Blockchain Mining: Optimal Resource Allocation

Yunming Xiao
yunming.xiao@u.northwestern.edu
Northwestern University

Sarit Markovich
s-markovich@northwestern.edu
Kellogg School of Management
Northwestern University

Aleksandar Kuzmanovic
akuzma@northwestern.edu
Northwestern University

ABSTRACT

Having enabled numerous applications, blockchains have attracted not only much attention, in the past decade, but also huge amount of resources: talent, capital, energy, etc. Focusing on the mining side of the market, in this paper, we aim at understanding how to efficiently use the resources mining and staking pools attract. We start with developing predictions about factors that increase the efficient allocation of pools' resources. We then test our predictions based on a general model for optimal resource allocation that we develop, as well as data we collected on pools' actual resource allocations. We find that pools can increase resource efficiency by mining for more blockchains as well as by increasing the frequency of resource re-allocation. Further, we enroll to mining pools as a miner to understand and comment on how pools can encourage their miners to increase the efficiency of their allocation. While our empirical investigation mostly focuses on the BTC family, we show that our theory and results are general and applicable to the Ethereum family as well as other proof-of-work (PoW) and proof-of-stake (PoS) chains.

ACM Reference Format:
Yunming Xiao, Sarit Markovich, and Aleksandar Kuzmanovic. 2022. Blockchain Mining: Optimal Resource Allocation. In *4th ACM Conference on Advances in Financial Technologies (AFT '22), September 19–21, 2022, Cambridge, MA, USA.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3558535.3559792

1 INTRODUCTION

Blockchains, in general, and specifically blockchain-based cryptocurrencies (cryptos) are disrupting financial markets – opening up the opportunity for various applications, *e.g.*, micropayments and smart contracts, which simplify, automate, and facilitate inclusion. Participation in blockchain-based crypto markets has grown tremendously, recently, both in terms of usage as well as in terms of the mining community that supports the blockchains' activity. From usage perspective, crypto daily trading volume nowadays averages around $120 billion after reaching a peak of $500 billion in May 2021. During this same period of time, miners' revenues grew from less than $0.5 million to more than $44.5 million – peaking at $70.7 million on May 1, 2021 and then again at $68.4 million on October, 2021.

The monetary rewards associated with adding blocks to blockchains like Bitcoin and Etheruem (a process known as "mining") attracts an enormous number of miners; resulting in individual mining becoming economically unviable. This, in turn, gave rise to mining and staking pools (pools). Mining pools distribute small parts of solving computational puzzles as part of the proof-of-work (PoW) process. Staking pools (or validators) attract crypto holders to lock their tokens with the pool as part of the proof-of-stake (PoS) process. Both types of pools then share the mining rewards with their users.

Under both, PoW and PoS, pools utilize a limited resource — hashrate in the case of PoW and the blockchain's native token in the case of PoS. Managing these resources efficiently is thus crucial for profitability. Moreover, PoW blockchains, such as Bitcoin, utilize enormous amounts of hashpower as miners attempt to solve computational puzzles associated with the PoW consensus mechanism. Many view this as a "waste of energy" that could be otherwise better allocated. Mining pools, thus, should guarantee the hashpower under their hands is used in the most efficient way.

In this paper, we aim at understanding how to increase the efficient use of scarce resources under different blockchain consensus mechanisms. Taking the Bitcoin family as an example, the hardware and software used for mining Bitcoin (BTC), can be also used to mine other cryptos (*e.g.*, Bitcoin Cash (BCH) and Bitcoin SV (BSV)) without much additional cost. This raises the question: how should mining pools allocate their hashrate across these different coins. While on the one hand some cryptos, *i.e.*, BTC, are approximately an order of magnitude more valuable than others (*e.g.*, BCH and BSV), on the other hand, many more miners compete over adding a block to the Bitcoin blockchain and thus the probability of winning a reward is much lower. Furthermore, the price of cryptos and the allocated hashrate fluctuate, often widely, over time. As we show below, dynamically re-distributing the hashrate increases the efficient use of the limited hashrate pools have.

To study how to increase resource allocation efficiency, we start with developing some predictions on factors affecting efficient resource allocation. In order to test our predictions, we develop a general model for the optimal resource allocation. We use the model to comprehensively evaluate our predictions on a large-scale data set that we collected over a period of two years on mining of the Bitcoin family. We analyze a number of system parameters in different scenarios. While our simulations focus on the Bitcoin family, we demonstrate that our analysis applies to the Ethereum and PoS blockchain protocols. Our key contributions and findings are as follows.

First, we show that pools' current hashrate allocation is largely sub-optimal; yet, varies across the different pools. As a result, a

This work was supported by NSF grant no. 1810582.

AFT '22, September 19–21, 2022, Cambridge, MA, USA

ACM ISBN 978-1-4503-9861-9/22/09…$15.00
https://doi.org/10.1145/3558535.3559792

lot of "money is left on the table." Specifically, most of the mining hashrate is dis-proportionally allocated to BTC, the most-highly priced crypto. This is likely the case since pools do not, in fact, control their hashrate. Rather, it is the individual miners that utilize the pool's service that decide how to allocate their hashrate. Interestingly, we find that a pool needs to only control a relatively small fraction of its hashrate to significantly increase its revenue. Surveying the leading pools, we learn that some pools, indeed, control such hash power.

Second, we use our data on PoW to validate our predictions. In particular, we find that (1) pools can increase resource efficiency by mining for a larger number of blockchains; (2) the more frequent the pool re-allocates its resources, the more efficient it can allocate and take advantage of its resources; and (3) the value from efficiently reallocating resources increases with the volatility of the blockchain's native token. To test this last prediction we examine times with disruptive events, which created price volatility. In particular, we study: (a) BSV price surge in January 2020; (b) BCH and BSV halving in April 2021; and (c) BTC halving in May 2021. We demonstrate that during such events, when price fluctuations is most pronounced, pools' reallocation strategies may be particularly valuable.

Third, we examine how pools can encourage their miners to increase the efficiency of their allocation. To this end, we investigate pools that offer smart-mining, where individual miners give the pool the right to allocate their hashrate in a way that would maximize the individual miner's compensation. We enroll as a miner in smart mining and examine how pools allocate our hashrate. In a robustness check, we show that pools do not seem to try and take advantage of miners' enrolled in smart mining for the pool's own benefit. We further test the robustness of our results by considering the fees pools charge from miners, as well as studying whether our results are driven by block transaction fees.

Fourth, we identify switching costs that miners may incur when switching from mining one coin to another. Specifically, there is an opportunity cost associated with the time it takes to switch, during which miners cannot mine and enjoy revenues. We assert that such opportunity costs of switching are relevant both for PoW and PoS blockchains. We show that as long as switching can be done within fairly moderate time scales, i.e., tens of seconds, dynamically switching across coins can still be profitable.

Finally, we show that our model and findings hold beyond the Bitcoin family and empirically demonstrate that our predictions hold for the Ethereum family as well. A good example of the importance of efficiently reallocating resources is the two recent failures of the Solana blockchain. Specifically, Solana (a PoS chain) had a 17 hours outage on September 14, 2021 and then again a 7 hours outage on April 30, 2022. During the outages, validators were unable to mine any SOL tokens. Moreover, as CoinDesk reported on September 14, 2021: "It is unclear what implications the Tuesday outage may have for Solana."[1] Such uncertainty typically results in high price volatility. Indeed, as reported in CoinDesk reported on May 1, 2022: "The outage contributed to a bloody, albeit brief, drawdown in SOL markets. Solana's native token crashed to a 24-hour low of $83.13 about three hours into the outage before climbing back toward $89..."[2] The effect on Solana's token price post the September 2021 failure was even larger. That is, network outages are associated with price volatility which, as we discuss below, represent times during which efficient reallocation of resource could be of great value to pools. Furthermore, as Solana was working on fixes to its issues, Bitcoin.com reported that: "Some have said that the fix could take anywhere between 24 hours to 48 hours to fix the issues."[3] This represents another time period where, unlike validators that only mine for Solana, validators that mine for other chains, in addition to Solana, could reallocate their capital away of Solana.

The rest of the paper is structured as follows. In Section 2, we provide some background on PoW and PoS pools as well as review the literature. Our predictions are summarized in Section 3. Then, in Section 4 we introduce our model, and in Section 5, the methodology for collecting data. We present the results in Section 6. We then analyze how pools can encourage efficient mining in Section 7, and discuss related issues in Section 8. We conclude in Section 9.

2 BACKGROUND AND RELATED WORK

In this section we provide some background on the PoW and PoS mechanisms, mining and staking pools, as well as review the literature.

Proof of Work and Mining Pools. Blockchain-based cryptocurrencies, such as Bitcoin, order and confirm transactions based on a consensus mechanism. The PoW consensus typically involves solving some computationally complex problems, which necessarily require a lot of computational resources. Solving the PoW puzzles is referred to as *mining*.

