How to Rationally Select Your Delegatee in PoS

Yuzhe Zhang, Qin Wang, Shiping Chen, Chen Wang
CSIRO Data61, Australia

ABSTRACT

This paper centers around a simple yet crucial question for everyday users: *How should one choose their delegated validators within proof-of-stake (PoS) protocols, particularly in the context of Ethereum 2.0?* This has been a long-overlooked gap, as existing studies have primarily focused on inter-committee (validator set) behaviors and activities, while neglecting the dynamic formation of committees, especially for individual stakeholders seeking reliable validators. Our study bridges this gap by diving into the delegation process (*normal users delegate their small-value tokens to delegatees who later act as validators*) before entering an actual consensus phase.

We propose a Bayesian model to quantify normal users' trust in delegatees, which we further incorporate into a game-theoretical model to simulate users' reactions against a set of critical factors identified through extensive research (including 10+ staking service providers as well as 30+ PoS blockchains). Our results reveal that users tend to choose their delegatees and utilize their tokens by carefully weighing the delegation cost, the behaviors of other users, and the reputation of delegatees, ultimately reaching a Nash equilibrium. Unfortunately, the collective trend significantly increases the likelihood of token concentration on a small number of delegatees.

KEYWORDS

Proof of Sake, Delegation, Game Theory, Ethereum 2.0

1 INTRODUCTION

After the Merge [1], Ethereum officially transitions from the *proof-of-work* (PoW) consensus mechanism to *proof-of-stake* (PoS) [2], marking the advent of its 2.0 version. This upgrade brings forth several enhancements, including reduced entry barriers, improved energy efficiency (99.5% [1]), and more robust crypto-economic incentives (e.g., inspiring DeFi protocols and staking services), catering to a wider spectrum of users. Beyond Ethereum, a multitude of industry-leading blockchains have embraced PoS as their consensus mechanism. These prominent PoS-based blockchain platforms include Cosmos, Polygon, Tezos, BNB Chain, Avalanche, Fantom, Cardano, Solana, Kuasama, Polkadot, Aptos, NEAR, Flow, Secret Network, SUI, Oasis, Kava, Band Protocol, and Casper Network. PoS-based blockchain shares have now increased to over 48% of the entire cryptocurrency market (#CoinMarketCap).

We focus our attention specifically on the Ethereum ecosystem due to its representativeness. To date (as of Oct. 2023), Ethereum has become the second-biggest cryptocurrency in terms of market capitalization (US$195,154,903,316, #CoinMarketCap) and the largest PoS blockchain platform. Ethereum boasts a significant number of registered validators, with 856,167 validators in total. Additionally, the platform has 150,622 actual depositors (#Dune[1]). The total staked deposits amount to 27,556,644 ETH (equiv. US$45B), representing an estimated staking ratio of 22.77%. This indicates a substantial portion of the Ethereum network's native cryptocurrency is being actively staked by validators.

Validators play a pivotal role in PoS-based Ethereum. A validator assumes responsibility for a wide array of processes throughout the consensus procedures, which includes managing stakes, coordinating committees, proposing blocks, validating neighboring blocks, and casting votes for finalization (as detailed in Sec.2). Consequently, becoming a validator offers various potential income streams, including staking annual reward rates (typically ranging from 3% to 9%), block rewards, transaction fees (with a consensus layer APR of up to 3.18%), and the ability to extract profits, such as *miner extractable value* (MEV [3], resulting in a net profit of US$672,889). Given these incentives, becoming a validator emerges as an attractive option for rational Ethereum participants.

The forgotten majority. However, becoming a validator demands a minimum of 32 ETH [2] (equiv. US$52,556), a threshold that can pose a significant barrier to entry for many participants. Based on our calculation, the ratio of active validators (150,622, #Dune) to the total cumulative users (246.52M, based on the hint of *unique Ethereum addresses*[2]) stands at a mere 0.61%. This indicates that only a very small fraction of normal users, even when accounting for users with multiple accounts, can afford the substantial deposit required to become validators. Notably, despite normal users forming the majority of this permissionless network, only a minuscule portion can actively participate in its maintenance. This lack of accessibility raises concerns as it appears that the broader market often overlooks the needs of these users. For many of these participants, the only viable option (more in Sec.2.2) to participate in the network is by delegating their stakes to eligible validators.

Delegation reliance. Existing ways of delegation for normal users are either *custodial*[3] (without private keys) or *non-custodial* (holding keys). Besides depositing stakes within exchanges, users prefer to delegate their stakes to one or multiple reputable existing validators, often referred to as stake service providers (Sec.2.3). This approach allows users to delegate only their tokens, rather than giving up full control of their accounts. Consequently, it has given rise to a category of services specializing in collecting small investments from normal users and distributing rewards based on their proportional holdings. These service providers typically charge fees, known as commission fees, as compensation for their services.

⚠ **User's decision?** Then, *how do users select their delegated validators?* This is a fundamental yet crucial question. Assuming a simple case, when a normal user Alice opens the wallet and decides to delegate her assets, she confronts a multitude of options. Should she choose the validator with the lowest service fees, the one with

[1] Data source from Dune: #1 https://dune.com/conelohan/ethereum-pos-and-merge-sql-challenge-twigblock, #2 https://dune.com/hildobby/eth2-staking. #3 https://dune.com/21co/ethereum-staking-and-withdrawals, #4 https://dune.com/chorus_one_research/ethereum-mev-data. [Oct. 2023].

[2] #YChart https://ycharts.com/indicators/ethereum_cumulative_unique_addresses.

[3] We skip custodial delegation due to its limited usage for CEXes (Sec.2.2).

the largest user base, the validator with the most substantial deposits, or perhaps one they are already familiar with? It's a complex decision influenced by various factors. To our surprise, despite the importance of this decision, our extensive research, including a thorough investigation detailed in Apx.C[4], revealed a significant gap in existing research. While a few recent studies [4, 5] have delved into the delegation process in *permissioned* blockchains like EOS and STEEM, there is a noticeable absence of focus on *permissionless* chains, including Ethereum 2.0. This leads us to explore their delegation patterns, processes, and potential impacts.

Contributions. We approach these goals through a series of efforts.

① **We are the first to identify the delegation problem in permissless blockchains** (Sec.2). Our motivation originates from our practical experiences participating in Ethereum PoS staking, as well as several other PoS-based projects (e.g., Oasis). We encountered difficulties when deciding which validator was most suitable for us. We embarked on a study from various aspects including PoS-based Ethereum (Sec.2.1), participant behaviors (Sec.2.2), staking services (Sec.2.3) and surrounding works (Apx.C). Our systematic studies enable us to develop a better understanding of the delegation process and identify several critical factors that significantly influence users' decisions, such as *provider's reputation*, *staking scale*, *depositor scale*, and *commission fees*. Here, we emphasize *permissionless* because, unlike permissioned blockchains, permissionless chains align more closely with Nakamoto's original idea of decentralization [6]. Additionally, such projects continue to dominate crypto markets (53%, CoinMarketCap) and boast large communities.

② **We develop a suite of models to maximize the simulation of the delegation process, suitable for both *permissioned* and *permissionless* settings** (Sec.3). In the context of PoS protocols, a delegation involves both a set of delegators A (equiv. users) and delegatees V (equiv. validators). Our model is designed to seamlessly accommodate both blockchain configurations without sacrificing its forward compatibility with previous definitions.

✎ *We introduce adjustability to set the validity of both delegators and delegatees, allowing us to simulate the dynamic joining and leaving of participants in permissionless settings.* In line with the common practice of configuring delegation models (e.g., [4, 5, 7]), we assume a finite set V, which has a fixed size (e.g., $|V| = m$) during an entire epoch for permissioned blockchains where the committee remains static. However, our game-theoretical model (i.e., the validator selection game (VSG) defined in Sec.4.1) can potentially capture the dynamic feature of permissionless settings in three ways:

◇ The VSG can be modified by setting that A and V are adjustable inter-epoch under constraint $P = A \cup V$ and adding the utility function of validators. Then, each participant in the game is able to exchange their role between a user and a validator, and the game becomes an evolutionary game [8].

◇ We can easily modify each user's strategy space ($\Sigma_i \in V \times \mathbb{R}_{\geq 0}$, Sec.4.1) by removing specific validators, such that we render the removed validators silent participants.

◇ We can easily extend this permission adaptability of silent mode to users (A) by simply setting a user's budget ($\mathbf{b}$, Dfn.1) to 0.

✎ *We introduce a novel metric, "trust," to emulate delegators' belief in validators' integrity information within the market (Sec.3.2).* This metric (T) is updated within a Bayesian probabilistic framework (see Fig.1), incorporating factors revealing validators' integrity, such as *brand reputation* and *rumors*, and the other users' *judgment* on such integrity factors that contribute to a validator's trustworthiness.

Notably, our proposed trust metric mainly reflects users' belief in validators' *intrinsic* motivation to leave the market, as well as other users' judgment on this motivation. In practice, validators might be motivated to leave the market by various *extrinsic* factors, e.g., their cost to run a client, their received number of delegation tokens, and so on. Our model is ready to be extended to take such extrinsic factors into consideration and to simulate validators' behavior.

✎ In accordance with this approach, *we have developed a Validator Selection Game (VSG) designed to replicate real-world user delegation scenarios* (Sec.2.2). By formally analyzing VSGs, we gain insight into how delegators compete to secure their utility in delegation scenarios. Our game (Dfn.1) takes into consideration all the factors mentioned earlier, including participants (A, V), user attributes like accuracy and error ($\mathbf{q}$, see Tab.5), budget ($\mathbf{b}$), strategy (Σ), validator characteristics such as integrity ($\mathbf{p}$), and external elements like commission fees ($\mathbf{c}$). This holistic approach covers a wide range of behaviors that users may exhibit throughout their engagement with the game. Furthermore, we have formulated a *utility* function (Dfn.2) based on this model, which accounts for all these behaviors.

