This article was downloaded by: [160.39.34.79] On: 11 February 2024, At: 14:50
Publisher: Institute for Operations Research and the Management Sciences (INFORMS)
INFORMS is located in Maryland, USA

Mathematics of Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Polynomial Voting Rules

Wenpin Tang, David D. Yao

To cite this article:
Wenpin Tang, David D. Yao (2024) Polynomial Voting Rules. Mathematics of Operations Research

Published online in Articles in Advance 07 Feb 2024

. https://doi.org/10.1287/moor.2023.0080

Full terms and conditions of use: https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2024, INFORMS

Please scroll down for article—it is on subsequent pages

With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.
For more information on INFORMS, its publications, membership, or meetings visit http://www.informs.org

https://pubsonline.informs.org/journal/moor

MATHEMATICS OF OPERATIONS RESEARCH
Articles in Advance, pp. 1–17
ISSN 0364-765X (print), ISSN 1526-5471 (online)

Polynomial Voting Rules

Wenpin Tang,[a,*] **David D. Yao**[a]

[a] Department of Industrial Engineering and Operations Research, Columbia University, New York, New York 10027
*Corresponding author
Contact: wt2319@columbia.edu, https://orcid.org/0000-0001-7228-1954 (WT); yao@columbia.edu (DDY)

Received: March 16, 2023
Revised: October 4, 2023
Accepted: January 7, 2024
Published Online in Articles in Advance: February 7, 2024

MSC2020 Subject Classification: Primary: 60C05, 91B08; secondary: 60F05, 60G42

https://doi.org/10.1287/moor.2023.0080

Copyright: © 2024 INFORMS

Abstract. We propose and study a new class of polynomial voting rules for a general decentralized decision/consensus system, and more specifically for the 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, whereas voter shares form a martingale process that converges to a Dirichlet distribution, their voting powers follow a supermartingale 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 aversion) in three categories, corresponding to the voter's utility being a supermartingale, a submartingale, and a martingale. For each category, we identify the voter's best strategy in terms of participation and trading.

Funding: W. Tang gratefully acknowledges financial support through the National Science Foundation [Grants DMS-2113779 and DMS-2206038] and through a start-up grant at Columbia University. D. D. Yao's work is part of a Columbia–City University/Hong Kong collaborative project that is supported by InnoHK Initiative, the Government of Hong Kong Special Administrative Region, and the Laboratory for AI-Powered Financial Technologies.

Keywords: cryptocurrency • economic incentive • fluid limit • phase transition • polynomial voting rules • proof-of-stake protocol • stability • urn models

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, in which the voters are often abstract entities (virtual) and the voting process automated, and the purpose of reaching consensus is often nonsocial and nonpolitical, such as to enhance the overall security of an industrial operation or infrastructure (Garcia-Molina [11], Lamport et al. [19]). Examples include cloud computing (Armbrust et al. [1], Dean and Ghemawat [7]), smart power grids (Huang and Baliga [13]), and more recently trading or payment platforms and exchanges built upon the blockchain technology (Nakamoto [21], Wood [30]).

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 [21],) and proof of stake (PoS; King and Nadal [14], Wood [30]). 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 the blockchain. Hence, whereas 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 existed long before the PoS protocol and is widely used in many applications, such as auctions and initial public offerings.

Downloaded from informs.org by [160.39.34.79] on 11 February 2024, at 14:50 . For personal use only, all rights reserved.

Downloaded from informs.org by [160.39.34.79] on 11 February 2024, at 14:50 . For personal use only, all rights reserved.

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, a 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, is 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 need 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 [22] famously proposed a square-root voting rule, around the time when the United Nations was founded shortly after World War II. According to Penrose, a world assembly such as the United Nations 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 [15]). The idea is that each voter is 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 quadratic voting rules mentioned earlier 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 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 being lower bounded because of network delay (see Shi [25, section 14.3]). Specifically, there is the principle of security:

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

Here, honest/dishonest power refers to the voting power of honest/dishonest bidders. The parameter γ is a user-defined security factor; for example, $\gamma = 2$ means honest power is expected to be twice as much as dishonest power, and 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 reduces 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, whereas dishonest bidders do not comply with the rules. As we illustrate (in the remarks following Theorem 2), slowing down the voting process enhances security. This is because decreasing voting power over time increases the block generation time, which mitigates network delay and makes the principle of security (1) easier to hold.

- Second, PoS blockchains suffer from malicious attacks known as nothing at stake (see, e.g., Deirmentzoglou et al. [8]). As pointed out in Bagaria et al. [2], for the PoS longest chain protocol, honest bidders focus exclusively on the longest chain, whereas dishonest bidders can work simultaneously on all existing blocks. They show 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 and (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 $t \to \infty$, whereas their voting powers follow a supermartingale process that decreases to zero over time (Theorem 2), 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 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 4). Notably, the same phenomenon is

Downloaded from informs.org by [160.39.34.79] on 11 February 2024, at 14:50 . For personal use only, all rights reserved.

demonstrated under the traditional voting rule ($\alpha = 0$); refer to Roşu and Saleh [24] and Tang [26]. 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 is recently studied in Roşu and Saleh [24]. Not only is our model more general in allowing any $\alpha \geq 0$, our results are also richer and sharper (Theorem 3). For instance, we quantify the level of risk sensitivity (or risk aversion) that results in three cases according to the voter's utility being a supermartingale, a submartingale, or a martingale. Each case leads to a best strategy for the voter, including nonparticipation (not to participate at all in the bidding) and buyout (buying as many stakes as are available), which are not considered in Roşu and Saleh [24]. Note that a buyout 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 [28]) on trading PoS stakes with volume constraint. See also Tang [27] for various problems (including transaction costs and voter's collective behavior) related to 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 1) and a fluid limit (Proposition 3).

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, and 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 earlier and (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 Poly(α) 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 and 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 $\overset{d}{\to}$ 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 toward zero as $b \to \infty$.
- $d_W(\mu, \nu)$ denotes the 1-Wasserstein distance between two probability distributions μ and ν. Refer to Villani [29, 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, are the participants in the decentralized system, in which they engage in rounds of bidding following a prespecified 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 stays fixed throughout the paper, and let $[K] := \{1, \ldots, K\}$ denote the set of all bidders.