The outcome of mining is two-fold. First, the result of the puzzle will be included in the next block, such that the puzzle of the next block depends on it. In this way, the blocks are chained, and the entire process keeps the transactions trustworthy and secure. Second, the *miner*, who adds the block to the blockchain, will earn a reward in the native token along with all the *transaction fees* in that block.

The high returns from Bitcoin block rewards, and later on from other blockchains' block rewards, attracted many miners to join the competition to solve the PoW puzzles. The intense competition across miners enticed innovation in hardware and specifically the development of Application Specific Integrated Circuits (ASICs). Given their high computational capacity and their power consumption optimization, ASICs quickly replaced CPU-based mining.

Still, nowadays it is virtually impossible for a miner to win a block reward with a single machine. The competition led to the formation of *mining pools*, which consist of many miners. Mining pools distribute small parts of the PoW to individual miners, often distributed across the world. This increases the pools' chance to win the reward. Individual miners who joined the pool are compensated proportionally to the hashrate they contribute, as we further discuss below. The mining pools may or may not host mining machines themselves. Usually mining pools keep a small share of the block rewards, as a fee for their services.

[1] See www.coindesk.com/markets/2021/09/14/solana-validators-ready-potential-restart-amid-blockchain-outage/; accessed on May 4, 2022.

[2] See www.coindesk.com/tech/2022/05/01/solana-goes-dark-for-7-hours-as-bots-swarm-candy-machine-nft-minting-tool/; accessed on May 4, 2022.

[3] See news.bitcoin.com/solana-block-production-stalls-for-hours-sol-holders-unable-to-transact-validators-deploy-a-fix/; accessed on May 4, 2022.

Proof of Stake and Staking Pools. Concerns about the waste of energy associated with the PoW consensus have given rise to the PoS consensus. The first functioning implementation of the proof of stake consensus was in 2012 [21]. Since then, PoS has enabled many real-world applications [13, 19, 30]. Under the PoS consensus, those who wish to propose new blocks to be added to the blockchain – a.k.a validators – must *stake* a certain amount of the blockchain's native token for a chance of being randomly selected for the task. The minimum staking amounts differ, depending on the blockchain in question. For each new block, a validator will be chosen – with probability increasing with their staking amount – to propose the new block and earn the rewards. Thus, the more tokens a validator stakes, the more blocks the validator will create, and the higher the validator's returns. Consequently, validators are encouraged and motivated to attract other token holders to delegate their tokens to the validator who would in return share the mining rewards with them. A penalty will be imposed on the validator who misconducts. In this way, PoS reduces the waste of computational power in the PoW's hashrate competition.

While validators do not compete over computational resources, they compete over token staking to increase their probability to propose the next block. Furthermore, many validators (e.g., Lido.io) mine for multiple PoS chains and thus must consider how to allocate the capital associated with staked tokens across the different chains. As such, our predictions and modeling below are applicable to both PoW and PoS pools.

Prior Work. The evolution of Bitcoin mining and in particular the switch from solo to pooled mining was studied in [29]. In [12], the authors collected over 800,000 Bitcoin nodes and presented the scale and geographic distribution of the nodes. Authors of [28] conducted the latest long-term measurement of Bitcoin mining pools– using blocks' information to estimate the hashrate. In comparison, we perform a finer-grained study on multiple cryptos. Mining pools were further analyzed in the context of game theory [6, 9, 22, 24, 26, 27, 31] and in similar related contexts [10]. In particular, it has been proven that there exists a singular equilibrium for the resource allocation between any two cryptos, driven by the rewards [9]. Furthermore, it has been shown that slow and cautious adjustment of resource allocation will lead to an equilibrium, whereas there could be oscillations otherwise. It is noteworthy that it is proved that in an equilibrium, the miners may not want to devote all the power given the difference of mining cost [27], and the miners or mining pools have no financial incentive to occupy over 50% of the hashrate of a crypto unless one of them has outstandingly low mining cost [31]. Pools' hashrate allocation was examined in [8, 11, 16], taking the perspective of finance and risks; rather than optimal allocation for profitability, as we do. To the best of our knowledge, our work is the first to analyze hashrate reallocation across multiple tokens with the same consensus basis at a fine-grained time scale.

There is a large literature on the PoS mechanism [7, 17], protocols [20, 23, 25], and the existing PoS market [18]. On the staking pools, Fanti et al. [14] demonstrated that existing mechanisms lead to poor equitability and hence result in wealth centralization. Moreover, a recent study suggests that a strategic miner with around one-third of the total stake can outperform an honest miner in PoS, while such thing cannot happen in PoW [15]. The same study also finds that the mining strategies in PoS are much richer than PoW. Nevertheless, the scope of their discussions is limited to within one cryptocurrency. To the best of our knowledge, no study has investigated the resource allocation across multiple PoS tokens, which we shed light on in this paper.

Ethics. Given that we rank the pools in terms of performance, we anonymized the actual names of the explored pools so as to guarantee that our findings do not affect the already competitive business environment. We do not report the specific size and hashrate of each pool as it may enable their indirect identification.

3 PREDICTIONS DEVELOPMENT

Mining and staking pools' profitability depends on the combined resources (*i.e.*, hashrate and native tokens) the pool attracts as well as on how the pool allocates these resources across the different blockchains it mines. For example, while the block reward of BTC is higher than the reward from mining BCH, the probability of winning a block (given a certain amount of hashrate) is lower for BTC as compared to BCH and BSV. This is because there are more miners competing over mining BTC than over mining BCH or BSV. Or in the case of staking, the block reward of Binance is much higher than that of Polygon, however, attracting enough BNB (the Binance native token) to become a validator requires much more resources relative to attracting enough of Polygon's native token. While easily transferring resources across different blockchains is not always possible and at times may be very expensive, there are "families" of coins for which such a switch is easy and almost costless. In the case of PoW, the BTC family is the largest of such families.

This suggests that important decisions mining and staking pools must make are how many blockchains to mine for, how to allocate its resources across these blockchains, and how frequently to reallocate these resources. Below we develop some predictions about the efficient use of pools' resources.

As mentioned above, blockchains that use the same type of consensus mechanisms may still differ in the reward value from mining as well as the probability of being the next to add a block or in the case of PoS, the probability of being able to be a validator and have the right to participate as a miner. For example, as of April 2022, Binance has only 20 validators yet announced that it plans to increase this number to 50. Other chains like Polygon and Solana have more than 100 and more than 1000 validators, respectively. Moreover, the price of the native token of some blockchains may be more volatile than others. This suggests that mining for multiple blockchains not only allows pools to diversify themselves but also to reallocate their resources as market conditions change, for example an increase in hashrate or the number of validators. Our first prediction thus is as follows:

Prediction 1: Pools can increase resource efficiency by actively mining and reallocating resources across a larger number of blockchains

Reallocating resources also depends on the frequency the pool reallocates its resources and the price volatility of the blockchain's native token. Specifically, if all prices are stable and hardly move,

pools need not reallocate resources after the initial choice of efficient allocation. In contrast, when prices are volatile, resources reallocation is more valuable and the more frequent the re-allocation, the more efficient the resource allocation. We, therefore, predict as follows:

Prediction 2: The more frequent a pool re-allocates its resources, the closer it can get to the efficient allocation of its resources

Prediction 3: The value from efficiently reallocating resources increases with the volatility of the blockchain's native token

In order to study the predictions above, we first present a model for efficient resource allocation.

4 A MODEL OF A POOL'S EFFICIENT RESOURCE ALLOCATION

Below, we present a model for the efficient resource allocation for a single pool. For concreteness, we take the resource to be hashrate. The model, however, is general and applies to PoS protocols as well.