③ **We both theoretically and practically analyze the dynamic delegation via game theory, elucidating the methods to attain a Nash Equilibrium (NE)** (Sec.4&Sec.5).

Theoretically, we investigate the existence and feature of Nash equilibria (NE) in Validator Selection Games (VSGs) under various conditions. This includes scenarios with a single validator (Sec.4.2.1), with multiple homogeneous validators (Sec.4.2.2), as well as with commission-free validators (Sec.4.2.3) in VSGs. Our analyses yield a series of proofs that demonstrate the existence of NE under specific conditions (Thm.1 to Thm.3).

Practically, we conduct a series of experiments to assess the performance of our game by varying parameter configurations (Sec.5). In simulations, we design an algorithm modeling that delegators noisily execute an optimization of their utility against the other delegators' behaviors, referred to as the *best response* (Dfn.4). Given the inherent uncertainty in these games, we focus on several key parameters that best elucidate our game. The collective results (Sec.5.2) demonstrate alignment with our theoretical analyses.

♠ **We further offer several key takeaways from our study.**

◇ A non-misled delegator is more inclined to trust a delegatee selected by a larger number of delegators and possessing a strong reputation. This represents a type of the *80-20 rule* [9] (a.k.a., Pareto Principle) existing in staking markets proved by our theory.

◇ A rational delegator makes choices regarding their delegatee and the number of tokens to delegate by carefully balancing their trust in the delegatees and the associated delegation costs.

◇ A large amount of tokens might be concentrated on a limited number of delegatees who have a good balance between high reputation and low delegation cost. This trend becomes more obvious if delegators are able to choose their best strategies more accurately, or repeatedly alter their strategies.

[4] **Abbreviation**: definition (Dfn.), equation (Eq.). theorem (Thm.), lemma (Lm.), appendix (Apx.), algorithm (Alg.), table (tab.), figure (Fig.), section (sec.).

2 UNDERSTANDING ETHEREUM 2.0

2.1 System Overview

Network assumption. Similar to its previous version, Ethereum 2.0 operates on a *partially synchronous model* [10], ensuring that messages will eventually be delivered, albeit with an unknown but finite upper-bound time delay.

Entities. Two types of participants are involved: (i) *normal users* (delegators in this work) engage in the blockchain network for simply staking (thus becoming stakeholders) or trading tokens. They typically access the network through lightweight clients, such as web browsers, wallets, or mobile apps. (ii) *validators* (delegatees) are either individual participants or groups representing normal users who delegate their tokens. They take on the responsibility of performing the consensus mechanism. In the context of PoW protocols, validators are referred to as miners. They play a pivotal role in maintaining the safety and liveness of blockchain systems [11].

Chain operation. Ethereum 2.0 operates on three key pillars: execution, consensus, and incentive.

Execution: Execution focuses on the responsibilities of validating and executing transactions. Notably, a recent shift known as *proposer-builder separation* (PBS) [12] has emerged, aiming to decouple the tasks of ordering transactions (the builder) from those proposing the block (the proposer). This step streamlines transaction packaging and block production.

Consensus (*Casper* [13]): Ethereum 2.0 utilizes a combination of two fundamental primitives: the fork choice rule *LMD GHOST* [14] and the finality gadget *Casper FFG* [15]. LMD GHOST[5] builds upon the heaviest chain rule akin to Nakamoto consensus [16] while simultaneously considering the latest message from each validator. Casper FFG introduces a gadget capable of adding finality to an underlying consensus protocol through the specification of epochs and checkpoints. This step finalizes block confirmations, resolves forks, and facilitates chain growth.

Incentive: It primarily defines the policies [17] for rewarding honest validators and imposing penalties on malicious ones [18].

2.2 Validator and Normal User

Validators becomes more important. Validators have a range of responsibilities: (i) periodically acting as block proposers to generate blocks after validating and ordering transactions from the mempool; (ii) continuously verifying the validity of blocks created by other validators and attesting them to the canonical chain; (iii) participating in consensus operations and voting for finalization; (iv) taking part in sync committees to ensure the network remains operational; and (v) managing stakes and distributing profits, with half of the rewards going to normal users (individual shareholders).

Compared to miners in PoW, validators play a more crucial role in PoS. PoS requires continuous *active participation* from over two-thirds of the validators to maintain blockchain progress. Validators do not incur the same tangible costs (e.g., electricity expenses in PoW), making it easier for them to accumulate stakes from normal users, thereby increasing their influence within the chain. As evidence, validators exercise control over more than 87% of the ETH on-chain assets (#Dune) in practice.

Benefits for validators. Ethereum validators have the opportunity to earn rewards from various sources. The first line is based on faithful consensus-related activities: validators can engage in the consensus process by proposing blocks (approx. 0.04 ETH per successful proposal [19]). They can also attest to blocks, including attesting to the source epoch, target epoch, and head block (0.00001 ETH per attestation). Meanwhile, a small portion of rewards can be contributed by participating sync committee process [19]. The second line is to report misbehaves: validators can report dishonest validators and receive whistle-blowing rewards [12]. Another line is to chase extra profits: validators can generate MEV profits by running MEV-boost services [20] (0.1 ETH per block [21]). The gross profit by arbitrage reaches US$2,297,234 as of Oct 2023 (#Dune).

Correspondingly, validators also face penalties for misconduct, such as producing two blocks for the same slot, which can result in their stakes being slashed (at least 1/32 of their staked ETH [22]).

Barriers in becoming validators. In Ethereum, becoming a validator to take responsibility for the entire network and earn revenues is required to deposit a minimum of 32 ETH, which is worth over US$53k (based on the price as of Oct 1st, 2023). This high barrier poses a significant challenge for ordinary users with limited cryptoassets. Our investigation reveals that only 0.61% of Ethereum addresses hold more than 32 ETH (equiv. US$52,556) in their accounts, which has also been mentioned in Sec.1.

How do users participate? Normal users join in the game by three ways (i) becoming a validator by depositing 32 ETH and maintaining a full node; (ii) delegating their tokens via staking service providers, where users only send tokens to the delegatee's address while still holding their account private keys (*non-custodial*); and (iii) delegating their tokens within platforms provided by service providers like centralized exchanges (e.g., Binance, Coinbase) whereas the private keys are owned by these providers. Our research is centered around the second route.

2.3 Staking Services

Following above, there are several ways of staking [23]: *solo-home staking*, delegating tokens to a *staking as a service* provider, *pooled staking* and creating accounts in *centralized exchanges* (CEXes).

We investigate two major types of such staking services in the market: *non-custodial* pooled stakings (cf. Tab.1, #Dune), and *custodial* stakings (Tab.3). We further provide their comparisons in Tab.4 and offer the mappings between non-custodial staking providers and current PoS blockchains (Tab.6, as demonstrated in Apx.B).

3 WARM-UP CONSTRUCTION

3.1 Delegation Modeling

We will focus on the delegation phase before validators enter the consensus procedures. Consider that in a finite set of participants P, each participant chooses to be either a *user* or a *validator*. P is therefore divided into two subsets: one consists of users $A = \{a_1, \dots, a_n\}$ ($|A| = n$), and the other of validators $V = \{v_1, \dots, v_m\}$ ($|V| = m$). Hence, we have $A \cup V = P$. Then, each *user* would like to choose

[5] LMD GHOST is a variant of the GHOST [16] (Greedy Heaviest-Observed Sub-Tree) rule that is based on each participant's most recent vote (LMD, latest message-driven).

Table 1: Staking services (*non-custodial*, Pooled)

Pools	Mini.	Browser	Wallet	GUI	Open-source	Audited	Bug bounty	Battle-tested	Trustless	Permissionless	Cons. diversity	Liqui. token	Validators	Mark. share	Comissions
Lido	Any	✓	✓	✓	✓	✓	✓	✓	✓	-	✓	✓	30.7k	32%	10%
Rocket pool	0.01	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	25k	3.1%	14%
Stakefish	0.1	✓	✓	✓	-	✓	-	✓	✓	-	✓	-	23k	2.7%	0.1ETH
Staked	Any	✓	-	✓	-	-	-	-	✓	✓	✓	-	21k	2.5%	-
P2P	Any	✓	-	✓	-	-	-	-	✓	✓	✓	-	8k	0.9%	10%
StakeWise	Any	✓	-	✓	✓	✓	✓	✓	✓	-	✓	✓	6k	0.7%	10%
StaFi	0.01	✓	-	✓	✓	✓	-	✓	-	✓	-	✓	5k	0.6%	10%
Ankr	Any	✓	✓	✓	✓	✓	✓	✓	-	-	-	✓	1.7k	0.2%	10%
RockX	Any	✓	-	✓	✓	✓	-	✓	✓	-	-	✓	0.13k	-	20%
StakingFacilities	Any	✓	-	✓	-	-	-	-	✓	✓	✓	-	-	-	10%
StakeWithUs	Any	✓	-	-	✓	✓	✓	✓	✓	-	✓	✓	-	-	10%
Stakin	Any	✓	-	✓	-	-	-	-	✓	✓	✓	-	-	-	10%
		Plugin			Features								Etereum 2.0		

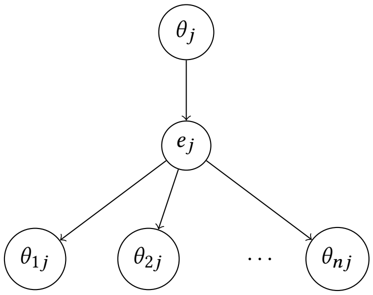

Figure 1: The influence relationship between θ_j, e_j and θ_{ij} follows this Bayesian network.

a *validator* among V to delegate to. If user a_i chooses to delegate to validator v_j, a_i would also decide a weight t_i attached to their delegation, which is called the number of *tokens* that a_i delegates to v_j. A *token profile* $\mathbf{t} = (t_1, \ldots, t_n)$ records all users' tokens that they delegate, and $\mathcal{T}$ is the set of all token profiles. We call a_i's delegation and the token attached to their delegation a_i's *delegation strategy*. Note that we assume that each user cannot delegate to multiple validators. Note also that for user a_i and validator v_j, $t_{ij} = 0$ indicates $d_i = 0$.