Time is discrete, indexed by $t = 0, 1, 2, \ldots$, and corresponds to the rounds of bidding mentioned. 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

$$\pi_{k,0} := \frac{n_{k,0}}{N}, \quad k \in [K]. \tag{2}$$

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

$$\pi_{k,t} := \frac{n_{k,t}}{N_t}, \quad k \in [K], \quad \text{with } N_t := \sum_{k=1}^{K} n_{k,t}. \tag{3}$$

Here, N_t is the total number of stakes at time t, and thus, $N_0 = N$. (We 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

$$\theta_{k,t} := \frac{n_{k,t}}{N_t^{1+\alpha}} = \frac{\pi_{k,t}}{N_t^{\alpha}}, \tag{4}$$

Downloaded from informs.org by [160.39.34.79] on 11 February 2024, at 14:50 . For personal use only, all rights reserved.

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 earlier, $\theta_{k,t}$ can also be viewed as bidder k's voting power at time t. (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 [24] and Tang [26].

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:

$$n_{k,t} = n_{k,t-1} + 1_{S_{k,t}}, \quad k \in [K]; \tag{5}$$

or simply,

$$n_{k,t} = \begin{cases} n_{k,t-1} & \text{with probability} \quad 1-\theta_{k,t-1}, \\ n_{k,t-1}+1 & \text{with probability} \quad \theta_{k,t-1}. \end{cases} \tag{6}$$

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

$$N_t = \begin{cases} N_{t-1} & \text{with probability} \quad 1-1/N_{t-1}^{\alpha}, \\ N_{t-1}+1 & \text{with probability} \quad 1/N_{t-1}^{\alpha}. \end{cases} \tag{7}$$

The counting process $(N_t, t \geq 0)$ specified in (7) evolves as a time-homogeneous Markov chain on $\{N, N+1, \dots\}$ in contrast with the Pólya urn (with a constant reward) in which N_t grows deterministically and linearly in t. As we see in the next section, the $\texttt{Poly}(\alpha)$ 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 $\boldsymbol{n}_t = (n_{1,t}, \dots, n_{K,t})$ 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 interarrival 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 earlier.

Proof. It is clear from the dynamics in (7) 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. □

This 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 is used to study the long-time behavior of bidder shares and reward rates in Section 2.2. But we first study in the next section 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 $\mathcal{F}_t$-submartingale, 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:*

$$\frac{N_t^{1+\alpha}}{t} \to 1+\alpha \quad \text{as } t \to \infty. \tag{8}$$

Proof.

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

Downloaded from informs.org by [160.39.34.79] on 11 February 2024, at 14:50 . For personal use only, all rights reserved.

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

$$\mathsf{E}N_t^{1+\alpha}\sim(1+\alpha)t \quad \text{as } t\to\infty. \tag{9}$$

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, $\mathsf{E}(N_{t+1}^{2(1+\alpha)}-N_t^{2(1+\alpha)})=(2(1+\alpha)+o(1))\mathsf{E}N_t^{1+\alpha}\sim 2(1+\alpha)^2t$ by (9). Then, we get $\mathsf{E}(N_t^{2(1+\alpha)})\sim(1+\alpha)^2t^2$ as $t\to\infty$. We proceed by induction. Assuming that $\mathsf{E}(N_t^{j(1+\alpha)})\sim(1+\alpha)^jt^j$ 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

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

By the method of moments (see, e.g., Billingsley [3, 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 (8). 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 of how N_t fluctuates 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 1 (Large Deviations for N_t). *Define a function $f_\alpha(\cdot)$,*

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

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 $\varepsilon>0$,*

$$\mathsf{P}(N_t<\lambda t^{\frac{1}{1+\alpha}})\le\exp\left(-(1-\varepsilon)f_\alpha(\lambda)\,t^{\frac{1}{1+\alpha}}\right) \quad \text{as } t\to\infty. \tag{10}$$

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

$$\mathsf{P}(N_t>\lambda t^{\frac{1}{1+\alpha}})\le\exp\left(-(1-\varepsilon)f_\alpha(\lambda)\,t^{\frac{1}{1+\alpha}}\right) \quad \text{as } t\to\infty. \tag{11}$$

Proof. Without loss of generality, assume that $N_0=1$. Note that $(N_t,\,t\ge0)$ 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

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

because 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

$$a_k:=\binom{t}{k}\frac{1}{(k!)^\alpha}\left(1-\frac{1}{(k+1)^\alpha}\right)^{t-k}. \tag{12}$$

Downloaded from informs.org by [160.39.34.79] on 11 February 2024, at 14:50 . For personal use only, all rights reserved.

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 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 these estimates yields

$$a_{\lambda\sqrt{t}} \sim (2\pi\lambda)^{-\frac{1+\alpha}{2}} t^{-\frac{1}{2}} \exp\left(-f_\alpha(\lambda)\, t^{\frac{1}{1+\alpha}}\right). \tag{13}$$

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 for 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 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 because, 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 Appendix A 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 1. *Under the* `Poly`(α) *voting rule, we have, for each* $\delta > 0$,

$$\mathsf{P}\left(|N_t - ((1+\alpha)t)^{\frac{1}{1+\alpha}}| > \delta t^{\frac{1}{1+\alpha}}\right) = \mathcal{O}(t^{-\frac{1}{1+\alpha}}) \quad \text{as } t \to \infty. \tag{14}$$

Proof. Note that $N_0 \le N_t \le t + N_0$, and by Theorem 1, 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 $\mathbb{E}(N_t^{2(1+\alpha)}) = (1+\alpha)^2 t^2 + \mathcal{O}(t^{\frac{1+2\alpha}{1+\alpha}})$, which implies $\mathsf{Var}(N_t^{1+\alpha}) = \mathcal{O}(t^{\frac{1+2\alpha}{1+\alpha}})$. Hence,

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

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 these estimates yields (14). □

Downloaded from informs.org by [160.39.34.79] on 11 February 2024, at 14:50 . For personal use only, all rights reserved.