Consider N coins that miners can freely switch across. Denote by H_i the hashrate devoted to coin i by pool P, and by H_i^o the total hashrate devoted to coin i by all *other* pools.[4] Then, pool P's fraction of total hashrate is $\frac{H_i}{H_i^o+H_i}$. Given that the probability of mining a block is proportional to the devoted hashrate, the fraction of total hashrate also represents the pool's expected (in the long run) share of the total reward of coin i. Let R_i denote coin i's block reward.[5] Pool P's reward from coin i is then $R_i \cdot \frac{H_i}{H_i^o+H_i}$. Pools' goal is to maximize revenues by efficiently allocating their hashrate across the different coins they mine (assuming the cost per unit of hashrate does not depend on the coin mined), *i.e.*,

$$\textbf{\textit{Objective:}} \quad \underset{H_i \in [1,N]}{\arg\max} \sum_{i=1}^{N} R_i \cdot \frac{H_i}{H_i^o + H_i}. \tag{1}$$

To solve the objective function, *i.e.,* to determine the optimal hashrate allocation for pool P, we apply the method of Lagrange multipliers as follows. Let H^p be the total hashrate of pool P. We then have the constraint

$$g(\mathbf{H}) = \sum_{i}^{N} H_i - H^p = 0. \tag{2}$$

Let $f(\mathbf{H})$ be the function after the arg max in Eq. 1. The Lagrange function $L(\mathbf{H}, \lambda)$ is then defined by

$$L(\mathbf{H}, \lambda) = f(\mathbf{H}) + \lambda g(\mathbf{H}). \tag{3}$$

It follows that the partial derivatives of L is:

$$\nabla_{H,\lambda} L(\mathbf{H}, \lambda) = \sum_{i}^{N} \left(\frac{R_i H_i}{(H_i^o + H_i)^2} - \lambda\right)\partial H_i + \left(-\sum_{i}^{N} H_i + H^p\right)\partial\lambda. \tag{4}$$

Next, we obtain the condition when Eq. 1 is optimal, *i.e.,* condition for achieving the optimal hashrate allocation, by solving $\nabla_{H,\lambda} L(\mathbf{H}, \lambda) = 0$, *i.e.,*

$$\begin{cases} \dfrac{R_i H_i}{(H_i^o + H_i)^2} - \lambda = 0, \\ -\displaystyle\sum_{i}^{N} H_i + H^p = 0. \end{cases} \tag{5}$$

This is equivalent to

$$\begin{cases} \dfrac{R_1 H_1}{(H_1^o + H_1)^2} = \dfrac{R_2 H_2}{(H_2^o + H_2)^2} = \ldots = \dfrac{R_i H_i}{(H_i^o + H_i)^2}, \\ \displaystyle\sum_{i}^{N} H_i = H^p. \end{cases} \tag{6}$$

Given that all the values are positive, Eq. 1 is at its optimal when:

$$\frac{\sqrt{R_1 H_1^o}}{H_1^o + H_1} = \frac{\sqrt{R_2 H_2^o}}{H_2^o + H_2} = \ldots = \frac{\sqrt{R_N H_N^o}}{H_N^o + H_N}. \tag{7}$$

The level of hashrate pool P should devote to coin i, H_i, is then

$$H_i = \frac{H^p \sqrt{R_i H_i^o} + \sqrt{R_i H_i^o} \sum_{j \neq i} H_j^o - H_i^o \sum_j \sqrt{R_j H_j^o}}{\sum_j R_j}. \tag{8}$$

5 DATA AND METHODOLOGY

We test our predictions on the BTC family. That is pools that mine one or more of the following tokens: BTC, BCH, BSV. Currently, Bitcoin (BTC) is the most valuable coin with the largest market share. Its PoW is SHA256 based.

Bitcoin and its forks. Several Bitcoin "forks" exist. Most prominently, Bitcoin Cash (BCH) forked from BTC in 2017. BCH later experienced a hard fork itself, which resulted in the creation of Bitcoin SV (BSV). All three coins have identical SHA256-based PoW algorithms, and they have the same settings for reward halving, *etc.* Given these similarities, all three coins can be mined with the same ASIC machines. Consequently, many mining pools support the mining of two or all of these three coins (BTC, BCH, and BSV). Furthermore, switching hashpower across them is easy. Still, we acknowledge that miners likely bear some costs when reallocating hashrate across the different coins. We account for these costs in Section 7.

In order to calculate the optimal hashrate allocation for pools, we first need to measure mining pools' hashrate. The two most popular methods to measure mining pools' hashrate are: 1) based on pools' block rewards [1, 29], and 2) the data reported by the pools. While the first method works well when averaged over a long period of time, it is not a good fit for our analysis as we are interested in exploring frequent hashrate reallocations. To this end, below we offer a third unique approach based on miners' compensation.

5.1 Methodology

In [1, 29], the authors estimate a pool's hashrate based on the pool's inter-block time relative to the default *network difficulty*; i.e., the minimum threshold required for the acceptance of a result of PoW. For BTC, BCH and BSV, the network difficulty is set to a value that results in the creation of a new block, on average, every T^{block} =

[4] Since, as we show below, pools honestly report their hashrate, we assume full information–pools know the hashrate allocation of all other pools.

[5] We assume that pools are myopic and convert all payments to $US based on current prices.

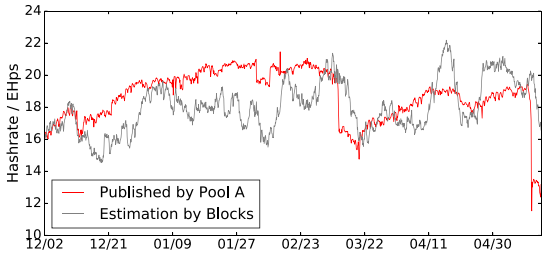

Figure 1: Pool A's estimated and published hashrate (Dec 2019 - May 2020).

10 minutes. Assuming a network difficulty of D and inter-block time T, the expected network hashrate H^{net} is given by:

$$H^{net} = \frac{T}{T^{block}} \cdot D \cdot 2^{32}. \tag{9}$$

A miner's or a pool's hashrate can be estimated by replacing the inter-block time with that of the actual blocks earned.

While, on average, this method reflects the hashrate of a mining pool, this measure is heavily influenced by the pool's "luck"–i.e., the realization of the probability of mining a block. Consequently, when considering a short period of time, the method results in a large variance between the estimated and actual hashrate. This is especially problematic when considering small mining pools. Figure 1 demonstrates this by presenting the discrepancy between the hashrate reported by pool A and the one calculated by method (1). Indeed, on average over a long period of time the two measures are similar. However, given that we want to investigate daily and at times hourly reallocations, we cannot use this measure for our analysis.

To this end, we have developed a process to validate the truthfulness of the public data published by the mining pools. Specifically, we joined the different pools as miners and took advantage of particular payment schemes pools utilize to pay their miners to calculate the pool's actual hashrate.

The most popular compensation methods are: Pay Per Share (PPS), Pay Per Share Plus (PPS+) and Pay Per Last N Share (PPLNS). Under PPS, mining pools pay miners based on the hashrate the miner devotes to the pool and the network difficulty. For example, given a period of time and network difficulty D, if the miner devoted H^m hashrate to the mining pool that was expected to result in N blocks mined, then if the reward from each block is R, the miner's income, I, is calculated as:

$$I = \frac{H^m}{D \cdot 2^{32}} \cdot N \cdot R. \tag{10}$$

Note, that this compensation is paid out regardless of whether the blocks were indeed rewarded to the pool. That is, under PPS, miners earn a guaranteed fixed payment and the mining pool bears the risk of bad luck. Consequently, one cannot infer the hashrate of pools based on the PPS compensation.

Under PPLNS, the mining pool shares the risk with the miners and pays miners based on the actual block rewards the pool gained. For example, consider a time period where overall N blocks were mined on the network and assume the pool received rewards from $M \leq N$ blocks. If a miner devoted an average H^m hashrate power to the mining pool and the pool's total hashrate is H^p, the miner will receive

$$I = \frac{H^m}{H^p} \cdot M \cdot R. \tag{11}$$

Note that miners know the hashrate they devote to the pool, H, and their income, I. Moreover, block rewarded to the pool, M, and the rewards, R, are publicly available. Assuming that pools pay miners honestly, as a miner, we can infer each pool's hashrate based on our compensation from each and every pool we joined.

In general, miners' compensation for mining BTC, BCH and BSV consists of two parts: 1) block reward; and 2) transaction fees associated with the block mined. An alternative compensation method is Pay Per Share Plus (PPS+), which combines the PPS and PPLNS methods. Under PPS+, which became one of the most common payment methods, the mining pool pays miners the base block reward as in PPS, and their proportional part of transaction fees associated with the blocks mined. The mining pools usually keep a small portion of the transaction fees as their own profit. Just like with PPLNS, given that each mined block's transaction fees are public information, as a miner compensated by PPS+, we can verify the pool's hashrate.

We joined six pools as miners and received compensation based on PPS or PPS+, depending on the compensation plans the pool offered.

5.2 Data Collection

We have used an ASIC machine of type AntMiner S9 SE to connect to different pools. It has an energy consumption of 1360W and can achieve around 17T hash computations per second for BTC, BCH or BSV mining. We started mining on Mar 18, 2020; and the data presented in Section 7 covers the period until May 18, 2020. We present data on later dates in Sections 8. We joined six major mining pools, whose sum hashrate exceeds half of the total hashrate devoted to BTC, BCH and BSV. We join each pool and mine BTC for one day at a time; iterating through the six mining pools continuously.

