POLYNOMIAL VOTING RULES

WENPIN TANG AND DAVID D. YAO

Abstract. We propose and study a new class of polynomial voting rules for a general decentralized decision/consensus system, and more specifically for the PoS (Proof of Stake) protocol. The main idea, inspired by the Penrose square-root law and the more recent quadratic voting rule, is to differentiate a voter's voting power and the voter's share (fraction of the total in the system). We show that while voter shares form a martingale process that converge to a Dirichlet distribution, their voting powers follow a super-martingale process that decays to zero over time. This prevents any voter from controlling the voting process, and thus enhances security. For both limiting results, we also provide explicit rates of convergence. When the initial total volume of votes (or stakes) is large, we show a phase transition in share stability (or the lack thereof), corresponding to the voter's initial share relative to the total. We also study the scenario in which trading (of votes/stakes) among the voters is allowed, and quantify the level of risk sensitivity (or risk averse) in three categories, corresponding to the voter's utility being a super-martingale, a sub-martingale, and a martingale. For each category, we identify the voter's best strategy in terms of participation and trading.

Key words: Cryptocurrency, economic incentive, fluid limit, phase transition, polynomial voting rules, Proof of Stake protocol, stability, urn models.
AMS 2020 Mathematics Subject Classification: 60C05, 60F05, 60G42, 91B08.

1. Introduction

Voting, in the traditional sense, refers to a set of rules for a community of individuals or groups ("voters") to reach an agreement, or to make a collective decision on some choices and ranking problems. In today's world, "voting" has become a ubiquitous notion that includes any decentralized decision-making protocol or system, where the voters are often abstract entities ("virtual"), and the voting process automated; and the purpose of reaching consensus is often non-social and non-political, such as to enhance the overall security of an industrial operation or infrastructure (Garcia-Molina (1982); Lamport et al. (1982)). Examples include cloud computing (Armbrust et al. (2009); Dean and Ghemawat (2008)), smart power grids (Huang and Baliga (2009)), and more recently, trading or payment platforms and exchanges built upon the blockchain technology (Nakamoto (2008); Wood (2014)).

At the core of a blockchain is the *consensus protocol*, which specifies a set of voting rules for the participants ("miners" or validators) to agree on an ever-growing log of transactions (the "longest chain") so as to form a distributed ledger. There are several existing blockchain protocols, among which the most popular are Proof of Work (PoW, Nakamoto (2008)) and Proof of Stake (PoS, King and Nadal (2012); Wood (2014)). In the PoW protocol, miners compete with each other by solving a hashing puzzle. The miner who solves the puzzle first receives a reward (a number of coins) and whose work validates a new block's addition to

Date: January 8, 2024.

the blockchain. Hence, while the competition is open to everyone, the chance of winning is proportional to a miner's computing power.

In the PoS protocol, there is a bidding mechanism to select a miner to do the work of validating a new block. Participants who choose to join the bidding are required to commit some "stakes" (coins they own), and the winning probability is proportional to the number of stakes committed. Hence, a participant in a PoS blockchain is actually a "bidder," as opposed to a "miner" — only the winning bidder becomes the miner validating the block. (Any participant who chooses not to join the bidding can be viewed as a bidder who commits zero stakes.) Needless to add, bidding exists long before the PoS protocol, and has been widely used in many applications, such as auctions and initial public offerings (IPOs).

Let's explore the PoS bidding mechanism a bit more formally. Suppose a voter (or bidder) k is in possession of $n_{k,t}$ votes (or stakes) at time t, an index that counts the rounds of voting or bidding in the protocol; and $N_t := \sum_k n_{k,t}$ is the total number of votes over all voters. Hence, voter k's share, fraction of the total, is $\pi_{k,t} := n_{k,t}/N_t$. Following a traditional voting rule, voter k's chance or probability of winning, which we call *voting power*, will be equal to $\pi_{k,t}$, voter k's share. Yet, this doesn't have to be the case. That is, any voter's voting power needs not be equal to the voter's share (of the system total). Indeed, there are often good reasons for the two to be different.

Historically, the English scholar Lionel Penrose famously proposed a square-root voting rule (Penrose (1946)), around the time when the United Nations was founded shortly after WWII. According to Penrose, a world assembly such as the UN should designate each country a number of votes that is proportional to the square root of its population. The obvious implication (which may or may not be what Penrose initially intended) is to limit the voting power of nations with very large populations. In the same spirit, the quadratic voting rule has attracted much attention in recent years (Lalley and Weyl (2018)). The idea is that each voter be given a budget (in dollars, for instance); the voter can cast multiple votes on any single or subset of choices or candidates on the ballot, with x votes (for any choice) costing x^2 dollars. Under both voting rules, the voting power is different from the voter's share or representation in the system, population in the first case, and the budget in the second case.

Inspired by these ideas, we propose a class of polynomial voting rules, denoted $\texttt{Poly}(\alpha)$, which grant every voter k a voting power that scales the voter's share $\pi_{k,t}$ by a factor $N_t^{-\alpha}$ for $\alpha \geq 0$. When $\alpha = 0$, this reduces to the traditional voting case of power=share, which is a linear rule. When $\alpha = 1$, the rule resembles the square-root or the quadratic voting rules mentioned above in spirit, in terms of decoupling voting power from a voter's share, but of course differs in both the application context and implementation schemes. As we will demonstrate, the general $\texttt{Poly}(\alpha)$ rule is a time change of the $\texttt{Poly}(0)$ rule, with the parameter α measuring how much the traditional $\alpha = 0$ rule is "slowed down," namely, the voting power is diminished over time.

There are (at least) two reasons to consider slowed-down voting schemes in blockchains.

- First, the block-generation time requires to be lower bounded due to network delay (see (Shi, 2020, Section 14.3)). Specifically, there is the principle of security:

$$(1-v)\cdot \text{honest power} > \gamma \cdot \text{dishonest power}, \quad \text{or} \quad v < 1 - \frac{\gamma \cdot \text{dishonest power}}{\text{honest power}}. \tag{1.1}$$

Here "honest/dishonest power" refers to the voting power of honest/dishonest bidders. The parameter γ is a user-defined *security factor*; e.g., $\gamma = 2$ means, honest power is expected to be twice as much as dishonest power; hence, γ measures how secure a distributed system is. When honest bidders broadcast their validation results, dishonest bidders may exploit network delay to attack; equivalently, network delay will reduce the honest power. Thus the term "$1 - v$" plays the role of a discount factor, with v proportional to network delay (the more severe network delay is, the smaller the discount $1 - v$ is, and hence the larger v is). Honest power is discounted also because honest bidders follow exactly the protocol while dishonest bidders do not comply with the rules. As we will illustrate (in the remarks following Theorem 6) slowing down the voting process enhances security. This is because decreasing voting power over time will increase the block generation time, which will mitigate network delay and make the principle of security (1.1) "easier" to hold.

- Second, PoS blockchains suffer from malicious attacks known as *Nothing at Stake* (see e.g. Deirmentzoglou et al. (2019)). As pointed out in Bagaria et al. (2019), for the PoS longest-chain protocol, honest bidders focus exclusively on the longest chain while dishonest bidders can work simultaneously on all existing blocks. They showed that the PoS longest-chain is less secure than its PoW counterpart, assuming both honest and dishonest parties have constant voting power over time. However, as dishonest bidders have more flexibility, it is (much) more likely that they win and get rewarded, and their advantage is only amplified over time. This makes "constant voting power" highly undesirable. There are two general approaches to solving this problem: (i) adjust the amount of reward over time, (ii) slow down the voting process; both are aimed at preventing dishonest bidders from overpowering honest bidders as time evolves.

Here is an overview of our main findings and results. We prove that under the $\texttt{Poly}(\alpha)$ voting rule, voter shares form a martingale process that converges to a Dirichlet distribution as $t \to \infty$, while their voting powers follow a super-martingale process that decreases to zero over time (Theorem 6); and for both limits we also explicitly characterize their rates of convergence. Thus, the $\texttt{Poly}(\alpha)$ voting scheme enhances security, preventing any voter or any group of voters from controlling the voting process and overpowering the system.

We further group the voters into two categories: large and small, according to the initial (time zero) votes they own relative to the total (N_0). When N_0 is large, which is the case in most applications, we show a phase transition in the stability of voter shares across the two categories (Proposition 7). Notably, the same phenomenon is demonstrated under the traditional voting rule ($\alpha = 0$), refer to Roşu and Saleh (2021); Tang (2022). Our result establishes that this phase transition is in fact *universal*, in the sense that it applies to all values of $\alpha(\geq 0)$.

We also study the scenario in which trading (of votes/stakes) among the voters (or "bidders") is allowed, motivated by PoS applications in cryptocurrency. For $\alpha = 0$, the trading scenario has been recently studied in Roşu and Saleh (2021). Not only our model is more general, in allowing any $\alpha \geq 0$, our results are also richer and sharper (Theorem 9). For instance, we quantify the level of risk sensitivity (or risk-averse) that results in three cases according to the voter's utility being a super-martingale, a sub-martingale or a martingale. Each case will lead to a best strategy for the voter, including "non-participation" (not to

participate at all in the bidding) and "buy out" (buying as many stakes as what is available), which are not considered in Roşu and Saleh (2021). Note that "buy-out" is a monopoly, and it is desirable to limit the number of stakes that any voter can acquire in a single round. This is studied in our subsequent paper Tang and Yao (2023) on trading PoS stakes with volume constraint. See also Tang (2023) for various problems (including transaction costs and voter's collective behavior) related to the PoS trading.

The key to our analysis relies on the study of the random process N_t (the total volume of votes/stakes at time t), which is a time-homogeneous Markov chain. We develop some asymptotic results for this Markov chain, including large-deviation bounds (Theorem 3) and a fluid limit (Proposition 5).

In the remainder of this paper there are two main sections. Section 2 studies the $\texttt{Poly}(\alpha)$ voting model from a general perspective, focusing on the associated stochastic processes, such as N_t, voter shares, and voting powers; their long-term behavior and limits, some of which are further characterized by concentration inequalities or large-deviation bounds. Section 3 concerns two aspects of the $\texttt{Poly}(\alpha)$ voting rule that are more closely associated with the application of PoS in cryptocurrency: (a) the evolution of bidder shares over time and the phase transition phenomenon mentioned above; (b) the issue of incentive and risk-sensitivity when trading is allowed. Concluding remarks and suggestions for further research are collected in Section 4.

2. The $\texttt{Poly}(\alpha)$ Voting Model

In this section, we develop a formal model for the $\texttt{Poly}(\alpha)$ voting rule, focusing on the stochastic processes associated with the model, their properties and limiting behavior.

First, here is a list of some of the common notation used throughout the paper.

- $\mathbb{N}_+$ denotes the set of positive integers, and $\mathbb{R}$ denotes the set of real numbers.
- $\overset{d}{=}$ denotes equal in distribution, and $\overset{d}{\longrightarrow}$ denotes convergence in distribution.
- $a = \mathcal{O}(b)$ means $\frac{a}{b}$ is bounded from above as $b \to \infty$; $a = \Theta(b)$ means $\frac{a}{b}$ is bounded from below and above as $b \to \infty$; and $a = o(b)$ or $b \gg a$ means $\frac{a}{b}$ decays towards zero as $b \to \infty$.
- $d_W(\mu, \nu)$ denotes the 1-Wasserstein distance between two probability distributions μ and ν. Refer to (Villani, 2009, Chapter 6).

We use C, C', C'' etc to denote generic constants (which may change from line to line).

The voters, referred to as *bidders* below, are the participants in the decentralized system, where they engage in rounds of bidding following a pre-specified voting rule (the "consensus protocol") so as to win more votes, or *stakes*. (The PoS protocol described in the Introduction provides a concrete instance to motivate the model here.) Let $K \in \mathbb{N}_+$ be the total number of bidders, which will stay fixed throughout the paper; and let $[K] := \{1, \dots, K\}$ denote the set of all bidders.

