biryukov, A., and Khovratovich, D. (2016) Equihash: Asymmetric Proof-of-Work Based on the Generalized Birthday Problem. https://eprint.iacr.org/2015/946
39. Bitcoin Private Whitepaper, https://btcprivate.org/whitepaper.pdf. Last accessed 10 August 2018.
40. Bitcoin Block 500439, Source: Blockchair Explorer, https://blockchair.com/bitcoin/block/500439. Last accessed 10 August 2018.
41. Ammous, S. (2018) The Bitcoin Standard: The Decentralized Alternative to Central Banking. John Wiley & Sons.
42. Let's talk about ASIC mining, Zcash Forum, https://forum.z.cash/t/let-s-talk-about-asic-mining/27353. Last accessed 10 August 2018.
43. Source: OnChainFX, https://onchainfx.com/asset/zclassic. Last accessed 10 August 2018.
44. Source: MiningSpeed Pool Monitor, https://pool.miningspeed.com. Last accessed 10 August 2018.

45. Source: PoW 51% Attack Cost, https://www.crypto51.app. Last accessed 10 August 2018.
46. Wong, J. I., Bitcoin cash could lead to bitcoin ”death spiral”, Quartz, https://qz.com/1127817/bitcoin-cash-bch-price-could-lead-to-bitcoin-death-spiral
47. Source: Blockchair, http://www.blockchair.com. Last accessed 10 August 2018.
48. Source: BCH Doublespend Monitor, http://doublespend.cash. Last accessed 10 August 2018.
49. Judmayer, A., Zamyatin, A., Stifter, N., Voyiatzis, A. G., and Weippl, E. (2017) Merged Mining: Curse or Cure? In Data Privacy Management, Cryptocurrencies and Blockchain Technology, 316-333. Springer, Cham.
50. Delayed Proof of Work Whitepaper, https://github.com/SuperNETorg/komodo/wiki/Delayed-Proof-of-Work-(dPoW)-Whitepaper
51. Casper Ethereum Github, https://github.com/ethereum/casper. Last accessed 10 August 2018.
52. Griffith, V., and Buterin, V. (2017) Casper the Friendly Finality Gadget, https://arxiv.org/pdf/1710.09437.pdf
53. Mazur, M. (2016) ECP1017, https://github.com/ethereumproject/ECIPs/pull/20/files. Last accessed 10 August 2018.
54. Sztorc, P. (2017) Fork Futures (via the Exchanges), http://www.truthcoin.info/blog/fork-futures/
55. Zamyatin, A., Stifter, N., Judmayer, A., Schindler, P., Weippl E., and Knottenbelt W. (Short Paper) A Wild Velvet Fork Appears! Inclusive Blockchain Protocol Changes in Practice. In Bitcoin ’18: Proceedings of the 5th Workshop on Bitcoin and Blockchain Research.
56. Gordon, S. (2018) Nano Team Target of Cryptocurrency Class Action Lawsuit, https://bitcoinmagazine.com/articles/nano-team-target-cryptocurrency-class-action-lawsuit/
57. Antonopoulos, A. M. (2017) Forkology: A Study of Forks, https://www.youtube.com/watch?v=rpeceXY1QBM.