

Estimating Profitability of Alternative Crypto-currencies

Danny Yuxing Huang
Princeton University

Kirill Levchenko and Alex C. Snoeren
University of California, San Diego

Abstract

Digital currencies have flourished in recent years, buoyed by the tremendous success of Bitcoin. These blockchain-based currencies, called *altcoins*, are associated with a few thousand to millions of dollars of market capitalization. Altcoins have attracted enthusiasts who enter the market by mining or buying them, but the risks and rewards could potentially be significant, especially when the market is volatile. In this work, we estimate the potential profitability of mining and speculating 18 altcoins using real-world blockchain and trade data. Using opportunity cost as a metric, we estimate the mining cost for an altcoin with respect to a more popular but stable coin. For every dollar invested in mining or buying a coin, we compute the potential returns under various conditions, such as time of market entry and hold positions. While some coins offer the potential for spectacular returns, many follow a simple bubble-and-crash scenario, which highlights the extreme risks — and potential gains — in altcoin markets.¹

1 Introduction

In its nine years of existence, Bitcoin [21] (BTC) has been tremendously popular, reaching a market capitalization of \$60 billion at the time of this writing. Its success has inspired the creation of many new digital currencies that borrow Bitcoin’s key design principles — a blockchain-based public ledger and a means of acquiring a stake in the currency computationally. Indeed, many of these digital currencies are direct clones of Bitcoin, created by tinkering with a few of the design parameters and coming up with a catchy name. Many coins, for instance, change the frequency with which blocks are generated,

shortening Bitcoin’s 10-minute block period to something more amenable to digital transactions. As an example, one of Bitcoin’s major competitors, Litecoin (LTC), produces a block every two and a half minutes, which the developers hoped would reduce transaction confirmation delays [7]. Coins also differ in the rate at which coins are generated by mining, as well as the hash function used to secure the network, among other variations.

Today, there are over 1,400 such currencies, collectively called *altcoins*. Unlike Bitcoin, which can be used as a medium of exchange, few merchants are known to accept altcoins. Some coins, like Auroracoin (AUR) and Dogecoin (DOGE), briefly dominated news headlines for their purported goals of supporting national economies [19] and facilitating digital payments [13], but few of these ends have yet materialized. The vast majority of altcoins appear to serve largely as speculative investment vehicles. They are often listed on altcoin exchanges, such as Poloniex or Bittrex [2, 8], which are open marketplaces where buyers and sellers exchange altcoins. The market capitalization and trade volume for a given altcoin can range from thousands to millions of dollars. Whether one believes in a coin’s merits or not, altcoins offer ample opportunities for currency speculation, especially given the volatile prices. As such, the potential risks and rewards can be significant.

In general, altcoins attract mindshare and market-share among digital currency enthusiasts. Among them are *miners*, who expend energy in finding hash collisions to computationally produce the digital assets in a process known as *mining*; and *speculators*, who take advantage of the price volatility of altcoins and profit from speculation.² An investor has the flexibility to choose to become a miner or a speculator at any point in time. To enter the altcoin market, an investor can mine an altcoin or buy it

¹This paper is an extended version (University of California, San Diego Technical Report CS2017-1019) based on a short paper of the same title published in the *Proceedings of Financial Cryptography and Data Security Conference*, 2018.

²We are aware that there are many ways to profit from altcoins, including gaming the mining protocol [18, 15] or trading altcoins as if they were penny stocks [12, 17]. It is beyond the scope of this paper to discuss all these ways. Furthermore, there may be other participants in the altcoin ecosystem that are not necessarily profit-driven; again, these participants are beyond our scope.

from the market. If she mines some units of an altcoin, she can further hold onto the coin units and sell them when the price rises, thereby speculating in addition to mining.³

This fluidity in role makes it difficult for us to retroactively analyze the profitability of investing in altcoins. For instance, it is a non-trivial task to identify activities of the same miners based on the blockchain of an altcoin, as a miner may use multiple wallet addresses. Also, even though there are techniques to identify wallet addresses of Bitcoin exchanges [20], identifying the wallet addresses for altcoin exchanges across the 1,400 altcoins would be a major effort. Even if we can identify the wallet addresses of altcoin exchanges, once a miner moves her mined coins into an exchange wallet, exchanges typically do not publicly disclose when the coins are sold and at what price. Tracing an individual investor's profitability through exchange data is therefore a significant challenge.

Given these issues, our goal is to develop a set of techniques to analyze the profitability of altcoin investment, from the collective perspectives of miners and speculators. Specifically, our analysis asks these questions: (1) For every \$1 of investment, what is the profitability of mining versus speculating an altcoin? (2) How does the potential profitability of mining/speculating vary across multiple coins? (3) What are the risks involved?

In this work, we use historical data to examine the risks and rewards of altcoin mining and speculation. Because many coins use the same hash function for the mining process, an altcoin miner has a choice of which currency to mine given the same mining hardware. Assuming he or she has no predilection for a particular coin, a miner will choose to mine a coin that offers the highest instantaneous reward, and will switch from currency to currency in response to changing market conditions. This creates competition among coins for mining power, and this allows us to quantify the relative mining risk and reward of an altcoin measured as the *opportunity cost* relative to a more stable currency. For example, mining \$1 worth of Dogecoin on January 1, 2014 required more than 7 billion hash computations on average; the same amount of effort could be spent mining Litecoin, which uses the same hash algorithm as Dogecoin, for an expected revenue of \$0.61. Regardless of the costs associated with operating the mining hardware, on that day, Dogecoin offered an additional \$0.39 of revenue compared with Litecoin, for doing the same 7 billion hash computations. By mining Dogecoin, the miner has given up the opportunity to earn Litecoin; thus, the opportunity cost of mining Dogecoin relative to Litecoin \$0.61. We

³Throughout this paper, we use "altcoin" or "coin" interchangeably to refer to the crypto-currency. In contrast, we use "units" to refer to individual units of reward as a result of mining an altcoin. Following the convention, we use "Bitcoin" to refer to the crypto-currency, and "bitcoin" to refer to units of Bitcoin.

use this opportunity cost to study the relative profitability of each coin as viewed by a profit-driven miner.

In addition to mining, another way to access the altcoin market is to buy coins at an exchange. We use market data covering 1,436 altcoins to measure the potential profitability of investing \$1 in each coin under different simulated environments. Due to market variability, altcoin speculation involves quite a bit of risk; however, we find well-timed investments can potentially net hefty gains — and some even predictable (modest) profits. In our analysis, we compare coins not only on their potential for profit but also the risks associated with speculating on the coin.

In summary, the contributions of this work are as follows. First, we develop a methodology for estimating a miner's investment in an altcoin by calculating her opportunity cost of mining a more stable alternative like Bitcoin or Litecoin. Second, we compare the profitability of mining and speculation of different coins through simulations under varying conditions. Using this entirely descriptive rather than prescriptive method, we find that miners who mine an altcoin immediately after it is listed on exchanges tend to enjoy higher potential returns than miners who mine on subsequent days. In contrast, speculators who buy an altcoin shortly after it is listed on exchanges are likely to generate lower returns than speculators who buy at a later point, and they also generate lower returns than miners in the same period.

The rest of this paper is organized as follows. Section 2 provides the technical background for subsequent analyses. Section 3 describes our dataset, and Section 4 explains our methodology. In Section 5, we estimate the opportunity cost of mining, followed by Section 6, where we use the opportunity costs and market prices to estimate the potential profitability of mining and trading. We discuss two altcoins in detail in Section 7 to highlight common observations across altcoins, along with exceptions to such patterns. Section 8 discusses limitations to our methodology, and Section 9 explores related work online and in the literature. Finally, we conclude our paper in Section 10.