Time is discrete, indexed by $t = 0, 1, 2, \dots$, and corresponds to the rounds of bidding mentioned above. Bidder k initially owns $n_{k,0}$ stakes. Let $N := \sum_{k=1}^{K} n_{k,0}$ denote the total number of initial stakes owned by all K bidders. The term *bidder share* refers to the fraction

of stakes each bidder owns. So the initial bidder shares ($\pi_{k,0}$, $k \in [K]$) are given by

$$\pi_{k,0} := \frac{n_{k,0}}{N}, \quad k \in [K]. \tag{2.1}$$

Similarly, $n_{k,t}$ denotes the number of stakes owned by bidder k at time $t \in \mathbb{N}_+$, and the corresponding share is

$$\pi_{k,t} := \frac{n_{k,t}}{N_t}, \quad k \in [K], \quad \text{with } N_t := \sum_{k=1}^{K} n_{k,t}. \tag{2.2}$$

Here N_t is the total number of stakes at time t, and thus $N_0 = N$. (We shall often refer to N_t as the "volume of stakes" or, simply, "volume.") Clearly, for each $t \geq 0$, ($\pi_{k,t}$, $k \in [K]$) forms a probability distribution on $[K]$.

In each period t, a single stake (or "reward") is distributed as follows: each bidder k receives the reward with probability

$$\theta_{k,t} := \frac{n_{k,t}}{N_t^{1+\alpha}} = \frac{\pi_{k,t}}{N_t^{\alpha}}, \tag{2.3}$$

and receives nothing with probability $1 - \theta_{k,t}$. Clearly, $\theta_{k,t}$ is bidder k's *reward rate*, as $1/\theta_{k,t}$ is the average number of rounds for bidder k to win an additional unit of stake. To the extent the reward is coupled with the voting mechanism outlined above, $\theta_{k,t}$ can also be viewed as bidder k's voting power at time t. (Below, we use the terms "reward rate" and "voting power" interchangeably if there is no ambiguity.) When $\alpha = 0$, the voting power $\theta_{k,t}$ coincides with the bidder share $\pi_{k,t}$, which is the Pólya urn framework in Roşu and Saleh (2021), and in Tang (2022).

Let $S_{k,t}$ be the random event that bidder k receives one unit of reward in period t. Thus, the number of stakes owned by each bidder evolves as follows,

$$n_{k,t} = n_{k,t-1} + 1_{S_{k,t}}, \quad k \in [K]; \tag{2.4}$$

or simply,

$$n_{k,t} = \begin{cases} n_{k,t-1} & \text{with probability} \quad 1 - \theta_{k,t-1}, \\ n_{k,t-1} + 1 & \text{with probability} \quad \theta_{k,t-1}. \end{cases} \tag{2.5}$$

Accordingly, the total number of stakes N_t evolves as follows, taking into account $\sum_{k=1}^{K} n_{k,t} = N_t$,

$$N_t = \begin{cases} N_{t-1} & \text{with probability} \quad 1 - 1/N_{t-1}^{\alpha}, \\ N_{t-1} + 1 & \text{with probability} \quad 1/N_{t-1}^{\alpha}. \end{cases} \tag{2.6}$$

The counting process (N_t, $t \geq 0$) specified in (2.6) evolves as a time-homogeneous Markov chain on $\{N, N+1, \ldots\}$, in contrast with the Pólya urn (with a constant reward) in which N_t grows deterministically and linearly in t. As we will see in the next subsection, the `Poly`(α) voting rule slows down the distribution of rewards, so the volume of stakes grows sublinearly. This is consistent with the volume growth in many cryptocurrencies such as Bitcoin and Ethereum.

Let $\boldsymbol{n}_t = (n_{1,t}, \ldots, n_{K,t})$ be the vector of bidder stakes at time t. An alternative (and useful) characterization of ($\boldsymbol{n}_t$, $t \geq 0$) is given as follows.

Proposition 1. *Let* (L_t, $t \geq 0$) *be a counting process with arrivals occurring at* $0 = T_0 < T_1 < \cdots$, *such that the inter-arrival times are independent, with* $T_{k+1} - T_k$, *for every* $k \geq 0$,

following a geometric distribution with success probability parameter $(N+k)^{-\alpha}$. *Define the process* $(\boldsymbol{l}_t,\, t \geq 0)$ *by*

$$\boldsymbol{l}_t = \boldsymbol{l}_{T_k} \quad \text{for } T_k \leq t < T_{k+1},$$

where $(\boldsymbol{l}_{T_k},\, k \geq 0)$ *is a copy of the Pólya urn process with* K *colors and* N *initial balls. Then, we have* $(\boldsymbol{n}_t,\, t \geq 0) \overset{d}{=} (\boldsymbol{l}_t,\, t \geq 0)$, *where* $\boldsymbol{n}_t$ *is the process of bidder stakes defined above.*

Proof. It is clear from the dynamics in (2.6) that the two counting processes $(N_t,\, t \geq 0)$ and $(L_t,\, t \geq 0)$ have the same distribution. Given $\boldsymbol{n}_t$, the probability that the next (unit of) stake goes to bidder k is $\frac{n_{k,t}}{N_t^{1+\alpha}} / \frac{1}{N_t^{\alpha}} = \frac{n_{k,t}}{N_t}$ by the craps principle. The connection to the Pólya urn process with K colors (voters) and N initial balls (stakes) is obvious. □

The above proposition implies that the Pólya urn is embedded in the process of stakes $(\boldsymbol{n}_t,\, t \geq 0)$ through a random time change $(N_t,\, t \geq 0)$. This fact will be used below to study the long-time behavior of bidder shares and reward rates in §2.2. But we first study in the next subsection how the issuance of rewards is slowed down under the $\texttt{Poly}(\alpha)$ voting rule.

2.1. **The volume** $(N_t,\, t \geq 0)$**.** Let $\mathcal{F}_t$ be the filtration generated by the random events $(S_{k,r} : k \in [K],\, r \leq t)$.

Proposition 2 (Long-time behavior of N_t)**.** *Under the* $\texttt{Poly}(\alpha)$ *voting rule, the following results hold:*

(i) The process $(N_t,\, t \geq 0)$ *is an* $\mathcal{F}_t$*-sub-martingale, and its compensator is*

$$A_t = \sum_{k \leq t-1} N_k^{-\alpha} \quad \text{for } t \geq 1.$$

(ii) There is the convergence in probability:

$$\frac{N_t^{1+\alpha}}{t} \longrightarrow 1+\alpha \quad \text{as } t \to \infty. \tag{2.7}$$

Proof. (i) It suffices to note that $\mathsf{E}(N_{t+1} \,|\, \mathcal{F}_t) = N_t + N_t^{-\alpha}$, for all $t \geq 0$.

(ii) Apply the method of moments by computing $\mathsf{E}(N_t^{(1+\alpha)j})$ for all j. For $j=1$, we have by definition

$$\begin{aligned}\mathsf{E}(N_{t+1}^{1+\alpha} - N_t^{1+\alpha} \,|\, N_t = x) &= (1+x)^{1+\alpha}\frac{1}{x^{\alpha}} + x^{1+\alpha}\left(1 - \frac{1}{x^{\alpha}}\right) - x^{1+\alpha} \\ &= 1 + \alpha + \mathcal{O}(x^{-1}) \quad \text{as } x \to \infty.\end{aligned}$$

It is clear that with probability one $N_t \to \infty$ as $t \to \infty$. As a result, $\mathsf{E}(N_{t+1}^{1+\alpha} - N_t^{1+\alpha}) \to 1+\alpha$ as $t \to \infty$ which yields

$$\mathsf{E}N_t^{1+\alpha} \sim (1+\alpha)t \quad \text{as } t \to \infty. \tag{2.8}$$

Next for $j=2$, we have

$$\mathsf{E}(N_{t+1}^{2(1+\alpha)} - N_t^{2(1+\alpha)} \,|\, N_t = x) = 2(1+\alpha)x^{1+\alpha} + \mathcal{O}(x^{\alpha}) \quad \text{as } x \to \infty.$$

Thus, $\mathsf{E}(N_{t+1}^{2(1+\alpha)} - N_t^{2(1+\alpha)}) = (2(1+\alpha) + o(1))\,\mathsf{E}N_t^{1+\alpha} \sim 2(1+\alpha)^2 t$ by (2.8). Then we get $\mathsf{E}(N_t^{2(1+\alpha)}) \sim (1+\alpha)^2 t^2$ as $t \to \infty$. We proceed by induction. Assuming that $\mathsf{E}(N_t^{j(1+\alpha)}) \sim$

$(1+\alpha)^j t^j$ as $t \to \infty$, we get

$$\mathsf{E}(N_t^{(1+\alpha)(j+1)} - N_t^{(1+\alpha)(j+1)}) = ((j+1)(1+\alpha) + o(1))\mathsf{E}(N_t^{j(1+\alpha)}) \sim (j+1)(1+\alpha)^{j+1} t^j,$$

which implies that $\mathsf{E}(N_t^{(1+\alpha)(j+1)}) \sim (1+\alpha)^{j+1} t^{j+1}$ as $t \to \infty$. Thus, we have

$$\mathsf{E}(N_t^{(1+\alpha)j}) \sim (1+\alpha)^j t^j \quad \text{as } t \to \infty, \quad j = 1, 2, \dots$$

By the method of moments (see e.g. (Billingsley, 1995, Section 30)), $N_t^{(1+\alpha)}/t$ converges in distribution, and thus in probability to $1+\alpha$. □

The proposition gives the growth rate of the volume of stakes: N_t grows as $((\alpha+1)t)^{\frac{1}{1+\alpha}}$ as $t \to \infty$. Part (i) suggests that $N_t \sim \sum_{k \le t-1} N_k^{-\alpha}$, which is consistent with the limit in (2.7). When $\alpha = 0$, N_t follows the (deterministic) linear growth of the Pólya urn model (with a constant reward). For $\alpha = 1$, N_t grows as $\sqrt{t}$.

Even more important is the question, how does N_t "fluctuate" around its growth trajectory $((\alpha+1)t)^{\frac{1}{1+\alpha}}$; specifically, how to establish large-deviation bounds on N_t? This is addressed in the next theorem, along with a corollary that confirms N_t's concentration around its growth trajectory, for large t.

Theorem 3 (Large deviations for N_t)**.** *Define a function $f_\alpha(\cdot)$,*

$$f_\alpha : \lambda \in (0, \infty) \ \mapsto \ (1+\alpha)\lambda \log \lambda - (1+\alpha)\lambda + \frac{1}{\lambda^\alpha} \in \mathbb{R}.$$

Let $\lambda_-(\alpha) < \lambda_+(\alpha)$ be the two roots of $f_\alpha(\cdot)$ on $(-\infty, \infty)$. Under the `Poly`(α) *voting rule, the following results hold:*

(i) For each $\lambda < \lambda_-(\alpha)$, and for any $\varepsilon > 0$,

$$\mathsf{P}(N_t < \lambda t^{\frac{1}{1+\alpha}}) \le \exp\left(-(1-\varepsilon) f_\alpha(\lambda)\, t^{\frac{1}{1+\alpha}}\right) \quad \text{as } t \to \infty. \tag{2.9}$$

(ii) For each $\lambda > \lambda_+(\alpha)$, and for any $\varepsilon > 0$,

$$\mathsf{P}(N_t > \lambda t^{\frac{1}{1+\alpha}}) \le \exp\left(-(1-\varepsilon) f_\alpha(\lambda)\, t^{\frac{1}{1+\alpha}}\right) \quad \text{as } t \to \infty. \tag{2.10}$$

Proof. Without loss of generality, assume that $N_0 = 1$. Note that $(N_t,\, t \ge 0)$ has increments $\{0, 1\}$, so there are $\binom{t}{k}$ paths ending at $N_t = k+1$ (one has to choose k upward steps "1" out of t steps). Moreover, the probability of each path ending at $N_t = k+1$ is upper bounded by

$$\frac{1}{(k!)^\alpha} \left(1 - \frac{1}{(k+1)^\alpha}\right)^{t-k},$$

since the k upward "1" steps contribute $1/k!$, and the remaining $t-k$ flat "0" steps have at most probability $\left(1 - \frac{1}{(k+1)^\alpha}\right)^{t-k}$. Thus,

$$\mathsf{P}(N_t \le m+1) \le \sum_{k \le m} a_k \quad \text{and} \quad \mathsf{P}(N_t > m) \le \sum_{k \ge m} a_k$$

where

$$a_k := \binom{t}{k} \frac{1}{(k!)^\alpha} \left(1 - \frac{1}{(k+1)^\alpha}\right)^{t-k}. \tag{2.11}$$