Given that we switch pools every day, we do not have continuous data for any pool but rather see windows of operation. Furthermore, since the pools only pay on a daily basis, the data we obtained from the pools does not meet our needs for the analysis of efficient allocation. Instead, we use the data to independently and accurately verify the hashrate reported by the pools can be trusted and used in our analysis. Specifically, pools report their hashrate every 1-3 minutes and there are different websites that aggregate this information. Pools, however, have incentives to over-report their hashrate. It is, therefore, crucial for us to first confirm that the hashrate reported by pools is truthful.

To this end, we collect the hashrate reported by the mining pools from miningpoolStats [4], a leading website for hashrate information that aggregates pools' information by APIs to the pools' websites. We have continuously collected the hashrate published by the mining pools and the block information since Dec 1, 2019.

We obtained data on blocks mined from blockchair.com [2]. The information includes block time, block rewards, transaction fees included with the block, and the guessed miner (mining pool). Though the guessed miner is not presented in every block, we find that for the major pools, the information is accurate comparing against the

Table 1: Hashrate verification of mining pools with PPS+ (Pool B, Pool C, and Pool D) and with PPS payment methods (all the pools). 95% Confidence Interval (CI) calculated for the series of differences of daily expected and actual incomes. All the units are BTC in 10^{-6}.

Pool	Expected Daily Income	Actual Daily Income	Differecen	95% CI
Transaction fees in rewarded blocks in addition to base block reward (PPS+)				
Pool B	9.146	9.135	−0.11	±0.239
Pool C	13.808	13.981	0.173	±0.345
Pool D	10.448	10.339	−0.109	±0.243
Base block reward (PPS)				
Pool A	289.49	286.15	−3.34	±5.36
Pool B	273.48	273.09	−0.39	±1.09
Pool C	283.10	281.44	−1.66	±2.32
Pool D	276.05	274.30	−1.75	±4.58
Pool E	274.90	274.19	0.71	±2.25
Pool F	293.45	290.34	−3.11	±3.76

data published by the pools. We derive the total hashrate of each network from the network difficulty stored in the block information. Finally, we collect the coins' market price from CoinGeko [3].

5.3 Do pools report truthfully?

To answer this question, we take all pools that paid us based on the PPS+ compensation method and compare our actual income to the income we would have expected to gain if we calculate our compensation based on the hashrate published by the pool. Specifically, Pool B, Pool C, and Pool D paid us based on the PPS+ compensation method. The expected income is calculated based on Eq. 11 where before the May 09, 2020 BTC halving, the base block reward was 12.5 coins, and became 6.25 after. Table 1 presents the results.

As the table shows, the discrepancy between our measure and the values published by the pools is within the 95% Confidence Interval (CI) (where CIs were calculated on the series of differences between the actual and expected pay). We, thus, consider the hashrate reported by the pools to match theoretical expectation. Note that the CI for the PPS+ reward is larger than that for PPS. This is driven by the fact that transaction fees are block specific and thus the fees we receive as a miner depend on the blocks the pools assigned to us, which is not observable. Also note that given that we mined for the same period of time and devoted the same level of hashrate to all pools, we can infer that the much higher compensation we received from Pool C implies that during the time that we conducted the experiment, Pool C had the best luck and Pool B got the least luck either in terms of blocks mined or in terms of transaction fees from blocks.

The hashrate verification above assumes that pools pay their miners honestly. In order to verify that this is indeed the case, we use the PPS compensation we received from all the pools. Table 1 shows the PPS income we received from the pools and compares it to the expected income based on Eq.10. As expected, the actual PPS income from all six pools is almost the same. Moreover, the difference between the actual and expected income is minimal. We conclude that the data reported by the six pools we study is reliable.

6 RESULTS

In this section, we examine the predictions we developed in Section 3. In general, pools' ability to optimize hashrate allocation is limited as, typically, it is the individual miners that choose which crypto to mine. Consequently, unless the pool "owns" some hashrate, hashrate allocation is not always in the control of the pool. Indeed, surveying the leading mining pools, we found that some of the pools control a portion of the mining hashrate, *i.e.*, are miners themselves. Moreover, as we further discuss below, recently, some pools started offering "smart mining" for their users where the pool automatically switches miners' hashrate from one coin to another, based on user-defined policies or, more often, the pool's decision.

For our analysis, we use pools' current allocation as the benchmark, and measure the value from our different hashrate allocations by comparing the difference between the value a pool would have earned were it to allocate optimally to that pool's value from its actual allocation. We analyze the hashrate allocation of six pools– all among the largest pools in terms of hashrate devoted to BTC, BCH and BSV. Pool A, Pool B, and Pool C mine all three cryptos, whereas Pool D, Pool E, and Pool F only mine BTC and BCH.

To standardize the data, we take each data point to be the average hashrate published over one hour (pools publish data every 1-3 minutes). Taking our unit of analysis to be at least an hour is a reasonable assumption given that the inter-block times of BTC, BCH or BSV varies between few minutes to half an hour and necessarily takes longer for individual pools. We apply the efficient hashrate allocation for that point of time. We consider the pool's revenue per block mined to be the coin's block reward multiplied by the current market price.[6]

6.1 Two vs. Three Chains

We analyze the value from reallocating across a larger number of blockchains by comparing the benefits from reallocating across two coins with the benefits from reallocating across three coins. Pools A, B, and C mine all three coins. Since the results for two-coins

[6]We present in Section 8.3 the results when block transaction fees are also incorporated in the pools' revenues and show that our results remain qualitatively unchanged.

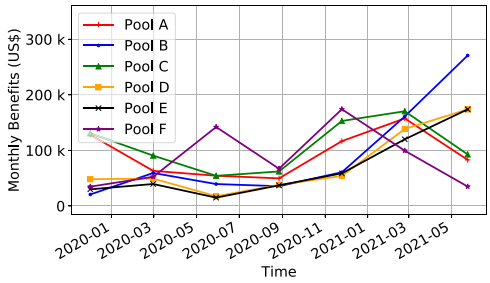

(a) Two coins.

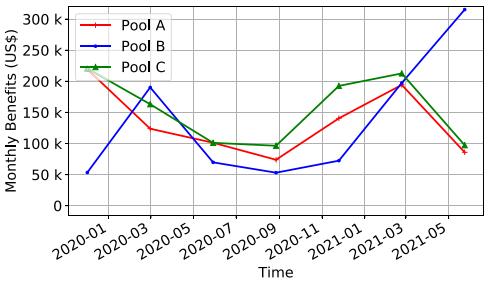

(b) Three coins.

Figure 2: Benefit over longer period of time.

reallocation between BTC and BCH are similar to those achieved when reallocating between BTC and BSV, we only present the results for BTC and BCH.

We first present the benefits from optimal hashrate reallocation over time. Figure 2 illustrates the monthly benefits for nearly two years, from December 2019 to August 2021, for the case of two-coins (Figure 2(a)) and three-coins (Figure 2(b)) reallocation. Each data point represents the average monthly benefits for the following three months. For instance, the last data point of June 2021 represents the average from June to August 2021. For all pools that mine all three coins, the benefit from reallocating across three coins are much larger than the potential benefits in the two-coins case. This supports our first prediction. Interestingly, the benefits of most pools are flat or decreasing till September 2020 when they start climbing up. This period of time highly overlaps with the period of time when the prices of many cryptocurrencies surged.

Since pools do not control all their hashrate, next we present the potential benefits from actively reallocating hashrate across multiple blockchains as a function of the percentage of hashrate the pool can transfer across the different coins. Figure 3(b) shows that the benefits from reallocating across three coins almost doubles the potential benefits relative to the two-coins case in Figure 3(a). This further supports our first prediction. Moreover, as the figure shows, the increased value from optimal hashrate reallocation differs substantially across the different pools. This likely implies that some pools have better control over some of their hashrate than others. Interestingly, pools need only to control a relatively small share of their hashrate in order to achieve the optimal allocation. For example, Pool B and Pool C need to control approximately 2% of their hashrate, in order to reach near-optimal performance. In general, controlling between 3% and 10% of a pool's hashrate enables optimal performance, and provides monthly benefits in the order of hundreds of thousands of USD, as shown on the y-axis. Still, as the figure shows, even if the pool does not control enough hashrate to achieve the optimal allocation, the benefits from three-coins reallocation is still slightly larger than the two-coins case. That is, in line with our first prediction, reallocating across more coins increases the potential benefits from hashrate reallocation.

6.2 Reallocation Frequency

The analysis above assumes that pools reallocate hashrate every hour. This might require too much logistics and thus raises the question of whether the returns from a less frequent optimization would be significantly lower than when reallocating every hour. Figure 3(c) and Figure 3(d) depict the monthly expected rewards as a function of the reallocation frequency between two cryptos and among three cryptos, respectively. Specifically, to simulate the case where a pool reallocates every x hours ($x \geq 1$), we keep the optimal hashrate set at time t constant till time $t + x$.

