http://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Vol. 67, No. 2, February 2021, pp. 661–672
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Evolution of Shares in a Proof-of-Stake Cryptocurrency

Ioanid Roşu,[a] Fahad Saleh[b]

[a] Finance Department, HEC Paris, 78351 Jouy-en-Josas, France; [b] School of Business, Wake Forest University, Winston-Salem, North Carolina 27109

Contact: rosu@hec.fr (IR); salehf@wfu.edu, https://orcid.org/0000-0003-1652-5189 (FS)

Received: April 26, 2020
Revised: June 14, 2020
Accepted: July 27, 2020
Published Online in Articles in Advance: November 6, 2020

https://doi.org/10.1287/mnsc.2020.3791

Abstract. Do the rich always get richer by investing in a cryptocurrency for which new coins are issued according to a proof-of-stake (PoS) protocol? We answer this question in the negative: Without trading, the investor shares in the cryptocurrency are martingales that converge to a well-defined limiting distribution and, hence, are stable in the long run. This result is robust to allowing trading when investors are risk neutral. Then, investors have no incentive to accumulate coins and gamble on the PoS protocol but weakly prefer not to trade.

History: Accepted by Kay Giesecke, finance.

Keywords: blockchain • cryptocurrency • asset allocation • martingale • Pólya urn • Dirichlet distribution

1. Introduction

In recent years, a large number of cryptocurrencies has emerged. A cryptocurrency is a type of electronic money for which the transaction log is based on a distributed ledger technology, such as blockchain.[1] A blockchain is a growing chain of records, called blocks, which are linked and secured using cryptography. Several protocols for achieving blockchain consensus exist, the most important being proof-of-work (PoW) and proof-of-stake (PoS). The PoW protocol requires agents to compete to update the blockchain by solving a computational puzzle so that success probabilities depend upon raw computational power. In the PoS protocol, the blockchain is updated by a randomly selected stakeholder, and the probability of an investor being drawn is equal to the investor's share, that is, the fraction of coins that the investor owns.[2]

The PoS protocol involves essentially no direct costs to the stakeholders. However, just as for the PoW protocol, the agent that updates the blockchain receives a coin reward. This reward feature of PoS has led critics across academia and the cryptocurrency press to argue that PoS induces wealth concentration. For example, Fanti et al. (2019, p. 43) argue that "PoS systems [lead] to a rich-get-richer effect, causing dramatic concentration of wealth." Similarly, one editorial in the cryptocurrency press argues that "the PoS model creates a centralizing effect where the rich will indefinitely get richer."[3] Thus, it natural to ask: What is the long-term evolution of investor shares in a cryptocurrency that uses a PoS protocol?

To answer this question, we consider a discrete-time, infinite-horizon model with several investors who can trade a risky asset called the cryptocurrency with units called coins. The PoS protocol requires that, before each trading time $t = 1,2,\ldots$, one investor is selected at random With probability given by the investor's share, that is, by the fraction of the total number of coins that the investor owns. Once selected, the investor receives new coins as a reward.

A key observation is that, when all investors are buy-and-hold, that is, when their trades are zero, the evolution of their shares is equivalent to a Pólya's urn problem (see Pemantle 2007 and the references therein). Indeed, consider an urn with balls of different colors, in which the number of colors corresponds to the number of investors. At each time t, a ball is extracted at random from the urn, which corresponds to an investor being selected at random by the PoS protocol with probability given by the investor's share. The ball is put back into the urn along with more balls of the same color, which corresponds to the selected investor receiving additional coins. Thus, the evolution of the fraction of balls of a given color in a Pólya's urn is the same as the evolution of investor shares in our context.

Our first result, which is standard in Pólya's urn problems, is that the share of an investor with a buy-and-hold strategy evolves according to a martingale. Intuitively, an investor with a large initial share (i.e., who is "rich") is more likely to receive the coin reward via the PoS protocol, but if the investor is not selected, the investor's share also decreases by a larger amount. As a result, the investor's share is not expected to increase or decrease, which is precisely the martingale condition. As an investor's share is bounded between zero and one, it possesses a well-defined limiting distribution with a mean equal to its initial value. This is the sense in which investor shares are stable in the long run.

Our second result is that, when all investors are buy-and-hold and the reward schedule is constant (normalized to one), the investor shares jointly approach a known distribution, called the Dirichlet distribution. In the case with only two investors, this reduces to a beta distribution. This case is sufficient to analyze the share of a particular investor because we can consider the aggregate holdings of the other investors as belonging to a single investor. Our analysis (see Section 3.3) shows that, if the coin rewards do not grow too fast, investor shares are stable in a stricter sense: They remain fairly close to the initial value. Moreover, we show that "poor" investors (i.e., those who start with a lower fraction of coins) end up with a more stable share distribution than rich investors.

Our third result is a trade irrelevance result that requires two additional assumptions: (i) Investors are risk neutral, and (ii) investors exit the model at an integrable random stopping time. Under these two assumptions, investors are indifferent between trading and a buy-and-hold strategy. Intuitively, when an investor buys more coins, there are two effects on the investor's utility. First, the purchase increases the probability that the investor receives a larger coin reward via the PoS protocol. At the same time, the additional coins lose in value because of the dilution effect. In equilibrium, the two effects exactly offset each other, and as a result, the investor is indifferent between trading and not trading. Moreover, with an infinitesimal trading cost, all investors would prefer not to trade, and thus, they would become "buy-and-hold" investors for which our first result applies.

Our paper contributes to the literature on the decentralization of blockchains, which mainly focuses on the PoW protocol and provides theoretical channels that drive PoW blockchains toward extreme centralization. Arnosti and Weinberg (2019) model PoW mining as a one-stage game in which miners simultaneously select hash rates. They find that asymmetries in hash-rate costs generate an extreme concentration of mining power. Alsabah and Capponi (2020) also establish an extreme concentration of mining power arising but in a model that incorporates R&D investment. They demonstrate that miners not investing sufficiently in R&D for mining equipment are driven out of the mining market. Neither of the aforementioned concentration channels arise in a PoS setting. The PoS analog of purchasing hash rate is to purchase PoS coins, but PoS coin prices do not vary across buyers within an efficient market, so asymmetric costs do not arise. Moreover, R&D investment is not relevant for the PoS setting.

Our paper also contributes to the literature on the economics of the PoS protocol, for example, Irresberger (2018), Fanti et al. (2019), and Saleh (2020). Irresberger (2018) provides an empirical analysis of coin concentration for three cryptocurrencies. The PoS protocols for these three cryptocurrencies vary in terms of their specific implementation, but Irresberger (2018) finds that coin concentration, measured by the Herfindahl index, does not deviate much from its starting value, barring sudden changes in network characteristics. Irresberger (2018) also provides a simulation analysis, indicating also that share concentration can be largely avoided in PoS blockchains. Saleh (2020) provides conditions under which a PoS protocol generates consensus among the investors and finds that a modest reward schedule helps to generate that consensus expediently. Similarly, Brown-Cohen et al. (2019) demonstrate security advantages from lower block rewards. Our results highlight further advantages for a modest reward schedule by focusing on stability in the wealth distribution.

Fanti et al. (2019) study the optimal reward function for a PoS cryptocurrency according to a "fairness" criterion; that is, they minimize the investor share variance over a given horizon, subject to a constraint regarding the number of coins distributed over that horizon. They show that a geometric reward is optimal in their context. Our analysis (see Appendix B) confirms their result, but because we are interested in the *limiting* evolution of investor shares, we show that, beyond the given horizon, the geometric reward produces a large and increasing variance of investor shares. This is not surprising as one should not expect exponentially increasing rewards to generate a stable share distribution over the long run. By contrast, our results show that the widely used constant reward function does generate stable share distributions in the limit.

Our results should not be interpreted as generally supportive of stake-based blockchain governance proposals. Tsoukalas and Falk (2020) examine stake-based voting for the purpose of crowdsourcing on blockchain. In such a setting, there exists a potential misalignment between the extent to which agents possess relevant information and the extent to which agents hold stake. Tsoukalas and Falk (2020) show that this misalignment may lead to suboptimal outcomes. Moreover, even when agents have the ability to endogenously acquire information, that ability does not lead to a first-best outcome because agents do not internalize the benefits to other users from their own information acquisition efforts. Our results apply only to the PoS protocol that is a special case of stake-based blockchain governance. Because PoS protocols specify publicly verifiable rules for updating the blockchain, asymmetric information has limited relevance for the evaluation of PoS protocols.

2. Environment

Time is discrete and infinite. There are two assets: (i) a risky cryptocurrency with units called coins

and (ii) a one-period-ahead risk-free asset. Trading in each asset takes place at each date $t \in \mathbb{N}_+ = \{1, 2, 3, \dots\}$. The time before trading begins is denoted by $t = 0$. The total supply of the cryptocurrency is distributed at $t = 0$ among a discrete population of investors, indexed by a set $\mathcal{I} = \{1, 2, \dots, I\}$ with $I \geq 2$.

2.1. Investor Shares

Denote by $n_{i,0} \in \mathbb{N}$ investor i's initial endowment in coins and by $n_{i,t}$ the number of coins owned by i after trading at t. The "investor share" is the fraction of coins that i owns at $t \in \mathbb{N}$:

$$\pi_{i,t} = \frac{n_{i,t}}{N_t}, \quad \text{with} \quad N_t = \sum_{i=1}^{I} n_{i,t}, \tag{1}$$

where N_t is the total (outstanding) number of coins at t.