We call a_i's choice of validator the *delegation* of a_i, denoted as $d_i \in V \cup \{0\}$. $d_i = v_j$ if a_i delegates to validator v_j, and if $d_i = 0$, a_i does not delegate to any validator. Then, a *delegation profile* $\mathbf{d} = (d_1, \ldots, d_n)$ records each user's delegation. Let $\mathcal{D}$ be all delegation profiles. Given a delegation profile $\mathbf{d}$, for each validator $v_j \in V$, let $d(v_j) = \{a_i \in A \mid d_i = v_j\}$ denote the set of users who delegate to v_j. We call $(\mathbf{d}, \mathbf{t})$ a *strategy profile*.

Users decide their delegation strategies based on two factors: (i) the commission cost by delegating to a validator, and (ii) whether a validator would stay in the market during the slot to validate the next block. Suppose that user a_i delegates t_{ij} tokens to validator v_j whose commission is $c_j \in [0, 1]$, then, a_i would spend $t_{ij}(1 + c_j)$ in total to accomplish the delegation.

For each validator $v_j \in V$, we have *a priori* probability of events $\theta_j = 1$ and $\theta_j = 0$, denoting that v_j would stay in the market during the slot, and not, respectively. The probability of event $\theta_j = 1$ is called the *integrity* of v_j, denoted as p_j, and we have that $\Pr(\theta_j = 0) = 1 - \Pr(\theta_j = 1) = 1 - p_j$. However, this priori information is not observable by the public, and instead, it is revealed by public *evidence* $e_j \in \{0, 1\}$ with noise, e.g., the reputation of validators' brands. $e_j = 1$ indicates that the evidence signal shows v_j would stay in the market, while $e_j = 0$ indicates that the signal shows v_j would leave. Then, we use $z_j = \Pr(e_j = 1 \mid \theta_j = 1) = \Pr(e_j = 0 \mid \theta_j = 0)$, i.e., the probability that v_j's evidence truthfully reveals a priori integrity of v_j, to denote the *quality* of the noisy evidence e_j. In other words, z_j implies how accurate e_j can reveal $\Pr(\theta_j)$.

3.2 A Probabilistic Model of Trust

Consider that evidence e_j of each validator v_j is observable to the public. However, users can only act according to the evidence with noise. Let $q_{ij} = \Pr(\theta_{ij} \mid e_j = 1)$ be the *accuracy* of user a_i's choice of validator v_j based on evidence e_j. This definition denotes the probability of event $\theta_{ij} : d_i = v_j$ conditioned on $e_j = 1$. On the other hand, let $\bar{q}_{ij} = \Pr(\theta_{ij} \mid e_j = 0)$ denote the *error* of a_i's choice of v_j based on evidence e_j. Note that we assume that users' accuracies and errors are independent, i.e., for each pair of users $a_i, a_{i'} \in A$ and each validator $v_j \in V$, $\Pr(\theta_{ij}, \theta_{i'j} \mid e_j) = \Pr(\theta_{ij} \mid e_j) \Pr(\theta_{i'j} \mid e_j)$.

We assume that this noise model satisfies the properties of a Bayesian network shown in Fig.1. That is, users' accuracies are not conditioned on a priori $\Pr(\theta_j)$. We remark that such probabilistic structure is also used in the literature of jury theorem, where voters' voting competencies are correlated due to their shared evidence, e.g., [24]. Then, for user a_i and validator v_j, given p_j, q_j and q_{ij}, if a_i delegates to v_j, a_i's *trust* on v_j's integrity is updated based on their behavior θ_{ij} as:

$$T_{ij} = \Pr(\theta_j = 1 \mid \theta_{ij}) = \frac{\mathrm{B}(a_i, \theta_j = 1)}{\mathrm{B}(a_i, \theta_j = 1) + \mathrm{B}(a_i, \theta_j = 0)} \tag{1}$$

$$= \frac{q_{ij} z_j p_j + \bar{q}_{ij}(1 - z_j) p_j}{q_{ij} z_j p_j + \bar{q}_{ij}(1 - z_j) p_j + q_{ij}(1 - z_j)(1 - p_j) + \bar{q}_{ij} z_j (1 - p_j)}, \tag{2}$$

where

$$\begin{aligned}\mathrm{B}(a_i, \theta_j = k) = &\Pr(\theta_{ij} \mid e_j = 1) \Pr(e_j = 1 \mid \theta_j = k) \Pr(\theta_j = k) + \\ &\Pr(\theta_{ij} \mid e_j = 0) \Pr(e_j = 0 \mid \theta_j = k) \Pr(\theta_j = k),\end{aligned}$$

such that $k \in \{0, 1\}$, by the Bayesian chain rule.

Equation 1 illustrates a user's trust on a validator based on their own delegation behavior. Assume that the delegation profile and token profile is observable to the public. This trust can then be updated based on the observation of the other users' behavior. Given a delegation profile $\mathbf{d}$, user a_i can update their trust on validator v_j based on the users who delegate to v_j. Formally, given validator v_j's integrity p_j, evidence e_j's quality z_j and each user a_k's accuracy q_{kj} and error $\bar{q}_{kj}$ (where $a_k \in A$), the trust of a_i on v_j is a function $T_{ij} : \mathcal{D} \to [0, 1]$, defined as

$$T_{ij}(\mathbf{d}) = \Pr(\theta_j = 1 \mid \wedge_{a_\ell \in d(v_j)} \theta_{\ell j}) \tag{3}$$

$$= \frac{\mathrm{B}(\{a_\ell \mid d_\ell = v_j\}, \theta_j = 1)}{\mathrm{B}(\{a_\ell \mid d_\ell = v_j\}, \theta_j = 1) + \mathrm{B}(\{a_\ell \mid d_\ell = v_j\}, \theta_j = 0)}, \tag{4}$$

where

$$
\begin{aligned}
&\mathrm{B}(\{a_\ell \mid d_\ell = v_j\}, \theta_j = 1)\\
=&\Pr(\wedge_{a_\ell \in d(v_j)} \theta_{\ell j} \mid e_j = 1) \Pr(e_j = 1 \mid \theta_j = 1) \Pr(\theta_j = 1) +\\
&\Pr(\wedge_{a_\ell \in d(v_j)} \theta_{\ell j} \mid e_j = 0) \Pr(e_j = 0 \mid \theta_j = 1) \Pr(\theta_j = 1)\\
=& \prod_{a_\ell \in d(v_j)} \Pr(\theta_{\ell j} \mid e_j = 1) \Pr(e_j = 1 \mid \theta_j = 1) \Pr(\theta_j = 1) +\\
& \prod_{a_\ell \in d(v_j)} \Pr(\theta_{\ell j} \mid e_j = 0) \Pr(e_j = 0 \mid \theta_j = 1) \Pr(\theta_j = 1)\\
=& \prod_{a_\ell \in d(v_j)} q_{\ell j} q_j p_j + \prod_{a_\ell \in d(v_j)} \bar{q}_{\ell j} (1 - q_j) p_j,
\end{aligned}
$$

and similarly

$$
\begin{aligned}
&\mathrm{B}(\{a_\ell \mid d_\ell = v_j\}, \theta_j = 0)\\
=&\Pr(\wedge_{a_\ell \in d(v_j)} \theta_{\ell j} \mid e_j = 1) \Pr(e_j = 1 \mid \theta_j = 0) \Pr(\theta_j = 0) +\\
&\Pr(\wedge_{a_\ell \in d(v_j)} \theta_{\ell j} \mid e_j = 0) \Pr(e_j = 0 \mid \theta_j = 0) \Pr(\theta_j = 0)\\
=& \prod_{a_\ell \in d(v_j)} q_{\ell j} (1 - z_j)(1 - p_j) + \prod_{a_\ell \in d(v_j)} \bar{q}_{\ell j} z_j (1 - p_j).
\end{aligned}
$$

Observe that all users who delegate to the same validator have the same trust in the validator.

In Sec.4, for the feasibility of theoretical analysis, we will be working on a simplified class of the above setting: (1) for each validator v_j, the evidence has perfect quality, i.e., $z_j = \Pr(e_j = 1 \mid \theta_j = 1) = \Pr(e_j = 0 \mid \theta_j = 0) = 1$; (2) each user a_i has the same accuracy/error on each validator v_j, i.e., $q = q_{ij}$ and $\bar{q} = \bar{q}_{ij}$ for all $a_i \in A$ and $v_j \in V$. Then, given a profile $\mathbf{d}$, we simplify the components of Equation 3 as: $\mathbf{B}(\{a_\ell \mid d_\ell = v_j\}, \theta_j = 1) = q^k p_j$ and $\mathbf{B}(\{a_\ell \mid d_\ell = v_j\}, \theta_j = 0) = \bar{q}^k (1 - p_j)$, where $k = |d(v_j)|$, and therefore, Eq.3 becomes

$$T_{ij}(\mathbf{d}) = \frac{q^k p_j}{q^k p_j + \bar{q}^k (1 - p_j)}. \tag{5}$$

We denote this subclass the *Homogeneous User and Perfect Evidence* (HUPE) trust.