As expected, the figure shows that reallocating every hour provides the largest potential benefits such that as the time window between reallocations gets longer, the potential benefits decrease. This result supports our second prediction. However, it is noteworthy that the decrease in the potential benefits due to less frequent hashrate reallocation flattens once re-optimization is performed every 12 hours or more. This suggests that pools that cannot optimize every half a day can re-optimize every other day or two without losing much.

6.3 Price Volatility

Next we examine how price volatility of the blokchain's native token affects the potential benefits from efficient hashrate allocation. To this end, we explore how exogenous shocks that have large effects on the relevant token price affect the potential benefits from efficient resource allocation. Specifically, during the period for which we have data, three important events occurred: (i) BSV's market price surpassed that of BCH; (ii) BCH and BSV block reward halving that happened two days apart; and (iii) BTC block reward halving. The three events are depicted in Figure 4, which shows the hashrate and market price for the three coins over the six months for which we have data. In the analysis below, we assume that the pools control 10% of their hashrate and can reallocate it every hour. We depict the *daily* benefits for the explored mining pools as a function of the time of the event, where Day 0 marks the time of the explored event.

BSV Surpasses BCH. On Jan 14, 2020, the market price of BSV surpassed that of its original BCH fork for the first time (BCH also went through a price surge that day). In order to capture the effect on potential benefits from efficient reallocation, we only study the pools that mine all 3 coins.

The results are depicted in Figure 5(a). Supporting our third prediction, the figure shows that benefits peak on Day 0, when the given event occurs and price volatility is at its highest. The peak is more pronounced for Pool A and Pool C, and is in the order of $20k. While the peak for Pool B is much smaller, it is still the case that the

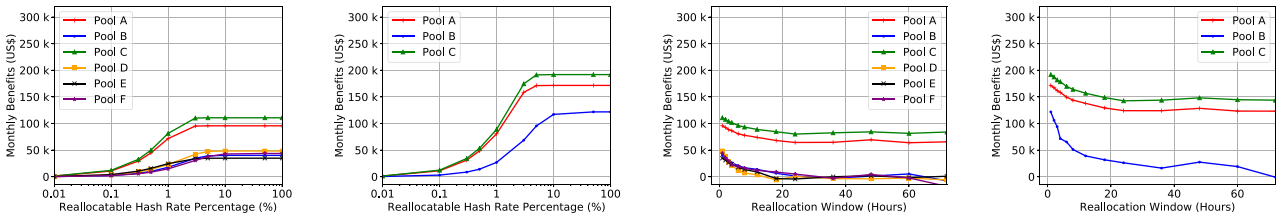

(a) BTC and BCH. (b) BTC, BCH and BSV. (c) BTC and BCH. (d) BTC, BCH and BSV.

Figure 3: Monthly average benefit from efficient hashrate allocation from Dec 2019 to May 2020. With transaction rewards.

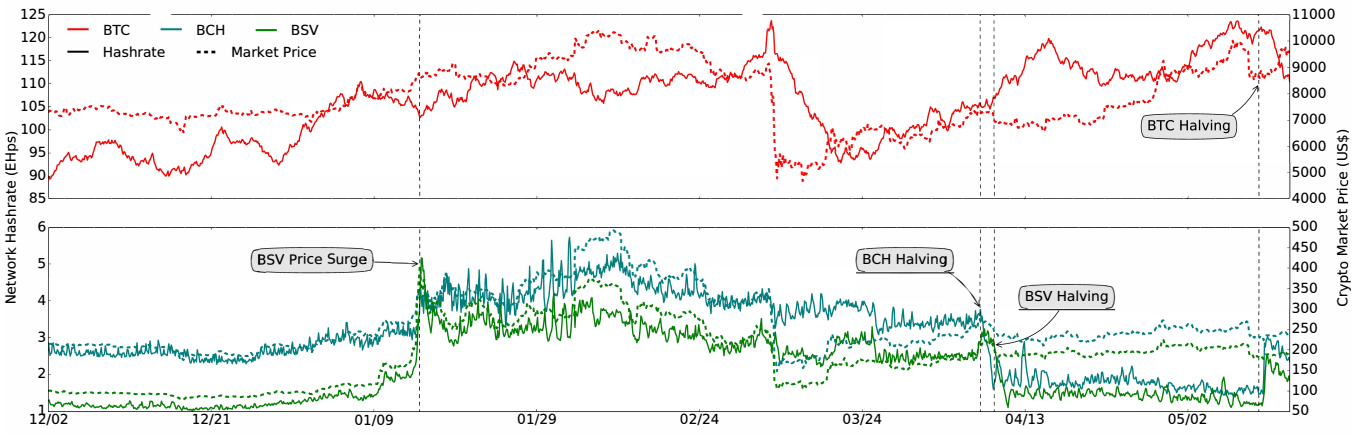

Figure 4: Total network hashrate and market prices for BTC, BCH and BSV; Dec 2019 to May 2020.

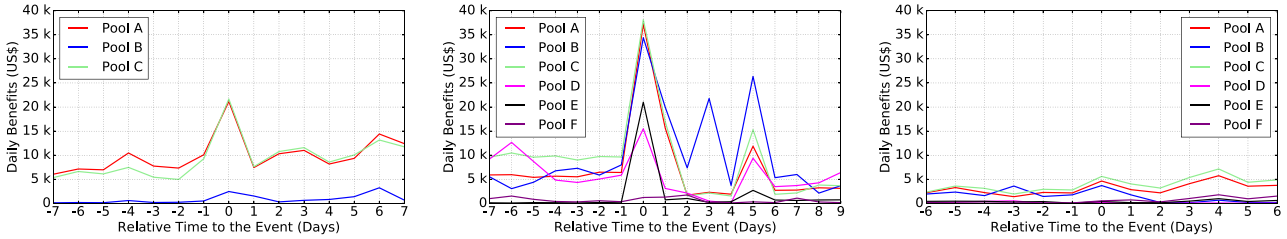

(a) BSV price surpasses BCH. (b) BCH and BSV halving. (c) BTC halving.

Figure 5: Daily benefit from optimal hashrate allocation for different events.

benefits from efficient allocation are larger when price volatility is higher–the comparison we are after. Note that the smaller effect for Pool B is consistent throughout our analysis (see previous figures) and suggests that Pool B is likely more efficient in its hashrate allocation and thus its benefits from more effectively re-allocating its hashrate are smaller. Interestingly, there is a small peak few days before and after the event. This is consistent with the fact that some in the BTC/BCH/BSV community have anticipated the event and prices of the three coins responded accordingly.[7]

BCH and BSV Halving. On Apr 8 and 10, 2020, the base block rewards of BCH and BSV halved from 12.5 to 6.25 coins per block. Figure 5(b) shows that the above events lead to a significant price drop of both BCH and BSV. As suggested by our third prediction, this price volatility indeed resulted in large increase in benefits from a more efficient hashrate allocation. In particular, daily benefits on Day 0 (i.e., Apr 8) reached $37k. In general, the figure shows that potential benefits for the different pools follow similar trends. The hashrate devoted to BCH and BSV experienced oscillations after the event, resulting in the oscillations in the benefits from more efficient allocations. This is likely the result of all pools re-allocating based on what other pools are doing. According to the figure, it took the market at least a week to get back to equilibrium allocation.

BTC Halving. On May 10, 2020, the base block reward of BTC halved from 12.5 to 6.25 coins per block. Given the massive hashrate

[7]While we only present the results for the pools that mine all three coins, the benefits for all other mining pools follow similar trends.

BTC attracts, one might expect the halving to trigger a significant "shock wave" in the mining ecosystem. Nevertheless, the results in Figure 5(c) show the smallest amounts of potential benefits (peak at around $5k on Day 0). This is likely due to the fact that all pools, and the market in general, have gone through several BTC halvings, and thus were able to anticipate the effect of the event. Indeed, BTC price does not change much in the day of the reward halving. We do observe a BTC price drop and BCH and BSV price hike few days after the halving. These then correspond to the benefit peaks we observe 3-4 days after the halving in Figure 5(c). As before, Figure 5(c) shows that Pool B adjusts quickly to the new situation, while Pool A and Pool C continue to sub-optimally favor BTC, in the days after the halving.

7 HOW CAN POOLS ENCOURAGE EFFICIENT RESOURCE ALLOCATION?

As mentioned above, pools' ability to efficiently allocate their hashrate is limited as they do not control most of their hashrate. The results above, however, suggest that pools "leave a lot of money on the table" by not reallocating efficiently. In this section we discuss pools' attempts at encouraging individual miners to more efficiently allocate their hashrate.

Given the large benefits from efficiently reallocating hashrate, one might wonder why individual miners do not reallocate their hashrate more frequently. To examine this, below we detail the process individual miners must follow in order to mine for the blockchain of their choice.