Standard analysis shows that there are $0 < k_1 < k_2$ such that a_k is nondecreasing on $[1, k_1)$ and $[k_2, t)$. As we will see, $k_1 \sim \lambda_-(\alpha)\, t^{\frac{1}{1+\alpha}}$ and $k_2 \sim \lambda_+(\alpha)\, t^{\frac{1}{1+\alpha}}$ as $t \to \infty$. The idea is to study the term a_k with $k = \lambda t^{\frac{1}{1+\alpha}}$ for $\lambda > 0$ as $t \to \infty$. By Stirling's formula,

$$\binom{t}{\lambda t^{\frac{1}{1+\alpha}}} \sim \frac{1}{\sqrt{2\pi\lambda}} t^{-\frac{1}{2(1+\alpha)}} \exp\left(\frac{\lambda\alpha}{1+\alpha} t^{\frac{1}{1+\alpha}} \log t + (\lambda - \lambda\log\lambda) t^{\frac{1}{1+\alpha}} + o\left(t^{\frac{1}{1+\alpha}}\right)\right),$$

$$\left(\lambda t^{\frac{1}{1+\alpha}}\right)! \sim \sqrt{2\pi\lambda}\, t^{\frac{1}{2(1+\alpha)}} \exp\left(\frac{\lambda}{1+\alpha} t^{\frac{1}{1+\alpha}} \log t + (\lambda\log\lambda - \lambda) t^{\frac{1}{1+\alpha}} \right),$$

and

$$\left(1 - \frac{1}{\lambda^\alpha t^{\frac{\alpha}{1+\alpha}}}\right)^{t - \lambda t^{\frac{1}{1+\alpha}}} = \exp\left(-\frac{t^{\frac{1}{1+\alpha}}}{\lambda^\alpha} + o\left(t^{\frac{1}{1+\alpha}}\right)\right).$$

Combining the above estimates yields

$$a_{\lambda\sqrt{t}} \sim (2\pi\lambda)^{-\frac{1+\alpha}{2}} t^{-\frac{1}{2}} \exp\left(-f_\alpha(\lambda)\, t^{\frac{1}{1+\alpha}}\right). \tag{2.12}$$

Note that $f'_\alpha(\lambda) = (1+\alpha)\log\lambda - \alpha\lambda^{-1-\alpha}$, which is increasing from $-\infty$ to ∞ on $[0,\infty)$. The unique stationary point of f_α on $[0,\infty)$ is achieved at λ_* such that $\lambda_*^{\alpha+1}\log\lambda_* = \frac{\alpha}{1+\alpha}$, so it is clear that $\lambda_* > 1$. We have

$$f_\alpha(\lambda_*) = (\alpha-1)\lambda_*^{-\alpha} - (1+\alpha)\lambda_* < 0.$$

Thus, the function $\lambda \to f_\alpha(\lambda)$ has two roots $\lambda_-(\alpha) < \lambda_+(\alpha)$ on $[0,\infty)$, and $f_\alpha > 0$ on $(0, \lambda_-(\alpha)) \cup (\lambda_+(\alpha), \infty)$. As a result, for each $\lambda < \lambda_-(\alpha)$ we have

$$\mathsf{P}(N_t < \lambda t^{\frac{1}{1+\alpha}}) \le C_\lambda t^{-\frac{1}{2}+\frac{1}{1+\alpha}} \exp\left(-f_\alpha(\lambda)\, t^{\frac{1}{1+\alpha}}\right) \le \exp\left(-(1-\varepsilon) f_\alpha(\lambda)\, t^{\frac{1}{1+\alpha}}\right) \text{ as } t \to \infty,$$

and each $\lambda > \lambda_+(\alpha)$ we have

$$\mathsf{P}(N_t > \lambda t^{\frac{1}{1+\alpha}}) \le C'_\lambda t^{\frac{1}{2}} \exp\left(-f_\alpha(\lambda)\, t^{\frac{1}{1+\alpha}}\right) \le \exp\left(-(1-\varepsilon) f_\alpha(\lambda)\, t^{\frac{1}{1+\alpha}}\right) \text{ as } t \to \infty.$$

□

The theorem above gives exponential deviation bounds for N_t when it is either sufficiently small (below $\lambda_-(\alpha) t^{\frac{1}{1+\alpha}}$), or sufficiently large (above $\lambda_+(\alpha) t^{\frac{1}{1+\alpha}}$). Note that there is a gap between the two bounding curves, since for each $\alpha > 0$, $\lambda_-(\alpha) < (1+\alpha)^{\frac{1}{1+\alpha}} < \lambda_+(\alpha)$. (For instance, for $\alpha = 1$, $\lambda_-(1) \approx 0.56 < \sqrt{2} < 2.51 \approx \lambda_+(1)$.) The gap is due to the combinatorial estimates in our proof, which may very well be improved. Refer to the Appendix 5.1 for a numerical procedure that shrinks the gap.

As a corollary, the volume of stakes N_t concentrates around $((1+\alpha)t)^{\frac{1}{1+\alpha}}$ for large t.

Corollary 4. *Under the* `Poly`(α) *voting rule, we have, for each* $\delta > 0$,

$$\mathsf{P}\left(|N_t - ((1+\alpha)t)^{\frac{1}{1+\alpha}}| > \delta t^{\frac{1}{1+\alpha}}\right) = \mathcal{O}(t^{-\frac{1}{1+\alpha}}) \quad \text{as } t \to \infty. \tag{2.13}$$

Proof. Note that $N_0 \le N_t \le t + N_0$, and by Theorem 3, we get $\lambda_1 t^{\frac{1}{1+\alpha}} \le N_t \le \lambda_2 t^{\frac{1}{1+\alpha}}$ with probability $1 - \exp(-C\lambda t^{\frac{1}{1+\alpha}})$ for some $\lambda_1, \lambda_2, C > 0$. Thus,

$$\mathbb{E}N_t^{-1} = \mathcal{O}(t^{-\frac{1}{1+\alpha}}) \quad \text{and} \quad \mathbb{E}N_t^{\alpha} = \mathcal{O}(t^{\frac{\alpha}{1+\alpha}}).$$

According to the proof of Proposition 2 (ii), we have

$$\mathbb{E}(N_{t+1}^{1+\alpha}-N_t^{1+\alpha}) = 1+\alpha+\mathcal{O}(\mathbb{E}N_t^{-1}) \quad \text{and} \quad \mathbb{E}(N_{t+1}^{2(1+\alpha)}-N_t^{2(1+\alpha)}) = 2(1+\alpha)\mathbb{E}N_t^{1+\alpha}+\mathcal{O}(\mathbb{E}N_t^{\alpha}).$$

Therefore, $\mathbb{E}N_t^{1+\alpha} = (1+\alpha)t + \mathcal{O}(t^{\frac{\alpha}{1+\alpha}})$ and $\mathbb{E}(N_t^{2(1+\alpha)}) = (1+\alpha)^2t^2 + \mathcal{O}(t^{\frac{1+2\alpha}{1+\alpha}})$, which implies $\mathsf{Var}(N_t^{1+\alpha}) = \mathcal{O}(t^{\frac{1+2\alpha}{1+\alpha}})$. Hence,

$$\mathsf{P}\left(|N_t^{1+\alpha} - (1+\alpha)t| > \delta t\right) = \mathcal{O}(t^{-\frac{1}{1+\alpha}}) \quad \text{for } t \to \infty.$$

Taking $\lambda < \lambda_-(\alpha)$, we have

$$\begin{aligned}
&\mathsf{P}\left(|N_t - ((1+\alpha)t)^{\frac{1}{1+\alpha}}| > \delta t^{\frac{1}{1+\alpha}}\right) \\
&\le \mathsf{P}\left(N_t < \lambda t^{\frac{1}{1+\alpha}}\right) + \mathsf{P}\left(|N_t - ((1+\alpha)t)^{\frac{1}{1+\alpha}}| > \delta t^{\frac{1}{1+\alpha}},\, N_t \ge \lambda t^{\frac{1}{1+\alpha}}\right) \\
&\le \exp(-C't^{\frac{1}{1+\alpha}}) + \mathsf{P}(|N_t^{1+\alpha} - (1+\alpha)t| > C''t),
\end{aligned}$$

for some $C', C'' > 0$ (depending on α, δ, λ). Combining the above estimates yield (2.13). □

Recall that the process $(N_t,\, t \ge 0)$ is a time-homogenous Markov chain. The path properties of a general time-homogenous Markov chain $(Z_t,\, t \ge 0)$ has long been studied since the work of Lamperti (1960, 1962, 1963). The basic idea is to study the recurrence, or transience of $(Z_t,\, t \ge 0)$ based on

$$m_1(x) = \mathsf{E}(Z_{t+1} - Z_t \,|\, Z_t = x) \quad \text{and} \quad m_2(x) = \mathsf{E}((Z_{t+1} - Z_t)^2 \,|\, Z_t = x).$$

For instance, if $\limsup_{x\to\infty} 2xm_1(x) + m_2(x) \le 0$ then $(Z_t,\, t \ge 0)$ is recurrent; and if $\liminf_{x\to\infty} 2xm_1(x) + m_2(x) > 0$ then $(Z_t,\, t \ge 0)$ is transient. The regime corresponding to $m_1(x) = o(1)$ is called the *Markov chain with asymptotic zero drift*, and features active research (see e.g. Denisov et al. (2016); Menshikov et al. (2017)). Specializing to the process $(N_t,\, t \ge 0)$, we have

$$m_1(x) = m_2(x) = \frac{1}{x^\alpha}.$$

Interestingly, there seem to be few results on the Lamperti's problem where both $m_1(x)$ and $m_2(x)$ decreases to zero as $x \to \infty$, except that $(N_t,\, t \ge 0)$ is transient. Observe that $(N_t,\, t \ge 0)$ is nondecreasing, and

$$\mathsf{Var}(N_{t+1} \,|\, N_t = x) = \left(1 - \frac{1}{x^\alpha}\right)^2 \frac{1}{x^\alpha} + \left(-\frac{1}{x^\alpha}\right)^2 \left(1 - \frac{1}{x^\alpha}\right),$$

where the "upward" contribution $\left(1 - \frac{1}{x^\alpha}\right)^2 \frac{1}{x^\alpha}$ is larger than the "downward" counterpart $\frac{1}{x^{2\alpha}}\left(1 - \frac{1}{x^\alpha}\right)$ as $x \to \infty$. In the similar spirit to Lamperti (1962), the asymptotic growth (2.7) hinges on a degenerate fluid approximation of the process $(N_t,\, t \ge 0)$, as stated in the following proposition.

Proposition 5 (Fluid limit of N_t)**.** *Under the* `Poly`(α) *voting rule, we have*

$$\left(\frac{N_{nu}}{n^{\frac{1}{1+\alpha}}},\, u \ge 0\right) \xrightarrow{d} (X_u,\, u \ge 0) \quad \text{as } n \to \infty \text{ in } \mathcal{C}[0,\infty), \tag{2.14}$$

where N_s for non-integer s is defined by the linear interpolation of the chain $(N_t,\, t \ge 0)$, and $X_u = ((1+\alpha)u)^{\frac{1}{1+\alpha}}$, $u \ge 0$ is the solution to the ordinary differential equation $dX_u = X_u^{-\alpha}dt$ with $X_0 = 0$.