Observations in the HUPE trust class. We first show that, based on this probabilistic model, when a user cannot access the others' delegation strategies, this user tends to delegate to a validator with higher integrity, if they are more likely to make correct decisions.

Lemma 1. *For each pair of validators $v_j, v_{j'} \in V$ ($j \neq j'$) with $p_j > p_{j'}$ and each user $a_i \in A$, $\Pr(\theta_{ij}) > \Pr(\theta_{ij'})$ if $q > \bar{q}$, but $\Pr(\theta_{ij}) > \Pr(\theta_{ij'})$ if $q < \bar{q}$.*

Proof. For a_i and v_j, we have that

$$
\begin{aligned}
\Pr(\theta_{ij}) &= \Pr(\theta_{ij} \mid e_j = 1) \Pr(\theta_j = 1) + \Pr(\theta_{ij'} \mid e_j = 0) \Pr(\theta_j = 0)\\
&= q p_j + \bar{q}(1 - p_j)\\
&= \bar{q} + p_j (q - \bar{q}).
\end{aligned}
$$

Similarly, we have that $\Pr(\theta_{ij'}) = \bar{q} + p_{j'}(q - \bar{q})$.
Then, if $q > \bar{q}$, we have that $\Pr(\theta_{ij}) > \Pr(\theta_{ij'})$ since $p_j > p_{j'}$. On the other hand, if $q < \bar{q}$, we have $\Pr(\theta_{ij}) < \Pr(\theta_{ij'})$. □

In practice, this may be because users can be correctly directed by the evidence, e.g., the reputation of validators.

Next, we demonstrate that when users possess a higher level of accuracy compared to errors, for a specific user group size, they exhibit more trust in validators with higher integrity.

Lemma 2. *Given a pair of delegation profiles $\mathbf{d}$ and $\mathbf{d}'$ and a pair of validators v_j and $v_{j'}$, such that $d(v_j) = d'(d_{j'})$ and $p_{j'} > p_j$, we have that for all $a_i \in d(v_j)$, $Tij'(\mathbf{d}') > T_{ij}(\mathbf{d})$ if $q > \bar{q}$.*

Proof. Let $|d(v_j)| = |d'(v_{j'})| = k$. We write the trusts of a_i on v_j and $v_{j'}$ in delegation profiles $\mathbf{d}$ and $\mathbf{d}'$ as:

$$T_{ij}(\mathbf{d}) = \frac{p_j q^k}{p_j q^k + (1 - p_j)\bar{q}^k} = \frac{p_j q^k}{p_j (q_k - \bar{q}^k) + \bar{q}^k}, \tag{6}$$

and

$$T_{ij'}(\mathbf{d}') = \frac{p_{j'} q^k}{p_{j'} (q_k - \bar{q}^k) + \bar{q}^k}. \tag{7}$$

Dividing Eq.7 by Eq.6 and we have that

$$T_{ij'}(\mathbf{d}')/T_{ij}(\mathbf{d}) = \frac{p_j p_{j'} (q^k - \bar{q}^k) + p_{j'} \bar{q}^k}{p_j p_{j'} (q^k - \bar{q}^k) + p_j \bar{q}^k}.$$

Since $p_{j'} > p_j$, and both $p_j p_{j'} (q^k - \bar{q}^k) + p_{j'} \bar{q}^k$ and $p_j p_{j'} (q^k - \bar{q}^k) + p_j \bar{q}^k$ are positive, we have that $T_{ij'}(\mathbf{d}')/T_{ij}(\mathbf{d}) > 1$ which implies that $T_{ij'}(\mathbf{d}') > T_{ij}(\mathbf{d})$. □

Our last observation on trust shows that more users delegating to a validator enhances the trust of users on the validator if users have higher accuracy than error, however, it undermines the trust if users have lower accuracy than error.

Lemma 3. *Given a pair of delegation profiles $\mathbf{d}$ and $\mathbf{d}'$, such that for validator v_j, $d'(v_j) = d(v_j) \cup \{a_k\}$ ($a_k \notin d(v_j)$), we have that for each $a_i \in d(v_j)$, $T_{ij}(\mathbf{d}') > T_{ij}(\mathbf{d})$ if $q > \bar{q}$, $T_{ij}(\mathbf{d}') < T_{ij}(\mathbf{d})$ if $q < \bar{q}$, and $T_{ij}(\mathbf{d}') = T_{ij}(\mathbf{d})$ if $q = \bar{q}$.*

Proof. We prove this lemma by showing that if $q > \bar{q}$, $T_{ij}(\mathbf{d})/T_{ij}(\mathbf{d}') < 1$, if $q < \bar{q}$, $T_{ij}(\mathbf{d})/T_{ij}(\mathbf{d}') > 1$, and if $q = \bar{q}$, $T_{ij}(\mathbf{d})/T_{ij}(\mathbf{d}') = 1$.

First, we notice that since each user's accuracy and error are identical for each validator, $T_{ij}(\cdot)$ solely depends on p_j and the number of users delegating to v_j. Let $|d(j)| = k$. We then have that

$$T_{ij}(\mathbf{d}) = \frac{p_j q^k}{p_j q^k + (1 - p_j)\bar{q}^k}, \tag{8}$$

and

$$T_{ij}(\mathbf{d}') = \frac{p_j q^{k+1}}{p_j q^{k+1} + (1 - p_j)\bar{q}^{k+1}}. \tag{9}$$

By dividing Eq.8 by Eq.9, we have that

$$T_{ij}(\mathbf{d})/T_{ij}(\mathbf{d}') = \frac{(1 - p_j)\bar{q}^{k+1} + p_j q^{k+1}}{(1 - p_j)\bar{q}^k q + p_j q^{k+1}}. \tag{10}$$

Hence, if $q > \bar{q}$, $T_{ij}(\mathbf{d})/T_{ij}(\mathbf{d}') < 1$, and therefore, $T_{ij}(\mathbf{d}') > T_{ij}(\mathbf{d})$. Reversely, if $q < \bar{q}$, $T_{ij}(\mathbf{d}') < T_{ij}(\mathbf{d})$. Specially, if $q = \bar{q}$, $T_{ij}(\mathbf{d}')/T_{ij}(\mathbf{d}) = 1$, which indicates $T_{ij}(\mathbf{d}') = T_{ij}(\mathbf{d})$. □

In other words, given the fact that each user is more likely to make a decision coincident with the observed evidence, i.e., $q > \bar{q}$, users better trust a validator if more users delegate to the validator, but users less trust the validator if fewer users delegate to the validator. A special case is that the trust on a validator is insensitive to the number of delegations if $q = \bar{q}$.

4 A GAME THEORETIC MODEL FOR VALIDATOR SELECTION

In this section, we define a game theoretical model in which each user chooses their delegation strategy in order to maximize their expected profit through the block validation slot.

In the rest of this section, we first provide the definition of the game in the general setting introduced in Sec.3, and then, we show theoretical analysis in subclasses under the HUPE trust class, e.g., the homogeneous belief class.

4.1 Validator Selection Game (VSG)

Definition 1 (VSG). *A validator selection game (VSG) is denoted as a tuple*

$$g = \langle A, V, \mathbf{q}, \mathbf{p}, \mathbf{z}, \mathbf{c}, \mathbf{b}, \Sigma, r, \mathbf{u} \rangle,$$

where A is a finite set of users and V a finite set of validators. $\mathbf{q} = ((q_{11}, \bar{q}_{11}), \ldots, (q_{nm}, \bar{q}_{nm})) \in \mathbb{R}_{\geq 0}^{n \times m \times 2}$ is a profile of users accuracies and errors on m validators. $\mathbf{p} = (p_1, \ldots, p_m) \in \mathbb{R}_{\geq 0}^m$ is a integrity profile. $\mathbf{z} = (z_1, \ldots, z_m) \in \mathbb{R}_{\geq 0}^m$ is a profile of evidence quality. $\mathbf{c} = (c_1, \ldots, c_m) \in \mathbb{R}_{\geq 0}^m$ is a commission profile. $\mathbf{b} = (b_1, \ldots, b_n) \in \mathbb{R}_{>0}^n$ is a budget profile, i.e., the number of tokens that each user can use. $\Sigma_i \in V \times \mathbb{R}_{\geq 0}$ is the strategy space of user a_i, where $(d_i, t_i) \in \Sigma_i$ denotes that user a_i delegates t_i tokens to validator v_j under the budget constraint $t_i(1 + c_j) \leq b_i$. User a_i abstains if $t_i = 0$, that is, they do not delegate. Finally, given the accuracy and error profile and the profit r (that the system returns when the next block is validated), $u_i : \mathcal{D} \times \mathcal{T} \longrightarrow \mathbb{R}$ is the utility function of user a_i.

We define the utility function as follows.

Definition 2 (Utility). *Given a strategy profile $(\mathbf{d}, \mathbf{t})$ and a profit r, for each user a_i, their utility function is:*

$$u_i(\mathbf{d}, \mathbf{t}) = \frac{rT_{id_i}(\mathbf{d})t_i}{\sum_{j=1}^n T_{jd_j}(\mathbf{d})t_j} - c_{d_i}t_i - (1 - T_{id_i}(\mathbf{d}))t_i, \tag{11}$$

where T is the trust function described in Sec.3.2.