Joining a mining pool involves some fixed costs and some ongoing costs. In particular, in order to join a mining pool, individual miners need to first configure their mining machine with a URL provided by the mining pool. Miners then need to create an account with a "worker name," and have the option to create a password to protect their privacy and secure them from malicious attackers. These are fixed costs that individual miners incur only once when joining a certain pool. As long as the miner keeps mining the same token, they bear no additional costs. However, the URLs for mining different coins are typically different, even when using the same mining pool. Consequently, if a miner wants to switch and mine a different coin, they must first reconnect to the mining pool. This, in turn, implies that the mining pool needs to *validate* the miner as an "incoming miner" before allocating actual mining tasks to the miner and starting to count the income.

We have tested our machine with all six mining pools investigated above, and found that the validation time varies from 3 to 5 minutes. The duration of the validation process, however, depends on the latency between the miner and the mining pool as well as the hash power of the miner, *i.e.,* a more powerful machine might be able to complete the validation task in a shorter time period. That is, the validation process of larger more sophisticated miners might be faster than ours. Regardless, miners cannot mine during the validation process and thus hashrate reallocation is associated with high opportunity costs. Note that since the pool cannot use the reconnecting miner's hashrate during that time as well, this is a cost for the pool as well; and thus an overall market inefficiency.

Below we examine two possible actions pools may take to help move the market toward more efficient hashrate allocation.

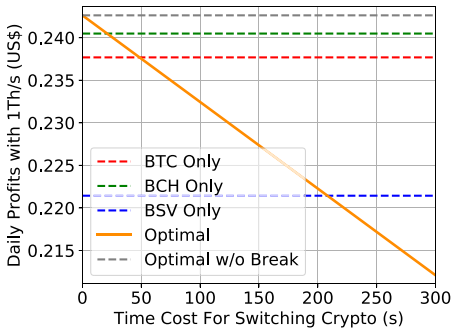

Figure 6: Daily optimal mining profit as a function of the opportunity cost of switching the mined coin.

7.1 Reducing Validation Time

The opportunity cost associated with switching from mining one coin to another depends on the duration of the re-validation process. Pools can, thus, encourage more switching by reducing the time, and thus the cost, individual miners bear when switching the mined coin. How fast should validation be? Figure 6 illustrates the daily value from efficient reallocation as a function of the opportunity cost associated with switching the mined coin. We present four extreme cases that do not involve an costs: *(i-iii)* mining only one coin – BTC, BCH, or BSV; and *(iv)* following our optimal reallocation with no re-validation time breaks. We compare these no-switching costs cases to the case where a miner follows the efficient reallocation and the time it takes to re-validate the miner's account varies from zero seconds to 300 seconds (presented on the axis in Figure 6.)

Since miners cannot generate revenue during the validation process, the line representing the value from following the efficient reallocation with validation breaks is decreasing linearly with the time it takes the pool to validate the miner. The slope of the line then depends on the number of times the miner needs to switch across coins (based on the efficient allocation algorithm). The graph is based on data between May and July 2021, during which, the optimal algorithm dictated switching coins, on average, 36 times.

As Figure 6 shows, if switching across coins takes more than 50 seconds, miners are better off only mining BTC than reallocating efficiently. The break-even duration is shorter for BCH – it is efficient to only mine BCH, if it takes more than 23 seconds to switch across coins. Yet, as far as only mining BSV, the break-even duration is much higher and get to about 3.5 *minutes*.

The results above explain why, recently, some pools started to provide a "one-click switch" function that allows miners to switch the mined coin within the same mining pool without the need to be re-validated [5].

7.2 Smart Mining

Another action pools can take to encourage more efficient allocation is for the pool to allocate the hashrate on the miner's behalf. Indeed, recently, some pools have started offering 'smart mining' features, or SHA256 mining, to their miners. Individual miners enrolled in smart-mining essentially give the pool the decision-making over how to allocate their hashrate. Pools then automatically reallocate

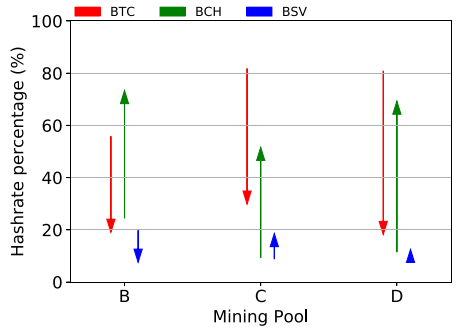

(a) Different pools

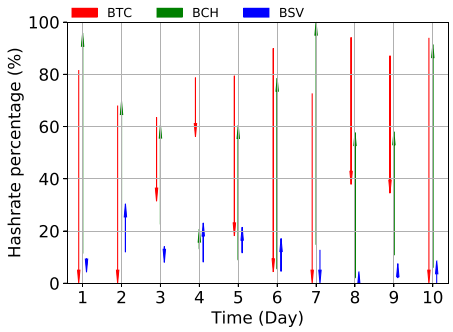

(b) Pool D

Figure 7: Smart mining hashrate allocation of our machine relative to our optimal model.

these miners' hashrate across BTC, BCH and BSV based on expected returns.

We examine the performance of the pools' smart mining offering by using our mining machine and signing up to the smart mining option. We join for a month three major pools, Pool B, Pool C, and Pool D that provide this service, and compare the results to our optimal allocation model. Note that while Pool D does not provide independent BSV mining, as discussed in Section 4, it does provide SHA256 mining among all three coins. We iterate one day at a time through the three mining pools continuously (i.e., mining 10 days for each mining pool) and, as before, derive the daily hashrate allocation based on our awarded income. Our mining machine is AntMiner S9 SE, which can achieve around 17T hash computations per second.

Figure 7(a) shows the difference between the actual allocation of our machine's hashrate by the different pools and our model's efficient reallocation. The arrows start at the actual allocation of our machine and point to our model's optimal allocation. Interestingly, even with smart mining, all three pools over-allocate hashrate to BTC and under-allocate hashrate to BCH, compared to our optimal model. Pool B's allocation is closest to our model; reflected in its relatively lower allocation to BTC at around 55% and higher BCH allocation at around 25%. In contrast, the other pools allocated more than 80% of our machine's hashrate to BTC and only around 10% to BCH.

The large differences between the pools' allocation and our model may be a result of the pools' different optimization model. Specifically, it is possible that pools try to maximize a different objective function and may care, for example, about their hashrate share within the BTC network. Alternatively, it is possible that our machine is too small to effectively reflect the logic behind the pools' strategy. In particular, currently, pools assign miners a task every one or two minutes. During each communication, each machine's hashrate is indivisible. Therefore, to achieve a specific hashrate allocation, pools must assign a specific number of miners to a specific coin, rather than assign each miner's machine to mine several different coins in proportion to the optimal allocation. As pools likely have greater reallocatable hashrate than what is needed to optimize the allocation, they likely randomly move only some of the smart-mining miners to other coins. As a result, our machine would not be always selected to the reallocation and we would not be able to see all reallocations.[8]

Note that Pool D's allocation to BSV is almost optimal. This is interesting as, given that Pool D does not support independent mining of BSV, its BSV allocation is coming only from smart-mining, so the pool can decide exactly how much to allocate to BSV. This supports the view that pools, in general, aim at efficient allocation, however, are limited in their ability to do so. We further examine this by closely studying Pool D's allocation during the 10 days of smart mining our machine devoted to Pool D. As Figure 7(b) shows the difference between the actual hashrate allocation of our machine when joining Pool D and the corresponding optimal allocation of our model for BSV is very close; and specifically, much closer as compared to BTC and BCH.

8 DISCUSSION

In this section, we discuss some potential extensions and robustness checks.

8.1 Pools' Compensation Fees

As mentioned above, pools charge miners a fee as compensation for the service they provide. Our analysis so far assumed that the pools' fees are based only on the block reward. Pools' fees, however, are typically more complex and may account for block transaction fees, as well as vary across the different coins mined. For instance, Pool B takes 4% of the block rewards but only 2% of the block transaction fee rewards, while Pool C charges 3% of the block rewards for BCH and BSV, whereas the fee for BTC is only 2.5%. These fee structures may result in a different efficient allocation for individual miners. To examine whether our analysis is robust to these fee structures, we simulate allocations based on our model accounting for each pool's actual fees.

Our results show little effect. Specifically, there is a slight shift (less than 1%) toward BTC relative to the results in Figure 7 for all three pools. Yet, our results remain qualitatively unchanged.

[8]To verify this, we later join smart-mining with two ASIC machines. We discuss the results in Section 8.2.