Recall that the process $(N_t, t \geq 0)$ is a time-homogenous Markov chain. The path properties of a general time-homogenous Markov chain $(Z_t, t \geq 0)$ has long been studied since the work of Lamperti [16–18]. The basic idea is to study the recurrence or transience of $(Z_t, t \geq 0)$ 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 $\limsup_{x\to\infty} 2xm_1(x) + m_2(x) \leq 0$, then $(Z_t, t \geq 0)$ is recurrent, and if $\liminf_{x\to\infty} 2xm_1(x) + m_2(x) > 0$, then $(Z_t, t \geq 0)$ 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. [9], Menshikov et al. [20]). Specializing to the process $(N_t, t \geq 0)$, we have

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

Interestingly, there seem to be few results on Lamperti's problem in which both $m_1(x)$ and $m_2(x)$ decrease to zero as $x \to \infty$ except that $(N_t, t \geq 0)$ is transient. Observe that $(N_t, t \geq 0)$ 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 a similar spirit as Lamperti [17], the asymptotic growth (8) hinges on a degenerate fluid approximation of the process $(N_t, t \geq 0)$ as stated in the following proposition.

Proposition 3 (Fluid Limit of N_t). *Under the* `Poly`(α) *voting rule, we have*

$$\left(\frac{N_{nu}}{n^{\frac{1}{1+\alpha}}}, u \geq 0\right) \xrightarrow{d} (X_u, u \geq 0) \quad \text{as } n \to \infty \text{ in } \mathcal{C}[0, \infty), \tag{15}$$

where N_s for noninteger s is defined by the linear interpolation of the chain $(N_t, t \geq 0)$ and $X_u = ((1+\alpha)u)^{\frac{1}{1+\alpha}}$, $u \geq 0$ 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 (15) 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 $K(\varepsilon) := \max(2X_T, \max_{n \leq n(\varepsilon)}(N_0 + [nT])/n^{\frac{1}{1+\alpha}})$. 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 $N^{n,T} := \left(N_{[nT]}/n^{\frac{1}{1+\alpha}}, 0 \leq t \leq T\right)$ is nondecreasing. Thus,

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

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 $0 \leq t_1 < \cdots < t_k$, the vector $(N_{[nt_1]}/n^{\frac{1}{1+\alpha}}, \ldots, N_{[nt_k]}/n^{\frac{1}{1+\alpha}})$ converges in probability to $(X_{t_1}, \ldots, X_{t_k})$, that is, 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 [4, chapter 2]). □

Note that the fluid limit in the proposition is different from the fluid limit in the literature of stochastic networks, in which it usually takes the form of a functional strong law of large numbers (FSLLN) 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 [6, section 6.1]. Here, the process $(N_t, t \geq 0)$ is nonrenewal; thus, the FSLLN limit does not apply, yet there is still the weak convergence, and the limit is still a deterministic function of time $(X_u, u \geq 0)$, explicitly characterized earlier. Another notable point is, in the FSLLN setting, both time and space are scaled by the same scaling factor n, whereas in (15), the time scaling remains the same, and the space scaling is by $n^{\frac{1}{1+\alpha}}$. But this is only because $(N_t, t \geq 0)$ 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 long-time behavior of $(\pi_{k,t}, k \in [K])$ and $(\theta_{k,t}, k \in [K])$. Recall that the Dirichlet distribution with parameters $(a_1, \ldots, a_K)$, which we denote by $\mathrm{Dir}(a_1, \ldots, a_K)$, has support on the standard simplex

Downloaded from informs.org by [160.39.34.79] on 11 February 2024, at 14:50 . For personal use only, all rights reserved.

$\{(x_1,\dots,x_K)\in\mathbb{R}_+^K:\sum_{k=1}^K x_k=1\}$ and has density

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

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 $(x_1,\dots,x_K)\overset{d}{=}\mathrm{Dir}(a_1,\dots,a_K)$, then for each $k\in[K]$, $x_k\overset{d}{=}\mathrm{Beta}(a_k,\sum_{j\neq k}a_j)$.

Theorem 2 (Long-Time Behavior). *Under the* `Poly`(α) *voting rule, we have the following limiting distributions.*

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

$$(\pi_{1,t},\dots,\pi_{K,t})\to(\pi_{1,\infty},\dots,\pi_{K,\infty})\quad \text{as } t\to\infty, \tag{16}$$

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

$$d_W(\pi_{k,t},\mathrm{Beta}(n_{k,0},N-n_{k,0}))=\mathcal{O}(t^{-\frac{1}{1+\alpha}})\quad \text{as } t\to\infty. \tag{17}$$

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

$$(1+\alpha)^{\frac{\alpha}{1+\alpha}}t^{\frac{\alpha}{1+\alpha}}\theta_{k,t}\xrightarrow{d}\mathrm{Beta}(n_{k,0},N-n_{k,0})\quad \text{as } t\to\infty. \tag{18}$$

Proof.

i. By (3) and (6), it is easily seen that, for each $k\in[K]$ and $t\ge 0$,

$$\mathsf{E}(\pi_{k,t+1}\mid\mathcal{F}_t)=\frac{n_{k,t}}{N_t}\left(1-\frac{1}{N_t^\alpha}\right)+\frac{n_{k,t}}{N_t+1}\frac{N_t-n_{k,t}}{N_t^{1+\alpha}}+\frac{n_{k,t}+1}{N_t+1}\frac{n_{k,t}}{N_t^{1+\alpha}}. \tag{19}$$

Recognizing the first term on the right side of (19), $\frac{n_{k,t}}{N_t}=\pi_{k,t}$, whereas all other terms sum up to zero, we conclude that $(\pi_{k,t},\,t\ge 0)$ is an $\mathcal{F}_t$-martingale. The convergence in (16) follows from the martingale convergence theorem (see, e.g., Durrett [10, section 4.2]). By Proposition 1, $(\boldsymbol{n}_t,\,t\ge 0)$ is a time-changed Pólya urn. So the limiting shares $(\pi_{1,\infty},\dots,\pi_{K,\infty})$ have the same distribution as that of the Pólya urn, which is $\mathrm{Dir}(n_{1,0},\dots,n_{K,0})$.