Intuitively, user a_i expects that they delegate to a validator a number of tokens that is the multiplication of the actual delegated token number t_i and the trust T_{id_i}, i.e., the probability that they believe the validator would not leave the market. Then, their utility is a proportion of the total profit r subtracting the commission and the expected loss of tokens (i.e., $(1 - T_{id_i}(\mathbf{d}))t_i$). The proportion is decided by the ratio of the user's expected delegating token number out of the entire expected token number in the pool.

Note that in this section, we assume that, users' accuracy is higher than their error, i.e., $q > \bar{q}$.

Remark 1. In the above definition, we assume that each user's token strategy space lies in the non-negative real space, i.e., $\Sigma_i \in V \times \mathbb{R}_{\geq 0}$ for all $a_i \in A$. That is, depending on the specific VSG setting, a user's token strategy might be any real number between 0 and their budget b_i. However, in practice, users' token strategy spaces are usually discrete, e.g., $\Sigma_i = V \times [b_i]$, where $[b_i] = \{0, 1, \ldots, b_i\}$. In this work, we consider continuous token strategy spaces mainly because of their simplicity for theoretical analysis. In Sec.5, we consider the more practical setting of discrete token strategy spaces.

To study users' behavior in VSGs, we will be considering the existence and the structure of the well-known game solution, the *Nash equilibrium* (NE).

Definition 3 (Nash equilibrium). *Given a VSG* g*, a strategy profile $(\mathbf{d}, \mathbf{t})$ is a Nash equilibrium if there is no user $a_i \in A$ and their strategy (d_i', t_i') such that $(d_i', t_i') \neq (d_i, t_i)$, and $u_i((\mathbf{d}_{-i}, d_i'), (\mathbf{t}_{-i}, t_i')) > u_i(\mathbf{d}, \mathbf{t})$, where $\mathbf{d}_{-i}$ and $\mathbf{t}_{-i}$ denotes that the other users than a_i take the delegation strategy and the token strategy in $\mathbf{d}$ and $\mathbf{t}$ respectively.*

We use the following example to illustrate the above definitions.

Example 1. Consider a VSG with two users $A = \{a_1, a_2\}$ and two validators $V = \{v_1, v_2\}$. Both users a_1 and a_2 have the same accuracy and error $(q, \bar{q}) = (0.8, 0.3)$ on both validators v_1 and v_2, and we consider both users' budgets are large. For validators, the commission profile is $\mathbf{c} = (0.2, 0.1)$, the integrity profile is $\mathbf{p} = (0.8, 0.6)$. When the next block is validated, users will receive a large amount of profit r.

We first compute both users' trusts on both validators based on all possible delegation profiles by Eq.5 as follows.

Table 2: Users' trusts on validators. T_{ij} denotes the trust of user a_i on validator v_j.

	$d_2 = v_1$	$d_2 = v_2$
$d_1 = v_1$	$T_{11} = 0.966, T_{21} = 0.966$	$T_{11} = 0.914, T_{22} = 0.8$
$d_1 = v_2$	$T_{12} = 0.8, T_{21} = 0.914$	$T_{12} = 0.914, T_{22} = 0.914$

We then study the utility each user can obtain by taking different strategies. We first investigate how users decide their number of delegating tokens $\mathbf{t} = (t_1, t_2)$ with their delegation profile fixed as $\mathbf{d} = (d_1, d_2)$. Let T_1 and T_2 denote the trust of a_1 and a_2 on their delegating validators, respectively. By Eq.11, we have that a_1 obtains utility of:

$$u_1(\mathbf{d}, \mathbf{t}) = \frac{rT_1t_1}{T_1t_1 + T_2t_2} - c_{d_1}t_1 - (1 - T_1)t_1. \tag{12}$$

Differentiating Eq.12 by t_1, we have that

$$\frac{du_1}{dt_1} = \frac{rT_2t_2T_1}{(T_1t_1 + T_2t_2)^2} - (c_{d_1} + 1 - T_1).$$

Therefore, we have that user a_1 optimizes $u_1(\mathbf{d}, \mathbf{t})$ by (1) $t_1^* = 0$, corresponding to utility of 0, if $\frac{rT_1}{T_2t_2} - (c_{d_1} + 1 - T_1) < 0$ (i.e., u_1 always decreases as t_1 increases), or (2) $t_1^* = \sqrt{\frac{rT_2t_2}{w_1T_1}} - \frac{T_2}{T_1}t_2$, by letting $\frac{du_1}{dt_1} = 0$, where $w_1 = (c_{d_1} + 1 - T_1)$. Then, in case (2), we have that the optimal utility a_1 can obtain is

$$u_1^*(\mathbf{d}, (t_1^*, t_2)) = r - 2\sqrt{\frac{rT_2t_2w_1}{T_1}} + \frac{T_2}{T_1}t_2w_1 = \left(\sqrt{\frac{T_2t_2w_1}{T_1}} - \sqrt{r}\right)^2. \tag{13}$$

Assume that both users enter the market by taking delegation profile $(d_1 = v_2, d_2 = v_2)$. Their trusts on v_2 are identical since they delegate to the same validator, and we let the trust be T. Then, having $\frac{du_1}{t_1} = \frac{du_2}{t_2} = 0$, we have that both users delegate $t =$

$\frac{r}{4w}$ tokens, otherwise, they have an incentive to alter their token strategy to achieve a higher utility.

Then, we show that strategy profile $((d_1 = v_2, d_2 = v_2), (t, t))$ is a NE. First notice that under this strategy profile, each user obtains utility of $\frac{r}{4} > 0$, which indicates that no user has an incentive to abstain. Then, we show that no user has an incentive to delegate to validator v_1. In particular, we show that a_1 has no incentive to deviate to $d'_1 = v_1$, and a similar reasoning can be developed to show that a_2 has no incentive to deviate from $(d_1 = v_2, d_2 = v_2)$ to $d'_2 = v_1$.

Consider strategy profile $((d_1 = v_2, d_2 = v_2), (t, t))$. If a_2 fixes their strategy (d_2, t) but a_1 changes from $d_1 = v_2$ to $d'_1 = v_1$, the trusts of a_1 on v_2 before the change and on v_1 after the change are identical (0.914) as shown in Tab.2, but the trust of a_2 on v_1 changes from 0.914 to 0.8. Considering a_1 bears a higher commission rate from delegating to c_2 to c_1, the only changed component in Equation 13 is w_1T_2, from 0.17 to 0.2288. Since $u_1(\mathbf{d}, \mathbf{t}) > 0$ in Equation 12, otherwise a_1 would rather abstain, we have that

$$\frac{rT_1}{T_2t_2} - w_1 > \frac{rT_1}{T_1t_1 + T_2t_2} - w_1 > 0,$$

which further indicates $\frac{T_2t_2w_1}{T_1} < r$.

This indicates that, according to Equation 13, user a_1 obtains a lower utility by changing from $d_1 = v_2$ to $d'_1 = v_1$ even by always taking the optimal token strategy because w_1T_2 increases. A similar argument can be developed for the user a_2, and therefore, $((d_1 = v_2, d_2 = v_2), (t, t))$ is a NE.

Observe that in the above NE, users do not delegate to v_1 who has a higher integrity. Instead, a rational user tries to reach a good balance between high reputation and low commission, such as to optimize their utility. □

Next, we investigate the existence and the feature of NE in VSGs under the HUPE trust class, which we call HUPE VSGs.

4.2 Equilibria in HUPE VSG

We show the existence and structure of NE in several subclasses of HUPE VSGs, through which we gain insight into how users delegate and use their tokens in markets under certain conditions.

4.2.1 *Single Validator VSG.*

We first study a subclass of the HUPE VSGs: the *single validator* VSGs, where there is only one validator in the market. Then, each user can choose to delegate to this validator with a number of tokens within their budget, or abstain. As follows, we show that there always exists an NE in each single validator VSG if a necessary condition on users' budgets holds.

Theorem 1. *In a single validator VSG, i.e., there is only one validator, there always exists a NE if for all user $a_i \in A$, their budget b_i satisfies $b_i \geq \frac{(n-1)r}{n^2(1+c-T)}(1+c)$, where $n = |A|$, r is the profit, c is the commission rate of the validator, and $T = \frac{q^np}{q^np+\bar{q}^n(1-p)}$ (p is the validator's integrity, and q and $\bar{q}$ are users' accuracy and error).*

Proof (Thm.1). We prove this theorem by showing that the strategy profile $(\mathbf{d}^*, \mathbf{t}^*) = (d_1 = \cdots = d_n = v, t_1 = \cdots = t_n = t = \frac{(n-1)r}{n^2(1+c-T)})$, i.e., each user delegates to the only validator v with t tokens, is an NE.

We first show that in strategy profile $(\mathbf{d}^*, \mathbf{t}^*)$, no user has the incentive to deviate by abstaining. Since each user delegates to v, all users' trusts on v are identical, and equal

$$T = \frac{q^np}{q^np + \bar{q}^n(1-p)}.$$

For an arbitrary user $a_i \in A$, their utility by taking the above strategy profile is:

$$u_i(\mathbf{d}^*, \mathbf{t}^*) = \frac{rTt}{\sum_{a_i \in A} Tt} - wt = \frac{r}{n} - wt = \frac{r}{n^2} \geq 0,$$

where $w = 1 + c - T$. If a_i abstains, they obtain utility of 0, which indicates that a_i prefers to take $(\mathbf{d}^*, \mathbf{t}^*)$ rather than to abstain.