Proof. Fix $T > 0$. It suffices to prove the weak convergence (2.14) on $[0, T]$. By Proposition 2 (ii), there is the convergence in probability $N_{[nT]}/n^{\frac{1}{1+\alpha}} \to X_T$ as $n \to \infty$. Given $\varepsilon > 0$, there is $n(\varepsilon) > 0$ such that for any $n > n(\varepsilon)$,

$$\mathsf{P}\left(N_{[nT]}/n^{\frac{1}{1+\alpha}} < 2X_T\right) > 1 - \varepsilon.$$

Let $K(\varepsilon) := \max(2X_T, \max_{n \le n(\varepsilon)}(N_0 + [nT])/n^{\frac{1}{1+\alpha}})$. We have $\mathsf{P}(N_{[nT]}/n^{\frac{1}{1+\alpha}} < K(\varepsilon)) > 1 - \varepsilon$ for each $n \in \mathbb{N}_+$. Note that for each $n \in \mathbb{N}_+$, the process $N^{n,T} := \left(N_{[nT]}/n^{\frac{1}{1+\alpha}},\, 0 \le t \le T\right)$ is nondecreasing. Thus,

$$\mathsf{P}\left(N^{n,T} \in [0, T] \times [0, K(\varepsilon)]\right) > 1 - \varepsilon.$$

So the sequence of processes $(N^{n,T},\, n \in \mathbb{N}_+)$ is tight. Moreover, for each $t \in [0, T]$, $N_{[nt]}/n^{\frac{1}{1+\alpha}}$ converges in probability to X_t as $n \to \infty$. Then for $0 \le t_1 < \cdots < t_k$, the vector $(N_{[nt_1]}/n^{\frac{1}{1+\alpha}}, \cdots, N_{[nt_k]}/n^{\frac{1}{1+\alpha}})$ converges in probability to $(X_{t_1}, \cdots, X_{t_k})$, i.e. the convergence in finite-dimensional distributions. The weak convergence follows readily from the tightness and the convergence in finite-dimensional distributions (see e.g. (Billingsley, 1999, Chapter 2)). □

Note that the fluid limit in the proposition is different from the fluid limit in the literature of stochastic networks, where it usually takes the form of a FSLLN (functional strong law of large numbers) concerning a renewal process and the associated counting process. In that setting, the convergence (to a deterministic function of time) is stronger – the almost sure convergence, and uniformly on $[0, T]$. Refer to (Chen and Yao, 2001, §6.1). Here, the process $(N_t,\, t \ge 0)$ is non-renewal, thus the FSLLN limit does not apply; yet there is still the *weak* convergence, and the limit is still a deterministic function of time $(X_u,\, u \ge 0)$, explicitly characterized above. Another notable point is, in the FSLLN setting, both time and space are scaled by the same scaling factor n whereas in (2.14) the time scaling remains the same, and the space scaling is by $n^{\frac{1}{1+\alpha}}$. But this is only because $(N_t,\, t \ge 0)$ grows in the order of $t^{\frac{1}{1+\alpha}}$ (Proposition 2 (ii)), whereas a renewal (counting) process grows linearly in t.

2.2. **Bidder shares and voting powers.** Here we study the evolution and the long-time behavior of $(\pi_{k,t},\, k \in [K])$ and $(\theta_{k,t},\, k \in [K])$. Recall that the Dirichlet distribution with parameters $(a_1, ..., a_K)$, which we denote by $\mathrm{Dir}(a_1, ..., a_K)$, has support on the standard simplex $\{(x_1, \ldots, x_K) \in \mathbb{R}_+^K : \sum_{k=1}^K x_k = 1\}$ and has density

$$f(x_1, \ldots x_K) = \frac{\Gamma(\sum_{k=1}^K a_k)}{\prod_{k=1}^K \Gamma(a_k)} \prod_{k=1}^K x_k^{a_k - 1},$$

where $\Gamma(z) = \int_0^\infty x^{z-1} e^{-x} dx$ is the Gamma function. For $K = 2$, the Dirichlet distribution reduces to the beta distribution, denoted as $\mathrm{Beta}(a_1, a_2)$. It is easily seen that if $(x_1, \ldots, x_K) \overset{d}{=} \mathrm{Dir}(a_1, \ldots, a_K)$ then for each $k \in [K]$, $x_k \overset{d}{=} \mathrm{Beta}(a_k, \sum_{j \ne k} a_j)$.

Theorem 6 (Long-time behavior)**.** *Under the* `Poly`(α) *voting rule, we have the following limiting distributions.*

(i) Bidder shares: the process $(\pi_{k,t},\, t \geq 0)$ *is an* $\mathcal{F}_t$*-martingale, and with probability one,*

$$(\pi_{1,t}, \ldots, \pi_{K,t}) \longrightarrow (\pi_{1,\infty}, \ldots, \pi_{K,\infty}) \quad \text{as } t \to \infty, \tag{2.15}$$

where $(\pi_{1,\infty}, \ldots, \pi_{K,\infty}) \stackrel{d}{=} \mathit{Dir}(n_{1,0}, \ldots, n_{K,0})$*. Moreover, for each* $k \in [K]$*,*

$$d_W\left(\pi_{k,t},\, \mathit{Beta}\left(n_{k,0},\, N - n_{k,0}\right)\right) = \mathcal{O}(t^{-\frac{1}{1+\alpha}}) \quad \text{as } t \to \infty. \tag{2.16}$$

(ii) Voting powers: the process $(\theta_{k,t},\, t \geq 0)$ *is an* $\mathcal{F}_t$*-super-martingale, and for* $\alpha > 0$*, with probability one,* $\theta_{k,t} \to 0$ *as* $t \to \infty$ *for each* $k \in [K]$*. Moreover, for each* $k \in [K]$*,*

$$(1+\alpha)^{\frac{\alpha}{1+\alpha}} t^{\frac{\alpha}{1+\alpha}} \theta_{k,t} \stackrel{d}{\longrightarrow} \mathit{Beta}\left(n_{k,0},\, N - n_{k,0}\right) \quad \text{as } t \to \infty. \tag{2.17}$$

Proof. (i) By (2.2) and (2.5), it is easily seen that for each $k \in [K]$ and $t \geq 0$,

$$\mathsf{E}(\pi_{k,t+1} \,|\, \mathcal{F}_t) = \frac{n_{k,t}}{N_t}\left(1 - \frac{1}{N_t^\alpha}\right) + \frac{n_{k,t}}{N_t+1}\,\frac{N_t - n_{k,t}}{N_t^{1+\alpha}} + \frac{n_{k,t}+1}{N_t+1}\,\frac{n_{k,t}}{N_t^{1+\alpha}}. \tag{2.18}$$

Recognizing the first term on the right side of (2.18), $\frac{n_{k,t}}{N_t} = \pi_{k,t}$, while all other terms sum up to zero, we conclude that $(\pi_{k,t},\, t \geq 0)$ is an $\mathcal{F}_t$-martingale. The convergence in (2.15) follows from the martingale convergence theorem (see e.g. (Durrett, 2019, Section 4.2)). By Proposition 1, $(\boldsymbol{n}_t,\, t \geq 0)$ is a time-changed Pólya urn. So the limiting shares $(\pi_{1,\infty}, \ldots, \pi_{K,\infty})$ have the same distribution as that of the Pólya urn, which is $\mathrm{Dir}(n_{1,0}, \ldots, n_{K,0})$.

Let $(\boldsymbol{n}_t^\dagger,\, t \geq 0)$ be the Pólya urn with $n_{k,0}^\dagger = n_{k,0}$, and $(\pi_{k,t}^\dagger, \ldots, \pi_{K,t}^\dagger)$ be the corresponding shares. Set $Z \stackrel{d}{=} \mathrm{Beta}\left(n_{k,0}, N - n_{k,0}\right)$. By Goldstein and Reinert (2013), we have for each $k \in [K]$,

$$d_W\left(\pi_{k,t}^\dagger,\, Z\right) = \mathcal{O}(t^{-1}) \quad \text{as } t \to \infty. \tag{2.19}$$

Taking $\lambda < \lambda_-(\alpha)$, we get

$$\begin{aligned}
d_W(\pi_{k,t},\, Z) &\leq \mathsf{P}(N_t < \lambda t^{\frac{1}{1+\alpha}}) + d_W\left(\pi_{k,t} 1_{N_t \geq \lambda t^{\frac{1}{1+\alpha}}},\, Z\right) \\
&\leq C\mathsf{P}(N_t < \lambda t^{\frac{1}{1+\alpha}}) + \sum_{s \geq \lambda t^{\frac{1}{1+\alpha}}} d_W((\pi_{k,t} | N_t = s), Z)\mathsf{P}(N_t = s) \\
&= C\mathsf{P}(N_t < \lambda t^{\frac{1}{1+\alpha}}) + \sum_{s \geq \lambda t^{\frac{1}{1+\alpha}}} d_W(\pi_{k,s-N}^\dagger, Z)\mathsf{P}(N_t = s) \\
&\leq C\exp(-C' t^{\frac{1}{1+\alpha}}) + C'' t^{-\frac{1}{1+\alpha}},
\end{aligned}$$

where the last inequality follows from Theorem 3 and (2.19). This yields the bound (2.16).
(ii) Applying the same derivation as in (2.18) but to $\theta_{k,t}$ instead, we have

$$\mathsf{E}(\theta_{k,t+1} \,|\, \mathcal{F}_t) = \frac{\pi_{k,t}}{N_t^\alpha}\left(1 - \frac{1}{N_t^\alpha}\right) + \frac{N_t \pi_{k,t}}{(N_t+1)^{1+\alpha}} \cdot \frac{1 - \pi_{k,t}}{N_t^\alpha} + \frac{N_t \pi_{k,t} + 1}{(N_t+1)^{1+\alpha}} \cdot \frac{\pi_{k,t}}{N_t^\alpha}.$$

The last two terms add up to $\frac{\pi_{k,t}}{N_t^\alpha (N_t+1)^\alpha} = \frac{\theta_{k,t}}{(N_t+1)^\alpha}$. Thus, we have

$$\mathsf{E}(\theta_{k,t+1} \,|\, \mathcal{F}_t) = \theta_{k,t}\left(1 - \frac{1}{N_t^\alpha} + \frac{1}{(N_t+1)^\alpha}\right) \leq \theta_{k,t},$$

i.e. $(\theta_{k,t},\, t \geq 0)$ is an $\mathcal{F}_t$-super-martingale for each k. Recall that $N_t^\alpha \theta_{k,t} = \pi_{k,t}$ so $\theta_{k,t} \leq N_t^{-\alpha}$ which converges to 0 with probability one. By (i), $N_t^\alpha \theta_{k,t}$ converges almost surely, and hence in distribution to Z. By Proposition 2, $N_t/((1+\alpha)t)^{\frac{1}{1+\alpha}}$ converges in probability to 1. We then apply Slutsky's theorem to get the convergence in (2.17). □

Several remarks are in order. Part (i) of Theorem 6 shows that the bidder shares form a martingale, and converges to a Dirichlet distribution (independent of α). This should be expected from the fact that the underlying bidder stakes $(\boldsymbol{n}_t,\, t \geq 0)$ is a time-changed Pólya urn; refer to Proposition 1. What's more revealing is the Wasserstein bound in (2.16) between a bidder's share and its limit. In fact, a matching lower bound can also be established (which we leave to the interested reader). Thus, the convergence rate of the bidder shares is exactly of order $t^{-\frac{1}{1+\alpha}}$. (Also refer to Proposition 7 below for further discussions on the stability of the bidder shares when the initial stakes $N := N_0$ is large.)

Part (ii) of the theorem implies that each bidder's voting power decays to zero at rate $t^{-\frac{\alpha}{1+\alpha}}$. Or, equivalently, the reward rate is slowed down: it takes a time of order $\Theta(t^{\frac{\alpha}{1+\alpha}})$ for any bidder to be rewarded a new (unit of) stake. This enhances security, so that no bidder can manipulate or control the bidding/voting process; while the level of decentralization remains unchanged. This also means the principle of security in (1.1) becomes "easier" to hold at large time t, since (due to the network delay) $v \propto N_t^{-\alpha} \downarrow 0$ as $t \to \infty$. On the other hand, if the reward is associated with transaction validation (which does not need to be), then the time required to validate a new block becomes uncontrolled in the long run. A possible remedy is to dynamically tune the parameter α over time, as detailed in the Appendix 5.2.