Let $(\boldsymbol{n}_t^\dagger,\,t\ge 0)$ be the Pólya urn with $n_{k,0}^\dagger=n_{k,0}$ and $(\pi_{k,t}^\dagger,\dots,\pi_{K,t}^\dagger)$ be the corresponding shares. Set $Z\overset{d}{=}\mathrm{Beta}(n_{K,0},N-n_{k,0})$. By Goldstein and Reinert [12], we have, for each $k\in[K]$,

$$d_W(\pi_{k,t}^\dagger,Z)=\mathcal{O}(t^{-1})\quad \text{as } t\to\infty. \tag{20}$$

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

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

where the last inequality follows from Theorem 1 and (20). This yields the bound (17).

ii. Applying the same derivation as in (19) but to $\theta_{k,t}$ instead, we have

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

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

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

that is, $(\theta_{k,t},\,t\ge 0)$ is an $\mathcal{F}_t$-supermartingale for each k. Recall that $N_t^\alpha\theta_{k,t}=\pi_{k,t}$, so $\theta_{k,t}\le N_t^{-\alpha}$, which converges to zero with probability one. By (i), $N_t^\alpha\theta_{k,t}$ converges almost surely and, hence, in distribution to Z. By Proposition 2,

Downloaded from informs.org by [160.39.34.79] on 11 February 2024, at 14:50 . For personal use only, all rights reserved.

$N_t/((1+\alpha)t)^{\frac{1}{1+\alpha}}$ converges in probability to one. We then apply Slutsky's theorem to get the convergence in (18). □

Several remarks are in order. Part (i) of Theorem 2 shows that the bidder shares form a martingale and converge to a Dirichlet distribution (independent of α). This should be expected from the fact that the underlying bidder stakes $(\boldsymbol{n}_t, t \geq 0)$ are a time-changed Pólya urn; refer to Proposition 1. What's more revealing is the Wasserstein bound in (17) 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 4 for further discussion on the stability of the bidder shares when the initial stakes $N := N_0$ are 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; the level of decentralization remains unchanged. This also means the principle of security in (1) becomes easier to hold at large time t because (because of 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 it 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 Appendix B.

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 sections: In Section 3.1, we study the evolution of bidder shares when $N := N_0$, the volume of initial stakes, is large. In Section 3.2, we study the additional feature of allowing the bidders to trade stakes among themselves, focusing on the issue of trading incentives (or the lack thereof). We remark that the results in both sections exhibit some type of phase transition 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 2(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 [24] that, for a large bidder with initial stake $n_{k,0} = \Theta(N)$, there is 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 [26] for small bidders (following the definition in part (ii) of the following corollary). Here, we focus on the ratio $\pi_{k,t}/\pi_{k,0}$. Because $\pi_{k,\infty} \overset{d}{=} \text{Beta}(n_{k,0}, N - n_{k,0})$, the results in Tang [26] hold. The following proposition is a refined version of Tang [26, theorem 2.1].

Proposition 4 (Phase Transitions of $\pi_{k,t}$). *Let $N_0 = N$ be the total number of initial stakes. Under the $\texttt{Poly}(\alpha)$ 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 $t = \infty$,*

$$\mathsf{P}\left(\left|\frac{\pi_{k,t}}{\pi_{k,0}} - 1\right| > \varepsilon\right) \leq \frac{1}{\varepsilon^2 f(N)}, \tag{21}$$

which converges to zero 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*

$$\pi_{k,\infty}/\pi_{k,0} \xrightarrow{d} \frac{1}{n_{k,0}}\gamma(n_{k,0}) \quad \text{as } N \to \infty, \tag{22}$$

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

$$d_W\left(\frac{\pi_{k,t}}{\pi_{k,0}}, \frac{1}{n_{k,0}}\gamma(n_{k,0})\right) \le C\left(N^3 t^{-\frac{1}{1+\alpha}} + \frac{1}{\sqrt{N}}\right). \tag{23}$$

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 (21).

ii. Note that

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

A careful application of Goldstein and Reinert [12] yields a refinement of (20): there is $C > 0$ such that $d_W(\pi_{k,t}^{\dagger}, Z) \le \frac{CN^3}{t}$. Adapting the argument in Theorem 2 yields

$$d_W\left(\frac{\pi_{k,t}}{\pi_{k,0}}, \frac{1}{\pi_{k,0}}\mathrm{Beta}(n_{k,0}, N-n_{k,0})\right) \le C' N^3 t^{-\frac{1}{1+\alpha}} \quad \text{for some } C' > 0. \tag{24}$$

Next, we claim that

$$d_W\left(\frac{1}{\pi_{k,0}}\mathrm{Beta}(n_{k,0}, N-n_{k,0}),\ \frac{1}{n_{k,0}}\gamma(n_{k,0})\right) \le \frac{C''}{\sqrt{N}} \quad \text{for some } C'' > 0, \tag{25}$$

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,

$$d_W(N\,\mathrm{Beta}(1, N-1), \gamma(1)) \le \mathbb{E}\left|\frac{NX}{X+X'} - X\right|. \tag{26}$$

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 [23]). Injecting into (26) yields the desired bound. Finally, combining the estimates (24) and (25) gives the bound (23). □

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 (21): that the share ratio $\pi_{k,t}/\pi_{k,0}$ concentrates at one and converges to one in probability when $N \to \infty$ for any $t \ge 1$ (including $t = \infty$). For small bidders, this is reversed: the concentration inequality in (21) becomes the anticoncentration inequality

$$\mathsf{P}\left(\left|\frac{\pi_{k,\infty}}{\pi_{k,0}} - 1\right| > \varepsilon\right) > c \quad \text{for } c > 0 \text{ independent of } \varepsilon, \tag{27}$$

implying volatility. The Wasserstein bound (23) 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 $t \ge N^{\frac{7}{2}(1+\alpha)}$. However, we do not know whether the N^3 dependence in (23) 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 [24] that, under certain conditions (which enforce some notion of risk neutrality), there is no incentive for any bidder to trade. Here, we extend that to the `Poly`(α) model, allowing α to take any nonnegative values. Furthermore, we allow a bidder-dependent risk-sensitivity (or risk-aversion) parameter δ_k and study the issue of incentive as it relates to δ_k.