2.2. Proof-of-Stake Protocol

Before trading at $t \in \mathbb{N}_+$, investor i is selected at random among the I investors with probability $\pi_{i,t-1}$. Once selected, an investor receives a deterministic reward of $R_t \geq 0$ coins (not necessarily an integer).[4] Denote by $S_{i,t}$ the event of i being selected at t, which is assumed independent of all other random variables. Define its indicator variable by $\mathbf{1}_{S_{i,t}}$, which is one if i is selected or zero otherwise. At each $t \in \mathbb{N}_+$, define the filtration that keeps track of the awards of the PoS protocol as the σ-algebra generated by the prices and indicator variables:

$$\mathcal{F}_t = \langle P_s, \mathbf{1}_{S_{i,s}} \rangle_{i \in \{1,\dots,I\},\, s \in \{1,\dots,t\}}. \tag{2}$$

2.3. Trading and Prices

The price of the cryptocurrency is an exogenous stochastic process $P_t > 0$, $t \in \mathbb{N}$. At each $t \in \mathbb{N}_+$, the order of events is as follows: (i) R_t coins are rewarded according to the PoS protocol, (ii) the price changes from P_{t-1} to P_t, and (iii) trading takes place at P_t.

A trading strategy of investor i is a process $\nu_i = (\nu_{i,t})_{t \in \mathbb{N}}$ adapted to the filtration $\mathcal{F}_t$, such that, for all $t \in \mathbb{N}_+$, the number of coins after trading, $n_{i,t}$, belongs to $[0, N_t]$, where

$$n_{i,t} = n'_{i,t} + \nu_{i,t}, \quad \text{with} \quad n'_{i,t} = n_{i,t-1} + R_t \mathbf{1}_{S_{i,t}}. \tag{3}$$

Define the total market capitalization of the cryptocurrency as

$$M_t = N_t P_t, \quad t \in \mathbb{N}. \tag{4}$$

3. Zero Trading

In this section, we analyze the case in which at least one investor never trades any coins after $t = 0$; that is, the investor's trading strategy satisfies $\nu_{i,t} = 0$ for all $t \in \mathbb{N}_+$. We call such investors with zero trades buy-and-hold. As the trades of all investors sum to zero, that is, $\sum_{j=1}^{I} \nu_{j,t} = 0$, the same is true about the trades of the non–buy-and-hold investors. Thus, if we aggregate the coins of all the non–buy-and-hold investors, they behave collectively as one buy-and-hold investor. In that case, we show that the investor shares are martingales that converge jointly to a limiting distribution, which we compute in closed form.

3.1. One Investor

Consider a buy-and-hold investor $i \in \mathcal{I} = \{1, 2, \dots, I\}$, who starts with an endowment of $n_{i,0}$ coins. Then, at each time $t \in \mathbb{N}_+$, the number of coins owned by investor i changes only if the investor is selected by the PoS protocol. Equation (3) implies that

$$n_{i,t} = n_{i,t-1} + R_t \mathbf{1}_{S_{i,t}}. \tag{5}$$

This setup is equivalent to a Pólya's urn problem:[5] Consider an urn with balls of I different colors and let $n_{t,i}$ be the number of balls of color i at t. At t, a ball is extracted at random from the urn (with probability $\pi_{i,t-1} = n_{i,t-1} / \sum_{j=1}^{I} n_{j,t-1}$), and it is put back into the urn along with R_t ball of the same color. Note that $R_t = 1$ in the original Pólya's urn problem, but the problem has since been adapted to include more general numbers (see Pemantle 2007). Thus, the number of balls of color i evolves as in Equation (5). A standard result in Pólya's urn problems is that the fraction of balls of color i follows a martingale. We prove this result in the context of our proof-of-stake model.

Proposition 1. *Suppose investor $i \in \mathcal{I}$ never trades any coins. Then the investor's share $\pi_{i,t}$ follows a martingale. Moreover, this martingale process has a well-defined limiting distribution, $\pi_{i,\infty}$, whose mean, $\mathsf{E}(\pi_{i,\infty})$, is equal to the initial share, $\pi_{i,0}$.*

To get intuition for this result, let $R_t = 1$ for all t.[6] Suppose there are, in total, 10 investors, each holding, initially, one coin. Thus, at $t = 0$, the number of coins outstanding is $N_0 = 10$, and the investor shares are all 1/10. At $t = 1$, one investor randomly receives the coin, and the number of coins outstanding increases by one; hence, $N_1 = 11$. Then, with probability 0.1, investor i's share increases by $9/(10 \times 11)$ (from 1/10 to 2/11), and with probability 0.9, investor i's share decreases by $1/(10 \times 11)$ (from 1/10 to 1/11), an amount that is nine times smaller.[7] Thus, the change in investor i's share has zero conditional mean, which is the martingale condition.

3.2. Multiple Investors

We now assume that all investors are buy-and-hold. Proposition 1 then implies that all investor shares are martingales. As all shares lie between zero and one, they are bounded martingales; hence, according to a classical

theorem by Doob (see Pemantle 2007), the investor shares converge in probability to a well-defined distribution on [0,1]. The next result, which is standard in Pólya's urn problems, identifies the limiting distribution as the Dirichlet distribution under the hypothesis of a constant reward schedule.

Let $\Gamma(z) = \int_0^\infty x^{z-1}e^{-x}\mathrm{d}x$ be the Gamma function, which, for positive integers, is the same as the factorial (i.e., $\Gamma(n) = (n-1)!$). Recall that the Dirichlet distribution with parameters $(a_1, \ldots, a_I)$ has support on the set $\{(x_1, x_2, \ldots, x_I) \in \mathbb{R}_+^I \mid \sum_{i=1}^I x_i = 1\}$ and has density function

$$f(x_1, \ldots, x_I) = C \prod_{i=1}^{I} x_i^{a_i - 1}, \quad \text{with} \quad C = \frac{\Gamma\left(\sum_{i=1}^I a_i\right)}{\prod_{i=1}^I \Gamma(a_i)}, \tag{6}$$

and when $I = 2$, the Dirichlet density reduces to the beta density on [0,1] with parameters (a_1, a_2):

$$f(x) = C\, x^{a_1 - 1}(1-x)^{a_2 - 1}, \quad \text{with} \quad C = \frac{\Gamma(a_1 + a_2)}{\Gamma(a_1)\Gamma(a_2)}. \tag{7}$$

When $a_1 = a_2 = 1$, the beta distribution on [0,1] is the uniform distribution.

Proposition 2. *Suppose there are no coin transactions among the I investors and the coin reward is $R_t = 1$. Then, the investor shares $\pi_{i,t}$ converge in distribution to a Dirichlet distribution with parameters $(n_{1,0}, \ldots, n_{I,0})$.*

The intuition of Proposition 2 is based on the martingale result of Proposition 1. One may think that investor shares are explosive. For example, if investor i is selected at t, i's share increases, and therefore, in the next period, investor i is more likely to be selected, and this can lead, via a "snowballing" effect, to larger and larger shares such that investor i's share converges in probability to one. This argument is wrong: A "richer" investor (i.e., an investor with a larger share at t) is indeed more likely to be selected than a "poorer" investor, but if the richer investor ends up not being selected, the investor's share would drop by more than the corresponding share decrease of a poorer investor. For a numeric example with two investors with different initial shares, see the discussion after Corollary 1. Formally, i's coin share at t changes by[8]

$$\pi_{i,t} - \pi_{i,t-1} = R_t \, \frac{\mathbf{1}_{S_{i,t}} - \pi_{i,t-1}}{N_t}. \tag{8}$$

Thus, i's share increases by $1 - \pi_{i,t-1}/N_t$ with probability $\pi_{i,t-1}$ and decreases by $\pi_{i,t-1}/N_t$ with probability $1 - \pi_{i,t-1}$. Moreover, as we show both numerically and analytically in Section 3.3, if the initial number of coins is large relative to the coin reward, the limiting Dirichlet distribution is concentrated around the initial shares.

3.3. Limiting Distribution and Stability

In this section, we examine the limiting distribution of investor shares from the perspective of a buy-and-hold investor. Thus far, we have considered an investor share to be stable if it does not change on average. Proposition 1 shows that the investor shares are martingales and, thus, are stable. Another way of defining stability is suggested by Proposition 2, which describes the limiting distribution of investor shares. An investor share is then stable in the stricter sense if the limiting distribution is tight around the initial share. We show that, when the number of coins outstanding is large relative to the coin reward, investors' shares are stable in the stricter sense as long as the coin reward does not increase too fast.[9]

3.3.1. Constant Reward Schedules. We consider first the case of a constant coin reward schedule with value normalized to one (i.e., $R_t = 1$). Consider a buy-and-hold investor, called investor 1. Denote by $N = N_0$ the initial number of coins outstanding. Let π_1 denote investor 1's initial share and $n_1 = \pi_1 N$ investor 1's initial number of coins. As the aggregate trade is zero, we can aggregate the other investors' holdings and obtain another buy-and-hold investor, called investor 2. Let π_2 denote investor 2's initial share and n_2 investor 2's initial number of coins. Then, $\pi_2 = 1 - \pi_1$ and $n_2 = \pi_2 N$, which implies that the initial coin holdings and shares of both investors are completely determined by π_1 (the initial share of investor 1) and N (the initial number of coins outstanding).

Proposition 2 implies that the share of investor 1, $\pi_{1,t}$, converges to a well-defined limiting density, $\pi_{1,\infty}$, which is the beta distribution with parameters (n_1, n_2). An implication of that result is that a larger initial number of coins N leads to a tighter distribution $\pi_{1,\infty}$. Figure 1 illustrates this result by showing the finite-sample density of investor 1's share for a variety of initial coin numbers. Each simulation involves 100,000 steps and 10,000 sample paths, and we assume a constant reward schedule with $R_t = 1$ for all t. In all cases, the initial share is 0.5. Figure 1 shows that the density of 1's share depends on the initial coin number: A larger initial coin number induces a tighter distribution around the initial share. As we move from the top left graph (with $N = 2$) to the bottom right graph (with $N = 2{,}000$), the coin share distribution tightens, but its center remains equal to the initial coin share.