3. Other Results, with PoS-Crypto Applications

In this section, we present more results associated with the $\texttt{Poly}(\alpha)$ model that are largely motivated by the application of PoS in cryptocurrency. There are two subsections: In the first one, §3.1, we study the the evolution of bidder shares when $N := N_0$, the volume of initial stakes, is large. In §3.2, we study the additional feature of allowing the bidders to trade stakes among themselves, focusing on the issue of trading incentive (or the lack thereof). We remark that the results in both subsections exhibit some type of phase transitions, and are independent of the parametric value of α, and in this sense, *universal*.

3.1. **Evolution of bidder shares and phase transitions.**

As explained in the introduction, one key feature of the $\texttt{Poly}(\alpha)$ model is that the reward rate, or the voting power (if the reward goes with validation work) $\theta_{k,t}$ of any bidder k is different from k's share $\pi_{k,t}$ of the total volume of stakes at t. We have seen from Theorem 6 (iii) that the reward rate or voting power is decreasing over time, which facilitates security. On the other hand, the evolution of the share $\pi_{k,t}$ over time, from its initial value $\pi_{k,0}$, in both absolute and relative terms, is an important issue for any individual bidder k.

In the classical Pólya urn setting, it is shown in Roşu and Saleh (2021) that for a large bidder with initial stake $n_{k,0} = \Theta(N)$, there is the stability in bidder share, in the sense that

$$\mathsf{P}(|\pi_{k,\infty} - \pi_{k,0}| > \varepsilon) \to 0 \quad \text{as } N \to \infty.$$

Furthermore, similar, albeit qualitatively different, results are revealed in Tang (2022), for small bidders (following the definition in part (ii) of the corollary below). Here we focus on

the ratio $\pi_{k,t}/\pi_{k,0}$. Since $\pi_{k,\infty} \overset{d}{=} \mathrm{Beta}(n_{k,0}, N - n_{k,0})$, the results in Tang (2022) hold. The following proposition is a refined version of (Tang, 2022, Theorem 2.1).

Proposition 7 (Phase transitions of $\pi_{k,t}$)**.** *Let $N_0 = N$ be the total number of initial stakes. Under the* `Poly`(α) *voting rule, we have*

(i) For $n_{k,0} = f(N)$ such that $f(N) \to \infty$ as $N \to \infty$ (i.e. $\pi_{k,0} \gg 1/N$), and for each $\varepsilon > 0$ sufficiently small and each $t \geq 1$ or $t = \infty$,

$$\mathsf{P}\left(\left|\frac{\pi_{k,t}}{\pi_{k,0}} - 1\right| > \varepsilon\right) \leq \frac{1}{\varepsilon^2 f(N)}, \tag{3.1}$$

which converges to 0 as $N \to \infty$.

(ii) For $n_{k,0} = \Theta(1)$ (i.e. $\pi_{k,0} = \Theta(1/N)$), there is the convergence in distribution

$$\pi_{k,\infty}/\pi_{k,0} \xrightarrow{d} \frac{1}{n_{k,0}}\gamma(n_{k,0}) \quad \text{as } N \to \infty, \tag{3.2}$$

where $\gamma(n_{k,0})$ is a Gamma random variable with density $x^{n_{k,0}-1}e^{-x}1_{x>0}/\Gamma(n_{k,0})$. Moreover, there is $C > 0$ (independent of t and N) such that

$$d_W\left(\frac{\pi_{k,t}}{\pi_{k,0}}, \frac{1}{n_{k,0}}\gamma(n_{k,0})\right) \leq C\left(N^3 t^{-\frac{1}{1+\alpha}} + \frac{1}{\sqrt{N}}\right). \tag{3.3}$$

Proof. (i) Conditioning on N_t and using the law of total variance, we get

$$\mathsf{Var}(\pi_{k,t}) = \frac{1 - N\mathsf{E}(N_t^{-1})}{N+1}\pi_{k,0}(1 - \pi_{k,0}).$$

It suffices to apply Chebyshev's inequality to get the bound (3.1).

(ii) Note that

$$\begin{aligned} d_W\left(\frac{\pi_{k,t}}{\pi_{k,0}}, \frac{1}{n_{k,0}}\gamma(n_{k,0})\right) &\leq d_W\left(\frac{\pi_{k,t}}{\pi_{k,0}}, \frac{1}{\pi_{k,0}}\mathrm{Beta}(n_{k,0}, N - n_{k,0})\right) \\ &\quad + d_W\left(\frac{1}{\pi_{k,0}}\mathrm{Beta}(n_{k,0}, N - n_{k,0}), \frac{1}{n_{k,0}}\gamma(n_{k,0})\right) \end{aligned}$$

A careful application of Goldstein and Reinert (2013) yields a refinement of (2.19): there is $C > 0$ such that $d_W(\pi_{k,t}^{\dagger}, Z) \leq \frac{CN^3}{t}$. Adapting the argument in Theorem 6 yields

$$d_W\left(\frac{\pi_{k,t}}{\pi_{k,0}}, \frac{1}{\pi_{k,0}}\mathrm{Beta}(n_{k,0}, N - n_{k,0})\right) \leq C' N^3 t^{-\frac{1}{1+\alpha}} \quad \text{for some } C' > 0. \tag{3.4}$$

Next we claim that

$$d_W\left(\frac{1}{\pi_{k,0}}\mathrm{Beta}(n_{k,0}, N - n_{k,0}), \frac{1}{n_{k,0}}\gamma(n_{k,0})\right) \leq \frac{C''}{\sqrt{N}} \quad \text{for some } C'' > 0, \tag{3.5}$$

which can be proved by elementary calculus. Here we provide a sketch of proof. Set $n_{k,0} = 1$ for simplicity. Let $X \sim \gamma(1)$, and let X' be the sum of $N - 1$ independent $\gamma(1)$ random variables, independent of X. By beta-gamma algebra, $\frac{X}{X+X'}$ has the same distribution as $\mathrm{Beta}(1, N-1)$. Thus,

$$d_W(N\,\mathrm{Beta}(1, N-1), \gamma(1)) \leq \mathbb{E}\left|\frac{NX}{X + X'} - X\right|. \tag{3.6}$$

By normal approximation, we have $\frac{X+X'}{N} = 1 + \frac{1}{\sqrt{N}}\mathcal{N}(0,1) + o(N^{-\frac{1}{2}})$ where $\mathcal{N}(0,1)$ is standard normal (see Rio (2009)). Injecting into (3.6) yields the desired bound. Finally, combining the estimates (3.4) and (3.5) gives the bound (3.3). □

The proposition reveals a phase transition in the stability of shares, and identifies large and small bidders, in terms of the size of their stakes, according to the categories in the two parts. A large bidder k is guaranteed to have stability, in the precise sense characterized in (3.1), that the share ratio $\pi_{k,t}/\pi_{k,0}$ concentrates at 1, and converges to 1 in probability when $N \to \infty$, for any $t \geq 1$ (including $t = \infty$). For small bidders, this is reversed: the concentration inequality in (3.1) becomes the anti-concentration inequality:

$$\mathsf{P}\left(\left|\frac{\pi_{k,\infty}}{\pi_{k,0}} - 1\right| > \varepsilon\right) > c \quad \text{for } c > 0 \text{ independent of } \varepsilon, \tag{3.7}$$

implying volatility. The Wasserstein bound (3.3) is new, and it indicates that the ratio $\pi_{k,t}/\pi_{k,0}$ approaches the limiting Gamma distribution with an $N^{-\frac{1}{2}}$ error for $t \geq N^{\frac{7}{2}(1+\alpha)}$. However, we do not know whether the "N^3 dependence" in (3.3) is tight so the ratio $\pi_{k,t}/\pi_{k,0}$ may mix at a faster rate.

3.2. **Participation and trading.** So far, we have not considered the possibility of allowing the bidders to trade stakes (among themselves). In the classical Pólya urn model ($\alpha = 0$), it is shown in Roşu and Saleh (2021) that under certain conditions (which enforce some notion of "risk neutrality"), there will be no incentive for any bidder to trade. Here, we extend that to the `Poly`(α) model, allowing α to take any non-negative values. Furthermore, we allow a bidder-dependent risk-sensitivity (or risk-averse) parameter δ_k, and study the issue of incentive as it relates to δ_k.

In the new setting of allowing trading, we need to modify the problem formulation presented at the beginning of §2. First, for each $k \in [K]$, let $\nu_{k,t}$ be the number of stakes that bidder k will trade at time t. Then, instead of (2.4), the number of stakes $n_{k,t}$ evolves as

$$n_{k,t} = \underbrace{n_{k,t-1} + 1_{S_{k,t}}}_{n'_{k,t}} + \nu_{k,t}, \tag{3.8}$$

i.e. $n'_{k,t}$ denotes the number of stakes bidder k owns in between time $t-1$ and t, excluding those traded in period t.

Note that $\nu_{k,t}$ will be up to bidder k to decide, as opposed to the random event $S_{k,t}$ which is exogenous; in particular, $\nu_{k,t}$ can be negative (as well as positive or zero). We will elaborate more on this below, but note that $\nu_{k,t}$ will be constrained such that after the updating in (3.8) $n_{k,t}$ will remain nonnegative.

Let $\{P_t,\, t \geq 0\}$ be the price process of each (unit of) stake, which is a stochastic process assumed to be independent of the randomness induced by the `Poly`(α) voting rule (specifically, the process $\{S_{k,t}\}$). Hence, we augment the filtration $\{\mathcal{F}_t\}_{t\geq 0}$ with that of the exogenous price process $\{P_t,\, t \geq 0\}$ to a new filtration denoted $\{\mathcal{G}_t\}_{t\geq 0}$. Note that the price process P_t is also assumed as exogenous in Roşu and Saleh (2021). This assumption need not be so far off, as the crypto's price tends to be affected by market shocks (such as macroeconomics, geopolitics, breaking news, etc) much more than by trading activities. So, here we isolate the price of each stake from any bidder's trading impact.

Let $b_{k,t}$ denote (units of) the risk-free asset that bidder k holds at time t, and $r_{\text{free}} > 0$ the risk-free (interest) rate. (Here, the risk-free asset is naturally the one that underlies the above price process.) As we are mainly concerned with the effect of exchanging stakes to each individual, we allow bidders to trade stakes only internally among themselves, but not risk-free assets between them. Hence, each bidder has to trade risk-free asset with a third-party instead of trading that with another bidder.

The decision for each bidder k at t is hence a tuple, $(\nu_{k,t}, b_{k,t})$. Moreover, there is a terminal time, denoted $T_k \in \mathbb{N}_+$ (i.e., $T_k \geq 1$ is integer valued), by which time bidder k has to sell all assets, including both any risk-free asset and any stakes owned at that time, and leave the system. T_k can either be deterministic or random. In the latter case, assume it has a finite expectation, and is either adapted to $\{\mathcal{G}_t\}_{t\geq 0}$, or independent of all other randomness (in which case augment $\{\mathcal{G}_t\}$ accordingly). We also allow bidder k to leave the system and liquidate prior to T_k at a stopping time τ_k relative to $\{\mathcal{G}_t\}_{t\geq 0}$. Thus, bidder k will also decide at which time τ_k to stop and exit. To simplify the notation, we abuse τ_k for $\tau_k \wedge T_k$, the minimum of τ_k and T_k.