Downloaded from informs.org by [160.39.34.79] on 11 February 2024, at 14:50 . For personal use only, all rights reserved.

Downloaded from informs.org by [160.39.34.79] on 11 February 2024, at 14:50 . For personal use only, all rights reserved.

In the new setting of allowing trading, we need to modify the problem formulation presented at the beginning of Section 2. First, for each $k \in [K]$, let $\nu_{k,t}$ be the number of stakes that bidder k trades at time t. Then, instead of (5), the number of stakes $n_{k,t}$ evolves as

$$n_{k,t} = \underbrace{n_{k,t-1} + 1_{S_{k,t}}}_{n'_{k,t}} + \nu_{k,t}, \tag{28}$$

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

Note that $\nu_{k,t}$ is 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 elaborate more on this as follows, but note that $\nu_{k,t}$ is constrained such that, after the updating in (28), $n_{k,t}$ remains nonnegative.

Let $\{P_t,\, t \ge 0\}$ be the price process of each (unit of) stake, which is a stochastic process assumed to be independent of the randomness induced by the $\texttt{Poly}(\alpha)$ voting rule (specifically, the process $\{S_{k,t}\}$). Hence, we augment the filtration $\{\mathcal{F}_t\}_{t\ge 0}$ with that of the exogenous price process $\{P_t,\, t \ge 0\}$ to a new filtration denoted $\{\mathcal{G}_t\}_{t\ge 0}$. Note that the price process P_t is also assumed as exogenous in Roşu and Saleh [24]. 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 preceding 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 a 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., $T_k \ge 1$ 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 $\{\mathcal{G}_t\}_{t\ge 0}$ 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 $\{\mathcal{G}_t\}_{t\ge 0}$. Thus, bidder k also decides 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, that is,

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

with

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

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 $\nu_{k,t} \ge -n'_{k,t}$ (following $n_{k,t} \ge 0$), that is, 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 ($n_{k,t} \le N_t$). Equation (C3) specifies how the assets are liquidated at the exit time τ_k: both ν_{k,τ_k} and b_{k,τ_k} are 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 $(\nu, b) := \{(\nu_{k,t}, b_{k,t}),\, 1 \le t \le \tau_k\}$. The objective of bidder k is

$$U_k^* := \max_{\tau_k, (\nu,b)} U_k := \max_{\tau_k, (\nu,b)} \mathsf{E}\left(\sum_{t=1}^{\tau_k} \delta_k^t c_{k,t} \right), \quad \text{subject to (C1), (C2), (C3);} \tag{29}$$

$\delta_k \in (0, 1]$ 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.

Downloaded from informs.org by [160.39.34.79] on 11 February 2024, at 14:50 . For personal use only, all rights reserved.

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:

$$\Pi_{k,0} := n_{k,0} P_0, \quad \text{and} \quad \Pi_{k,t} := \delta_k^t n'_{k,t} P_t - \sum_{j=1}^{t-1} \delta_k^j \nu_{k,j} P_j, \quad t \geq 1; \tag{30}$$

$n'_{k,t+1}$ follows (28). Note that the two terms that define $\Pi_{k,t}$ are the discounted present values, respectively, of k's pretrading 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 1. *Under the* `Poly`(α) *voting rule, along with the trading specified earlier, we have*

$$\mathsf{E}(\Pi_{k,t+1} | \mathcal{G}_t) - \Pi_{k,t} = \delta_k^{t+1} \pi_{k,t} \mathsf{E}(M_{t+1} | \mathcal{G}_t) - \delta_k^t \pi_{k,t} M_t. \tag{31}$$

Proof. First, by (28) and (6), along with $\pi_{k,t} = n_{k,t}/N_t$, we have

$$\mathsf{E}(n'_{k,t+1} | \mathcal{F}_t) = n_{k,t}(1 + N_t^{-(1+\alpha)}) = \frac{n_{k,t}}{N_t}(N_t + N_t^{-\alpha}) = \pi_{k,t} \mathsf{E}(N_{t+1} | \mathcal{F}_t). \tag{32}$$

Next, from (30), we have

$$\Pi_{k,t+1} - \Pi_{k,t} = \delta_k^{t+1} n'_{k,t+1} P_{t+1} - \delta_k^t n'_{k,t} P_t - \delta_k^t \nu_{k,t} P_t, \quad t \geq 1. \tag{33}$$

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

$$\begin{aligned} \mathsf{E}(n'_{k,t+1} P_{t+1} | \mathcal{G}_t) &= \mathsf{E}(\mathsf{E}(n'_{k,t+1} | \mathcal{F}_t) P_{t+1} | \mathcal{G}_t) \\ &\overset{(32)}{=} \pi_{k,t} \mathsf{E}(N_{t+1} P_{t+1} | \mathcal{G}_t) = \pi_{k,t} \mathsf{E}(M_{t+1} | \mathcal{G}_t). \end{aligned}$$

This, along with (33) yields the desired expression in (31) 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 (29). To see this, summing up both sides of (C1) and (C3) over t (along with $b_{k,0} = 0$ in (C2)), we have

$$\sum_{t \leq \tau_k} \delta_k^t c_{k,t} = \sum_{t \leq \tau_k} \delta_k^t c_{k,t} = \delta_k^{\tau_k} n'_{\tau_k} P_{\tau_k} - \sum_{t=1}^{\tau_k - 1} \delta_k^t \nu_{k,t} P_t + \sum_{t=1}^{\tau_k - 1} \delta_k^t [(1 + r_{\text{free}})\delta_k - 1] b_{k,t}. \tag{34}$$