Corollary 1 helps us formalize this result.

Corollary 1. *Consider a buy-and hold investor with initial share π_1, and let N be the initial number of coins outstanding. The variance of the limiting distribution of investor 1's share is*

$$\mathrm{Var}(\pi_{1,\infty}) = \frac{\pi_1(1-\pi_1)}{N+1}, \tag{9}$$

which is increasing in π_1 if $\pi_1 < 0.5$ and is decreasing in N.

Thus, as illustrated by Figure 1, when N increases, the variance of the limiting distribution decreases. Moreover, in the limit when N approaches infinity, the variance of the limiting distribution converges to zero; that is, it becomes very tight around its mean. This implies that investor 1's share is stable in the stricter sense.

Corollary 1 is also useful in comparing investors with different initial shares. For example, consider a 1% investor with $\pi_1 = 0.01$ and a 10% investor with $\pi_1 = 0.1$. Then, the variance ratio of the limiting distribution of the two investors is

$$\frac{\mathrm{Var}\left(\pi_{1,\infty}^{1\%}\right)}{\mathrm{Var}\left(\pi_{1,\infty}^{10\%}\right)} = 0.11. \tag{10}$$

Thus, the investor shares are less stable for richer investors in the sense that their limiting distribution is less tight around its mean.[10] The intuition for why a poorer investor has a more stable share follows from the discussion after Proposition 1. Suppose initially there are $N = 100$ coins, and there are two investors: a 10% investor (with 10 coins) and a 1% investor (with one coin). Then, for the 10% investor, with probability 0.1, the investor's share increases from 10/100 to 11/101 (by 90/10,100), and with probability 0.9, the investor's share decreases from 10/100 to 10/101 (by 10/10,100). For a 1% investor, with probability 0.01, the investor's share increases from 1/100 to 2/101 (by 99/10,100), and with probability 0.99, the investor's share decreases from 1/100 to 1/101 (by 1/10,100). Although it is true that the share of the 1% investor almost doubles in rare cases (with 1% probability), most of the time (with 99% probability), the investor's share remains very close to the initial value. Thus, the share of the poorer investor is more stable in the long run.

3.3.2. General Reward Schedules. We now consider general reward schedules: nonincreasing, in Proposition 3, and increasing, in Proposition 4. We have already established that, regardless of the type of reward schedule, investor shares are martingales and, hence, are stable. In this section, we examine whether investor shares are stable in a stricter sense, that is, whether eventual deviations from the mean are unlikely when the initial number of coins outstanding is very large. Thus, investor 1's share is stable in the stricter sense if the limiting distribution $\pi_{1,\infty}$ satisfies $\lim_{N\to\infty}\mathrm{P}(|\pi_{1,\infty} - \pi_1| \geq \varepsilon) = 0$ for any $\varepsilon > 0$.[11]

Proposition 3. *Suppose there are no coin transactions among the I investors, and the coin reward is nonincreasing, that is, $R_{t+1} \leq R_t$ for all t. Then, investor 1's limiting share distribution satisfies $\lim_{N\to\infty}\mathrm{P}(|\pi_{1,\infty} - \pi_1| \geq \varepsilon) = 0$ for any $\varepsilon > 0$.*

Figure 1. (Color online) Simulated Limiting Distribution of the Investor Share

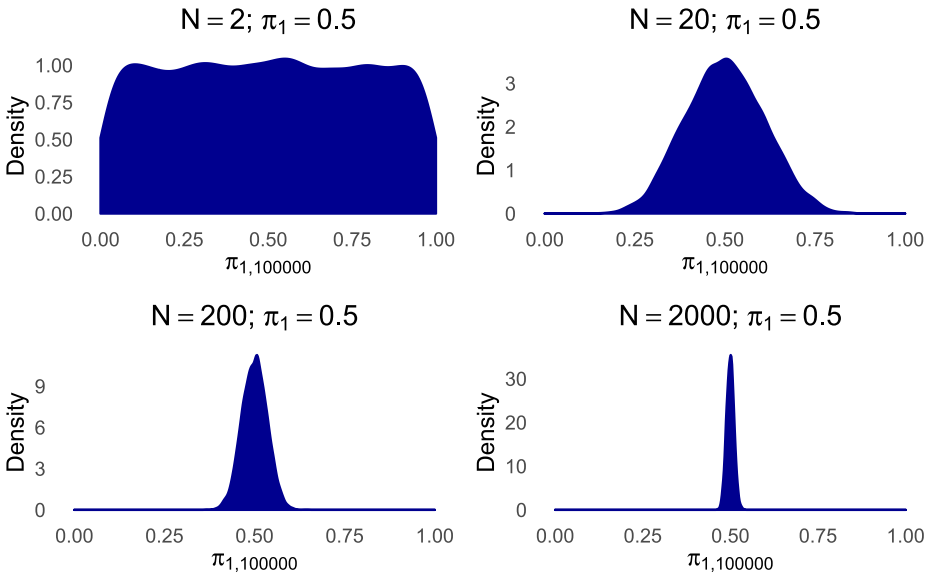

Notes. Let N be the total initial number of coins and π_1 investor 1's initial share. The graphs depict the density of investor 1's share after 100,000 steps. Each graph is generated from 10,000 sample paths and assumes a constant coin reward of one coin; that is, $R_t = 1$ for all $t \in \mathbb{N}_+$. The figure is generated using the R Statistical Software with a random seed of 100.

Thus, the investor shares are stable in the stricter sense if the reward schedule is nonincreasing, for example, if it is constant (already discussed in Section 3.3.1) or decreasing.

Proposition 3 also implies that the limiting distribution of investors' shares depends not only on the initial shares but also on the initial number of coins.[12] To illustrate this point, we consider a streak of five straight rewards for investor 1 under the following two cases: (i) $n_1 = n_2 = 1$ and (ii) $n_1 = n_2 = 1{,}000$. In each case, the initial shares are the same: $\pi_1 = \pi_2 = 0.5$. For exposition, we let $R_t = 1$ for all t. In case (i), the streak occurs with probability $(1/2)\times(2/3)\times(3/4)\times(4/5)\times(5/6) = 1/6$. In case (ii), the same streak occurs with probability $(1/2)\times(1{,}001/2{,}001)\times(1{,}002/2{,}002)\times(1{,}003/2{,}003)\times(1{,}004/2{,}004) \approx 1/32 < 1/6$. We first note that the streak is roughly five times more likely in case (i) than in case (ii). Moreover, after the streak, investor 1 possesses more than 85% of the coins in case (i) but still approximately 50% of the coins in the second case. In general, when the initial coin number is low, as in case (i), streaks occur with higher probability and have a more dramatic impact on the shares. Thus, the limiting distribution is more spread out when the initial coin number is low.

In Proposition 4, we examine the stability of investor shares if the reward schedule is an increasing function of the number of coins outstanding.

Proposition 4. *Suppose there are no coin transactions among the I investors, and the coin reward is increasing with the number of coins outstanding such that $R_t = \rho N_{t-1}^{\gamma}$, where $\rho, \gamma > 0$ are two constants. Then, investor 1's limiting share distribution satisfies $\lim_{N\to\infty} \mathsf{P}(|\pi_{1,\infty} - \pi_1| \geq \varepsilon) = 0$ for any $\varepsilon > 0$ if and only if $\gamma < 1$.*

Thus, investor shares are stable in the stricter sense even when the reward schedule is increasing in the number of coins outstanding as long as this increase is not too fast (i.e., as long as $\gamma < 1$). If, however, the increase is sufficiently fast (i.e., $\gamma \geq 1$), then the reward schedule exhibits exponential growth, and the probability of large deviations from the initial share does not vanish as the initial number of coins outstanding approaches infinity.

Overall, Propositions 3 and 4 clarify the extent to which arbitrary PoS implementations generate stability of investor shares.

4. A Trading Irrelevance Result

In this section, we strengthen our results by showing that, in equilibrium, investors should be indifferent about how much they trade or whether they trade at all. This trading irrelevance result is not obvious ex ante. For example, one may argue that, in this environment, an investor has an incentive to amass the cryptocurrency in order to increase the investor's probability of getting even more coins from the PoS protocol. We show, however, that this intuition is incorrect and that doing nothing is weakly preferred to hoarding.

4.1. Investor Preferences

To analyze the equilibrium behavior of investors, we need to supplement the assumptions in Section 2 with a description of investor preferences. We, thus, assume that all investors are risk neutral and impatient; that is, they discount each period by multiplying their expected payoff with a constant parameter $\delta \in (0,1]$, called the investor impatience. As in Biais et al. (2019), we assume that investor i incurs a liquidity shock at a random time $\tau_i > 0$ when the investor must sell all assets, consume the proceeds, and exit the model. The exit time τ_i is independent of all other variables and satisfies $\mathsf{E}(\tau_i) < \infty$. With an abuse of notation, we augment the filtration to include the liquidity shock time; that is, we redefine $\mathcal{F}_t$ to be $\langle \mathcal{F}_t, \tau_1, \ldots, \tau_I \rangle$ such that the investor can condition the investor's strategy on the time of the liquidity shock.