2 Background on Altcoins

Based on data we collected from exchanges, at least 1,400 altcoins — i.e., currencies derived from Bitcoin — were mined and traded after Bitcoin was released in 2008. Other crypto-currencies, like Ripple [9] and Ethereum [6], were developed completely independently. The vast majority, however, are adapted from Bitcoin, with a variety of minor differences. For example, the developers of LTC were dissatisfied with Bitcoin's average block rate of 10 minutes; thus, LTC has a block rate of 2.5 minutes in an apparent effort to speed up trans-

Table 1: Overview of the coins that we study.

Coin	Blockchain (Days)	Trade (Days)	(a) Overall Statistics						(b) Opportunity Cost Analysis			
			Market Cap (\$)	Trade Vol (\$)	V_1 (%)	V_7 (%)	HF	Pf	Anlyns (Days)	Market Cap (\$)	Opp Cost (\$)	Corr
BTC	2,711	2,128	9,129,946,498	30,750,001,190	0.00	0.81	SH	W	2,128	8,112,097,758	1,764,083,690	1.00
LTC	1,691	1,122	182,338,574	1,542,156,841	-0.16	-0.96	Sc	W	1,122	106,266,194	179,931,265	1.00
DOGE	919	915	26,938,503	236,217,401	-0.52	-2.48	Sc	A	276	17,476,815	43,769,744	0.45
PPC	1,375	1,045	8,511,846	187,487,304	-0.39	-1.76	SH	S	1,375	8,129,715	4,137,808	0.98
AUR	852	825	2,625,706	20,455,696	-0.95	-4.73	C	W	827	2,828,901	1,357,342	0.87
DGC	1,101	1,045	270,634	4,756,876	-0.50	-4.01	C	W	563	171,734	831,855	0.96
VIA	678	677	125,390	4,426,677	-1.16	-4.84	Sc	A	160	403,478	240,058	0.94
UNO	950	888	371,620	4,204,613	-0.15	-1.63	SH	A	566	491,855	436,511	0.62
RPC	878	686	5,561	938,403	-1.90	-6.96	Sc	W	691	4,323	276,416	0.93
ARG	1,079	858	8,630	733,145	-1.12	-5.61	Sc	W	892	7,540	91,499	0.90
EFL	797	796	152,200	372,766	-0.11	-1.04	Sc	W	797	162,355	77,047	0.91
WBB	464	454	192,722	317,571	-1.83	-5.74	Sc	W	464	164,643	61,213	0.86
CURE	754	744	436,351	250,942	-0.61	-4.60	SH	S	754	421,326	37,761	0.73
XJO	976	701	4,368	183,473	-0.85	-6.25	SH	W	789	3,519	114,994	0.97
BTA	383	379	32,229	30,163	-0.84	-3.38	Sc	W	383	31,293	8,020	0.93
HAM	694	497	14,166	25,042	-1.20	0.59	SH	S	575	12,698	5,575	0.70
SWING	269	230	3,091	23,354	-1.55	-5.40	SH	S	269	4,102	3,441	0.84
TROLL	177	57	22,144	2,639	0.05	0.85	Sc	S	58	24,776	280	0.70
VCN	378	246	3,390	982	1.61	7.38	SH	W	259	4,316	852	0.82
DOT	775	331	17,656	917	-1.31	-5.46	Sc	W	346	4,974	3,198	0.78

action processing. Furthermore, BTC uses SHA-256 as the proof-of-work (PoW) function. As Bitcoin miners were increasingly using dedicated hardware like ASICs to mine bitcoins, LTC used Scrypt as PoW, in an apparent attempt to discourage the use of ASICs in mining [23]. In short, these “cloned” coins, like LTC, form the basis of our study, as their shared structure admits straightforward comparisons.

We consider two logical participants in an altcoin ecosystem — although any single person or enterprise may play either or both roles: miners, who provide the computational resources to generate altcoins by mining; and speculators, who simply buy and sell altcoins on the open market. A typical coin’s lifetime starts when the developer releases the coin’s client to the public. A miner downloads the client, connects to the coin’s peer-to-peer network, and mines the coin with CPU, GPU, or dedicated hardware. Once there is sufficient mining power toward the coin, or once there is enough interest in the community, an exchange may decide to list the coin, in a process similar to initial public offerings. At this point, a public market for the coin exists. Miners can sell their altcoin units, and speculators can start buying the coins.

Because we want to measure the profitability of miners and speculators, we need a way to quantify the costs and revenues of each. Our speculation analysis relies on trading data. To analyze mining profitability, we compare the resources required to mine one coin against another, more popular base currency with the same technical attributes. Here, two attributes of a coin are important:

Types of proof: Each coin defines a mechanism to allow a miner to prove that they have successfully mined a block. Bitcoin, for instance, uses proof-of-work (PoW). Miners with higher computational power are more likely

to be rewarded. To compensate for the time-varying hashing capabilities of the active miners, PoW coins constantly adjust their *difficulty*, which dictates the expected number of hashes to mine a block, in an attempt to keep the expected rate at which coins are mined constant. The cost of PoW mining, therefore, is a time-varying function unique to each miner, dependent on that particular miner’s ability to acquire and operate the requisite computational resources, often referred to as the mining gear.

In contrast, altcoins employing *proof-of-stake* (PoS) reward miners based on the number of the coin units they already possess. Typically, a miner computes a very small number of hashes on the order of every second. For the purpose of this study, the computational resources required for PoS mining are effectively zero. For a typical PoS coin, the first sequence of blocks are mined using PoW to bootstrap the currency, and subsequent blocks are mined with PoS, or they alternate between PoS and PoW.

Lastly, a coin may also be mined using *auxiliary proof-of-work* (AuxPoW). Also known as merged-mining, this approach allows a coin to be generated alongside a parent coin (which are typically popular coins such as Bitcoin or Litecoin), while the miner computes hashes only for the parent coin. In this way, miners need not dedicate their mining power to a particular AuxPoW coin; they can mine the AuxPoW for free while generating the parent coin. A typical AuxPoW coin starts with only PoW blocks in the blockchain, and it eventually switches to AuxPoW blocks.

Hash functions: Mining a PoW or AuxPoW block requires searching for collisions in some prefix of a fixed hash function’s output. (The hash function is selected by the coin’s developer.) For Bitcoin, that hash function is

defined to be SHA-256 for all blocks; for Litecoin, it is Scrypt [22]. For most altcoins, each block is similarly mined using the same hash function, either SHA-256, Scrypt, or some other hash functions. Some altcoins, however, such as Auroracoin (AUR), allow a portion of their blocks to be mined with one hash function, and another portion of their blocks to be mined with a different hash function [1].

3 Data sets

In this section, we describe our blockchain and trade datasets. We explain how we collected the data, and we provide an overview of both dataset.

3.1 Blockchain

We obtained the blockchain data for 153 altcoins from the CryptoID website, which hosts and displays blockchain data in a way similar to blockchain.info [5]. Additionally, we also obtained the blockchains of Litecoin (LTC) and Bitcoin (BTC) by running the clients ourselves. By checking against altcoin mining pools such as ZPool and PoolSwitch, we were able to determine the hash algorithms and types of proofs for every block in 42 of the coins. We have also manually sampled blocks from CryptoID, ZPool and PoolSwitch and confirmed that they are consistent with one another. Out of these 42 coins, 18 coins, along with BTC and LTC, exclusively use SHA-256 or Scrypt as the hash function in parts of their blockchains. As we can compare these coins against BTC (based on SHA-256) and LTC (based on Scrypt), our study focuses on the 18 altcoins.

Table 1(a) summarizes the coins by hash functions (Column “HF”) and types of proofs (Column “PF”) as of November 8, 2016. For hash functions, “SH” is a shorthand for SHA-256, “Sc” for Scrypt, and “C” for Combined, where some blocks are mined with SHA-256, and some blocks are mined with Scrypt (e.g. AUR). For types of proofs, “W” is a shorthand for PoW, “S” for PoS, and “A” for AuxPoW. We note that any of these properties may change at any time. For example, BTC has consistently been a PoW coin that is mined with SHA-256. In contrast, for the first 827 days in the 852-day blockchain in our dataset, Auroracoin (AUR) is a proof-of-work coin that could be mined with Scrypt. Afterwards, blocks can be mined with SHA-256, Scrypt, or some other hash functions.