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

$$U_k(\tau_k, \nu, b) = \mathsf{E}[\Pi_{k,\tau_k}(\nu)] + \mathsf{E}\left(\sum_{t=1}^{\tau_k - 1} \delta_k^t [(1 + r_{\text{free}})\delta_k - 1] b_{k,t} \right); \tag{35}$$

hence, the RHS is separable: the first term depends on (ν) only, whereas 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 3) along with b being nonnegative, part of the feasibility in (C2). In this case, we 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 (29). A quick word on the parameter r_{cryp} that appears prominently in Theorem 3. Simply put, it is the rate (expected rate of return) associated with each stake (e.g., a unit of some cryptocurrency); that is, it is the counterpart of r_{free}, the rate for the risk-free asset. We elaborate more on the two rates after proving the theorem.

In the theorem, two strategies are singled out: the buyout 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 nonparticipation 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 nonparticipation 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 buyout strategy clearly belongs to the feasible class.

Downloaded from informs.org by [160.39.34.79] on 11 February 2024, at 14:50 . For personal use only, all rights reserved.

Theorem 3 (Buyout Strategy vs. Nonparticipation). *Assume the following two conditions:*

$$\text{(a)}\ \ \delta_k(1+r_{\text{free}}) \le 1 \quad \textit{and} \quad \text{(b)}\ \ \mathsf{E}(M_{t+1}|\mathcal{G}_t) = (1+r_{\text{cryp}})M_t. \tag{36}$$

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

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

In addition, all three parts of the following hold:

i. *If $\delta_k(1+r_{\text{cryp}}) \le 1$, then any feasible strategy provides no greater utility for bidder k than the nonparticipation strategy; that is, $U_k^* \le n_{k,0}P_0$.*

ii. *If $\delta_k(1+r_{\text{cryp}}) \ge 1$, then any feasible strategy provides no greater utility for bidder k than the buyout strategy. In this case, bidder k buys all available stakes at time 1 and participates in the bidding process until the terminal time T_k.*

iii. *If $\delta_k(1+r_{\text{cryp}}) = 1$, then, bidder k is indifferent between the nonparticipation and buyout strategies with any exit time, both of which provide no less utility than any feasible strategy. In other words, all strategies achieve the same utility (which is $\Pi_{k,0}=n_0P_{k,0}$).*

Moreover, when $\delta_k=\delta:=(1+r_{\text{cryp}})^{-1}$ for all k, then no bidder has any incentive to trade. Consequently, the long-term behaviors (of N_t, $\pi_{k,t}$, and $\theta_{k,t}$) characterized in Propositions 2 and 4 and Theorem 2 hold.

Proof. That $U_k^*=\max_{\tau_k,\nu}\mathsf{E}(\Pi_{k,\tau_k})$ (with $b_{k,t}$ being set to zero for all t) under condition (a) in (36) is already established in the discussions following (35). So it suffices to prove the three parts (i)–(iii).

i. Applying the given condition (b) in (36), along with the assumed inequality $\delta_k(1+r_{\text{cryp}})\le 1$, to the RHS of Equation (31) makes it ≤ 0; that is, $\{\Pi_{k,t}\}$ is a $\mathcal{G}_t$-supermartingale, implying $\mathsf{E}(\Pi_{k,\tau_k}) \le \Pi_{k,0}$. Because $\Pi_{k,0}$ is independent of ν, we have

$$U_k^* = \max_{\tau_k,\nu}\mathsf{E}(\Pi_{k,\tau_k}) \le \Pi_{k,0} = n_{k,0}P_0, \tag{37}$$

ii. With the assumed inequality $\delta_k(1+r_{\text{cryp}})\ge 1$, $\{\Pi_{k,t}\}$ now becomes a $\mathcal{G}_t$-submartingale, and hence, the inequality

$$\mathsf{E}(\Pi_{k,T_k}) \ge \mathsf{E}(\Pi_{k,\tau_k}) \ge \Pi_{k,0} = n_{k,0}P_0. \tag{38}$$

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

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

which is linear in ν_{k,T_k-1}. By assumed condition (b) in (36), we have

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

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 $\delta_k(1+r_{\text{cryp}})=1$, $\{\Pi_{k,t}\}$ is a $\mathcal{G}_t$-martingale; hence, the inequality in (37) now holds as equality, that is, $U_k^* = \Pi_{k,0} = n_{k,0}P_0$.

Thus, all strategies lead to the optimal utility, including any feasible strategy (in particular, the no-trading strategy) and the nonparticipation strategy.

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

In what remains of this section, we make a few remarks on Theorem 3, in particular, to motivate and explain its required conditions. First, the rate r_{cryp} is determined by condition (b), the second equation in (36). As such, it should be distinct from r_{free}, the latter being associated with a risk-free asset. For all practical purposes, we can assume $r_{\text{crpt}} \ge r_{\text{free}}$ even though this is not assumed in the theorem. When this relation does hold, then condition (a) becomes superfluous in cases (i) and (iii).

Downloaded from informs.org by [160.39.34.79] on 11 February 2024, at 14:50 . For personal use only, all rights reserved.

Second, the discount factor δ_k in the utility objective in (29), a parameter that measures bidder k's sensitivity toward risk, plays a key role in characterizing phase transitions in terms of $\delta_k(1+r_{\text{cryp}})$. In case (i), the inequality $\delta_k \le 1/(1+r_{\text{cryp}})$ implies bidder k is seriously risk averse, and this is reflected in k's nonparticipation 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 nonparticipation strategy provides less (no greater) utility for bidder k than the no-trading strategy and certainly no greater utility than the buyout strategy. In case (iii), the inequality becomes an equality $\delta_k = 1/(1+r_{\text{cryp}})$, and $\{\Pi_{k,t}\}$ becomes a martingale. Consequently, bidder k is indifferent between nonparticipation and participation and, in the latter case, indifferent to all (feasible) strategies, including the buyout (and no-trading) strategy. Indeed, the equality $\delta_k = 1/(1+r_{\text{cryp}})$ 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 [24], there is a single rate r_{free} or, equivalently, $r_{\text{cryp}} = r_{\text{free}}$ is assumed, which seems difficult to justify because, in most applications (cryptocurrency in particular) r_{cryp} is 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 [24] is limited to the martingale case only, reaching the same conclusion as our case (iii), that all feasible strategies, buyout included, yield the same (expected) utility. As there is no stopping decision and supermartingale or submartingale cases in Roşu and Saleh [24], nonparticipation does not come up at all, and neither do notions such as risk aversion or risk seeking.