Formally, investor i's strategy is a triple $(v_{i,t}, b_{i,t}, c_{i,t})_{t\in\mathbb{N}_+}$ of processes adapted to the filtration $\mathcal{F}_t$, where $v_{i,t}$ is the number of coins traded at t, $b_{i,t}$ is the end-of-period-t holding of the risk-free asset, and $c_{i,t}$ is the consumption at t. If $r_{t-1,t}$ denotes the one-period-ahead risk-free rate from $t-1$ to t, then investor i's strategy solves the maximization problem

$$
\begin{aligned}
U_i = \max_{v_{i,t}, b_{i,t}, c_{i,t}} \mathsf{E}\left(\sum_{t=1}^{\tau_i} \delta^t c_{i,t}\right) \quad \text{such that:}\\
c_{i,t} + b_{i,t} + v_{i,t} P_t = (1 + r_{t-1,t}) b_{i,t-1},\\
0 \leq n'_{i,t} + v_{i,t} \leq N_t,\\
c_{i,\tau_i} = n'_{i,\tau_i} P_{\tau_i} + (1 + r_{\tau_i - 1,\tau_i}) b_{i,\tau_i - 1},\\
b_{i,\tau} = n_{i,\tau} = 0 \quad \text{for } \tau \geq \tau_i,
\end{aligned} \tag{11}
$$

where $n'_{i,t}$ is the number of coins owned by i after the random coin reward but before trading at t, and $n_{i,t} = n'_{i,t} + v_{i,t}$ is the number of coins owned by i after trading at t. The first constraint, $c_{i,t} + b_{i,t} + v_{i,t} P_t = (1 + r_{t-1,t}) b_{i,t-1}$, is the standard budget constraint. The second constraint, $0 \leq n_{i,t} \leq N_t$, is that investor shares are bounded by zero and one. The third constraint, $c_{i,\tau_i} = n'_{i,\tau_i} P_{\tau_i} + (1 + r_{\tau_i - 1,\tau_i}) b_{i,\tau_i - 1}$, is that i liquidates i's holdings at τ_i. The fourth constraint is that i stops trading after the exit time τ_i.

Proposition 5 describes the equilibrium utility of an investor. We assume the following necessary equilibrium conditions for a risk-neutral economy:

$$
r_{t,t+1} = \frac{1}{\delta} - 1, \tag{12}
$$

$$
\mathsf{E}_t(M_{t+1}) = (1 + r_{t,t+1}) M_t, \tag{13}
$$

where Equation (12) determines the endogenous risk-free rate and Equation (13) arises as an intratemporal condition across the risk-free asset and the cryptocurrency.

Proposition 5 (Trading Irrelevance). *If the conditions* (12) *and* (13) *are satisfied, any trading strategy* $\nu_i = (\nu_{i,t})_{t\in\mathbb{N}_+}$ *provides the same expected utility for investor* i *at* $t = 0$*:*

$$U_i = n_{i,0}P_0. \tag{14}$$

Proposition 5 establishes the main result of this section: that, under certain conditions, investors are indifferent to how much they trade. Intuitively, when an investor buys more coins at t, there are two effects on the investor's utility: First, the purchase increases the probability that the investor receives a larger coin reward via the PoS protocol. At the same time, the additional coins lose in value because of the dilution effect. In equilibrium, the two effects exactly offset each other, and as a result, the investor is indifferent between trading and not trading. Furthermore, if investors face an infinitesimal trading cost, they prefer not to trade and, thus, become identical to the buy-and-hold investors of Section 3.[13]

5. Conclusion

We analyze the evolution of investor shares in a model of a cryptocurrency for which new coin issuance follows a PoS protocol. This problem closely parallels the evolution of color shares in a Pólya's urn. As in that literature, the shares of coins owned by buy-and-hold investors are bounded martingales and, therefore, have a limiting distribution. Thus, investor shares are stable in the long run. With a constant reward normalized to one coin, the limiting share distribution for buy-and-hold investors can be computed in closed form: It is a Dirichlet distribution and, in the case of two investors, a beta distribution. Further, we show that, when coin rewards are not increasing too fast, the investor shares are stable in a stricter sense: They remain fairly close to the initial value. Moreover, poor investors (i.e., those who start with a lower fraction of coins) end up with a more stable share distribution than rich investors.

By analyzing the optimal strategies of investors who are not necessarily buy-and-hold, we obtain a trading irrelevance result: Investors are indifferent between trading and being buy-and-hold. Thus, our results regarding the evolution of shares for buy-and-hold investors are robust to the case when trading is allowed.

Our results are counter to the intuition of some in the cryptocurrency press that say investors have an incentive to amass coins in order to increase the probability of getting even more coins under the PoS protocol. In our framework, we show that this intuition is incorrect and that, under plausible assumptions, the PoS protocol does not lead to wealth accumulation and the rich getting richer but rather to stable investor shares.

Acknowledgments

The authors thank Kay Giesecke (the editor), an anonymous associate editor, three anonymous referees, Bruno Biais, Lin William Cong, Nicolae Garleanu, Franz Hinzen, Felix Irresberger, Kose John, Evgeny Lyandres, Katya Malinova, and Gerry Tsoukalas for valuable comments.

Appendix A. Proofs of Results

Proof of Proposition 1. If $t \in \mathbb{N}_+$, the number of coins owned by i at $t-1$ is $n_{i,t-1} = N_{t-1}\pi_{i,t-1}$. At t, investor i receives R_t coins if selected, that is, if $\mathbf{1}_{S_{i,t}} = 1$. Therefore, $n_{i,t} = n_{i,t-1} + R_t\mathbf{1}_{S_{i,t}}$, and i's share evolves according to

$$\pi_{i,t} = \frac{N_{t-1}\pi_{i,t-1} + R_t\mathbf{1}_{S_{i,t}}}{N_t}. \tag{A.1}$$

The total number of coins satisfies $N_t = N_{t-1} + R_t$; therefore, the investor share satisfies.

$$\pi_{i,t} - \pi_{i,t-1} = R_t\,\frac{\mathbf{1}_{S_{i,t}} - \pi_{i,t-1}}{N_t}. \tag{A.2}$$

The event $S_{i,t}$ of i being selected at t has probability $\pi_{i,t-1}$ and is independent from everything else.[14] The expected change in investor share based on the information available at $t-1$ is

$$\mathsf{E}_{t-1}(\pi_{i,t} - \pi_{i,t-1}) = \mathsf{E}_{t-1}\left(\frac{R_t}{N_t}\right)\mathsf{E}_{t-1}(\mathbf{1}_{S_{i,t}} - \pi_{i,t-1}) = 0. \tag{A.3}$$

This implies that $\pi_{i,t}$ is a martingale process. As it is also a bounded process, the martingale convergence theorem implies that $\pi_{i,t}$ has a well-defined limit, which we denote by $\pi_{i,\infty}$. Moreover, by the bounded convergence theorem, $\mathsf{E}(\pi_{i,\infty}) = \mathsf{E}(\lim_{t\to\infty}\pi_{i,t}) = \lim_{t\to\infty}\mathsf{E}(\pi_{i,t}) = \pi_{i,0}$. □

Proof of Proposition 2. The proof is standard: See Pemantle (2007) and the references therein. Denote by $a_i = n_{i,0}$ the initial number of coins owned by $i \in \mathcal{I} = \{1,\dots,I\}$ and by $m_{i,T}$ the (random) number of coins received by i after T periods. Clearly, $n_{i,T} = a_i + m_{i,T}$. We need to compute the joint probability that $m_{i,T}$ equals some integer m_i. As one coin is gained in each period by one of the investors, we have $\sum_{i=1}^{I} m_{i,T} = T$. Denote the time indices when i receives one coin by $t_{i,1} < t_{i,2} < \dots < t_{i,m_i} \in \{1,2,\dots,T\}$. The joint probability of these sequences of times occurring is

$$\begin{aligned}\prod_{i=1}^{I}\left(\frac{a_i}{N_{t_{i,1}-1}}\,\frac{a_i+1}{N_{t_{i,2}-1}}\cdots\frac{a_i+m_i-1}{N_{t_{i,m_i}-1}}\right)&= \frac{\prod_{i=1}^{I}[a_i(a_i+1)\cdots(a_i+m_i-1)]}{N_0\cdots N_{T-1}}\\ &= \frac{\prod_{i=1}^{I}\frac{(a_i+m_i-1)!}{(a_i-1)!}}{\frac{(N_0+T-1)!}{(N_0-1)!}} = \frac{\prod_{i=1}^{I}\frac{\Gamma(a_i+m_i)}{\Gamma(a_i)}}{\frac{\Gamma(N_0+T)}{\Gamma(N_0)}}.\end{aligned} \tag{A.4}$$

Note that this probability does not depend on the particular sequences of times $t_{i,k}$; hence, the probability that $m_{i,T} = m_i$ for $i \in \mathcal{I}$ is the term in (A.4) multiplied by the number of times in which we can partition T coins into I subsets with m_i elements each. This number is $\binom{T}{m_1,\dots,m_I} = T!/\prod_{i=1}^I m_i! = \Gamma(T+1)/\prod_{i=1}^I \Gamma(m_i+1)$. We get

$$\begin{aligned}\mathsf{P}(m_{i,T} = m_i) &= \frac{\prod_{i=1}^I \frac{\Gamma(a_i+m_i)}{\Gamma(a_i)}}{\frac{\Gamma(N_0+T)}{\Gamma(N_0)}} \frac{\Gamma(T+1)}{\prod_{i=1}^I \Gamma(m_i+1)} \\ &= C\,\frac{\Gamma(T+1)}{\Gamma(T+N_0)} \prod_{i=1}^I \frac{\Gamma(m_i+a_i)}{\Gamma(m_i+1)}, \end{aligned} \tag{A.5}$$

where $C = \Gamma(N_0)/\prod_{i=1}^I \Gamma(a_i)$, as in (6), and $N_0 = \sum_{i=1}^I a_i$. Note that the formula (A.5) assumes that $\sum_{i=1}^I m_i = T$. Thus, if we want the formula to be true in general, we must also include the term $\mathbf{1}_{\sum m i = T}$.