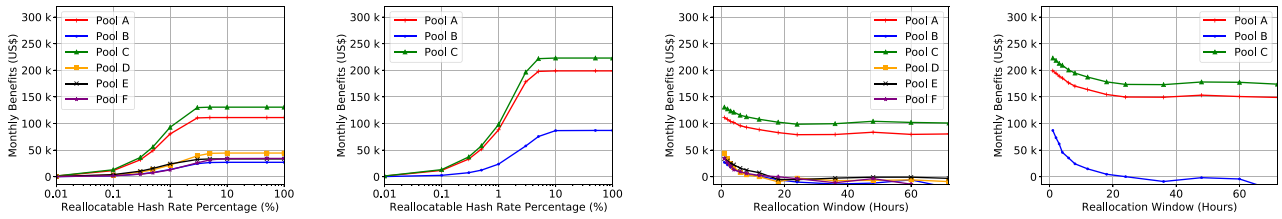

(a) BTC and BCH. (b) BTC, BCH and BSV. (c) BTC and BCH. (d) BTC, BCH and BSV.

Figure 8: Monthly average benefit from efficient hashrate allocation from Dec 2019 to May 2020. Without transaction rewards.

8.2 Smart Mining Task Assignment

Pools that offer smart-mining may, in theory, allocate the miner's hashrate in a way that optimizes the pool's revenues rather than the miner's. While we cannot directly observe whether pools take their own or the individual miner's perspective, we can test whether two identical miners' hashrate is allocated differently. In particular, one can imagine that if a pool is trying to maximize its own revenues, it may take advantage of the aggregated hashrate of all individual miners using the smart-mining feature and allocate it in a way that maximizes its own profit. In this case, it is likely that identical individual miners would be assigned different mining tasks, *i.e.,* to mine different coins, based on the allocation that maximizes the pool's revenues.

We examine this by joining the smart mining of each pool with two machines with identical hashrate that were associated with two different accounts at the same time. The measurement last two months for each pool.

Our results show that for all pools over the entire period of testing, the task assignments of our two identical miners were the same. This suggests that the mining pools are assigning the same task to all miners enrolled in smart mining, without utilizing this share of hashrate to facilitate optimal allocation for themselves.

To further examine pools' hashrate assignment for miners enrolled in smart-mining, we test whether different pools use the same optimization algorithm. To this end, we join two different pools with two identical machines at the same time and test whether our miners are assigned the same tasks. We find that the assigned mining tasks from Pool B, Pool C, and Pool D are not entirely the same. This suggests that the different pools have different optimization algorithms that they follow. In terms of performance, while some pools performed better on some days, none of the pools consistently outperformed others.

8.3 Transaction Fee Rewards

The analysis above assumes that miners' and pools' compensation are based on the block rewards as well as the block transaction fee rewards. During the period we study, however, block transaction fees have experienced several surges, which did not always correspond to surges in prices. Figure 9 presents prices (solid lines) and average block transaction fees (dotted lines) for BTC, BCH, and BSV.[9] As the figure shows, the block transaction fee rewards

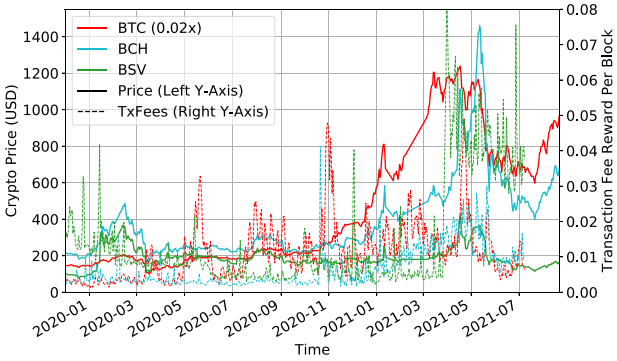

Figure 9: Histogram of prices and transaction reward per block (Dec 2019 - Aug 2021).

do not follow any trend. More importantly, they may, at times, be much more volatile than prices of the tokens. This may raise a concern that some of our results are driven by the volatility and/or magnitude of the block transaction fees.

We have performed our analysis based on only the block rewards; *i.e.,* excluding all block transaction fees. As Figure 8 shows, our results are robust and hold in this case as well. This suggests that mining pools could benefit from efficiently reallocating hashrate even without considering/predicting the block transaction fee rewards.

8.4 Ethereum Mining

In order to study the generality of our results, we study hashrate allocation for the Ethereum (ETH) family; the second largest family of cryptocurrency. The ETH family is composed of two main tokens Ethereum and Ethereum Classic (ETC), where ETC is a hard fork of ETH. The family is mined-by-GPU rather than ASIC. We joined the six largest mining pools that mine ETH and ETC as a miner.

As with the BTC family, we first confirmed that pools report and pay honestly. We then applied our efficient allocation model to test our second prediction.[10] Since inter-block times for ETH and ETC are much shorter than for the BTC family (~13 seconds vs ~10 minutes), we shorten the unit of analysis from an hour to 10 minutes, *i.e.,* we take each data point to be the average hashrate published over 10 minutes. Similarly to Figure 3(a) and Figure 3(c) in Section 6.2, we present in Figure 10(a) the monthly expected rewards as a function of the percentage of hashrate the pool can

[9] In order to make values comparable in a graph, we divide prices and block transaction fees for BTC by 50.

[10] Given that there are at most two coins in the family, we do not test our first prediction.

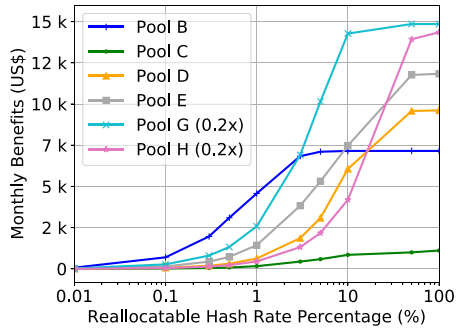

(a) ETH and ETC. Reallocation window is 10 minutes.

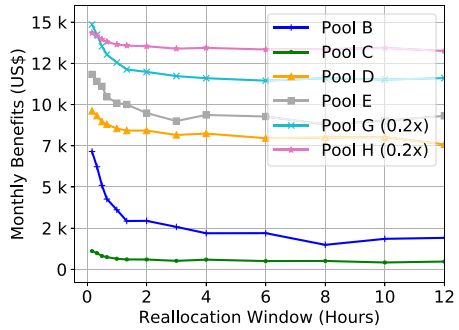

(b) ETH and ETC.

Figure 10: Monthly average benefit from optimal hashrate allocation of ETH and ETC from Dec 2019 to May 2020. With transaction rewards.

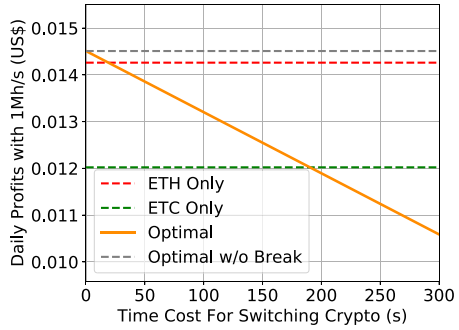

Figure 11: Daily optimal mining profit as a function of the opportunity cost for switching the mining coin.

transfer across the different coins, and in Figure 10(b) the monthly expected rewards as a function of the reallocation frequency.[11]

As expected, Figure 10(b) shows that reallocating every 10 minutes provides the largest potential benefits such that as the time window between reallocations gets longer, the potential benefits decrease. This result supports our second prediction. However, it is noteworthy that the decrease in the potential benefits due to less frequent hashrate reallocation flattens once re-optimization is performed every 2 hours or more. This suggests that pools that cannot optimize every 10 minutes can re-optimize every half day without great loss of efficiency.

Similarly to Section 7.1, we investigate the opportunity costs associated with switching from mining one coin to another. As Figure 11 shows, if switching across coins takes more than 19 seconds, miners are better off only mining ETH than reallocating efficiently. Yet, as far as only mining ETC, the break-even duration is about 3 minutes (190 seconds).

[11] In order to make values comparable in a graph, we divide benefits for Pool G and Pool H by 5.

9 CONCLUSION

This paper proposes a method to increase the efficiency of resource usage under different blockchain consensuses, PoW and PoS. We develop predictions on factors that can increase the efficiency of resource allocation and test these based on a general model for optimal resource allocation, and a unique, large-scale, data set we collected on mining pools of the Bitcoin family. Given that our theoretical model requires the knowledge of mining pools' hashrates, we present a novel hashrate measurement methodology, which we use to demonstrate that mining pools' hashrates can be indirectly, yet cheaply, accurately, and scalably measured. We show that resource efficiency can be improved by (*i*) mining for a larger number of chains, and (*ii*) re-allocating the hashrate more frequently. In addition, the value from resource efficiency increases as the price volatility of the native token increases. Next, we examined how pools can encourage the efficient hashrate allocation of their individual miners. We discuss the cost individual miners may bear when switching the mined coin. Specifically, there is an opportunity cost associated with the time it takes to switch from mining one coin to another, as miners can not mine and enjoy revenues during this time. We calculate the threshold switching time above which miners are better off mining a single token with no switching and show that this threshold varies between $\tilde{2}5$ seconds for BCH to more than 3.5 minutes for BSV.