The last point we emphasize is that the two conditions in (36) play very different roles. As evident from the proof of Theorem 3, condition (b) makes $\{\Pi_{k,t}\}$ a supermartingale, submartingale, or 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 (29), $\{\Pi_{k,t}\}$ 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 $t \ge 1$ and applicable to all three cases in Theorem 3. 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 submartingale case, and even there, condition (a) pins down the fact that to participate (even without trading) is better than nonparticipation.

Note that Theorem 3 can be readily extended. For instance, the rates $r_{\text{cryp}}(t)$ and $r_{\text{free}}(t)$ can vary over time. In this case, it suffices to modify the conditions in case (i) to

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

the conditions in case (ii) to

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

and the conditions in case (iii) to

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

Then, Theorem 3 continues to hold. We can also include a processing cost $\kappa > 0$ that any bidder selected by the $\texttt{Poly}(\alpha)$ mechanism pays 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 (36) is modified to $\mathsf{E}(M_{t+1}|\mathcal{G}_t) = (1+r_{\text{cryp}})M_t + \kappa$.

4. Conclusions

We propose 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, in which the voters are the bidders (competing for rewards or validation of new blocks), we show this decoupling enhances security, a key objective of the PoS protocol. Specifically, we prove that, whereas bidder shares form a martingale process that converges to a Dirichlet distribution, each bidder's voting power is a supermartingale 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

Downloaded from informs.org by [160.39.34.79] on 11 February 2024, at 14:50 . For personal use only, all rights reserved.

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 has no incentive to participate in the bidding process or, if participating, forgoes trading.

In the introduction, we mention 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, keeping the reward constant. It is possible to pursue a combination of both approaches, that is, adjusting the size of reward dynamically over time in the same manner as adjusting α (for the latter, refer to Appendix B). In another direction, it is also possible to study the trading problem in Section 3.2 under a suitable market impact model, in which the price process P_t is impacted by trading activities; for instance, a mean-field PoS model with linear impact (and transaction costs).

Acknowledgments

The authors thank the anonymous referees for helpful suggestions that improved the presentation of the paper.

Appendix A. Improvement on $\lambda_{\pm}(\alpha)$

Theorem 1 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 also implies 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 1. 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 multiscale 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 $\left(1-\frac{1}{\lambda\sqrt{t}}\right)^{t/2-\lambda\sqrt{t}}$, whereas 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''):

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

By equating the two coefficients before $\sqrt{t}$ in (A.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 preceding equation yields

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

Figure A.1. (Color online) Volume of stakes N_t with $N_0 = 5$ and $\alpha = 1$ (quadratic voting).

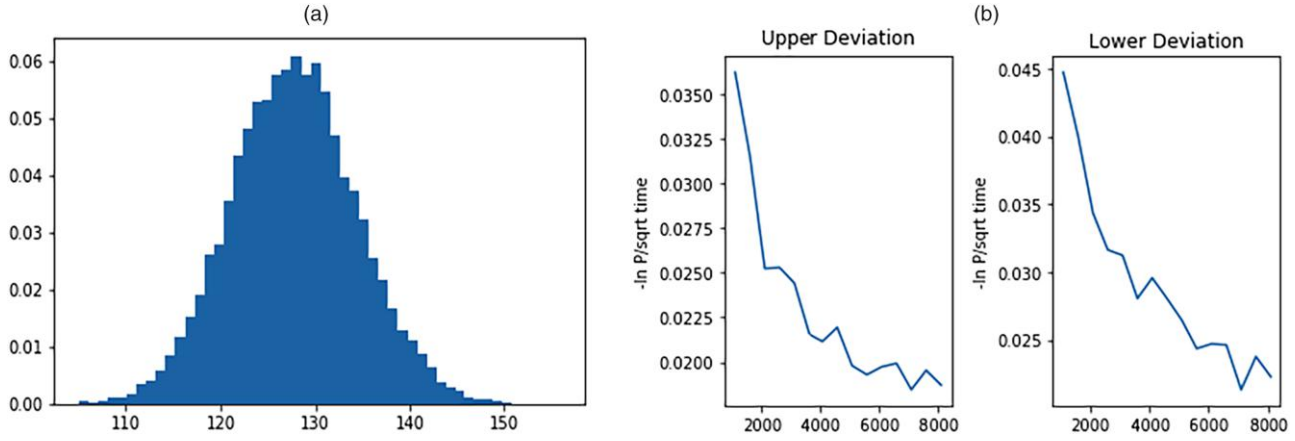

Notes. (a) Histogram of N_{8000} on MC simulation of 20,000 samples. (b) x-axis: $t \in \{1{,}000, 1{,}500, \ldots, 8{,}000\}$; y-axis: $-\ln \mathsf{P}(N_t > \sqrt{2.2t})/\sqrt{t}$ (left) and $-\ln \mathsf{P}(N_t < \sqrt{1.8t})/\sqrt{t}$ (right) on MC simulation of 20,000 samples.

The coefficient before $\sqrt{t}$ is

$$f(\lambda) = 2\lambda \log \lambda - 2\lambda + \frac{1}{2\lambda} + \frac{1}{2(1-\theta)\lambda}, \tag{A.3}$$

where θ is specified by (A.2). By injecting Expression (A.2) into (A.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 eventually gets 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 A.1.