We introduce the following notation when n is large: $x_n \approx y_n$, which, by definition, means $\lim_{n\to\infty} \frac{x_n}{y_n} = 1$. Stirling's formula is $n! = \Gamma(n+1) \approx \sqrt{2\pi n}(n/e)^n$, which implies $\Gamma(n+\alpha)/\Gamma(n+\beta) \approx n^{\alpha-\beta}$. Using this approximation, Equation (A.5) implies $\mathsf{P}(m_{i,T} = m_i) \approx C \times (\prod_{i=1}^I m_i^{a_i-1})/T^{N_0-1}$.

Consider I divisions of the interval [0,1] with points of the form $x_i^{(m)} = m/T$, $m \in \{0,1,\dots,T\}$. For each of the I divisions, the distance between two consecutive points is $\Delta x_i = x_i^{(m)} - x_i^{(m-1)} = 1/T$. Note that i's share at T is $\pi_{i,T} = (a^0 + m_{i,t})/(N_0 + T) \approx m_{i,T}/T = x_i^{(m_{i,T})}$. Thus, setting $\pi_{i,T} = x_i$ implies $m_{i,T} \approx x_i T$. Therefore, the joint probability that i's share at t equals x_i is

$$\begin{aligned}\mathsf{P}(\pi_{i,T} = x_i) &\approx C\,\frac{\prod_{i=1}^I (x_i T)^{a_i-1}}{T^{N_0-1}} \mathbf{1}_{\sum x i = 1} = \frac{C}{T^{I-1}} \prod_{i=1}^I x_i^{a_i-1} \mathbf{1}_{\sum x i = 1} \\ &\approx C\,\Delta x_1 \cdots \Delta x_{I-1} \left(\prod_{i=1}^{I-1} x_i^{a_i-1}\right) \left(1 - \sum_{i=1}^{I-1} x_i\right)^{a_I - 1}. \end{aligned} \tag{A.6}$$

We, thus, obtain the density function of the Dirichlet distribution, which finishes the proof. □

Proof of Corollary 1. Define $n_1 = \pi_1 N$ as the initial number of coins of investor 1 and $n_2 = N - n_1 = (1-\pi_1)N$ the remaining number of coins. Then, Proposition 2 implies that the limiting distribution of investor 1's share is a beta distribution with parameters n_1 and n_2. Its variance is then

$$\mathsf{Var}(\pi_{1,\infty}) = \frac{n_1 n_2}{(n_1+n_2)^2(n_1+n_2+1)} = \frac{\pi_1(1-\pi_1)}{N+1}. \tag{A.7}$$

The rest of the proof is straightforward. □

Before proving Propositions 3 and 4, we prove several useful lemmas.

Lemma A.1. *The conditional variance at t of investor 1's share at $t+1$ is*

$$\mathsf{Var}_t(\pi_{1,t+1}) = \left(\frac{R_{t+1}}{N_{t+1}}\right)^2 \pi_{1,t}(1-\pi_{1,t}). \tag{A.8}$$

Proof. Equation (A.1) implies that

$$\pi_{1,t+1} = \frac{N_t \pi_{1,t}}{N_{t+1}} + \frac{R_{t+1}}{N_{t+1}} \mathbf{1}_{S_{1,t+1}}. \tag{A.9}$$

As the coin reward R_{t+1} is deterministic, conditional on the information at t, we have

$$\mathsf{Var}_t(\pi_{1,t+1}) = \left(\frac{R_{t+1}}{N_{t+1}}\right)^2 \mathsf{Var}_t(\mathbf{1}_{S_{1,t+1}}). \tag{A.10}$$

As $\mathsf{Var}_t(\mathbf{1}_{S_{1,t+1}}) = \pi_{1,t}(1-\pi_{1,t})$, the proof of (A.8) is complete. □

Lemma A.2. *The unconditional variance of investor 1's share at $t+1$ is*

$$\mathsf{Var}(\pi_{1,t+1}) = a_{t+1}\pi_{1,0}(1-\pi_{1,0}), \tag{A.11}$$

where the sequence a_t satisfies

$$a_1 = \left(\frac{R_1}{N_1}\right)^2, \qquad a_{t+1} = a_t + \left(\frac{R_{t+1}}{N_{t+1}}\right)^2 (1-a_t). \tag{A.12}$$

Proof. We proceed by induction. Lemma A.1 establishes the base case $t = 0$. Let $t \in \mathbb{N}_+$. A standard formula of conditional expectations implies

$$\mathsf{Var}(\pi_{1,t+1}) = \mathsf{Var}(\mathsf{E}_t(\pi_{1,t+1})) + \mathsf{E}(\mathsf{Var}_t(\pi_{1,t+1})). \tag{A.13}$$

As $\pi_{1,t}$ is a martingale, $\mathsf{E}_t(\pi_{1,t+1}) = \pi_{1,t}$. Using the formula (A.8) for $\mathsf{Var}_t(\pi_{1,t+1})$, we compute

$$\begin{aligned}\mathsf{Var}(\pi_{1,t+1}) &= \mathsf{Var}(\pi_{1,t}) + \left(\frac{R_{t+1}}{N_{t+1}}\right)^2 \mathsf{E}(\pi_{1,t}(1-\pi_{1,t})) \\ &= \mathsf{Var}(\pi_{1,t}) + \left(\frac{R_{t+1}}{N_{t+1}}\right)^2 (\pi_{1,0}(1-\pi_{1,0}) - \mathsf{Var}(\pi_{1,t})). \end{aligned} \tag{A.14}$$

By induction, $\mathsf{Var}(\pi_{1,t}) = a_t \pi_{1,0}(1-\pi_{1,0})$; hence, we obtain

$$\mathsf{Var}(\pi_{1,t+1}) = \left(a_t + \left(\frac{R_{t+1}}{N_{t+1}}\right)^2 (1-a_t)\right) \pi_{1,0}(1-\pi_{1,0}). \tag{A.15}$$

Thus, $\mathsf{Var}(\pi_{1,t+1}) = a_{t+1}\pi_{1,0}(1-\pi_{1,0})$, which completes the induction step. □

Lemma A.3. *Let $\nu > 0$ and $\theta_n \geq 0$ for all $n \in \mathbb{N}_+$ be some real constants. Define the sequence α_n by*

$$\alpha_1 = \left(\frac{\theta_1}{\nu + \theta_1}\right)^2, \quad \textit{and}$$

$$\alpha_{n+1} = \alpha_n + \left(\frac{\theta_{n+1}}{\nu + \sum_{k=1}^{n+1} \theta_k}\right)^2 (1-\alpha_n). \tag{A.16}$$

Then, for all $n \in \mathbb{N}_+$, $\alpha_n \in [0,1]$ and $\alpha_n \leq \alpha_{n+1}$.

Proof. By induction, we prove that $\alpha_n \in [0,1]$ and $\alpha_n \leq \alpha_{n+1}$. As $\nu > 0$ and $\theta_1 \geq 0$, the case $n = 1$ follows from $\alpha_1 = [\theta_1/(\nu+\theta_1)]^2$. We now assume the induction hypothesis. Clearly, $\theta_{n+1}/(\nu + \sum_{k=1}^{n+1}\theta_k) \in [0,1]$; therefore, the induction hypothesis $\alpha_n \in [0,1]$ implies that $\alpha_{n+1} = \alpha_n + [\theta_{n+1}/\nu + \sum_{k=1}^{n+1}\theta_k)]^2(1-\alpha_n)$ belongs to [0,1]. Equation (A.16) also implies that $\alpha_n \leq \alpha_{n+1}$, which completes the induction step. □

Lemma A.4. *Let $N>0$ and assume that the sequence R_t is positive and nonincreasing; that is, $R_t \geq R_{t+1} \geq 0$ for all $t \in \mathbb{N}$. Define the sequence a_t by*

$$a_1 = \left(\frac{R_1}{N+R_1}\right)^2,$$

$$a_{t+1} = a_t + \left(\frac{R_{t+1}}{N+\sum_{n=1}^{t+1} R_n}\right)^2 (1-a_t). \tag{A.17}$$

Then, $a_t \leq \frac{R_1}{N}$ for all $t \in \mathbb{N}_+$.

Proof. We extend the sequence a_t at $t=0$ by $a_0 = 0$. By summing from $n=1$ to $n=t+1$ the differences $a_{n+1}-a_n$ computed from Equation (A.17), we obtain

$$a_{t+1} - a_0 = \sum_{n=1}^{t+1} \left(\frac{R_n}{N+\sum_{m=1}^{n} R_m}\right)^2 (1-a_{n-1}). \tag{A.18}$$

Lemma A.3 implies that $a_n \in [0,1]$. As $a_0 = 0$, we obtain

$$a_{t+1} \leq \sum_{n=1}^{t+1} \left(\frac{R_n}{N+\sum_{k=1}^{n} R_k}\right)^2. \tag{A.19}$$

As $R_t \geq R_{t+1} \geq 0$ for all t, we have

$$\frac{R_t}{N+\sum_{k=1}^{n} R_k} - \frac{R_1}{N+tR_1} = \frac{N(R_t-R_1) + R_1\left(\sum_{k=1}^{t}(R_t-R_k)\right)}{\left(N+\sum_{k=1}^{n} R_k\right)(N+tR_1)} \leq 0. \tag{A.20}$$

Together, Equations (A.19) and (A.20) imply

$$a_{t+1} \leq \sum_{n=1}^{t+1} \left(\frac{R_1}{N+nR_1}\right)^2. \tag{A.21}$$

Let $f(x) = (\frac{R_1}{N+xR_1})^2$, twhich is a strictly decreasing function of $x \in \mathbb{R}_+$. Then, the right-hand side term in Equation (A.21) can be interpreted as a Riemann sum for the integral $\int_0^{t+1} f(x)\mathrm{d}x$. Therefore, we have

$$a_{t+1} \leq \int_0^{\infty} \left(\frac{R_1}{N+R_1 x}\right)^2 \mathrm{d}x = \frac{R_1}{N}. \tag{A.22}$$