3.2 Trade

We obtained daily trade data of the 20 coins (18 altcoins plus LTC and BTC) from CryptoCoinCharts, which aggregates daily altcoin trades across 1,436 altcoins over

57 exchanges since 2010 [4], along with blockchain.info, which records daily BTC prices since 2011. We also scraped a well-known exchange, Cryptsy (now defunct), to verify altcoin prices on CryptoCoinCharts were the same as those on Cryptsy. The 57 exchanges include some of the most well-known altcoin exchanges like Cryptsy and Poloniex, and we believe that they have processed a representative sample of all altcoin trades, although we make no claims of completeness. While our volume data is therefore a lower bound on actual trade volume, we assume that any markets we do not capture offer roughly similar prices on a day-to-day basis. For each coin, our dataset includes the exchange(s) where the coin was listed, along with the daily trade volumes at each exchange. For each altcoin c , we compute the mean price on a given day t as $p(c,t) = v(BTC,t)/v(c,t) * p(BTC,t)$, where $v(BTC,t)$ is the trade volume of BTC against c at t (as altcoins are typically traded against Bitcoin), $v(c,t)$ is the trade volume of c against BTC at t , and $p(BTC,t)$ is Bitcoin’s daily USD price at t .

Table 1(a) presents trade-related statistics for the 18 altcoins we study, along with BTC and LTC. In particular, we show the lengths of trading activities in the “Trade” column, which counts the number of days from when a coin is first listed at an exchange to when it is last traded at some exchange. We contrast this value with the lengths of blockchains (“Blockchain” column) — the number of days from when the first block was mined to the last block in the data. In all cases, a coin is listed some time after the first block is mined.

The market capitalization, shown in the “Market Cap” column, reflects the market value of all units of an altcoin as of the last day of that coin in our dataset [11]. We calculate it as $(p \cdot q)$, where q is the total number of coin units ever mined up to the last day in our dataset (which differs across altcoins), and p is the USD price of the altcoin as of the last day the coin was listed in our dataset. In comparison, to compute the trade volume of an altcoin, we sum up the total number of bitcoins traded against the altcoin every day. Using daily BTC-USD exchange rates, we calculate the equivalent US dollar value for the bitcoins. Summing up these US dollars, we obtain the “Trade Vol” column. Typically, a high trade volume is strongly correlated with a large market capitalization, with the Pearson correlation coefficient of 0.9995. The total amount of trade for all 1,436 altcoins in the CryptoCoinCharts dataset is \$38.5 billion. Bitcoin accounts for the majority of the trade with a volume of \$30.8 billion. Of the remaining \$7.7 billion of trade, the 18 altcoins we study, plus LTC, account for \$2.0 billion of trade volume.

One feature of altcoin markets is the price volatility. To compare volatility across different altcoins, we examine the percentage difference between daily prices over time. We define volatility V_d as the median percentage

Table 2: Mining continuously for 7 and 30 days. All units are in percentages unless otherwise stated.

Coin	(a) 7 Days of Mining					(b) 30 Days of Mining				
	1st Day r	$E(r)$	$\sigma(r)$	P	$T_{r \geq 0}$ (Days)	1st Day r	$E(r)$	$\sigma(r)$	P	$T_{r \geq 0}$ (Days)
ARG	4.41	2.61	7.22	76.28	8	-0.51	0.71	1.50	80.71	0
AUR	10.17	0.63	2.25	61.35	19	1.37	0.13	0.35	62.31	18
BTA	6.67	2.01	5.18	69.33	14	0.99	0.45	1.07	68.97	13
CURE	6.23	-6.68	5.92	14.54	18	0.72	-1.54	1.01	8.79	8
DGC	3.24	1.16	2.79	61.66	111	0.77	0.27	0.57	64.89	207
DOGE	70.63	4.22	10.51	100.00	N/A	8.88	0.76	1.31	100.00	N/A
DOT	10.48	18.29	12.12	99.21	14	2.94	4.92	1.15	100.00	N/A
EFL	16.48	2.00	2.07	89.29	30	1.61	0.47	0.34	94.89	18
HAM	13.49	-2.19	6.43	18.82	46	3.10	-0.51	1.30	14.89	65
PPC	0.09	-1.13	1.37	16.50	1	0.02	-0.26	0.17	3.40	4
RPC	10.52	0.74	3.08	59.54	6	0.82	0.17	0.46	59.85	32
SWING	2.74	-2.46	3.77	19.09	3	-0.04	-0.53	0.55	21.83	0
TROLL	-0.01	4.37	1.31	98.08	0	0.87	0.99	0.06	100.00	N/A
UNO	8.44	-5.43	5.52	20.99	79	1.26	-1.34	1.23	15.98	72
VCN	56.98	-2.22	11.77	27.18	34	7.40	-0.75	1.87	30.23	25
VIA	-1.68	2.76	2.09	95.07	0	0.71	0.67	0.33	100.00	N/A
WBB	6.21	6.58	4.45	97.87	83	0.75	1.54	0.78	100.00	N/A
XJO	3.03	-5.51	4.00	4.48	15	0.18	-1.28	0.75	0.83	4

difference between prices on Day i and Day $(i + d - 1)$ for all possible i values within the trading period. A larger absolute value of V_d indicates a higher level of volatility. For instance, the Euro/USD exchange rate in the past year is associated with $V_1 = V_7 = 0.00\%$ [24]. In the same period, Google’s stock price (GOOG) has $V_1 = 0.08\%$ and $V_7 = 0.67\%$ [25]. For altcoins, we present the V_d values in Columns “ V_1 ” and “ V_7 ” in **Table 1(a)**. BTC and LTC, for example, have amongst the lowest absolute values for V_1 and V_7 .

4 Methodology

Our goal is to estimate the potential profitability of mining and speculating across different altcoins. Using the data we collected in Section 3, we discuss a methodology to estimate the cost of mining (Section 4.1) with the concept of opportunity costs. Using the cost estimate, we describe a method to estimate the profitability of mining, as well as speculating in Section 4.2.

4.1 Estimating Cost of Mining

For miners to decide whether to mine a particular coin through PoW, they first need to estimate the cost of mining. However, the precise value is difficult to calculate. The fixed cost of mining, such as investing in mining equipment, can vary across coins, depending which hash functions coins use. For instance, while coins based on SHA-256 are best mined with custom ASICs, hash functions such as X11 are designed to be ASIC-resistant and should be computed on general CPUs or GPUs [14]. We also need to account for variable costs, such as the cost of upgrading to faster mining equipment, as well as the

(significant) energy costs of operating the gear that differ widely across geographical regions.

An alternative to using direct cost is the *opportunity cost*. The opportunity cost of an activity is the revenue lost by engaging that activity *rather than the best alternative*. We can apply the idea to altcoin-mining. In particular, for two coins based on the same hash function, the costs of mining one versus the other are nearly equal, because the underlying unit of work, computing a hash, is the same. Thus, other than some initial software set-up cost, computing a million hashes of SHA-256 toward Bitcoin costs the same as computing a million hashes of SHA-256 toward XJO, another SHA-256-based currency.

Thus, the opportunity cost of mining XJO is the revenue a miner can expect from mining another SHA-256 currency *instead* of XJO over the same time period. To be a meaningful concept for the miner, this alternative revenue should be something a miner can reasonably expect to receive *a priori*. In other words, to say that a miner chose to mine A units of XJO rather than receive D US Dollars for mining another currency, the miner must be certain that he could get D US Dollars by choosing the alternative *before* choosing one or the other. In our comparisons, we use the least volatile alternative currencies with the highest trade volumes. For SHA-256-based coins, this is Bitcoin; for Scrypt-based coins, this is Litecoin. We call currency whose opportunity cost we are computing (e.g. XJO) the *target currency*, and Bitcoin or Litecoin the *base currency*.