Then, we show that no user can obtain a higher utility by unilaterally altering their token strategy. Assume that user a_i alters his token strategy to $t + x$, such that $t + x \in [0, b_i]$. Let the altered strategy profile be $(\mathbf{d}^*, \mathbf{t}' = (\mathbf{t}^*_{-i}, t_i = t + x))$, where $\mathbf{t}^*_{-i}$ denotes that all users except for a_i take $\mathbf{t}^*$. The utility of a_i becomes:

$$u_i(\mathbf{d}^*, \mathbf{t}') = \frac{r(t+x)}{nt+x} - w(t+x).$$

Differentiating $u_i(\mathbf{d}^*, \mathbf{t}')$ with respect to x, we have

$$\frac{du_i(\mathbf{d}^*, \mathbf{t}')}{dx} = \frac{(n-1)rt}{(nt+x)^2} - w.$$

Substitute t by $\frac{(n-1)r}{n^2w}$, and we obtain

$$\frac{du_i(\mathbf{d}^*, \mathbf{t}')}{dx} = \left(\frac{(n-1)^2r^2 - [(n-1)r + nwx]^2}{n^2w}\right) / \left(\frac{(n-1)r + nwx}{nw}\right)^2.$$

Observe that when $x = 0$, $\frac{du_i(\mathbf{d}^*,\mathbf{t}')}{dx} = 0$, and $\frac{du_i((\mathbf{d}^*,\mathbf{t}'))}{dx} > 0$ when $x < 0$ and $\frac{du_i(\mathbf{d}^*,\mathbf{t}')}{dx} < 0$ when $x > 0$. That is, a_i cannot increase their utility from $u_i(\mathbf{d}^*, \mathbf{t}^*)$ by choosing $x \neq 0$.
This completes the proof. □

Thm.1 illustrates that in a single validator VSG, users reach a NE by delegating to the validator with an identical number of tokens if their budgets admit. Observe that when users' budgets are high, they do not delegate as many as possible. Instead, compared to the NE token strategy shown in Thm.1, if a user's delegating token number is at a lower level, they can improve utility by increasing the delegating number, and if at a high level, they can improve by decreasing the number. Then, users' strategies eventually converge to the NE given that all users are rational and selfish.

4.2.2 *Homogeneous Validator VSG.*

We next consider another subclass of the HUPE VSGs, the *homogeneous validator VSG*, where each validator has the same integrity and commission rate. The following theorem shows that an NE can always be guaranteed in a homogeneous validator VSG under certain conditions.

Theorem 2. *In a homogenous validator VSG, there always exists an NE, if for each user $a_i \in A$, their budget b_i satisfies $b_i \geq \frac{(n-1)r}{n^2(1+c-T)}(1+c) \geq 1+c$, where $n = |A|$, r is the profit, c is the commission rate of the validator, and $T = \frac{q^np}{q^np+\bar{q}^n(1-p)}$ (p is the validator's integrity, and q and $\bar{q}$ are users' accuracy and error).*

Proof (Thm.2). We prove by showing that the strategy profile $(\mathbf{d}, \mathbf{t})$, where for each $a_i \in A$, $d_i = v_j$ and $t_i = t = \frac{(n-1)r}{n^2(1+c-T)}$ is a NE, where $v_j \in V$ and $T = \frac{q^n p}{q^n p + \bar{q}^n (1-p)}$. Note that $t \geq 1$ by the condition of the theorem. For an arbitrary user $a_i \in A$, we reason by three exclusive aspects: **(i)** a_i cannot obtain a higher utility by abstaining; **(ii)** a_i cannot obtain a higher utility by only altering their token strategy; and **(iii)** a_i cannot obtain a higher utility by delegating to another validator.

***Case*-(i).** Taking $(\mathbf{d}, \mathbf{t})$, a_i obtains utility:

$$u_i(\mathbf{d}, \mathbf{t}) = \frac{r}{n} - wt = \frac{r}{n^2} > 0,$$

where $w = 1 + c_j - T$. If a_i abstains, they obtain utility of 0, which is lower than $u_i(\mathbf{d}, \mathbf{t})$, which implies that a_i would prefers $(\mathbf{d}, \mathbf{t})$ over abstention.

***Case*-(ii).** By the proof of Thm.1, we have that this claim holds.

***Case*-(iii).** Assume that a_i delegate to validator $v_{j'}$ (this includes the case $v_j = v_{j'}$) with token strategy t', forming profile $(\mathbf{d}', \mathbf{t}') = ((\mathbf{d}_{-i}, d'_i = v_{j'}), (\mathbf{t}_{-i}, t'))$, and we have that a_i obtains utility:

$$u_i(\mathbf{d}', \mathbf{t}') = \frac{rT_{ij'}t'}{(n-1)T^{(n-1)}t + T_{ij'}t'} - w_{j'}t',$$

where $T^{(n-1)}$ is the trust of the other users than a_i on v_j, $T_{ij'}$ is the trust of a_i on $v_{j'}$, and $w_{j'} = 1 + c_{j'} - T_{ij'}$. We have the derivative of $u_i(\mathbf{d}', \mathbf{t}')$ with respect to t' is:

$$\frac{du_i(\mathbf{d}', \mathbf{t}')}{dt'} = \frac{r(n-1)T^{(n-1)}tT_{ij'}}{((n-1)T^{(n-1)}t + T_{ij'}t')^2} - w_{ij'}.$$

Making the above derivative 0, we have that the optimal token strategy for a_i is

$$t' = \sqrt{\frac{r(n-1)T^{(n-1)}t}{w_{ij'}T_{ij'}}} - \frac{(n-1)T^{(n-1)}t}{T_{ij'}}. \tag{14}$$

Then, we have that the optimal utility a_i can obtain by delegating to $v_{j'}$ is:

$$u_i^* = r - 2\sqrt{\frac{r(n-1)T^{(n-1)}w_{ij'}}{T_{ij'}}} + \frac{w_{ij'}(n-1)T^{(n-1)}}{T_{ij'}}.$$

Let $X = \sqrt{\frac{(n-1)T^{(n-1)}w_{ij'}}{T_{ij'}}}$, and we have that $u_i^* = (X - \sqrt{r})^2$.
Compare the two cases: (1) $v_{j'} = v_j$ and (2) $v_{j'} \neq v_j$. From (1) to (2), $w_{ij'}$ increases because $T_{ij'} < T$. We also have that, from (1) to (2), $T^{(n-1)}/T_{ij'}$ increases from 1 to larger than 1 because $T^{(n-1)} = T_{ij'}$ if $v_{j'} = v_j$, and $T^{(n-1)} > T_{ij'}$ if $v_{ij'} \neq v_j$ by Lm.3. Therefore, X increases from (1) to (2).
Since $u_i(\mathbf{d}', \mathbf{t}') > 0$, otherwise a_i prefers abstain, we have that

$$\frac{rT_{ij'}}{(n-1)T^{(n-1)}t} - w_{ij'} > \frac{rT_{ij'}}{(n-1)T^{(n-1)}t + T_{ij'}t'} - w_{ij'} > 0.$$

Since by the theorem's condition, $t \geq 1$, we have that

$$\frac{(n-1)T^{(n-1)}w_{ij'}}{T_{ij'}} < r.$$

This indicates that, since X increases from (1) to (2), $u_i^* = (X - \sqrt{r})^2$ decreases. Thus, a_i does not have an incentive to change from $(\mathbf{d}, \mathbf{t})$ to delegating to another validator.
This completes the proof. □

4.2.3 Commission-free VSG.

Lastly, we consider the subclass of commission-free VSG where the validator's commission rate is 0. We show that if commission-free VSG satisfies certain conditions, there always exists an NE.

Theorem 3. *In a commission-free VSG, a NE is always guaranteed to exist, if, for each user $a_i \in A$, their budget satisfies $b_i \geq \frac{(n-1)r}{n^2(1-T)} \geq 1$, where $T = \frac{q^n p^*}{q^n p^* + \bar{q}^n (1-p^*)}$ such that p^* is the maximal integrity among all validators.*

Proof (Thm.3). We prove this theorem by showing that strategy profile $(\mathbf{b}, \mathbf{t})$ such that for each user $a_i \in A$, $b_i = v^*$ ($v^* \in V$ is a validator with the maximal integrity, i.e., $v^* \in arg \max_{v_j \in V} p_j$) and $t_i = \frac{(n-1)r}{n^2(1-T)}$ is a NE. Intuitively, $(\mathbf{b}, \mathbf{t})$ is such a strategy profile that each user delegates the above number of tokens to the same validator who has the maximal integrity. We illustrate that each user a_i cannot obtain a higher utility by unilaterally altering their strategy in three exhaustive ways: **(i)** to abstain; **(ii)** to only alter their token strategy; and **(iii)** to delegate to another validator.

***Case*-(i).** Taking strategy $(\mathbf{d}, \mathbf{t})$, a_i obtains utility

$$u_i(\mathbf{d}, \mathbf{t}) = \frac{r}{n^2} > 0.$$

Thus, a_i has no incentive to deviate from $(\mathbf{d}, \mathbf{t})$ to abstaining otherwise a_i obtains a lower utility of 0.

***Case*-(ii).** We can reason this case by the same argument in the proof of Thm.1.

***Case*-(iii).** Assume that a_i delegates to validator $v_{j'}$ (either $v_{j'} = v_j$ or $v_{j'} \neq v_j$). We obtain that, the maximal utility that a_i can achieve by taking the optimal token strategy t' computed by Eq.14 is

$$u_i^* = \left(\sqrt{\frac{(1-T_{ij'})(n-1)T^{(n-1)}}{T_{ij'}}} - \sqrt{r}\right)^2,$$

where $T^{(n-1)}$ is, based on strategy profile $((\mathbf{d}_{-i}, d_i = v_{j'}), (\mathbf{t}_{-i}, t'))$, the trust of all users other than a_i on v^*, and $T_{ij'}$ is the trust of a_i on $v_{j'}$.
Compare two cases: (1) $v_{j'} = v_j$, and (2) $v_{j'} \neq v_j$. From (1) to (2), we have that $T_{ij'}$ decreases by Lm.3 and Lm.2, since $p_{j'} \leq p_j$. We let $X = \sqrt{\frac{(1-T_{ij'})(n-1)T^{(n-1)}}{T_{ij'}}}$, and thus, we have that X increases from (1) to (2), since $T_{ij'}$ decreases, and $fracT^{(n-1)}T_{ij'}$ increases from 1 to larger than 1 from (1) to (2) due to $T^{(n-1)} > T_{ij'}$ by Lm.3.