This completes the proof. □

Lemma A.5. *Let $\rho > 0$ and $\gamma \in [0,1)$. For all $t \in \mathbb{N}_+$, let $R_t = \rho N_{t-1}^{\gamma}$ and $N_t = N + \sum_{n=1}^{t} R_n$. Define the sequence a_t as in Equation (A.12). Then, for all $t \in \mathbb{N}_+$,*

$$a_{t+1} \leq \frac{\rho}{1-\gamma} N^{\gamma-1}. \tag{A.23}$$

Proof. As in the proof of Lemma A.4, we obtain (see Equation (A.19))

$$a_{t+1} \leq \sum_{n=1}^{t+1} \left(\frac{R_n}{N_n}\right)^2. \tag{A.24}$$

If we define $\Delta N_n = N_n - N_{n-1}$, we have $\Delta N_n = R_n$, which is an increasing sequence. As $R_{n+1} = \rho N_n^{\gamma}$, we compute

$$a_{t+1} \leq \sum_{n=1}^{t+1} \frac{R_n R_{n+1}}{N_n^2} \leq \rho \sum_{n=1}^{t+1} \frac{\Delta N_n}{N_n^{2-\gamma}}. \tag{A.25}$$

Let $g(x) = \rho x^{\gamma-2}$, which is a strictly decreasing function of $x \in \mathbb{R}_+$. Then, the rightmost term in Equation (A.25) can be interpreted as a lower Riemann sum for the integral $\int_N^{t+1} g(x)\mathrm{d}x$. Therefore, we have

$$a_{t+1} \leq \rho \int_N^{\infty} \frac{\mathrm{d}x}{x^{2-\gamma}} \leq \frac{\rho}{1-\gamma} N^{\gamma-1}, \tag{A.26}$$

which completes the proof. □

Lemma A.6. *The condition*

$$\lim_{N\to\infty} \mathsf{Var}(\pi_{1,\infty}) = 0 \tag{A.27}$$

implies that, for any $\varepsilon > 0$,

$$\lim_{N\to\infty} \mathsf{P}(|\pi_{1,\infty} - \pi_1| \geq \varepsilon) = 0. \tag{A.28}$$

Proof. Chebyshev's inequality states that any random variable X with finite mean μ and variance σ^2 satisfies $\mathsf{P}(|X-\mu| \geq k\sigma) \leq 1/k^2$. In our case, let $X = \pi_{1,\infty}$, and denote its variance by $\sigma_{1,\infty}^2$. Fix $\varepsilon > 0$, and let $k = \varepsilon/\sigma_{1,\infty}$. Chebyshev's inequality then implies that $\mathsf{P}(|\pi_{1,\infty} - \pi_1| \geq \varepsilon) \leq \sigma_{1,\infty}^2/\varepsilon^2$. As $\lim_{N\to\infty}\sigma_{1,\infty} = 0$, condition (A.28) follows. □

Proof of Proposition 3. Denote by $\mu_{1,\infty}$ and $\sigma_{1,\infty}$, respectively, the mean and standard deviation of the limiting distribution $\pi_{1,\infty}$. Lemma A.2 implies that $\mathsf{Var}(\pi_{1,t+1}) = a_{t+1}\pi_1(1-\pi_1)$, where π_1 is the initial share of investor 1. Lemma A.4 implies that $a_{t+1} \leq \frac{R_1}{N}$, where N is the initial number of coins outstanding (see Equation (A.22)). Therefore, $\lim_{N\to\infty}\sigma_{1,\infty} = 0$. Moreover, Proposition 1 implies that $\mu_{1,\infty} = \pi_1$. Then, Lemma A.6 completes the proof. □

Proof of Proposition 4. We first consider the case $\gamma \in [0,1)$. Lemma A.5 implies that $a_{t+1} \leq \rho N^{\gamma-1}/1-\gamma$. Using the same proof as for Proposition 3, one shows that, for any $\varepsilon > 0$, $\lim_{N\to\infty}\mathsf{P}(|\pi_{1,\infty} - \pi_1| \geq \varepsilon) = 0$.

Consider the case $\gamma \geq 1$. Proposition 1 shows that there exists a well-defined random variable $\pi_{1,\infty} = \lim_{t\to\infty}\pi_{1,t}$, and let $\mu_{1,\infty} = \pi_1$ be its mean and $\sigma_{1,\infty}$ its standard deviation. We show that there exists some $\varepsilon > 0$ such that

$$\lim_{N\to\infty} \mathsf{P}(|\pi_{1,\infty} - \pi_1| \geq \varepsilon) > 0. \tag{A.29}$$

Lemma A.2 shows that $a_t = \mathsf{Var}(\pi_{1,t})/[\pi_1(1-\pi_1)]$, where a_t is the sequence in Equation (A.12). Let $a_\infty = \mathsf{Var}(\pi_{1,\infty})/[\pi_1(1-\pi_1)]$. The bounded convergence then implies that $a_\infty = \lim_{t\to\infty} a_t$.

The inequality $\gamma \geq 1$ implies that the ratio $R_{t+1}/N_{t+1} = \rho N_t^{\gamma}/(N_t + \rho N_t^{\gamma}) = \rho N_t^{\gamma-1}/(1+\rho N_t^{\gamma-1})$ converges when $N \to \infty$ to a number that is at least equal to $\rho/(1+\rho)$. Lemma A.3 implies that the sequence a_t is weakly increasing and bounded by zero and one; hence, $a_\infty \in [0,1]$. Equation (A.12) implies that $a_{t+1} - a_t = (R_{t+1}/N_{t+1})^2(1-a_t)$. As $a_{t+1} - a_t$ converges to zero, it follows that $\lim_{N\to\infty} a_t = 1$. Thus, $a_\infty = 1$, which implies that $\mathsf{Var}(\pi_{1,\infty}) = \pi_1(1-\pi_1)$.[15] For any $\varepsilon > 0$, we compute

$$\begin{aligned}\pi_1(1-\pi_1) &= \mathsf{E}\big((\pi_{1,\infty} - \pi_1)^2\big) \\ &= \mathsf{E}\big((\pi_{1,\infty} - \pi_1)^2 \mathbf{1}_{|\pi_{1,\infty}-\pi_1| \geq \varepsilon}\big) \\ &\quad + \mathsf{E}\big((\pi_{1,\infty} - \pi_1)^2 \mathbf{1}_{|\pi_{1,\infty}-\pi_1| < \varepsilon}\big) \\ &\leq \mathsf{P}(|\pi_{1,\infty} - \pi_1| > \varepsilon) \\ &\quad + \varepsilon^2(1 - \P(|\pi_{1,\infty} - \pi_1| > \varepsilon)),\end{aligned} \tag{A.30}$$

where, for the last inequality, we use the fact that $|\pi_{1,\infty} - \pi_1| < 1$ almost surely as investor shares are bounded by zero and one. Choose a number $\varepsilon \in (0, \sqrt{\pi_1(1-\pi_1)})$. Rewriting Equation (A.30), we obtain

$$\mathsf{P}(|\pi_{1,\infty} - \pi_1| > \varepsilon) \geq \frac{\pi_1(1-\pi_1) - \varepsilon^2}{1-\varepsilon^2} > 0. \tag{A.31}$$

Taking $\liminf_{N\to\infty}$ on both sides completes the proof. □

Before proving Proposition 5, we prove a lemma that computes the expected utility gain from coin issuance. Recall that, at $t+1 \in \mathbb{N}_+$, the order of events is as follows: (i) R_{t+1} coins are rewarded to investor i with probability $\pi_{i,t} = n_{i,t}/N_t$, thus increasing investor i's ownership of coins from $n_{i,t}$ to $n'_{i,t+1}$; (ii) the price changes exogenously from P_t to P_{t+1}; and (iii) investor i trades $v_{i,t+1}$ coins at P_{t+1}.

Lemma A.7. *Investor i's expected utility gain from coin issuance at $t+1 \in \mathbb{N}_+$ satisfies*

$$\mathsf{E}_t\left(n'_{i,t+1}P_{t+1} - n_{i,t}P_t\right) = \pi_{i,t}(\mathsf{E}_t(M_{t+1}) - M_t). \tag{A.32}$$

Proof. Equation (3) implies that $n'_{i,t+1} = n_{i,t} + R_{t+1}\mathbf{1}_{S_{i,t+1}}$, where $S_{i,t+1}$ is the event of i being selected at $t+1$ with probability $\pi_{i,t} = n_{i,t}/N_t$. The equality $N_{t+1} = N_t + R_{t+1}$ then implies that $\mathsf{E}_t(n'_{i,t+1}) = n_{i,t}N_{t+1}/N_t$. (Here, we use that R_{t+1} is deterministic.) Equation (4) implies that $P_t = M_t/N_t$ and $P_{t+1} = M_{t+1}/N_{t+1}$. As the event $S_{i,t+1}$ is independent from P_{t+1}, we compute

$$\mathsf{E}_t\left(n'_{i,t+1}P_{t+1}\right) = n_{i,t}\mathsf{E}_t(M_{t+1})/N_t = \pi_{i,t}\mathsf{E}_t(M_{t+1}). \tag{A.33}$$

Also, $n_{i,t}P_t = \pi_{i,t}M_t$, which, together with (A.33), proves (A.32). □

Proof of Proposition 5. For any trading strategy $\nu_i = (\nu_{i,t})_{t\in\mathbb{N}_+}$, define a process $\Pi_{i,t}$ by

$$\Pi_{i,0} = n_{i,0}P_0, \qquad \Pi_{i,t} = \delta^t n'_{i,t}P_t - \sum_{s=1}^{t-1}\delta^s \nu_{i,s}P_s \quad \text{if } t \in \mathbb{N}_+. \tag{A.34}$$