Formally, we define the opportunity cost of mining a unit of a currency on a given day as follows. First, we determine the expected number of hashes, H , necessary to mine a unit of the target currency based on the difficulty

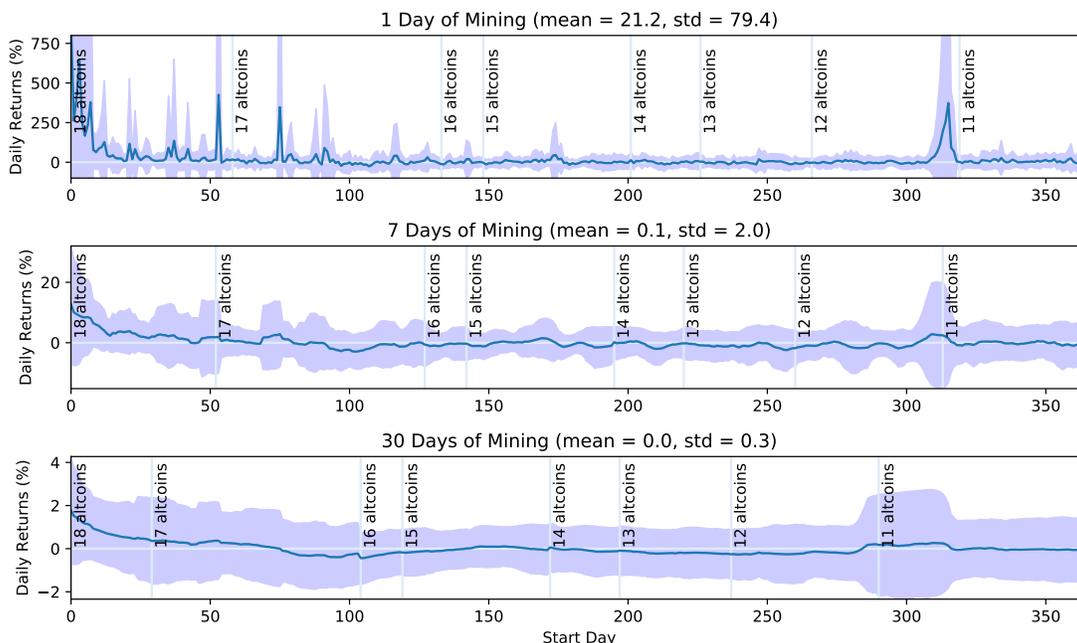


Figure 1: Mining a random altcoin.

of mining the currency that day. Next, we determine the expected number of units of Bitcoin (for SHA-256 currencies) or Litecoin (for Scrypt currency) that could be mined with H hashes. Finally, we convert this expected number of bitcoins or litecoins to US Dollars at the exchange rate that day. Thus, the opportunity cost of mining a unit of currency X is $\text{OppCost}_X = D_X \cdot R_B / D_B$, where D_X is the expected number of hashes required to mine a unit of X based on the day's difficulty, D_B is the expected number of hashes required to mine a unit of the base currency (Bitcoin or Litecoin) based on the day's difficulty, and R_B is the exchange rate of the base currency, in US Dollars per unit of the currency, on that day. We compute D_X and D_B based on the blockchain data of the target and base currencies, while we obtain R_B from our trade data.

An alternative way to understand opportunity cost is to consider an investor who would like to mine \$100 worth of currency X , say one based on SHA-256. He can approach a miner mining Bitcoin and offer him \$100, plus a nominal fee, to spend the same number of hashes it would take to mine \$100 worth of Bitcoin to mine X instead. The miner gets \$100 plus a fee, the investor gets \$100/ OppCost_X units of X . In an efficient market, the fee paid to the miner is negligible, so OppCost_X is the cost of acquiring a unit of X by mining. Even though such an efficient market for hashing power may not ex-

ist today, opportunity cost is a powerful *analysis tool* for comparing the relative profitability of mining a currency versus buying it outright.

4.2 Estimating Profitability

Our goal is to estimate the profitability of both miners and speculators for different coins. However, computing the profitability of individual actors is difficult. First, it is hard to identify the miners based on wallet addresses on the blockchain. Moreover, even if we did so, our trade data is anonymous and independent of the blockchain. More vexingly, we do not even have lot information: we do not know which coin units were sold when. The trade data only presents daily aggregates of how many coin units were bought and sold, and at what price. Instead, we focus on the potential profitability of a dollar invested, either through mining (i.e., \$1 of opportunity cost) or speculating (e.g., literally spending \$1 to buy into the market of a given altcoin).

We separate the act of altcoin investment into two logical roles: (a) the creation of the coin units (i.e. mining), and (b) continuing to hold the coin units and selling at the opportune moment (i.e. speculation). Specifically, we restrict mining to the act of dedicating computational resources to the coin's blockchain to obtain a reward, and then presume the miner sells immediately (if possi-

Table 3: A speculator that holds for 7 and 30 days. All units are in percentages unless otherwise specified.

Coin	(a) 7 Days of Speculating					(b) 30 Days of Speculating				
	1st Day r	$E(r)$	$\sigma(r)$	P	$T_{r<0}$ (Days)	1st Day r	$E(r)$	$\sigma(r)$	P	$T_{r<0}$ (Days)
ARG	10.72	-0.32	5.56	39.02	0	-0.48	-0.41	2.90	31.48	26
AUR	6.69	-0.22	5.87	41.38	0	3.39	-0.52	2.46	40.22	0
BTA	-26.24	0.35	6.87	47.31	11	-11.71	0.56	3.06	54.35	31
BTC	-1.97	0.44	2.51	55.17	6	-0.72	0.42	1.43	57.59	18
CURE	-3.74	-0.45	3.76	41.90	22	-6.30	-0.40	1.77	40.19	47
DGC	-5.33	0.00	4.44	41.26	10	-1.90	-0.06	2.21	37.04	10
DOGE	86.27	-0.06	4.53	39.06	0	5.24	-0.06	1.48	37.84	0
DOT	-17.47	-0.37	20.13	41.09	2	-18.67	-1.68	7.79	45.28	23
EFL	17.07	0.02	4.60	47.95	0	0.21	-0.18	2.06	47.83	0
HAM	-19.35	0.46	7.15	51.05	8	-1.12	0.35	2.35	64.24	1
LTC	-4.12	0.08	3.54	45.45	6	-0.34	0.05	1.63	39.55	6
PPC	-1.66	0.14	3.08	44.25	6	-0.16	0.11	1.60	46.99	3
RPC	-1.99	-1.13	4.89	39.08	1	-1.46	-1.25	2.24	28.50	88
SWING	1.61	-0.52	5.20	42.28	0	-1.04	-0.66	2.05	48.59	51
TROLL	-1.93	-0.69	4.38	51.92	1	-2.25	-0.16	1.22	51.72	13
UNO	3.34	-0.09	3.74	47.45	0	-2.08	-0.10	1.39	44.65	20
VCN	-43.26	0.80	7.99	58.97	5	-13.06	1.18	2.44	73.26	6
VIA	2.72	-0.21	3.53	44.74	0	1.49	-0.39	1.43	33.19	0
WBB	-20.82	0.17	5.77	41.97	22	-7.41	0.25	2.41	53.05	17
XJO	7.36	-0.82	3.37	35.68	0	-2.37	-0.86	1.58	29.90	25

ble). Hence, any gain or loss a miner experiences is due entirely to intra-day arbitrage between the valuation of the currency he chooses to mine and the alternative base currency he could have mined instead (i.e. opportunity cost). The mined coin units are logically transferred from miners to speculators, who would then hold the units and sell them at the opportune moment (e.g. when the price rises).