Since $u_i\sqrt{\frac{(1-T_{ij'})(n-1)T^{(n-1)}}{T_{ij'}}} > 0$, otherwise a_i would prefer to abstain, we have that

$$\frac{rT_{ij'}}{(n-1)T^{(n-1)}t} - (1 - T_{ij'}) > 0.$$

By condition $t > 1$, we further obtain that $X < \sqrt{r}$. Therefore, we have that u_i^* decrease from (1) to (2), which completes the proof. □

We can observe that, by Thm.2 and Thm.3, in a homogeneous VSG or a commission-free VSG, when users' budgets are high enough, the strategies of users may converge to such a structure

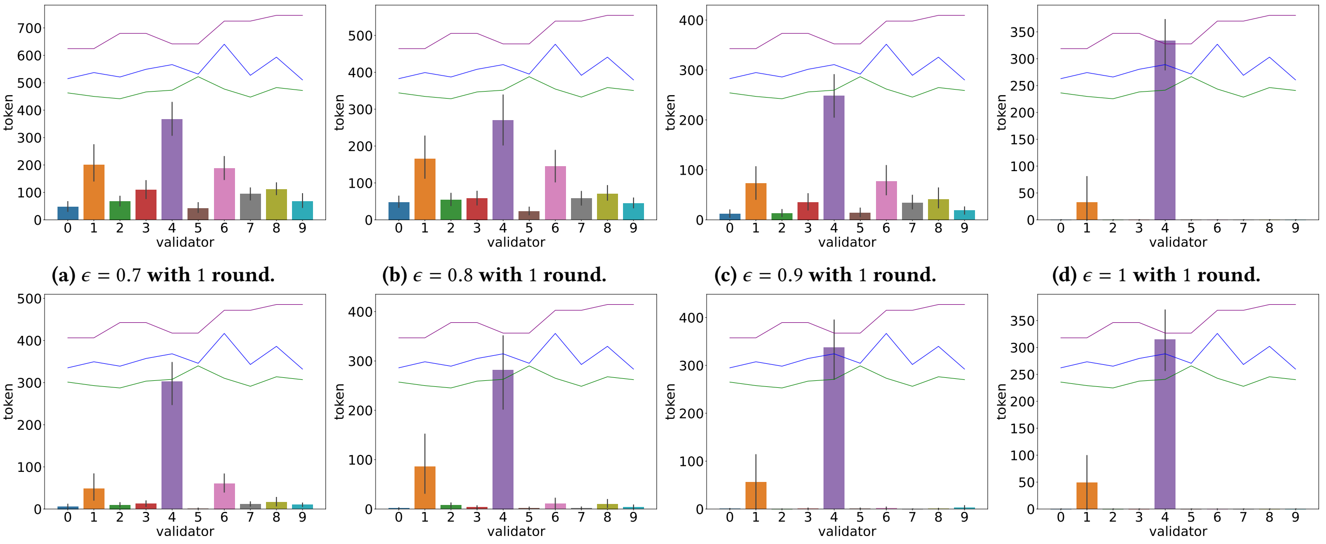

(a) ϵ = 0.7 **with** 1 **round.** **(b)** ϵ = 0.8 **with** 1 **round.** **(c)** ϵ = 0.9 **with** 1 **round.** **(d)** ϵ = 1 **with** 1 **round.**

(e) ϵ = 0.7 **with** 2 **round.** **(f)** ϵ = 0.8 **with** 2 **round.** **(g)** ϵ = 0.9 **with** 2 **round.** **(h)** ϵ = 1 **with** 2 **round.**

Figure 2: *Bars*: the average values of each validator's received delegation tokens in 20 simulation instances, varying ϵ and RL (i.e., the round limit). *Lines*: — validator integrities p_j, — validator evidences $\Pr(b_j = 1)$, — validator commission rates c_j.

that all users delegate to the same validator, who has the maximal integrity in the commission-free VSG, and users delegate with the same token strategy.

We configure that this is a reasonable result in VSGs. In a market, there should exist an optimal validator who has a good trade-off between their reputation and commission. Then, users would tend to delegate to this validator, and this validator is more trusted as more users delegate to them.

5 EXPERIMENT

Relaxing the assumptions in the above theoretical analysis, we now empirically investigate users' behavior by computer simulations. In the empirical study, we consider users can only take token strategies in a discrete space, which coincides with the practice. The assumptions on users' accuracies and errors are relaxed, such that misled users (i.e., $\exists a_i \in A$ and $\exists v_j \in V$, $q_{ij} < 0.5$ or $\bar{q}_{ij} > 0.5$) exist. Additionally, we also consider that validators' integrities are not revealed perfectly, i.e., $\exists v_j \in V$, the evidence quality (recall Equation 3) of v_j satisfies $q_j < 1$.

With the above practical settings, we use a noisy model to simulate that a set of users, in turn, choose their *best response strategy* against the current strategy profile, in a random *round-robin* order.

Definition 4 (Best Response). *Given a VSG and a strategy profile* $(\mathbf{d}, \mathbf{t})$*, for each user* $a_i \in A$*, their best response strategy against* $(\mathbf{d}, \mathbf{t})$ *is the strategy* (d_i', t_i') *by which* a_i *can obtain the highest utility assuming that the other users do not change their strategies in* $(\mathbf{d}, \mathbf{t})$*. Formally, given a strategy profile* $(\mathbf{d}, \mathbf{t})$*, any strategy* $(d_i', t_i') \in arg\max_{(d_i,t_i)\in\Sigma_i} u_i(((\mathbf{d}_{-i}, d_i), (\mathbf{t}_{-i}, t_i)))$ *is* a_i*'s best response strategy.*

We design a so-called ϵ-Greedy Best Response Dynamics (Alg.1, with abbreviation GBRD) to simulate users' behavior. In GBRD, we initiate the strategy profile by assigning a validator to each user at random, with a randomly sampled integer between 0 and $1/(1+c_j)$ of their budget as the token strategy, where c_j is the commission rate of each user's randomly assigned validator. The initiated strategy profile is denoted as $(\mathbf{d}^0, \mathbf{t}^0)$. Then, in each round, following a randomly generated order σ (i.e., a random permutation of all users in A), against the current strategy profile, each user in turn chooses their best response strategy by an ϵ-greedy manner. That is, given a strategy profile $(\mathbf{d}, \mathbf{t})$ and a user $a_i \in A$, let $R_i^{\text{best}}(\mathbf{d}, \mathbf{t}) = arg\max_{(d',t')\in\Sigma_i} u_i((\mathbf{d}_{-i}, d'), (\mathbf{t}_{-i}, t'))$ be the set of best response strategies of a_i, and $R_i^{\text{better}}(\mathbf{d}, \mathbf{t}) = \{(d', t') \in \Sigma_i \mid u_i((\mathbf{d}_{-i}, d'), (\mathbf{t}_{-i}, t')) > u_i(\mathbf{d}, \mathbf{t})\} \setminus R_i^{\text{best}}$ be the set of strategies that can improve a_i's utility against $(\mathbf{d}, \mathbf{t})$ except for those in R_i^{best}. Then, in GBRD, at each round, each user uniformly at random chooses a strategy from R_i^{best} with probability $\epsilon \in [0, 1]$, and uniformly at random chooses one from R_i^{better} with probability $1 - \epsilon$, in the turn of σ. After a round, if for each user, the change of their utility is less than a ratio θ of their utility of the last round, or the number of rounds reaches the limit RL, the algorithm terminates and returns the strategy profile output by the last round.

5.1 Experiment Settings

We conduct simulations with 10 validators and 200 users. The integrity and evidence quality of each validator, and the accuracy and error of each user are randomly generated by Gaussian distributions with a standard deviation of 0.1, and the means of integrity, evidence quality, accuracy, and error are 0.7, 0.8, 0.6, 0.5, respectively. Especially, each generated value of integrity, evidence quality, and accuracy is forced in the interval [0.5, 1]. The budget of each user is generated by a Gaussian distribution of $N(70, 15)$, and the profit is $r = 30$. The commission rate of each validator $v_j \in V$ is generated as $c_j = (p_j - 0.5)/3 + \delta$, where δ is a Guassian error following

Algorithm 1 ϵ-Greedy Best Response Dynamics

```
Initialization: $\sigma$, $\epsilon$, $s = 0$, $(\mathbf{d}^0, \mathbf{t}^0)$, $RL$, $\theta$.
Iteration:
    Round $s$:
    • Step1: $k = 1$, and $(\mathbf{d}^{s+1}, \mathbf{t}^{s+1}) = (\mathbf{d}^s, \mathbf{t}^t)$.
    • Step2: For $a = \sigma(k)$, compute $u_a(((\mathbf{d}^{s+1}_{-a}, d_k), (\mathbf{t}^{s+1}_{-a}, t_k)))$
      for all $(d_k, t_k) \in \Sigma_a$.
    • Step3: With probability of $\epsilon$, $a$ samples a strategy from
      $R^{\text{best}}(\mathbf{d}^{s+1}, \mathbf{t}^{s+1})$ uniformly at random, and with probability
      of $(1 - \epsilon)$, $a$ samples a strategy from $R^{\text{better}}(\mathbf{d}^s, \mathbf{t}^s)$.
      Replace $(d^{s+1}_a, t^{s+1}_a)$ by the sampled strategy.
    • Step4: If $k < n$, $k \leftarrow k + 1$ and go to Step2, else to Step5.
    • Step5: If $|u_i(\mathbf{d}^{s+1}, \mathbf{t}^{s+1}) - u_i(\mathbf{d}^s, \mathbf{t}^s)| \leq \theta$ or $s + 1 = RL$,
      terminate, else $s \leftarrow s + 1$ and go back to Step1.
Return: $(\mathbf{d}^{s+1}, \mathbf{t}^{s+1})$
```

$N(0, 0.01)$, and c_j is forced positive. Each commission rate is in the range $[0, 0.2]$ with high probability. Lastly, θ is set as 0.01.