Let $t \in \mathbb{N}$. As $n_{i,t} = n'_{i,t} + \nu_{i,t}$, we compute

$$\Pi_{i,t+1} - \Pi_{i,t} = \delta^{t+1}n'_{i,t+1}P_{t+1} - \delta^t n_{i,t}P_t. \tag{A.35}$$

Equations (A.33) and (A.35) imply that

$$\mathsf{E}_t(\Pi_{i,t+1}) - \Pi_{i,t} = \pi_{i,t}\left(\delta^{t+1}\mathsf{E}_t(M_{t+1}) - \delta^t M_t\right). \tag{A.36}$$

Equations (12) and (13) then imply that Π_t is a martingale. Moreover, the dominated convergence theorem implies that $\mathsf{E}(\Pi_{i,\tau_i}) = \Pi_{i,0} = n_{i,0}P_{i,0}$.[16] Finally, the budget constraint, $c_{i,\tau_i} = n'_{i,\tau_i}P_{\tau_i} + (1 + r_{\tau_i-1,\tau_i})b_{i,\tau_i-1}$ and $b_{i,\tau_i} = 0$ from (11) imply $\mathsf{E}(\sum_{t=1}^{\tau_i}\delta^t c_{i,t}) = \mathsf{E}(\delta^{\tau_i}n'_{i,\tau_i}P_{\tau_i} - \sum_{t=1}^{\tau_i-1}\delta^t\nu_{i,t}P_t) = \mathsf{E}(\Pi_{i,\tau_i}) = \Pi_{i,0} = n_{i,0}P_{i,0}$, which completes the proof. □

Appendix B. Reward Functions

Thus far in this paper, we have considered only the limiting distribution of investor shares that arise from constant rewards. As constant rewards are widely used in practice, our analysis is sufficient to address practical questions regarding the evolution of investor shares in cryptocurrencies with a PoS protocol. Nevertheless, the theoretical literature argues that other reward functions may produce less concentrated distributions of investor shares. For example, Fanti et al. (2019) show that a geometric reward function minimizes investor share variance over a finite horizon, subject to a constraint regarding the number of coins distributed over that horizon.

In this section, we compare the evolution of investor shares corresponding to both constant and geometric reward functions. As we are interested in the limiting distribution of investor shares, we consider the evolution of investor shares beyond the finite horizon set in Fanti et al. (2019).

We, thus, define the geometric reward as the reward function that minimizes investor share variance for a given horizon $\tilde{T}$ and analyze the evolution of shares over $T \geq \tilde{T}$. Fanti et al. (2019) show that the geometric reward in period t is of the form

$$R_t = N \times \left(\left(\frac{1+\tilde{R}}{N}\right)^{t/\tilde{T}} - \left(\frac{1+\tilde{R}}{N}\right)^{(t-1)/\tilde{T}}\right), \tag{B.1}$$

where N is the initial number of coins in circulation, $\tilde{T}$ is the number of periods, and $\tilde{R} = \sum_{t=1}^{\tilde{T}} R_t$ is the free parameter that determines the total reward distributed over the first $\tilde{T}$ periods.

Figure B.1 depicts simulated investor share distributions for both geometric and constant rewards over three horizons: $T \in \{1{,}000, 10{,}000, 25{,}000\}$. The geometric reward is computed as in Equation (B.1) for the following parameter values: $N = 100$, $\tilde{T} = 1{,}000$, and $\tilde{R} = 1{,}000$.

Our results are consistent with Fanti et al. (2019) at a horizon T equal to the parameter $\tilde{T}$. Indeed, the geometric reward produces a lower variance than that of the constant reward over that horizon. Nonetheless, the same geometric reward applied over longer horizons ($T = 10{,}000$ and $T = 25{,}000$) produces a significantly higher variance than the constant reward over the same horizons. Moreover, Although the investor share variance of the constant reward varies only modestly from $T = 1{,}000$ to $T = 25{,}000$, the investor share variance of the geometric reward rises dramatically from $T = 1{,}000$ to $T = 25{,}000$. Thus, although a geometric reward minimizes investor share variance over a finite horizon subject to a requirement regarding the number of coins disbursed over that horizon, such a reward function generates large variances when applied over large horizons.

Determining the optimal reward function lies beyond the scope of our analysis. Nakamoto (2008) proposes rewards within Bitcoin to achieve security, but PoS introduces a fundamentally different context in which rewards may play a different role than generating security. For example, Brown-Cohen et al. (2019) highlight that the "absence of rewards [...] achieves the same formal [security] guarantees" as having rewards for particular extant PoS protocols. Similarly, Saleh (2020) demonstrates that smaller rewards help PoS overcome the nothing-at-stake problem.

Determining the optimal reward function for a PoS protocol remains an active research area. Our paper does not aim to contribute to that area. Instead, we highlight that, contrary to conventional wisdom, the widely used constant

Figure B.1. (Color online) Investor Share Distributions Under Constant and Geometric Reward Functions

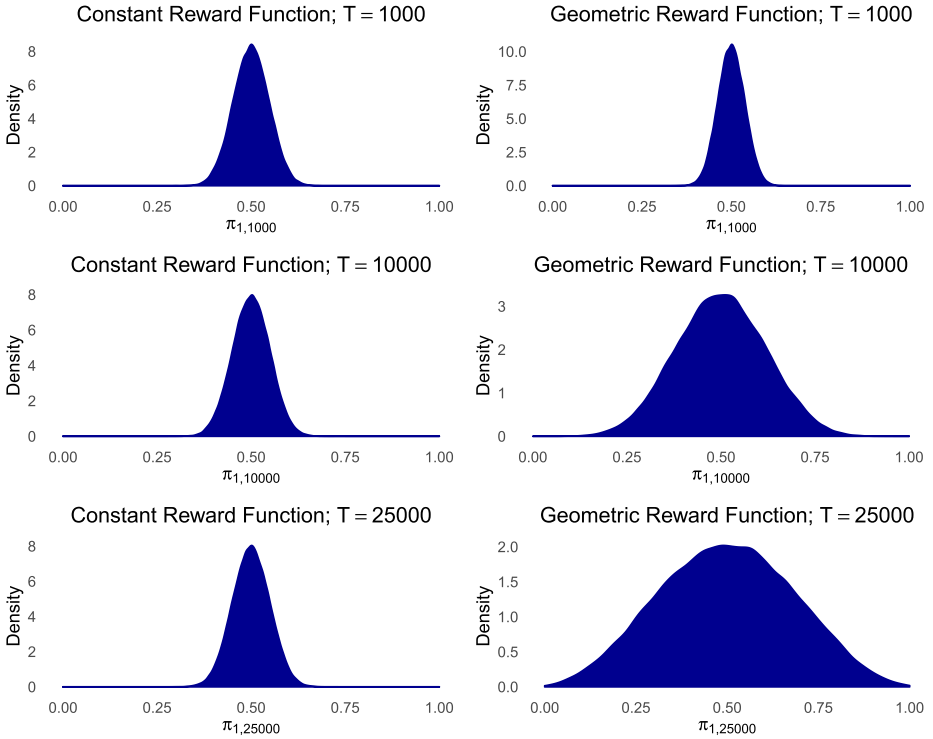

Notes. For an initial total number of coins $N = 100$ and an initial share of investor 1 equal to $\pi_1 = 0.5$, the plots show the density of investor 1's share after T steps. Each graph is generated from 10,000 sample paths. We select a geometric reward schedule that minimizes investor share variance over $T = 1{,}000$ steps. The figure is generated using the R Statistical Software with a random seed of 100.

reward function does not induce wealth concentration, and a geometric reward function produces significant wealth concentration in the limit.

Appendix C. Mining Costs

In this section, we generalize our results to the case in which investors must pay an explicit mining cost $\kappa > 0$ every time they get selected at t to validate a block. As in Section 4, investor i's strategy is a triple $(v_{i,t}, b_{i,t}, c_{i,t})_{t\in\mathbb{N}_+}$, where $v_{i,t}$ is the number of coins traded at t, $b_{i,t}$ is the end-of-period-t holding of the risk-free asset, and $c_{i,t}$ is the consumption at t. If $r_{t-1,t}$ denotes the one-period-ahead risk-free rate from $t-1$ to t, then investor i's strategy solves the maximization problem

$$\begin{aligned} U_i = \max_{v_{i,t},b_{i,t},c_{i,t}} \mathsf{E}\left(\sum_{t=1}^{\tau_i} \delta^t c_{i,t}\right) \quad &\text{such that:} \\ c_{i,t} + b_{i,t} + v_{i,t}P_t + \kappa \mathbf{1}_{S_{i,t}} &= (1 + r_{t-1,t})b_{i,t-1}, \\ 0 \le n'_{i,t} + v_{i,t} &\le N_t, \\ c_{i,\tau_i} = n'_{i,\tau_i}P_{\tau_i} + (1 + r_{\tau_i-1,\tau_i})b_{i,\tau_i-1} &- \kappa \mathbf{1}_{S_{i,t}}, \\ b_{i,\tau} = n_{i,\tau} = 0 \quad &\text{for } \tau \ge \tau_i. \end{aligned} \tag{C.1}$$

This is the same problem as in Equation (11) except for the additional term $\kappa \mathbf{1}_{S_{i,t}}$, which is the mining cost that is paid if investor i is selected at t.