Let $c_{k,t}$ denote the (free) cash flow (or, "consumption") of bidder k at time t, i.e.,

$$c_{k,t} = (1 + r_{\text{free}})b_{k,t-1} - b_{k,t} - \nu_{k,t}P_t, \quad \forall 1 \leq t < \tau_k; \tag{C1}$$

with

$$b_{k,0} = 0,\ b_{k,t} \geq 0, \quad 0 \leq n_{k,t} = n'_{k,t} + \nu_{k,t} \leq N_t, \quad \forall 1 \leq t < \tau_k; \tag{C2}$$

and

$$c_{k,\tau_k} = (1 + r_{\text{free}})b_{k,\tau_k-1} + n'_{k,\tau_k}P_{\tau_k}, \quad \text{and } \nu_{k,\tau_k} = b_{k,\tau_k} = 0. \tag{C3}$$

Observe that the equation in (C1) simply defines what's available for "consumption" in period t. It is simply an accounting or budget constraint on the cash flow. The requirements in (C2) are all in the spirit of disallowing shorting, on both components of the decision, the free asset $b_{k,t}$ and the traded stakes $\nu_{k,t}$. In particular the latter is constrained such that $\nu_{k,t} \geq -n'_{k,t}$ (following $n_{k,t} \geq 0$) i.e., bidder k cannot sell more than what's in possession at t; it also ensures that no bidder can own a number of stakes beyond the current total volume ($n_{k,t} \leq N_t$). (C3) specifies how the assets are liquidated at the exit time τ_k: both ν_{k,τ_k} and b_{k,τ_k} will be set at zero, and all remaining stakes n'_{k,τ_k} liquidated (cashed out at P_{τ_k} per unit).

Denote bidder k's decision (process) or "strategy" as τ_k and $(\nu, b) := \{(\nu_{k,t}, b_{k,t}),\ 1 \leq t \leq \tau_k\}$. The objective of bidder k is

$$U_k^* := \max_{\tau_k,(\nu,b)} U_k := \max_{\tau_k,(\nu,b)} \mathsf{E}\left(\sum_{t=1}^{\tau_k} \delta_k^t c_{k,t}\right), \quad \text{subject to (C1), (C2), (C3);} \tag{3.9}$$

where $\delta_k \in (0, 1]$ is a discount factor, a given parameter measuring the risk sensitivity of bidder k. Clearly, bidder k's objective is to maximize a utility that is just the present value of k's total cash flow cumulated up to T_k.

We need to introduce two more processes that are related and central to understanding the dynamics of the system in the presence of trading. The first one is $\{M_t,\ t \geq 1\}$, where $M_t := N_t P_t$ denotes the market value of the volume of stakes at time t. The second one is $\{\Pi_{k,t},\ t \geq 0\}$, for each bidder k, defined as follows:

$$\Pi_{k,0} := n_{k,0}P_0, \quad \text{and} \quad \Pi_{k,t} := \delta_k^t n'_{k,t}P_t - \sum_{j=1}^{t-1} \delta_k^j \nu_{k,j}P_j, \quad t \geq 1; \tag{3.10}$$

where $n'_{k,t+1}$ follows (3.8). Note the two terms that define $\Pi_{k,t}$ are the discounted present values, respectively, of k's pre-trading stakes ($n'_{k,t}$) and of the return from k's trading (cumulated up to $t-1$).

The connection between $\{M_t\}$ and $\{\Pi_{k,t}\}$ is presented in the following lemma, which reveals that their incremental gains (per time period) are proportional: each increment of $\Pi_{k,t}$ is a $\pi_{k,t}$ fraction of the corresponding increment of M_t. In other words, $\pi_{k,t}$ not only represents bidder k's share of the total volume of stakes, it also represents k's share of the system's market value, with or without trading.

Lemma 8. *Under the* `Poly`(α) *voting rule, along with the trading specified above, we have*

$$\mathsf{E}(\Pi_{k,t+1} \,|\, \mathcal{G}_t) - \Pi_{k,t} = \delta_k^{t+1} \pi_{k,t} \mathsf{E}(M_{t+1} \,|\, \mathcal{G}_t) - \delta_k^t \pi_{k,t} M_t. \tag{3.11}$$

Proof. First, by (3.8) and (2.5), along with $\pi_{k,t} = n_{k,t}/N_t$, we have

$$\mathsf{E}(n'_{k,t+1} \,|\, \mathcal{F}_t) = n_{k,t}(1 + N_t^{-(1+\alpha)}) = \frac{n_{k,t}}{N_t}(N_t + N_t^{-\alpha}) = \pi_{k,t} \mathsf{E}(N_{t+1} \,|\, \mathcal{F}_t). \tag{3.12}$$

Next, from (3.10), we have

$$\Pi_{k,t+1} - \Pi_{k,t} = \delta_k^{t+1} n'_{k,t+1} P_{t+1} - \delta_k^t n'_{k,t} P_t - \delta_k^t \nu_{k,t} P_t, \quad t \geq 1. \tag{3.13}$$

Furthermore, as the price process $(P_t,\, t \geq 0)$ is independent of $\mathcal{F}_t$, we have

$$\begin{aligned} \mathsf{E}(n'_{k,t+1} P_{t+1} \,|\, \mathcal{G}_t) &= \mathsf{E}(\mathsf{E}(n'_{k,t+1} \,|\, \mathcal{F}_t) P_{t+1} \,|\, \mathcal{G}_t) \\ &\overset{(3.12)}{=} \pi_{k,t} \mathsf{E}(N_{t+1} P_{t+1} \,|\, \mathcal{G}_t) = \pi_{k,t} \mathsf{E}(M_{t+1} \,|\, \mathcal{G}_t). \end{aligned}$$

This, along with (3.13) yields the desired expression in (3.11), along with $n_{k,t} = n'_{k,t} + \nu_{k,t}$, $n_{k,t} = \pi_{k,t} N_t$, and $M_t = N_t P_t$. □

The process $\{\Pi_{k,t}\}$ also connects to the utility U_k in (3.9). To see this, summing up both sides of (C1) and (C3) over t (along with $b_{k,0} = 0$ in (C2)), we have

$$\sum_{t \leq \tau_k} \delta_k^t c_{k,t} = \sum_{t \leq \tau_k} \delta_k^t c_{k,t} = \delta_k^{\tau_k} n'_{\tau_k} P_{\tau_k} - \sum_{t=1}^{\tau_k - 1} \delta_k^t \nu_{k,t} P_t + \sum_{t=1}^{\tau_k - 1} \delta_k^t \left[(1 + r_{\text{free}}) \delta_k - 1 \right] b_{k,t}. \tag{3.14}$$

Observe that the first two terms on the RHS are equal to Π_{k,τ_k}, so we can rewrite the above as follows (after taking expectations on both sides), emphasizing the exit time τ_k and the strategy (ν, b),

$$U_k(\tau_k, \nu, b) = \mathsf{E}\left[\Pi_{k,\tau_k}(\nu)\right] + \mathsf{E}\left(\sum_{t=1}^{\tau_k - 1} \delta_k^t \left[(1 + r_{\text{free}}) \delta_k - 1 \right] b_{k,t} \right); \tag{3.15}$$

hence, the RHS above is *separable*: the first term depends on (ν) only while the second term, the summation, on (b) only. Furthermore, the second term is ≤ 0 provided $(1 + r_{\text{free}})\delta_k \leq 1$ (which is the condition (a) assumed in Theorem 9 below), along with b being non-negative, part of the feasibility in (C2). In this case, we will have $U_k \leq \mathsf{E}(\Pi_{k,\tau_k}(\nu))$, which implies $U_k^* \leq \max_{\tau_k, \nu} \mathsf{E}(\Pi_{k,\tau_k}(\nu))$, with equality holding when $b_{k,t} = 0$ for all $t = 1, \ldots, \tau_k$.

We are now ready to present the main result regarding the utility maximization problem in (3.9). A quick word on the parameter r_{cryp} that will appear prominently in Theorem 9 below. Simply put, it is the rate (expected rate of return) associated with each stake (e.g., a unit of some cryptocurrency), i.e., it is the counterpart of r_{free}, the rate for the risk-free asset. We

will elaborate more on the two rates after proving the theorem. In the theorem, two strategies are singled out: the "buy-out" strategy, in which bidder k buys up all stakes available at time 1, and then participates in the bidding process until the end; and the "non-participation" strategy, in which bidder k turns all $n_{k,0}$ stakes into cash, and then never participates in either bidding or trading for all $t \geq 1$. Note that the non-participation strategy is executed at $\tau_k = 0$; as such, it complements the feasible class, which is for $\tau_k \geq 1$ and presumes participation. The buy-out strategy clearly belongs to the feasible class.

Theorem 9 (Buy-out strategy versus non-participation)**.** *Assume the following two conditions:*

$$\text{(a)}\ \ \delta_k(1+r_{free}) \leq 1 \quad \text{and} \quad \text{(b)}\ \ \mathsf{E}(M_{t+1}\,|\,\mathcal{G}_t) = (1+r_{cryp})M_t. \tag{3.16}$$

Then, under the `Poly`(α) *voting rule, the following results will hold.*

First, with condition (a), the maximal utility U_k^* is achieved by setting $b_{k,t} = 0$ for all $t = 1, \ldots, T_k$; i.e., $U_k^* = \max_\nu \mathsf{E}(\Pi_{k,T_k})$.

In addition, all three parts of the following will hold.

- *(i) If $\delta_k(1+r_{cryp}) \leq 1$, then any feasible strategy will provide no greater utility for bidder k than the non-participation strategy, i.e., $U_k^* \leq n_{k,0}P_0$.*
- *(ii) If $\delta_k(1+r_{cryp}) \geq 1$, then any feasible strategy will provide no greater utility for bidder k than the buy-out strategy. In this case, bidder k will buy all available stakes at time 1, and participate in the bidding process until the terminal time T_k.*
- *(iii) If $\delta_k(1+r_{cryp}) = 1$, then, bidder k is indifferent between the non-participation and the buy-out strategy with any exit time, both of which will provide no less utility than any feasible strategy. In other words, all strategies achieve the same utility (which is $\Pi_{k,0} = n_0 P_{k,0}$).*

Moreover, when $\delta_k = \delta := (1+r_{cryp})^{-1}$ for all k, then no bidder will have any incentive to trade. Consequently, the long-term behaviors (of N_t, $\pi_{k,t}$ and $\theta_{k,t}$) characterized in Proposition 2, Theorem 6 and Proposition 7 will hold.

Proof. That $U_k^* = \max_{\tau_k,\nu} \mathsf{E}(\Pi_{k,\tau_k})$ (with $b_{k,t}$ being set to 0 for all t), under condition (a) in (3.16), has already been established in the discussions following (3.15). So, it suffices to prove the three parts (i)-(iii).

(i) Applying the given condition (b) in (3.16), along with the assumed inequality $\delta_k(1 + r_{\rm cryp}) \leq 1$, to the RHS of the equation in (3.11) will make it ≤ 0; i.e.,, $\{\Pi_{k,t}\}$ is a $\mathcal{G}_t$-supermartingale, implying $\mathsf{E}(\Pi_{k,\tau_k}) \leq \Pi_{k,0}$. Since $\Pi_{k,0}$ is independent of ν, we have

$$U_k^* = \max_{\tau_k,\nu} \mathsf{E}(\Pi_{k,\tau_k}) \leq \Pi_{k,0} = n_{k,0}P_0, \tag{3.17}$$

(ii) With the assumed inequality $\delta_k(1 + r_{\rm cryp}) \geq 1$, $\{\Pi_{k,t}\}$ now becomes a $\mathcal{G}_t$-sub-martingale; and hence, the inequality below,

$$\mathsf{E}(\Pi_{k,T_k}) \geq \mathsf{E}(\Pi_{k,\tau_k}) \geq \Pi_{k,0} = n_{k,0}P_0. \tag{3.18}$$