We estimate the profitability of miners through simulation. Using historical blockchain and trade data, we simulate the investment of \$1 worth of opportunity cost in mining across different altcoins, start dates, and durations. For speculators, we use a similar simulation, varying the time when an investor enters the market by buying \$1 worth of an altcoin’s units, as well as the holding positions of the investor. In this way, we can compute the profitability of mining/speculating depending on the participant’s strategy. This profitability analysis, albeit retrospective, assesses the risks and rewards for each coin and across multiple coins.

To simplify our analysis, we assume that any action that a miner or speculator takes is sufficiently insignificant so that it is unable to affect the market price. Specifically, a miner spends \$1 of opportunity cost on mining an altcoin and sells the mined units on the same day. Similarly, each speculator purchases a \$1 worth of a given altcoin and sells all units some time later. An altcoin market typically has a trade volume much higher than \$1, such that \$1 of mining or speculating is unlikely to change the price significantly.

5 Estimating Cost of Mining

In this section, we approximate the cost of PoW mining by computing the opportunity cost. We are aware that

some coins start as simple SHA-256/Scrypt PoW cryptocurrencies, but later they change the hash functions or types of proof. To this end, we consider the history of the currency up to the day of change, so that we can use BTC or LTC as the base currency for calculating opportunity costs.

Table 1(b) presents the opportunity cost of mining. For each currency, we show the the length of the blockchain — since the beginning of the chain — for which mining involves PoW blocks based on SHA-256 or Scrypt; this length is our analysis period (“Anlys” column). For coins like WBB and EFL, the values in the “Anlys” column are the same as the actual lengths of the respective blockchains, because the entire blockchain accepts PoW mining with SHA-256/Scrypt. In contrast, DOGE blocks are based on Scrypt in the first 276 days; afterwards, DOGE allows AuxPoW mining. As a result, we only consider the first 276 days of the 919-day blockchain for DOGE.

We include PoS coins in our opportunity cost analysis as long as they also allow PoW mining. For every PoS currency we consider in **Table 1(b)**, the majority of the altcoin units were mined with PoW. For instance, only 1.5% of PPC units were mined with PoS, and 0.3% of CURE units were with PoS. As the computational cost of mining PoS coins is negligible, we effectively consider these altcoins as PoW and compute their opportunity cost accordingly.

Focusing on the analysis period, we compute the total opportunity cost of mining (“Opp Cost” column). The total opportunity cost reflects the amount of sunk cost into a given altcoin by all the miners.⁴ For example, DOGE miners would have potentially made more than

⁴The opportunity costs of mining BTC and LTC are simply the revenue of selling mined BTC/LTC on the same day.

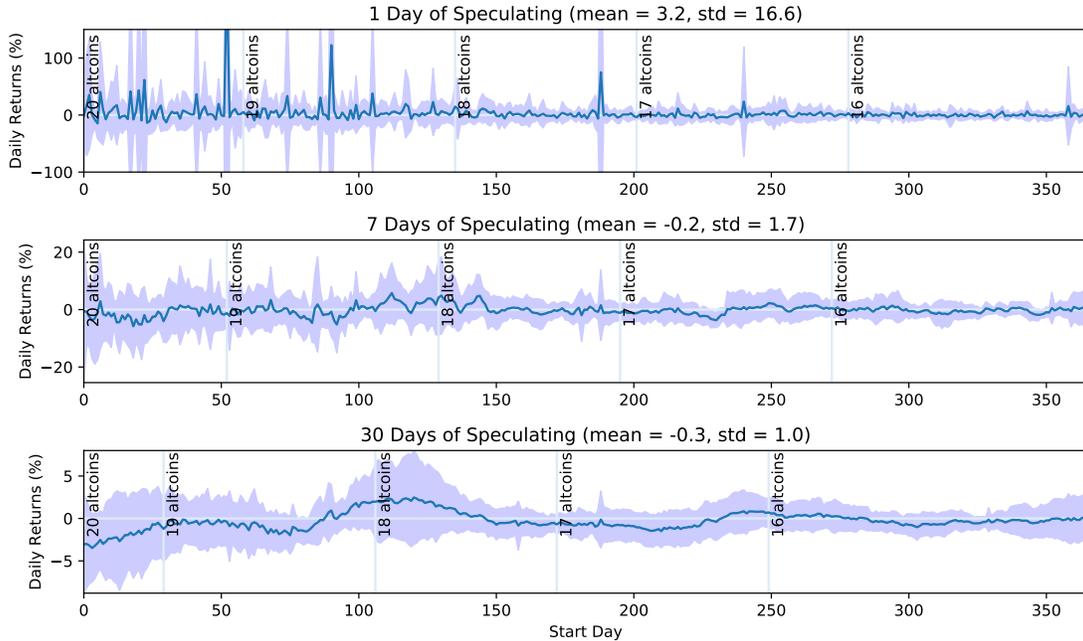


Figure 2: Speculating a random coin

\$43.8 million if they had been mining LTC instead. We contrast the total opportunity cost with the market capitalization, as shown in the “Market Cap” column in **Table 1(b)**. For a given altcoin, we compute the market capitalization as $(p \cdot q)$, where p is the price of the altcoin on the last day in the analysis period, and q is the total number of coin units mined during this period. The market capitalization reflects the value of all the coin units at the end of the analysis period. If this value is greater than the total opportunity cost, then all the coin units are worth more than what miners have collectively invested at the end of the analysis period. For example, miners invested \$4.1 million worth of opportunity cost in mining PPC, while PPC’s market capitalization is almost double this amount. Potentially, PPC could be overvalued and miners could make a profit.

We argue that opportunity cost is an effective estimate of the actual mining cost, as it is closely correlated to the market price. One possible explanation is that as more hype is created around a coin, the market price increases, which in turn attracts more miners. This increases the difficulty of mining and also the opportunity cost, so the opportunity cost goes up along with the price. Conversely, a coin that has attracted a significant amount of hashing capacity has a high difficulty and thus opportunity cost. Thus, miners expect to sell the mined coin units at a higher price. In general, the opportunity cost for mining a unit of a coin is correlated with the coin’s market price on the same day. The “Corr” column in **Ta-**

ble 1(b) shows the Pearson correlation coefficient for the market price and the unit opportunity cost of a given altcoin on the same day. Across the 18 altcoins in the table (except BTC and LTC), 12 altcoins are associated with a correlated coefficient > 0.80 . As a comparison, the Dow Jones Industrial Average and the S&P 500 Index between May 2012 and 2017 are correlated with a coefficient of 0.99 [16].

6 Estimating Profitability

In this section, we compute the profitability of miners and speculators. In our model, a miner always expends \$1 worth of opportunity cost in mining per day and sells the mined coin units on the same day. In contrast, a speculator purchases \$1 of an altcoin’s units on one day, and he sells all of the coin units later at a different price. Since our dataset does not offer details on exactly what individual miners or speculators did in the past, we can only use the historical data to simulate the behaviors of hypothetical miners and speculators. In this way, we compute the respective profitability under various conditions.