Proposition C.1 describes the equilibrium utility of an investor. We assume the following necessary equilibrium conditions for a risk-neutral economy:

$$r_{t,t+1} = \frac{1}{\delta} - 1, \tag{C.2}$$

$$\mathsf{E}_t(M_{t+1}) - \kappa = (1 + r_{t,t+1})M_t, \tag{C.3}$$

where Equation (C.2) determines the endogenous risk-free rate, and Equation (C.3) arises as an intratemporal condition across the risk-free asset and the cryptocurrency. Note that the cryptocurrency is a traded asset in a risk-neutral economy, and thus, it must have an expected return equal to the risk-free rate after mining costs. If the cryptocurrency did not provide at least such an expected return after mining costs, then investors would not hold the asset in equilibrium. Vice versa, if the cryptocurrency provided a higher expected return after mining costs then investors would have infinite demand for the asset, thus unraveling the equilibrium.

Proposition C.1 (Trading Irrelevance with Mining Costs)**.** *If the conditions* (C.2) *and* (C.3) *are satisfied, any trading strategy* $v_i = (v_{i,t})_{t\in\mathbb{N}_+}$ *provides the same expected utility for investor* i *at* $t = 0$*:*

$$U_i = n_{i,0}P_0. \tag{C.4}$$

Proposition C.1 shows that our irrelevance result arises even with mining costs. The intuition for this result is similar to the intuition for Proposition 5. Mining costs do

not affect this intuition because prices adjust to account for the mining costs.

Proof of Proposition C.1. For any trading strategy $\nu_{i,t}$, define a process $\Pi'_{i,t}$ by

$$\Pi'_{i,0} = n_{i,0}P_0,$$
$$\Pi'_{i,t} = \delta^t\left(n'_{i,t}P_t - \kappa\mathbf{1}_{S_{i,t}}\right) - \sum_{s=1}^{t-1}\delta^s\left(\nu_{i,s}P_s + \kappa\mathbf{1}_{S_{i,s}}\right). \tag{C.5}$$

Let $t\in\mathbb{N}$. As $n_{i,t} = n'_{i,t} + \nu_{i,t}$, we compute

$$\Pi'_{i,t+1} - \Pi'_{i,t} = \delta^{t+1}\left(n'_{i,t+1}P_{t+1} - \kappa\mathbf{1}_{S_{i,t+1}}\right) - \delta^t n_{i,t}P_t. \tag{C.6}$$

Equations (A.33) and (A.35) imply that

$$\mathsf{E}_t\left(\Pi'_{i,t+1}\right) - \Pi'_{i,t} = \pi_{i,t}\left(\delta^{t+1}(\mathsf{E}_t(M_{t+1}) - \kappa) - \delta^t M_t\right). \tag{C.7}$$

Equations (C.2) and (C.3) then imply that Π_t is a martingale. Moreover, the dominated convergence theorem implies that $\mathsf{E}(\Pi'_{i,\tau_i}) = \Pi'_{i,0} = n_{i,0}P_{i,0}$.[17] Finally, the budget constraint, $c_{i,\tau_i} = n'_{i,\tau_i}P_{\tau_i} + (1+r_{\tau_i-1,\tau_i})b_{i,\tau_i-1} - \kappa\mathbf{1}_{S_{i,\tau_i}}$ and $b_{i,\tau_i} = 0$ from (C.1) imply $\mathsf{E}(\sum_{t=1}^{\tau_i}\delta^t c_{i,t}) = \mathsf{E}(\delta^{\tau_i}(n'_{i,\tau_i}P_{\tau_i} - \kappa\mathbf{1}_{S_{i,\tau_i}}) - \sum_{t=1}^{\tau_i-1}\delta^t(\nu_{i,t}P_t + \kappa\mathbf{1}_{S_{i,s}})) = \mathsf{E}(\Pi'_{i,\tau_i}) = \Pi'_{i,0} = n_{i,0}P_{i,0}$, which completes the proof. □

Endnotes

[1] As of March 28, 2019, Cryptoslate lists 2,128 cryptocurrencies, out of which 835 have their own blockchain ledger and are sometimes called "coins," and the rest are called "tokens."

[2] Cryptoslate lists 402 coins with a PoS protocol, for example, Nxt, BlackCoin, and Wave, and it lists 531 coins with a PoW protocol, for example, Bitcoin and Ethereum. Some coins are hybrid and have both PoW and PoS protocols, for example, Peercoin. Irresberger et al. (2020) provide further details regarding the prevalence of various protocols among public blockchains.

[3] "Proof of Work vs Proof of Stake," *CoinGeek*, May 28, 2018.

[4] In general, reward schedules vary widely and include zero rewards (e.g., Nxt), constant rewards (e.g., Blackcoin), decreasing rewards (e.g., Bitcoin), and increasing rewards (e.g., EOS).

[5] See Pemantle (2007) and the references therein.

[6] We prove Proposition 1 for any deterministic coin reward R_t, but the same proof works when the coin reward R_t is random as long as the event of being selected at time t is independent from R_t.

[7] See the proof of Proposition 1. Note that Equation (A.2) implies that investor i's share change at t is $\pi_{i,t} - \pi_{i,t-1} = [\mathbf{1}_{S_{i,t}} - 1/10]/N_t$, and the probability of $S_{i,t}$ is 0.1.

[8] See the proof of Proposition 1 in Appendix A.

[9] It is not the aim of this paper to determine an optimal reward schedule but rather to analyze the concentration of coin shares by taking a reward schedule as given.

[10] Note that, according to Corollary 1, this result is true only if $\pi_1 < 0.5$. If investor 1 is "super-rich" (i.e., $\pi_1 \geq 0.5$), then investor 1's limiting distribution becomes tighter around its mean as investor 1 becomes richer. In practice, however, even a large investor is unlikely to own more than 50% of all coins, so we restrict ourselves to the case $\pi_1 < 0.5$.

[11] Another definition of stability is $\lim_{N\to\infty}\mathsf{Var}(\pi_{1,\infty}) = 0$, which implies $\lim_{N\to\infty}\mathsf{P}(|\pi_{1,\infty} - \pi_1| \geq \varepsilon) = 0$ for any $\varepsilon > 0$ (see Lemma A.6). In Appendix A, we show that Propositions 3 and 4 are true under this alternative notion of stability.

[12] The initial number of coins determines the initial shares but not vice versa. Thus, our results imply that the limiting share distribution depends not only on the initial shares, but also on the initial numbers of coins.

[13] In Appendix C, we show that trading irrelevance holds also in the presence of a mining cost.

[14] The proof of this proposition can be extended to random R_t and N_t as long as these variables are independent from the event $S_{i,t}$.

[15] Note that $\sigma_{1,\infty} > 0$ is constant, and therefore, condition (A.27) is not satisfied.

[16] Let $X_i = \sum_{t=0}^{\tau_i}\delta^t N_t P_t$. Then, $|\Pi_{i,t}| \leq X_i$ for all $t\in\mathbb{N}$, and $E(|X_i|) = (1 + E(\tau_i))N_0P_0 < \infty$. Thus, $\Pi_{i,0} = \lim_{t\to\infty}\mathsf{E}(\Pi_{i,t\wedge\tau_i}) = \mathsf{E}(\Pi_{i,\tau_i})$.

[17] Let $X_i = \sum_{t=0}^{\tau_i}\delta^t N_t P_t + \kappa/(1-\delta)$. Then, $|\Pi'_{i,t}| \leq X_i$ for all $t\in\mathbb{N}$, andxref $E(|X_i|) = (1+E(\tau_i))N_0P_0 + \kappa/(1-\delta) < \infty$.
Thus, $\Pi'_{i,0} = \lim_{t\to\infty}\mathsf{E}(\Pi'_{i,t\wedge\tau_i}) = \mathsf{E}(\Pi'_{i,\tau_i})$.

References

Alsabah H, Capponi A (2020) Pitfalls of bitcoin's proof-of-work: R&D arms race and mining centralization. Working paper, Columbia University, New York, NY.

Arnosti N, Weinberg SM (2019) Bitcoin: A natural oligopoly. Blum A, ed. *10th Innovations in Theoretical Computer Science, ITCS 2019 [5]*, Leibniz International Proceedings in Informatics, Vol. 124 (Schloss Dagstuhl-Leibniz-Zentrum fur Informatik GmbH, Dagstuhl Publishing, Wadern, Germany).

Biais B, Bisière C, Bouvard M, Casamatta C (2019) The blockchain folk theorem. *Rev. Financial Stud.* 32(5):1662–1715.

Brown-Cohen J, Narayanan A, Psomas C, Weinberg SM (2019) Formal barriers to longest-chain proof-of-stake protocols. *Proc. 2019 ACM Conf. Econom. Comput. (EC '19)* (Association for Computing Machinery, New York), 459–473.

Fanti GC, Kogan L, Oh S, Ruan K, Viswanath P, Wang G (2019). Compounding of wealth in proof-of-stake cryptocurrencies. Goldberg I, Moore T, eds. *Financial Cryptography and Data Security* (FC 2019), Lecture Notes in Computer Science, vol. 11598 (Springer, Cham, Switzerland).

Irresberger F (2018) Coin concentration of proof-of-stake blockchains. Working paper, Durham University, UK.

Irresberger F, John K, Saleh F (2020) The public blockchain ecosystem: An empirical analysis. Working paper, New York University Stern, New York.

Nakamoto S (2008) Bitcoin: A peer-to-peer electronic cash system. White paper. Accessed September 20, 2020, https://bitcoin.org/bitcoin.pdf.

Pemantle R (2007) A survey of random processes with reinforcement. *Probab. Surveys* 4:1–79.

Saleh F (2020) Blockchain without waste: Proof-of-stake. *Rev. Financial Stud.*, ePub ahead of print July 7, https://academic.oup.com/rfs/advance-article-abstract/doi/10.1093/rfs/hhaa075/5868423?redirectedFrom=fulltext.

Tsoukalas G, Falk BH (2020) Token-weighted crowdsourcing. *Management Sci.*, ePub ahead of print August 3, https://doi.org/10.1287/mnsc.2019.3515.