To identify the optimal trading strategy $\{\nu^*_{k,j}\}_{j\le T_k-1}$, we use backward induction (dynamic programming). To optimize ν_{k,T_k-1}, observe

$$\begin{aligned}
&\mathsf{E}(\delta_k^{T_k} n'_{k,T_k} P_{T_k} - \delta_k^{T_k-1}\nu_{k,T_k-1}P_{T_k-1} \,|\, \mathcal{G}_{T_k-1}) \\
&= \delta_k^{T_k}(n'_{k,T_k-1}+\nu_{k,T_k-1})\left(1+N_{T_k-1}^{-\alpha-1}\right)\mathsf{E}(P_{T_k}|\mathcal{G}_{T_k-1}) - \delta_k^{T_k-1}\nu_{k,T_k-1}P_{T_k-1} \\
&= \delta_k^{T_k} n'_{k,T_k-1}N_{T_k-1}^{-1}\mathsf{E}(N_{T_k}P_{T_k}|\mathcal{G}_{T_k-1}) + \delta_k^{T_k-1}\left(\delta_k N_{T_k-1}^{-1}\mathsf{E}(N_{T_k}P_{T_k}|\mathcal{G}_{T_k-1}) - P_{T_k-1}\right)\nu_{k,T_k-1} \\
&= \delta_k^{T_k} n'_{k,T_k-1}N_{T_k-1}^{-1}\mathsf{E}(M_{T_k}|\mathcal{G}_{T_k-1}) + \delta_k^{T_k-1}\left(\delta_k N_{T_k-1}^{-1}\mathsf{E}(M_{T_k}|\mathcal{G}_{T_k-1}) - P_{T_k-1}\right)\nu_{k,T_k-1}
\end{aligned}$$

which is linear in $\nu_{k,T_{k-1}}$. By assumed condition (b) in (3.16), we have

$$\delta_k N_{T_k-1}^{-1}\mathsf{E}(M_{T_k}|\mathcal{G}_{T_k-1}) - P_{T_k-1} \ge (\delta_k(1+r_{\rm cryp})-1)\,P_{T_k-1} \ge 0.$$

Thus, $(\nu^*_{k,T_k-1}\,|\,\mathcal{G}_{T_k-1}) = N_{T_k-1} - n'_{k,T_k-1}$, following the (binding) constraint in (C2). That is, bidder k's optimal strategy at the penultimate time T_k-1 is to buy all available stakes at that time. Going backward, we have $(\nu^*_{k,j}\,|\,\mathcal{G}_j) = N_{k,j} - n'_{k,j}$ for $j \ge 1$. Thus, the optimal trading strategy is $\nu^*_{k,1} = N_1 - n'_{k,1}$, $\nu^*_{k,2} = \cdots = \nu^*_{k,T_k-1} = 0$.

(iii) Under the assumed equality $\delta_k(1+r_{\rm cryp}) = 1$, $\{\Pi_{k,t}\}$ is a $\mathcal{G}_t$-martingale; hence, the inequality in (3.17) now holds as equality, i.e., $U_k^* = \Pi_{k,0} = n_{k,0}P_0$. Thus. all strategies lead to the optimal utility, including any feasible strategy (in particular, the no-trading strategy) and the non-participation strategy.

The "moreover" part of the theorem is immediate. □

In what remains of this section, we make a few remarks on Theorem 9, in particular, to motivate and explain its required conditions. First, the rate $r_{\rm cryp}$ is determined by condition (b), the second equation in (3.16). As such, it should be distinct from $r_{\rm free}$, the latter being associated with a risk-free asset. For all practical purpose, we can assume $r_{\rm crpt} \ge r_{\rm free}$, even though this is not assumed in the theorem. When this relation does hold, then condition (a) will become superfluous in cases (i) and (iii).

Second, the discount factor δ_k in the utility objective in (3.9), a parameter that measures bidder k's sensitivity towards risk, plays a key role in characterizing phase transitions in terms of $\delta_k(1+r_{\rm cryp})$. In case (i), the inequality $\delta_k \le 1/(1+r_{\rm cryp})$ implies bidder k is seriously risk-averse; and this is reflected in k's non-participation strategy. In case (ii), the inequality holds in the opposite direction, implying bidder k is lightly risk-averse or even a risk taker. Accordingly, k's strategy is to aggressively sweep up all the available stakes to reach monopoly, and participate (but not trade) until the terminal time. Also note in this case, the non-participation strategy will provide less (no greater) utility for bidder k than the "no-trading" strategy, and certainly no greater utility than the buy-out strategy. In case (iii), the inequality becomes an equality $\delta_k = 1/(1+r_{\rm cryp})$, and $\{\Pi_{k,t}\}$ becomes a martingale. Consequently, bidder k is indifferent between non-participation and participation, and in the latter case, indifferent to all (feasible) strategies, including the buy-out (and the no-trading) strategy. Indeed, the equality $\delta_k = 1/(1+r_{\rm cryp})$ is both necessary and sufficient for the no-trading strategy. This equality also has the effect to force all participating bidders to have the same risk sensitivity.

In contrast, in Roşu and Saleh (2021), there is a single rate $r_{\rm free}$, or equivalently. $r_{\rm cryp} = r_{\rm free}$ is assumed, which seems difficult to justify, since in most applications (cryptocurrency

in particular) r_{cryp} will be significantly larger than r_{free}. Moreover, there is also a single risk sensitivity for all bidders, which is set at $\delta = 1/(1 + r_{\text{free}})$. Thus, Roşu and Saleh (2021) is limited to the martingale case only, reaching the same conclusion as our case (iii), that all feasible strategies, buy-out included, yields the same (expected) utility. As there is no stopping decision and super- or sub-martingale cases in Roşu and Saleh (2021), non-participation does not come up at all, neither do notions like risk-averse or risk-seeking.

The last point we want to emphasize is that the two conditions in (3.16) play very different roles. As evident from the proof of Theorem 9, condition (b) makes $\{\Pi_{k,t}\}$ a super- or sub-martingale or a martingale, according to bidder k's risk sensitivity as specified by the inequalities and equality applied to δ_k (along with r_{free}) in the three cases. Yet, to solve the maximization problem in (3.9), $\{\Pi_{k,t}\}$ needs to be connected to the utility; and this is the role played by condition (a), under which it is necessary (for optimality) to set $b_{k,t} = 0$ for all $t \geq 1$, and applicable to all three cases in Theorem 9. In this sense, condition (a) alone solves half of the maximization problem, the $b_{k,t}$ half of the strategy. In fact, it's more than half, as the optimal ν strategy is only needed in the sub-martingale case; and even there, condition (a) pins down the fact that to participate (even without trading) is better than non-participation.

Note that Theorem 9 can be readily extended. For instance, the rates $r_{\text{cryp}}(t)$ and $r_{\text{free}}(t)$ can vary over the time. In this case, it suffices to modify the conditions in case (i) to

$$\left(1 + \sup_{t<T_k} r_{\text{cryp}}(t)\right)\delta_k \leq 1 \quad \text{and} \quad \left(1 + \sup_{t<T_k} r_{\text{free}}(t)\right)\delta_k \leq 1,$$

the conditions in case (ii) to

$$\left(1 + \inf_{t<T_k} r_{\text{cryp}}(t)\right)\delta_k \geq 1 \quad \text{and} \quad \left(1 + \sup_{t<T_k} r_{\text{free}}(t)\right)\delta_k \leq 1,$$

and the conditions in case (iii) to

$$\delta_k = (1 + r_{\text{cryp}})^{-1} \quad \text{and} \quad \sup_{t<T_k} r_{\text{free}}(t) \leq r_{\text{cryp}}, \quad \text{with } r_{\text{cryp}} \text{ being constant.}$$

Then, Theorem 9 will continue to hold. We can also include a processing cost $\kappa > 0$ that any bidder selected by the $\texttt{Poly}(\alpha)$ mechanism will pay to receive the reward. (This corresponds to the "mining" cost to validate the block.) In this case, the budget constraint (C1) is modified by adding a term $-\kappa 1_{S_{k,t}}$ to the right side of the equation, and the same applies to the liquidation constraint (with t replaced by T_k). Condition (b) in (3.16) is modified to $\mathsf{E}(M_{t+1} \,|\, \mathcal{G}_t) = (1 + r_{\text{cryp}})M_t + \kappa$.

4. Conclusions

We have proposed in this study a new $\texttt{Poly}(\alpha)$ voting rule that is more general than the traditional voting rule (which is linear, corresponding to $\alpha = 0$). More importantly, the $\texttt{Poly}(\alpha)$ voting rule distinguishes voting power from voter share, and hence decouples the two.

Applying the $\texttt{Poly}(\alpha)$ voting rule to the PoS protocol, where the voters are the bidders (competing for "rewards," or validation of new blocks), we show this decoupling will enhance security, a key objective of the PoS protocol. Specifically, we prove that while bidder shares

form a martingale process that will converge to a Dirichlet distribution, each bidder's voting power is a super-martingale that decreases to zero over time. For both limiting results, we explicitly characterize their rate of convergence as well. Furthermore, we show a phase transition in the stability of bidder shares in terms of each bidder's initial share relative to the total in the system. We also study the issue of a bidder's risk sensitivity when trading is allowed, and provide conditions under which a bidder will have no incentive to participate in the bidding process, or, if participate, will forgo trading.

In the Introduction, we mentioned two general approaches to enhance security in the PoS protocol: adjust the amount of reward over time and slow down the voting process; and the current study focuses on the latter while keeping the reward constant. It is possible to pursue a combination of both approaches, i.e., adjusting the size of reward dynamically over time in the same manner as adjusting α (for the latter, refer to §5.2 below). In another direction, it is also possible to study the trading problem in §3.2 under a suitable market impact model, where the price process P_t will be impacted by trading activities; for instance, a mean-field PoS model with linear impact (and transaction costs).

Acknowledgments: We thank anonymous referees for helpful suggestions which improve the presentation of the paper. W. Tang gratefully acknowledges financial support through NSF grants DMS-2113779 and DMS-2206038, and through a start-up grant at Columbia University. David Yao's work is part of a Columbia-CityU/HK collaborative project that is supported by InnoHK Initiative, The Government of the HKSAR and the AIFT Lab.

5. Appendix

5.1. **Improvement on $\lambda_\pm(\alpha)$.** Theorem 3 proves large-deviation bounds on N_t. However, it does not cover the whole range. It remains open to prove such bounds in the range $(\lambda_-(\alpha), \lambda_+(\alpha))$; and once proved, the result will also imply the almost sure convergence of $N_t/t^{\frac{1}{1+\alpha}}$ as $t \to \infty$.

Here we provide a way to (slightly) improve the values of $\lambda_\pm(\alpha)$ in Theorem 3. To simplify the presentation, we consider $\alpha = 1$ (quadratic voting rule) with $\lambda_-(1) \approx 0.56$ and $\lambda_+(1) \approx 2.51$. The idea relies on a multi-scale analysis by splitting the interval $[0, t]$ into $[0, t/2]$ and $[t/2, t]$, and the goal is to upper bound $\mathbb{P}(N_t = \lambda\sqrt{t})$ for $\lambda > 0$. In the sequel, we neglect the polynomial factors and only focus on the exponential terms. Note that

$$\mathsf{P}(N_t = \lambda\sqrt{t}) = \sum_{k \le \lambda\sqrt{t}} \binom{t/2}{k} \binom{t/2}{\lambda\sqrt{t} - k} \frac{1}{(\lambda\sqrt{t})!} \underbrace{\left(1 - \frac{1}{k}\right)^{t/2-k} \left(1 - \frac{1}{\lambda\sqrt{t}}\right)^{t/2+k-\lambda\sqrt{t}}}_{(a'')}.$$

Next we split the range of $k \le \lambda\sqrt{t}$ into $S_1 := \{k \le a\sqrt{t}\} \cup \{k \ge (\lambda - a)\sqrt{t}\}$, and $S_2 := \{a\sqrt{t} < k < (\lambda - a)\sqrt{t}\}$ with $a < \frac{\lambda}{2}$. For $k \in S_1$, we simply bound the term (a) by $\left(1 - \frac{1}{\lambda\sqrt{t}}\right)^{t/2-\lambda\sqrt{t}}$ while for $k \in S_2$ we bound the term (a'') by $\left(1 - \frac{1}{(\lambda-a)\sqrt{t}}\right)^{t/2-(\lambda-a)\sqrt{t}} \left(1 - \frac{1}{\lambda\sqrt{t}}\right)^{t/2-a\sqrt{t}}$.