Pools can also encourage the efficient use of individual miners' resources by offering to re-allocate the hashrate for them. Pools offering such features – smart-mining – seem to, indeed, optimize compensation for miners rather than take advantage of the miners' hashrate to optimize their own revenues.

We have verified the robustness of our results by considering the block transaction fees, the fees that pools charge as compensation for the service they provide, and enrolling in smart-mining as two independent miners. We further demonstrated the generality of our results by collecting data and testing our results on the Ethereum family.

Finally, we stress that our predictions and results are applicable to PoS pools. In particular, scarcity of resources, and thus competition over these resources, exists under both the PoW and the PoS consensuses. In the case of PoS, validators need to attract users to stake their native token with them, as well as allocate their capital across different coins. We provide some examples that highlight the relevance of our results to PoS pools.

For future work, we highlight two directions worth pursuing. First, it is unclear what would be the effects of the proposed optimization policies on the global participation and prices of the blockchains, if deployed widely. The question is if such policies can potentially exacerbate price volatility and create network instability. Even if such events are unlikely to happen, it appears useful to understand the conditions that may lead to such outcomes. Second, while for the near future, block rewards will still be the main return to the miners, it remains an open question if and how the optimization will change when this assumption no longer holds true.

REFERENCES

[1] Bitcoinwiki: Bitcoin difficulty, 2020. https://en.bitcoin.it/wiki/Difficulty/.

[2] BlockChair, 2020. https://blockchair.com/.

[3] CoinGecko API, 2020. https://www.coingecko.com/en/api/.

[4] Mining Pool Stats, 2020. https://miningpoolstats.stream/.

[5] What is the One-Click Switch and how to use it?, 2020. https://support.viabtc.com/hc/en-us/articles/900001529046-What-is-the-One-Click-Switch-and-how-to-use-it-.

[6] S. Abramova, P. Schöttle, and R. Böhme. Mixing coins of different quality: A game-theoretic approach. In *Financial Cryptography and Data Security - FC 2017 International Workshops, WAHC, BITCOIN, VOTING, WTSC, and TA, Sliema, Malta, April 7, 2017, Revised Selected Papers*, pages 280–297, 2017.

[7] I. Bentov, C. Lee, A. Mizrahi, and M. Rosenfeld. Proof of activity: Extending bitcoin's proof of work via proof of stake [extended abstract]y. *SIGMETRICS Perform. Evaluation Rev.*, 42(3):34–37, 2014.

[8] G. Bissias, B. N. Levine, and D. Thibodeau. Using economic risk to model miner hash rate allocation in cryptocurrencies. In *Data Privacy Management, Cryptocurrencies and Blockchain Technology - ESORICS 2018 International Workshops, DPM 2018 and CBT 2018, Barcelona, Spain, September 6-7, 2018, Proceedings*, pages 155–172, 2018.

[9] G. Bissias, B. N. Levine, and D. Thibodeau. Greedy but cautious: Conditions for miner convergence to resource allocation equilibrium. *CoRR*, abs/1907.09883, 2019.

[10] P. Chatzigiannis, F. Baldimtsi, I. Griva, and J. Li. Diversification across mining pools: Optimal mining strategies under pow. *CoRR*, abs/1905.04624, 2019.

[11] L. W. Cong, Z. He, and J. Li. Decentralized mining in centralized pools. Technical report, National Bureau of Economic Research, 2019.

[12] J. A. D. Donet, C. Pérez-Solà, and J. Herrera-Joancomartí. The bitcoin P2P network. In *Financial Cryptography and Data Security - FC 2014 Workshops, BITCOIN and WAHC 2014, Christ Church, Barbados, March 7, 2014, Revised Selected Papers*, pages 87–102, 2014.

[13] A. Dorri, S. S. Kanhere, R. Jurdak, and P. Gauravaram. Lsb: A lightweight scalable blockchain for iot security and anonymity. *Journal of Parallel and Distributed Computing*, 134:180–197, 2019.

[14] G. Fanti, L. Kogan, S. Oh, K. Ruan, P. Viswanath, and G. Wang. Compounding of wealth in proof-of-stake cryptocurrencies. In I. Goldberg and T. Moore, editors, *Financial Cryptography and Data Security - 23rd International Conference, FC 2019, Frigate Bay, St. Kitts and Nevis, February 18-22, 2019, Revised Selected Papers*, volume 11598 of *Lecture Notes in Computer Science*, pages 42–61. Springer, 2019.

[15] M. V. X. Ferreira and S. M. Weinberg. Proof-of-stake mining games with perfect randomness. In P. Biró, S. Chawla, and F. Echenique, editors, *EC '21: The 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, 2021*, pages 433–453. ACM, 2021.

[16] B. Fisch, R. Pass, and A. Shelat. Socially optimal mining pools. In *Web and Internet Economics - 13th International Conference, WINE 2017, Bangalore, India, December 17-20, 2017, Proceedings*, pages 205–218, 2017.

[17] P. Gazi, A. Kiayias, and D. Zindros. Proof-of-stake sidechains. In *2019 IEEE Symposium on Security and Privacy, SP 2019, San Francisco, CA, USA, May 19-23, 2019*, pages 139–156. IEEE, 2019.

[18] F. Irresberger, K. John, P. Mueller, and F. Saleh. The public blockchain ecosystem: An empirical analysis. *NYU Stern School of Business*, 2021.

[19] D. Ivan. Moving toward a blockchain-based method for the secure storage of patient records. In *ONC/NIST Use of Blockchain for Healthcare and Research Workshop. Gaithersburg, Maryland, United States: ONC/NIST*, pages 1–11, 2016.

[20] A. Kiayias, A. Russell, B. David, and R. Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol. In J. Katz and H. Shacham, editors, *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part I*, volume 10401 of *Lecture Notes in Computer Science*, pages 357–388. Springer, 2017.

[21] S. King and S. Nadal. Ppcoin: Peer-to-peer crypto-currency with proof-of-stake. *self-published paper, August*, 19(1), 2012.

[22] Y. Lewenberg, Y. Bachrach, Y. Sompolinsky, A. Zohar, and J. S. Rosenschein. Bitcoin mining pools: A cooperative game theoretic analysis. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2015, Istanbul, Turkey, May 4-8, 2015*, pages 919–927, 2015.

[23] W. Li, S. Andreina, J. Bohli, and G. Karame. Securing proof-of-stake blockchain protocols. In J. García-Alfaro, G. Navarro-Arribas, H. Hartenstein, and J. Herrera-Joancomartí, editors, *Data Privacy Management, Cryptocurrencies and Blockchain Technology - ESORICS 2017 International Workshops, DPM 2017 and CBT 2017, Oslo, Norway, September 14-15, 2017, Proceedings*, volume 10436 of *Lecture Notes in Computer Science*, pages 297–315. Springer, 2017.

[24] X. Liu, W. Wang, D. Niyato, N. Zhao, and P. Wang. Evolutionary game for mining pool selection in blockchain networks. *IEEE Wirel. Commun. Lett.*, 7(5):760–763, 2018.

[25] F. Saleh. Blockchain without waste: Proof-of-stake. *The Review of financial studies*, 34(3):1156–1190, 2021.

[26] A. Spiegelman, I. Keidar, and M. Tennenholtz. Game of coins. *CoRR*, abs/1805.08979, 2018.

[27] J. Sun, P. Tang, and Y. Zeng. Games of miners. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '20, Auckland, New Zealand, May 9-13, 2020*, pages 1323–1331, 2020.

[28] C. Wang, X. Chu, and Q. Yang. Measurement and analysis of the bitcoin networks: A view from mining pools. *CoRR*, abs/1902.07549, 2019.

[29] L. Wang and Y. Liu. Exploring miner evolution in bitcoin network. In *Passive and Active Measurement - 16th International Conference, PAM 2015, New York, NY, USA, March 19-20, 2015, Proceedings*, pages 290–302, 2015.

[30] W. Wang, D. Niyato, P. Wang, and A. Leshem. Decentralized caching for content delivery based on blockchain: A game theoretic perspective. In *2018 IEEE International Conference on Communications, ICC 2018, Kansas City, MO, USA, May 20-24, 2018*, pages 1–6. IEEE, 2018.

[31] G. Yamamoto, A. Laszka, and F. Kojima. Equilibrium of blockchain miners with dynamic asset allocation. *CoRR*, abs/2006.08016, 2020.