Appendix B. Control of Voting Powers

As proved in Theorem 2, 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) increase (and are 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

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

which is even much longer than the 10-minute block time of Bitcoin. (The block time is 10 seconds in Ethereum; see, e.g., Buterin [5].)

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 $\kappa^{1+\alpha_1^{-1}} \ldots$ and so on.

Here, $\kappa, \alpha_0, \alpha_1, \ldots$ are user-defined hyperparameters. 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 $\ldots$ and so on.

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

- Apply the $\texttt{Poly}(1)$ scheme up to round $5^2 \approx 4$ minutes.
- Apply 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 $\ldots$ 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 $\ldots$ and so on.

Downloaded from informs.org by [160.39.34.79] on 11 February 2024, at 14:50 . For personal use only, all rights reserved.

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 depends on the choice of $\{\alpha_k\}$).

References

[1] Armbrust M, Fox A, Griffith R, Joseph AD, Katz RH, Konwinski A, Lee G, Patterson DA, Rabkin A, Stoica I (2009) Above the clouds: A Berkeley view of cloud computing. Technical Report No. UCB/EECS-2009-28, University of California, Berkeley, CA.
[2] Bagaria V, Dembo A, Kannan S, Oh S, Tse D, Viswanath P, Wang X, Zeitouni O (2019) Proof-of-stake longest chain protocols: Security vs predictability. Preprint, submitted October 5, https://arxiv.org/abs/1910.02218.
[3] Billingsley P (1995) *Probability and Measure*, Wiley Series in Probability and Mathematical Statistics, 3rd ed. (John Wiley & Sons, Inc., New York).
[4] Billingsley P (1999) *Convergence of Probability Measures*, Wiley Series in Probability and Statistics, 2nd ed. (John Wiley & Sons, Inc., New York).
[5] Buterin V (2014) Toward a 12-second block time. Accessed January 25, 2024, https://blog.ethereum.org/2014/07/11/toward-a-12-second-block-time.
[6] Chen H, Yao DD (2001) *Fundamentals of Queueing Networks*, Applications of Mathematics, vol. 46 (Springer-Verlag, New York).
[7] Dean J, Ghemawat S (2008) MapReduce: Simplified data processing on large clusters. *Comm. ACM* 51(1):107–113.
[8] Deirmentzoglou E, Papakyriakopoulos G, Patsakis C (2019) A survey on long-range attacks for proof of stake protocols. *IEEE Access* 7:28712–28725.
[9] Denisov D, Korshunov D, Wachtel V (2016) At the edge of criticality: Markov chains with asymptotically zero drift. Preprint, submitted December 5, https://arxiv.org/abs/1612.01592.
[10] Durrett R (2019) *Probability—Theory and Examples* (Cambridge University Press, Cambridge, UK).
[11] Garcia-Molina H (1982) Elections in a distributed computing system. *IEEE Trans. Comput.* 31(1):48–59.
[12] Goldstein L, Reinert G (2013) Stein's method for the beta distribution and the Pólya-Eggenberger urn. *J. Appl. Probab.* 50(4):1187–1205.
[13] Huang AQ, Baliga J (2009) FREEDM System: Role of power electronics and power semiconductors in developing an energy internet. *21st Internat. Sympos. Power Semiconductor Devices ICs* (IEEE, Piscataway, NJ), 9–12.
[14] King S, Nadal S (2012) Ppcoin: Peer-to-peer crypto-currency with proof-of-stake. Accessed January 25, 2024, https://decred.org/research/king2012.pdf.
[15] Lalley SP, Weyl EG (2018) Quadratic voting: How mechanism design can radicalize democracy. *AEA Papers Proc.* 108:33–37.
[16] Lamperti J (1960) Criteria for the recurrence or transience of stochastic process. I. *J. Math. Anal. Appl.* 1(3–4):314–330.
[17] Lamperti J (1962) A new class of probability limit theorems. *J. Math. Mech.* 11(5):749–772.
[18] Lamperti J (1963) Criteria for stochastic processes. II. Passage-time moments. *J. Math. Anal. Appl.* 7(1):127–145.
[19] Lamport L, Shostak R, Pease M (1982) The Byzantine generals problem. *ACM Trans. Programming Languages Systems* 4(3):382–401.
[20] Menshikov M, Popov S, Wade A (2017) Lyapunov function methods for near-critical stochastic systems. *Non-Homogeneous Random Walks*, Cambridge Tracts in Mathematics, vol. 209 (Cambridge University Press, Cambridge, UK), 382.
[21] Nakamoto S (2008) Bitcoin: A peer-to-peer electronic cash system. *Decentralized Bus. Rev.* 21260.
[22] Penrose LS (1946) The elementary statistics of majority voting. *J. Roy. Statist. Soc.* 109(1):53–57.
[23] Rio E (2009) Upper bounds for minimal distances in the central limit theorem. *Ann. Inst. Henri Poincare Probab. Statist.* 45(3):802–817.
[24] Roşu I, Saleh F (2021) Evolution of shares in a proof-of-stake cryptocurrency. *Management Sci.* 67(2):661–672.
[25] Shi E (2020) Foundations of distributed consensus and blockchains. Accessed January 25, 2024, http://elaineshi.com/docs/blockchain-book.pdf.
[26] Tang W (2022) Stability of shares in the proof of stake protocol—Concentration and phase transitions. Preprint, submitted June 5, https://arxiv.org/abs/2206.02227.
[27] Tang W (2023) Trading and wealth evolution in the proof of stake protocol. Preprint, submitted August 3, https://arxiv.org/abs/2308.01803.
[28] Tang W, Yao DD (2023) Trading under the proof-of-stake protocol—A continuous-time control approach. *Math. Finance* 33(4):979–1004.
[29] Villani C (2009) Old and new. *Optimal Transport*, Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 338 (Springer-Verlag, Berlin), 976.
[30] Wood G (2014) Ethereum: A secure decentralised generalised transaction ledger. Ethereum Project Yellow Paper 151.

Downloaded from informs.org by [160.39.34.79] on 11 February 2024, at 14:50 . For personal use only, all rights reserved.