Consequently,

$$\mathsf{P}(N_t = \lambda\sqrt{t}) \le \underbrace{\left(\sum_{k\in S_1}\binom{t/2}{k}\binom{t/2}{\lambda\sqrt{t}-k}\right)\frac{1}{(\lambda\sqrt{t})!}\left(1-\frac{1}{\lambda\sqrt{t}}\right)^{t/2-\lambda\sqrt{t}}}_{(b'')}$$
$$+\underbrace{\left(\sum_{k\in S_2}\binom{t/2}{k}\binom{t/2}{\lambda\sqrt{t}-k}\right)\frac{1}{(\lambda\sqrt{t})!}\left(1-\frac{1}{(\lambda-a)\sqrt{t}}\right)^{t/2-(\lambda-a)\sqrt{t}}\left(1-\frac{1}{\lambda\sqrt{t}}\right)^{t/2-a\sqrt{t}}}_{(c'')}.$$

Using Stirling's formula, we get exponential bounds for the terms (b'') and (c''):

$$\begin{aligned}(b'') &\sim \exp\left((-\lambda\log 2 + 2\lambda - a\log a - (\lambda-a)\log(\lambda-a) - \lambda\log\lambda - \frac{1}{\lambda})\sqrt{t}\right),\\ (c'') &\sim \exp\left((2\lambda - 2\lambda\log\lambda - \frac{1}{2\lambda} - \frac{1}{2(\lambda-a)})\sqrt{t}\right).\end{aligned} \tag{5.1}$$

By equating the two coefficients before $\sqrt{t}$ in (5.1), we have

$$-\lambda\log 2 + 2\lambda - a\log a - (\lambda-a)\log(\lambda-a) - \lambda\log\lambda - \frac{1}{\lambda}) = 2\lambda - 2\lambda\log\lambda - \frac{1}{2\lambda} - \frac{1}{2(\lambda-a)}.$$

By letting $a = \theta\lambda$ with $\theta < \frac{1}{2}$, the above equation yields

$$\lambda = \sqrt{\frac{\theta}{2(1-\theta)(\log 2 + \theta\log\theta + (1-\theta)\log(1-\theta))}}. \tag{5.2}$$

and the coefficient before $\sqrt{t}$ is

$$f(\lambda) = 2\lambda\log\lambda - 2\lambda + \frac{1}{2\lambda} + \frac{1}{2(1-\theta)\lambda}, \tag{5.3}$$

where θ is specified by (5.2). By injecting the expression (5.2) into (5.3), f is a function of θ. It is easy to see that $f(\theta)$ has only one root on $(0, 1/2)$ which is approximately 0.1575, and $\lambda_-(1)$ is improved numerically to from 0.56 to 0.60. Similarly, the value of $\lambda_+(1)$ is improved numerically from 2.51 to 2.44.

We can continue this procedure, for instance to split $[0,t]$ into $[0,t/3]$, $[t/3,2t/3]$ and $[2t/3,t]$, and so on to get better and better numerical values of $\lambda_-(1)$ and $\lambda_+(1)$. However, it is not clear whether this approach will eventually get all the way to the threshold $\sqrt{2} \approx 1.41$. We conjecture that the exponential deviation holds right off the threshold $(1+\alpha)^{\frac{1}{1+\alpha}}$, which is supported by the numerical experiments; refer to Figure 1.

5.2. **Control of voting powers.** As proved in Theorem 6, the reward rate $\theta_{k,t}$ decays at rate $\Theta(t^{-\frac{\alpha}{1+\alpha}})$. If the reward is associated with the validation of a new block, then the duration between two consecutive validations (called "block time" below) will increase (and uncontrolled) over time. For instance, set $\alpha = 1$ (quadratic voting rule), and $T = 10^7$ seconds (≈ 4 months). Then, the duration required to see the next block at time T is approximately

$$10 \text{ seconds } \times (10^7/10)^{\frac{1}{2}} = 10^4 \text{ seconds } \approx 3 \text{ hours},$$

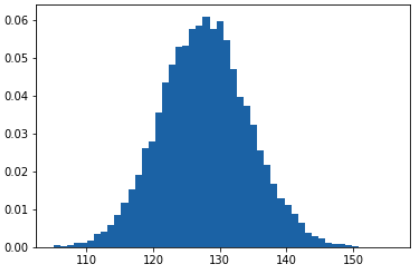

(A) Histogram of N_{8000} on MC simulation of 20000 samples.

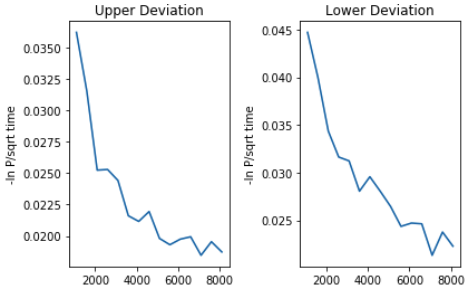

(B) x-axis: $t \in \{1000, 1500, \dots, 8000\}$; y-axis: $-\ln \mathsf{P}(N_t > \sqrt{2.2\,t})/\sqrt{t}$ (left) and $-\ln \mathsf{P}(N_t < \sqrt{1.8\,t})/\sqrt{t}$ (right) on MC simulation of 20000 samples.

FIGURE 1. Volume of stakes N_t with $N_0 = 5$ and $\alpha = 1$ (quadratic voting).

which is even much longer than the 10 minutes block time of Bitcoin. (The block time is 10 seconds in Ethereum, see e.g. Buterin (2014).)

One possible (and practical) solution is to dynamically "tune" the parameter α over time. Specifically, let κ denote a threshold for the expected number of rounds of bidding/voting between two validated blocks. Then,

- set $\alpha = \alpha_0 > 0$, and apply the $\texttt{Poly}(\alpha_0)$ scheme up to round $\kappa^{1+\alpha_0^{-1}}$;
- set $\alpha = \alpha_1 < \alpha_0$, and apply the $\texttt{Poly}(\alpha_1)$ scheme up to round $\kappa^{1+\alpha_1^{-1}}$ ... and so on.

Here $\kappa, \alpha_0, \alpha_1, \dots$ are user-defined hyper-parameters. To illustrate, by setting $\kappa = 50$ rounds (≈ 10 minutes in Ethereum) and $\alpha_k = (1+k)^{-1}$ for $k \geq 0$,

- Apply the $\texttt{Poly}(1)$ scheme up to round $50^2 \approx 7$ hours;
- Apply the $\texttt{Poly}(1/2)$ scheme up to round $50^3 \approx 2$ weeks;
- Apply the $\texttt{Poly}(1/3)$ scheme up to round $50^4 \approx 2$ years;
- Apply the $\texttt{Poly}(1/4)$ scheme up to round $50^5 \approx 100$ years ... and so on.

Similarly, by setting $\kappa = 5$ rounds (≈ 1 minutes in Ethereum),

- Apply the $\texttt{Poly}(1)$ scheme up to round $5^2 \approx 4$ minutes;
- Apply the the $\texttt{Poly}(1/2)$ scheme up to round $5^3 \approx 20$ minutes; $\cdots$
- Apply the $\texttt{Poly}(1/10)$ scheme up to round $5^{11} \approx 15$ years ..., and so on.

It is also possible to tune the parameter α at random time points adaptive to the reward rate. That is,

- Set $\alpha = \alpha_0 > 0$, and apply the $\texttt{Poly}(\alpha_0)$ scheme up to round k_0 where k_0 is the first time by which no new block is validated in κ rounds;
- Set $\alpha = \alpha_1 < \alpha_0$, and apply the $\texttt{Poly}(\alpha_0)$ scheme up to round k_1 where k_1 is the first time by which no new block is validated in κ rounds since then ... and so on.

Note that in either case the process of stakes is a time-changed Pólya urn, so the results in Section 3 continue to hold (except that the convergence rate will depend on the choice of $\{\alpha_k\}$).

References

M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. H. Katz, A. Konwinski, G. Lee, D. A. Patterson, A. Rabkin, and I. Stoica. Above the clouds: A Berkeley view of cloud computing. 2009. Technical Report UCB/EECS-2009-28. Available at `https://www2.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.pdf`.

V. Bagaria, A. Dembo, S. Kannan, S. Oh, D. Tse, P. Viswanath, X. Wang, and O. Zeitouni. Proof-of-stake longest chain protocols: Security vs predictability. 2019. arXiv:1910.02218.

P. Billingsley. *Probability and measure*. Wiley Series in Probability and Mathematical Statistics. John Wiley & Sons, Inc., New York, third edition, 1995.

P. Billingsley. *Convergence of probability measures*. Wiley Series in Probability and Statistics: Probability and Statistics. John Wiley & Sons, Inc., New York, second edition, 1999.

V. Buterin. Toward a 12-second block time. 2014. Available at `https://blog.ethereum.org/2014/07/11/toward-a-12-second-block-time`.

H. Chen and D. D. Yao. *Fundamentals of queueing networks*, volume 46 of *Applications of Mathematics*. Springer-Verlag, 2001.

J. Dean and S. Ghemawat. Mapreduce: simplified data processing on large clusters. *Commun. ACM*, 51(1):107–113, 2008.

E. Deirmentzoglou, G. Papakyriakopoulos, and C. Patsakis. A survey on long-range attacks for proof of stake protocols. *IEEE Access*, 7:28712–28725, 2019.

D. Denisov, D. Korshunov, and V. Wachtel. At the edge of criticality: Markov chains with asymptotically zero drift. 2016. arXiv:1612.01592.

R. Durrett. *Probability—theory and examples*. Cambridge University Press, 2019.

H. Garcia-Molina. Elections in a distributed computing system. *IEEE Trans. Comput.*, 31 (01):48–59, 1982.

L. Goldstein and G. Reinert. Stein's method for the beta distribution and the Pólya-Eggenberger urn. *J. Appl. Probab.*, 50(4):1187–1205, 2013.

A. Q. Huang and J. Baliga. FREEDM System: Role of power electronics and power semiconductors in developing an energy internet. In *21st International Symposium on Power Semiconductor Devices & IC's*, pages 9–12, 2009.

S. King and S. Nadal. Ppcoin: Peer-to-peer crypto-currency with proof-of-stake. 2012. Available at `https://decred.org/research/king2012.pdf`.

S. P. Lalley and E. G. Weyl. Quadratic voting: How mechanism design can radicalize democracy. In *AEA Papers and Proceedings*, volume 108, pages 33–37, 2018.

J. Lamperti. Criteria for the recurrence or transience of stochastic process. I. *J. Math. Anal. Appl.*, 1:314–330, 1960.

J. Lamperti. A new class of probability limit theorems. *J. Math. Mech.*, 11:749–772, 1962.

J. Lamperti. Criteria for stochastic processes. II. Passage-time moments. *J. Math. Anal. Appl.*, 7:127–145, 1963.

L. Lamport, R. Shostak, and M. Pease. The Byzantine Generals Problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, 1982.

M. Menshikov, S. Popov, and A. Wade. *Non-homogeneous random walks*, volume 209 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 2017. Lyapunov function methods for near-critical stochastic systems.

S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Decentralized Business Review*, page 21260, 2008.

L. S. Penrose. The elementary statistics of majority voting. *J. R. Stat. Soc.*, 109(1):53–57, 1946.

E. Rio. Upper bounds for minimal distances in the central limit theorem. *Ann. Inst. Henri Poincaré Probab. Stat.*, 45(3):802–817, 2009.

I. Roşu and F. Saleh. Evolution of shares in a proof-of-stake cryptocurrency. *Manag. Sci.*, 67(2):661–672, 2021.

E. Shi. *Foundations of Distributed Consensus and Blockchains*. 2020. Available at `http://elaineshi.com/docs/blockchain-book.pdf`.

W. Tang. Stability of shares in the Proof of Stake protocol – concentration and phase transitions. 2022. arXiv:2206.02227.

W. Tang. Trading and wealth evolution in the Proof of Stake protocol. 2023. arXiv:2308.01803.

W. Tang and D. D. Yao. Trading under the proof-of-stake protocol–a continuous-time control approach. *Math. Finance*, 33(4):979–1004, 2023.

C. Villani. *Optimal transport*, volume 338 of *Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 2009. Old and new.

G. Wood. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum Project Yellow Paper*, 151:1–32, 2014.

Department of Industrial Engineer and Operations Research, Columbia University.
Email address: `wt2319@columbia.edu`

Department of Industrial Engineer and Operations Research, Columbia University.
Email address: `yao@columbia.edu`