6.1 Miners

We start by considering the profitability of miners. Using the unit opportunity cost that we computed in the previous section, we construct a simulator in which a miner

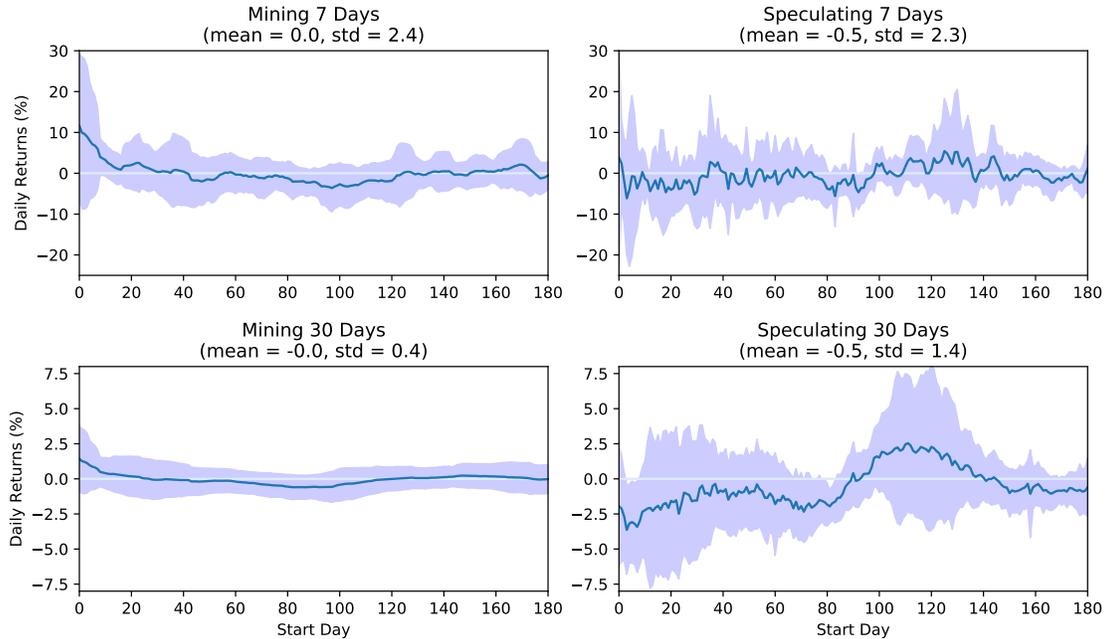


Figure 3: Comparing mining and speculating for the same set of 14 altcoins.

continuously mines over a duration of d days, starting on Day i . Every day, he invests \$1 worth of opportunity cost in mining. At the end of each day, he sells all the coin units mined on that day. At the end of the d days, his total revenue would be v dollars. We vary i between 1 to $N - d$, where N is the length of the trading period; for instance, $i = 1$ means that our simulated miner starts mining on the day when an altcoin is first listed on an exchange. We also vary d for $d = 1, 7, 30$ days. For all these i and d combinations, we compute the daily average returns, r , by solving for r in this equation: $d(1 + r)^d = v$.

Table 2(a) shows the result of our first simulation, in which a miner continuously mines for $d = 7$ days. Our analysis covers 18 altcoins during the same period as shown in **Table 1(b)**. The “1st Day r ” column shows the value of r if the miner starts mining a coin on Day 1 and continues until Day 7, selling any mined units on the same day. In contrast, the “ $E(r)$ ” column computes the expected/mean r if the miner starts mining on a random day. The standard deviation is shown in the “ $\sigma(r)$ ” column. The larger the standard deviation, the higher the risk of investment if a miner starts on a random day. For instance, mining DOT on a random day results in an expected daily return of $18.29 \pm 12.12\%$, which potentially presents a higher reward and risk than mining AUR, with an expected daily return of $0.63 \pm 2.25\%$. We stress that the r value is computed based on a simulated investment of \$1 worth of opportunity cost. In total, DOT has \$917 of trade volume and \$3,198 of total opportunity

cost (**Table 1(b)**). Even though its r value is high, the amount of actual profit extracted is likely to be limited. In contrast, a PPC miner can generate an expected return of $-1.13 \pm 1.37\%$. Given that PPC is associated with millions of dollars worth of trade volume and total opportunity cost, an actual miner has the potential to suffer significant losses.

Another way to measure risk is to compute the probability, P , of achieving a positive r if a miner starts mining an altcoin for 7 days on a random start day. As shown in the “ P ” column, miners of altcoins like DOGE and WBB will earn positive returns on any random start day. XJO miners, by contrast, will earn positive returns on a random start day with a probability of only 4.48%.

We observe that of the 18 altcoins in the table, the 1st Day r values are higher than the corresponding $E(r)$ values in 14 altcoins, which suggests that miners of these 14 coins obtain higher returns if they start mining as soon as an altcoin is listed on an exchange, relative to the average case (i.e. $E(r)$). One possible reason is that when an altcoin is first listed, the amount of mining capacity it has attracted is still on the rise, as there is friction for miners to reconfigure their equipment to mine a new-to-market altcoin; thus the opportunity cost of mining tends to be low in the beginning. Furthermore, the market price is often high when an altcoin is listed — a period typically associated with hype. The gap between high price and low opportunity cost creates a potential for miners to profit during this period.

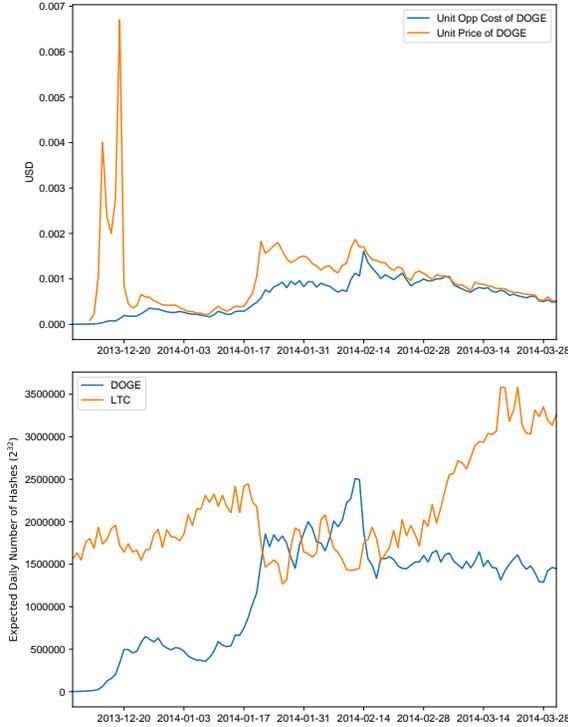


Figure 4: Top: The opportunity cost of mining a unit of DOGE, compared with the market price. Bottom: The expected number of hashes per day for DOGE and its base currency LTC. The x -axis starts on the same day as DOGE’s first block.

In fact, miners can potentially profit during the first few consecutive days after an altcoin is listed on exchanges. Column “ $T_{r \geq 0}$ ” shows the number of consecutive days since Day 1, such that a miner who starts mining on one of these days and continues for 7 days will not encounter a negative r . For example, an ARG miner who starts mining on any day of the first 8 days will receive $r \geq 0$; if she starts mining on Day 9, she would receive $r < 0$ for the first time (although she could still obtain positive returns subsequently). For TROLL and VIA, the 1st Day r is already negative, so $T_{r \geq 0} = 0$. For DOGE, all r values are positive regardless of the start day; thus $T_{r \geq 0} = N/A$. Finally, we repeat the simulation above, changing $d = 30$ days. We show the results in **Table 2(b)**. Again, for 14 out of the 18 altcoins, mining on Day 1 will yield a higher r than the expected case. The expected returns are lower for $d = 30$ than $d = 7$ in 11 of the altcoins, while all the $\sigma(r)$ values are smaller.

So far, we have examined the daily average returns for individual altcoins. This approach assumes that a miner already knows which altcoin to mine. Our next simulation departs from such an assumption, and it instead looks at a case where a miner randomly picks one of the

18 altcoins in **Table 2** and start mining on Day i . Again, $i = 1$ means the miner starts on the same day when a coin is listed on an exchange. The miner will stick to mining this altcoin for d days, devoting \$1 of opportunity cost of mining every day, and selling all mined units at the end of each day. Since the miner picks a coin at random for each i , our goal is to compute the distribution of daily returns for these 18 coins for given i and d values.

First, suppose we set $d = 1$ day. The top chart in **Figure 1** shows the result of our new simulation. The x -axis shows the start day of mining, i (relative to the first day of listing at an exchange for each coin). The y -axis shows the distribution of daily average returns, r . The solid line represents the expected (i.e. mean) r for picking a random coin and starting to mine on Day i for $d = 1$ day. The band above and below the solid line indicates the standard deviation of r . For i between 0 and 57, a miner can randomly mine one of the 18 altcoins on Day i . At its peak, the mean r is $885.1 \pm 2,781.0\%$ on Day $i = 0$. The expected r value decreases over time. Between $i = 58$ and $i = 132$, the miner can only pick one of 17 altcoins, as we do not have the trading data for one of the altcoins beyond 57 days. In general, the mean of the expected r values between $i = 0$ and $i = 365$ is $21.2 \pm 79.4\%$ (i.e. the expected returns for a miner who picks a random coin on a random start day).

The middle and bottom charts in **Figure 1** show the results for $d = 7$ and 30 days. As d increases, the expected r and its standard deviation both decrease. Effectively, with lower expected returns, the risks are potentially lower, too. Furthermore, as d increases, a persistent pattern is that the returns tend to be higher when i is low, regardless of the d values. This implies that a miner who picks a random altcoin and mines it shortly after the altcoin is listed is likely to receive higher returns than later.

6.2 Speculators

In contrast to miners who acquire altcoins through mining, speculators acquire altcoins by buying from exchanges. Also, while miners mine and sell on the same day, speculators buy coins on one day and sell them later. We design similar simulations to measure the potential profitability for speculators. Specifically, we require that a speculator enter the market on a random day i , buy \$1 worth of coins, hold them the next $d - 1$ days, sell all the coins at the end of the d -day period for a total of v dollars. Again, we compute the average daily return by solving for r in this equation: $1 \cdot (1 + r)^d = v$.

Table 3 shows the results of our simulation. Similar to **Table 2**, **Table 3** shows the returns on the 1st Day, as well as the expected r values for $d = 7$ and 30 days. However, the trend is the opposite. Across the 18 alt-

coins, plus LTC and BTC, a speculator who enters the market on Day 1 receives *lower* returns than the average case for 12 of the coins for $d = 7$; for 30 days, we observe the same trend across 16 of the coins. In contrast to the $T_{r \geq 0}$ metric in **Table 2**, we compute $T_{r < 0}$ here in **Table 3**, which counts the number of consecutive days since Day 1, such that if a speculator enters the market on one of these days for a given d value, she will receive $r < 0$. For instance, for $d = 7$ days, if a CURE speculator enters the market on any day between 1 and 22, she will receive negative returns; on Day 23, she will receive positive returns for the first time.

In addition to analyzing the returns for individual altcoins, we examine the case where a speculator picks an altcoin at random. **Figure 2** shows the result. Similar to **Figure 1**, **Figure 2** shows that as d increases from 1 Day, 7 Days, to 30 Days, the mean of the expected r values decreases, and so does the standard deviation; in other words, as the holding time increases, the potential returns and risks decrease. Contrary to **Figure 1**, **Figure 2** shows that a speculator who picks a random altcoin and enters the market soon after the altcoin is listed on exchanges is likely to receive *lower* returns than if she enters the market later.

To compare the returns between mining and speculating, we identify 14 out of the 18 coins, such that all of them can be involved in mining and speculation for $d = 7, 30$ days and $i = 0, \dots, 180$ days. Again, we assume that an investor randomly picks one of the 14 coins on Day i , enters the market either by mining or buying, and exits d days later. We show the results in **Figure 3**. For both $d = 7$ and 30 days, if an investor who picks a random altcoin decides to enter the market early, her expected returns will be higher if she becomes a miner. For $d = 30$, if an investor decides to become a speculator, she can potentially receive higher returns in the best case (e.g. more than 6% around $i = 120$) than the best-case returns in mining (i.e. less than 5% at $i = 0$). However, the risk of speculation is also higher, as indicated by the larger standard deviation for the expected r value.

7 Case Studies

Using WBB and DOGE as examples, we illustrate two observations that are common across other altcoins. First, as shown in Section 6, many coins are associated with high r values. While the price of coins tend to closely correlate with the opportunity cost, there are periods in which the price is higher than the opportunity cost, thus creating the potential for profitability. Second, coins like DOGE attract a significant amount hashing power, to a point where it is comparable to LTC. We show that using LTC as the base currency for estimating the opportunity cost is associated with decreased correlation be-

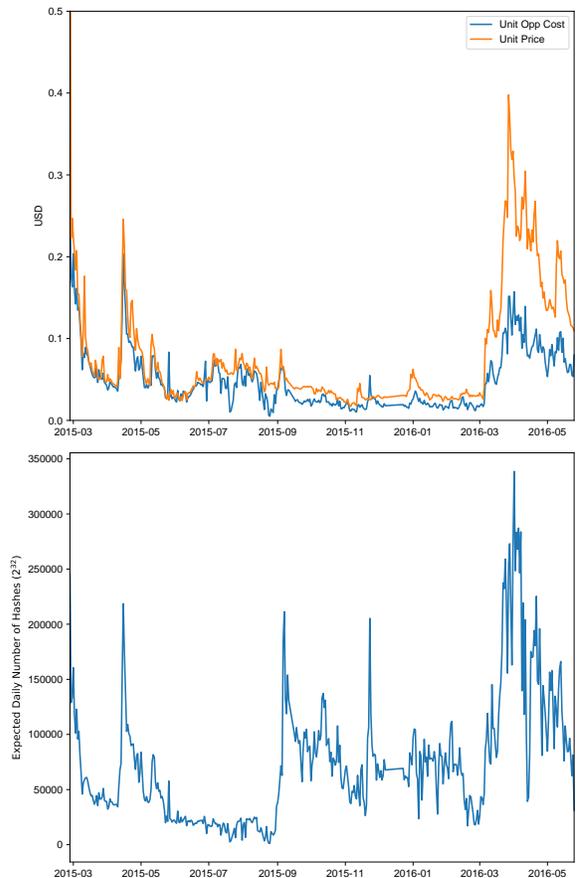


Figure 5: Top: The opportunity cost of mining a unit of WBB, compared with the market price. Bottom: The expected number of hashes per day. The x -axis starts on the same day as WBB’s first block.

tween the price and the opportunity cost. This suggests that the resultant opportunity cost may not be an accurate estimate for the actual mining cost.

To visualize these observations, we plot two graphs for each coin: (a) the unit price against the unit opportunity cost every day, which would allow us to visualize the correlation, or the lack thereof, between price and opportunity cost; and (b) the expected number of hashes across the network over the same time frame. Although we are unable to obtain the actual number of hashes, the expected number of hashes, derived from the difficulty value, can approximate the relative number of hashing power at different points in time for a given coin, or across multiple coins. It is an indicator of the relative supply of hashing power from miners.

7.1 WBB

Between March to May 2016, the unit price of WBB is consistently higher than the unit opportunity cost, as shown in the top chart of **Figure 5**. A miner during this period is likely to be profitable. For example, on April 26, 2016, the unit opportunity cost is 8.3 cents, while the unit price is 16.9 cents. In other words, expending the same energy on LTC-mining (i.e. the base currency) would produce only 8.3 cents of revenue; mining WBB would double the revenue.

Given the potential profitability, one would expect the market to react to close the gap between the price and the opportunity cost. Intuitively, a higher profitability would draw more miners, which would increase the difficulty of mining and thus the opportunity cost of mining. Miners would keep coming until the market reaches equilibrium, when the price is equal to the opportunity cost.

This intuitive behavior is partly observed in WBB, as shown in the bottom chart of **Figure 5**. As the gap between the price and opportunity cost widens beginning of March 2016, the expected daily number of hashes also begins to rise, which, in turn, causes the opportunity cost to rise. However, this rise in hash computations is not high enough to further drive up the opportunity cost and close the price-cost gap. We see this behavior not only in WBB, but also in coins like AUR and VIA, where the price is consistently higher than the opportunity cost during certain periods.

7.2 DOGE

DOGE is similar to WBB, in that around December 2013, the price is consistently higher than the opportunity cost, as shown in **Figure 4**. The expected number of daily hashes also rises, but not enough to drive up the opportunity cost and close the price-cost gap.

Furthermore, there is another observation unique to DOGE. Between Jan 17 and Feb 14, 2014, the price is higher than the opportunity cost, and the expected number of hashes of DOGE surpasses that of LTC, the base currency for DOGE. Throughout our analysis, we have always used LTC as the base currency for computing opportunity costs, as we assume LTC is associated with a higher trade volume and hash rate than a typical altcoin, and that miners are more likely to view LTC as an alternative when deciding to mine a Script-based coin.

This assumption is invalidated during the Jan-Feb period of DOGE. As the expected number of hashes toward DOGE-mining increases, the expected number of hashes toward LTC-mining decreases by similar amounts. It is likely that the hashing power originally dedicated to LTC-mining was shifted to DOGE. As a consequence, using LTC to compute the opportunity cost results in an

inaccurate estimate of cost and a decreased correlation between the price and opportunity cost.

8 Discussion

In this section, we discuss a number areas where our methodology might be limited. First, our choice of base currencies are BTC and LTC, as they have the highest trade volumes and their prices are relatively stable. However, they are by no means the only candidates for base currencies. DOGE and PPC, for instance, trail behind BTC and LTC in terms of total trade volumes, but they, too, enjoy relatively stable prices. Furthermore, as we show in **Figure 4**, there were more hashes toward mining DOGE than LTC for close to a month. Using LTC as the base currency is associated with a decreased correlation between the price and the opportunity cost. During this period, an alternative analysis for Script-based altcoins might use DOGE rather than LTC as the base currency. In general, using coins other than BTC and LTC as the base currencies may potentially produce different opportunity costs and present a different analysis of the market characteristics.

Second, the opportunity cost equates each hash of mining coin X to mining coin Y , provided X and Y both use the same hash function. In reality, such hash-to-hash comparison may be inaccurate. For example, a miner could be still computing hashes for an old block when a new block is announced. Depending on the network latency, it may take some time for the miner to receive the new block announcement. In the meantime, hashes are wasted on the old (“stale”) blocks.

Third, in characterizing various altcoin markets, we estimate the profitability of mining or speculating by assuming \$1 of investment. While this is a sufficiently small amount that would be unlikely to affect the market size, the resultant r values that we compute are not necessarily applicable to much larger investments. Even a few hundred dollars of investment may be enough to swing the price in small markets with only thousands of dollars of trade volume. Even in markets with much higher trade volumes, the actual amount of liquidity may be significantly less. A participant could be trading with herself to create an illusion of supply and demand in the market. Given the data we have, there is no way to rule out such “shadow” trades, thus making it difficult to estimate the true supply and demand.

Finally, in converting altcoin values into equivalent USD values, we first compute the value of altcoins in BTC and subsequently convert BTC into USD. While typically the markets with the highest trade volumes tend to be between altcoins and BTC, this is by no means the only way to sell altcoins. Many exchanges such as Cryptsy allowed speculators to trade altcoins with more

well-known altcoins such as PPC, LTC, and DOGE. Our analysis assumes that the market is perfectly efficient and that altcoin prices are the same regardless the trading pair. In relatively small markets with low trade volumes, however, this assumption may not be true.

9 Related Work

In this paper, we discuss the profitability of mining and speculating altcoins. There are existing online tools that help miners estimate the profitability of mining, such as Coinwarz and WhatToMine [3, 10]. Typically, these tools ask miners to enter the hash rate of their mining equipment, along with the cost of electricity in the respective geographic regions. While these tools are useful for miners who know their cost of operation, it is difficult to generalize across all miners. In addition, there are a few studies in the literature that empirically analyze the behaviors of miners. Most of the work focuses on theoretical behaviors of miners — in particular how they could potentially game the system and produce profits [18, 15]. Our measurement work complements these studies, as we combine simulation with real-world data in a way that allows us to extrapolate behaviors of actual miners under various conditions.

For speculation, we contrast our work against existing literature on penny stocks trading [12, 17]; in many ways, altcoin markets are similar to penny stock markets in terms of volatility and size. Many of the speculative behaviors in the penny stocks literature can also be observed in the world of altcoins. While much of the existing work focuses on the criminal aspects of penny stock scams, our work sidesteps the issue of intentionality — i.e. what speculators or miners do with the coins — and, instead, focuses on the economic incentives for each party, thus leaving the reader room to explore the implications of any gains and losses on certain altcoins.

10 Conclusion

In this work, we compare the profitability of mining versus speculation for 18 altcoins. By comparing against BTC and LTC, we use opportunity cost to estimate the miners' effort in the 18 coins, and we design simulations to estimate the daily returns per \$1 of investment, either through mining or speculating, under various conditions. These simulations show that a miner who starts mining shortly after an altcoin is listed can potentially earn higher returns than the average case, whereas a speculator who enters the market shortly after an altcoin is listed on exchanges might potentially earn lower returns. We also show that returns from mining a random altcoin tend to be lower with smaller standard deviations — less risky

— than from speculation.

References

- [1] Auroracoin. <http://auroracoin.is>.
- [2] Bittrex. <https://bittrex.com/>.
- [3] Coinwarz. <http://www.coinwarz.com>.
- [4] CryptoCoinCharts. <https://cryptocoincharts.info/>.
- [5] CryptoID. <http://chainz.cryptoid.info>.
- [6] Ethereum. <https://www.ethereum.org>.
- [7] Litecoin. <https://litecoin.org>.
- [8] Poloniex. <https://poloniex.com/>.
- [9] Ripple. <https://ripple.com>.
- [10] WhatToMine. <https://whattomine.com>.
- [11] Definition of Market Capitalisation. *Financial Times*, <http://lexicon.ft.com/Term?term=market-capitalisation>, 2017.
- [12] T. Bouraoui. Stock spams: An empirical study on penny stock market. *International Review of Business Research Papers*, 5:292–305, 2009.
- [13] M. J. Casey. PayPal to Allow More Merchants to Accept Bitcoin; Partnership Includes Bitcoin Payment Processors BitPay, Coinbase And GoCoin. *Wall Street Journal (Online)*, September 2014.
- [14] E. Duffield and D. Diaz. Dash: A Privacy-Centric Cryptocurrency. <https://github.com/dashpay/dash/wiki/Whitepaper>, 2015.
- [15] I. Eyal and E. G. Sirer. Majority is not enough: Bitcoin mining is vulnerable. In *International Conference on Financial Cryptography and Data Security*, pages 436–454. Springer, 2014.
- [16] Federal Reserve Bank of St. Louis. FRED Economic Data. <https://fred.stlouisfed.org/categories/32255?tg=gen>, 2017.
- [17] J. P. Fraedrich. Signs and signals of unethical behavior. In *Business Forum*, volume 17, page 13. California State University, Los Angeles, School of Business and Economics, 1992.
- [18] A. Kiayias, E. Koutsoupias, M. Kyropoulou, and Y. Tselekounis. Blockchain mining games. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 365–382. ACM, 2016.
- [19] E. Mack. National Bitcoin Alternative Auroracoin Launches To Save Iceland's Economy. *For*, <https://goo.gl/MPBzz6>, 2014.
- [20] S. Meiklejohn, M. Pomarole, G. Jordan, K. Levchenko, D. McCoy, G. Voelker, and S. Savage. A Fistful of Bitcoins: Characterizing Payments Among Men with No Names. In *Proceedings of the ACM Internet Measurement Conference*, 2013.
- [21] S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Consulted*, 1(2012):28, 2008.
- [22] C. Percival. Stronger key derivation via sequential memory-hard functions. *The BSD Conference (BSDCan)*, pages 1–16, 2009.
- [23] C. Percival and S. Josefsson. The scrypt password-based key derivation function. Technical report, 2016.
- [24] Quandl. Currency Exchange Rates - USD / EUR. <https://www.quandl.com/data/CUR/EUR>, 2017.
- [25] Yahoo. GOOG Historical Prices. <https://finance.yahoo.com/quote/GOOG/history?p=GOOG>, 2017.