We run all simulation instances under the same initiative setting of validator integrities, evidence qualities and commission rates, and user accuracies, errors, and budgets. For each parameter, we run 20 simulation instances. The following results are the average of values output by 20 instances.

We conduct the experiment in Python 3.9, and run the simulation on a MacBook Pro with an Apple M1 Pro chip.

5.2 Results

We first observe that, according to the noisy model, information error exists in the market: validators' integrities (purple line) and evidence (blue line) do not have the same trend, suggesting that a validator with high integrity may not necessarily have high evidence (i.e., a good reputation).

Our results also show that almost all users delegate: averagely more than 99% users delegate in each parameter setting. Among those who delegate, averagely, each user only uses less than 10% of their tokens, with a range from 2.7% to 9.9% in all parameter settings. Users delegate fewer tokens when they are able to specify their optimal strategies more accurately, i.e., corresponding to a higher ϵ or a higher round number RL. We configure that it is because we use a small profit $r = 30$, which leads the users to choose a low level of token strategies such that the users would not lose too much due to the commission and the risk of trust.

Generally, observe that in all figures in Fig.2, users' delegation tokens tend to concentrate on a small part of validators. However, most of the validators receiving a large amount of tokens are not among those with high integrity or high evidence, except for validator 6 in results of $\epsilon = 0.7, 0.8$ and 0.9 (i.e., Fig.2a, 2b, 2c). Instead, tokens are concentrated on validators with a good balance of high reputation and low commission rate, e.g., validators 4 and 1.

Fig.2a, 2b, 2c and 2d illustrate to which validators the users delegate their tokens when they in turn noisily choose their best response strategy only once ($RL = 1$ in Alg.1), by varying the noise indicator ϵ from 0.7 to 1. Observe that as ϵ becomes higher, tokens are more concentrated, especially on validators 4 and a relatively small amount on validator 1. This coincides with our configuration that, as ϵ becomes higher, users are more accurate in delegating to the best validator with the optimal token strategy with respect to the utility defined in Equation 11. Though a large amount of tokens are also delegated to validators 1 and 6 in smaller ϵ's (Fig.2a, 2b, 2c), those amounts gradually decrease as ϵ becomes larger. We can conclude that when users' rationality is high (i.e., corresponding to higher ϵ), they can better converge to strategy profiles where they delegate to the validators with the best balance of good reputation and low commission. We can also conclude that, in this randomized simulation instance, validator 4 should be the one with the best balance of good reputation and low commission for users, and validator 1 also shows a good trade-off on these two attributes.

A similar trend can be observed in Fig.2e, 5b, 5c, 5d, where users go through 2 rounds of iteration in Alg.1, varying ϵ also from 0.7 to 1. However, we can also observe that for each ϵ, tokens are better concentrated in figures of 2 rounds (the second row) than figures of 1 rounds (the first row). This is because, by each iteration, users can gradually improve their strategy. More users will thereby delegate their tokens to the optimal validators. This further supports the conclusion that validators 4 and 1 are the best choices for users in terms of good reputation and low commission rate.

6 CONCLUSION

This paper analyses the delegation process as regular users choose their validators, offering models to quantify virtual trust in delegatees. Through game-theoretical simulations, we explore user behavior while considering vital factors from active staking services and PoS blockchains. Our findings indicate that users make decisions that lead to a Nash equilibrium by weighing delegation costs, other users' actions, and delegatee reputation. However, this trend heightens the risk of token concentration among a few delegatees.

⚠ **Open problems for NEXT.** We present several future plans.

① As this work majorly focuses on the behaviors of normal users (stakers), we could extend our game theoretical model by incorporating validators as players. Validators may choose a strategy between staying in or leaving the market, based on their potential profit decided by extrinsic factors, such as their cost to run a client and their accrual of delegation tokens. Then, validators with lower integrity might be easily motivated by their potential profit to leave the market. This brings the users more concerns on judging validators' leaving risk than only considering the trust in validators' intrinsic motivation. Our configuration is that this setting may prevent users from concentrating their tokens on a small number of validators, such that the distribution of delegation tokens is more equal among validators.

② We initially model the subjective parameters, such as a user's accuracy, as static within the framework of the game. However, for a more realistic representation, we can introduce a dynamic learning process to account for users' potential growth and accelerated learning during each evolutionary phase. This extension would involve incorporating a dynamic accuracy parameter into the model to better align with practical scenarios.

③ As an extended part of experiments, we plan to perform a comparative analysis that involves evaluating the simulated data generated by our model alongside real-world data collected from publicly available sources related to stakers and staking providers.

REFERENCES

[1] Ethereum roadmap: The Merge. *https://ethereum.org/en/roadmap/merge/*, 2022.

[2] Proof of stake (PoS) Ethereum. *https://ethereum.org/en/developers/docs/consensus-mechanisms/pos/*, 2023.

[3] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *IEEE Symposium on Security and Privacy (SP)*, pages 910–927. IEEE, 2020.

[4] Li Chao, Palanisamy Balaji, Xu Runhua, Duan Li, Liu Jiqiang, and Wang Wei. How hard is takeover in DPoS blockchains? understanding the security of coin-based voting governance. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2023.

[5] Chao Li, Runhua Xu, and Li Duan. Liquid democracy in DPoS blockchains. In *Proceedings of the ACM International Symposium on Blockchain and Secure Critical Infrastructure (BSCI@AsiaCCS)*, pages 25–33, 2023.

[6] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Decentralized business review*, 2008.

[7] Yuzhe Zhang and Davide Grossi. Power in liquid democracy. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 35, pages 5822–5830, 2021.

[8] Seonggeun Kim and Sang-Geun Hahn. Mining pool manipulation in blockchain network over evolutionary block withholding attack. *IEEE Access*, 7:144230–144244, 2019.

[9] Richard Koch. *The 80/20 Principle: The Secret of Achieving More with Less: Updated 20th anniversary edition of the productivity and business classic.* Hachette UK, 2011.

[10] Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. Consensus in the presence of partial synchrony. *Journal of the ACM (JACM)*, 35(2):288–323, 1988.

[11] Juan Garay, Aggelos Kiayias, and Nikos Leonardos. The Bitcoin backbone protocol: Analysis and applications. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, pages 281–310. Springer, 2015.

[12] Lioba Heimbach, Lucianna Kiffer, Christof Ferreira Torres, and Roger Wattenhofer. Ethereum's proposer-builder separation: Promises and realities. *arXiv preprint arXiv:2305.19037*, 2023.

[13] Vitalik Buterin, Diego Hernandez, Thor Kamphefner, Khiem Pham, Zhi Qiao, Danny Ryan, Juhyeok Sin, Ying Wang, and Yan X Zhang. Combining GHOST and Casper. *arXiv preprint arXiv:2003.03052*, 2020.

[14] Joachim Neu, Ertem Nusret Tas, and David Tse. Ebb-and-flow protocols: A resolution of the availability-finality dilemma. In *IEEE Symposium on Security and Privacy (SP)*, pages 446–465. IEEE, 2021.

[15] Vitalik Buterin and Virgil Griffith. Casper the friendly finality gadget. *arXiv preprint arXiv:1710.09437*, 2017.

[16] Yonatan Sompolinsky and Aviv Zohar. Accelerating Bitcoin's transaction processing: Fast money grows on trees, not chains. *Cryptology ePrint Archive*, 2013.

[17] Vitalik Buterin, Daniël Reijsbergen, Stefanos Leonardos, and Georgios Piliouras. Incentives in Ethereum's hybrid Casper protocol. *International Journal of Network Management*, 30(5):e2098, 2020.

[18] Natale Umberto. Analysing Ethereum cryptoeconomics: The validator's perspective. *https://docs.google.com/document/d/1r640UQOm2z-Q9nsJzqBq3BVgCtTL1_Yc7WnPp4jEBgk/edit*, 2023.

[19] Edgington Ben. Upgrading Ethereum: A technical handbook on Ethereum's move to proof of stake and beyond. *https://eth2book.info/capella/*, 2023.

[20] MEV-boost in a nutshell. *https://boost.flashbots.net/*, 2023.

[21] Thalman Benjamin. Ethereum: A deep dive into new ETH rewards dynamics. *https://figment.io/insights/ethereum-a-deep-dive-into-new-eth-rewards-dynamicsethereum-a-deep-dive-into-new-eth-rewards-dynamics/*, 2023.

[22] Edgington Ben. Upgrading Ethereum: Slashing. *https://eth2book.info/capella/part2/incentives/slashing/*, 2023.

[23] Ethereum. Ethereum offical: Earn rewards while securing ethereum. *https://ethereum.org/en/staking/#how-to-stake-your-eth*, 2023.

[24] Franz Dietrich and Christian List. A model of jury decisions where all jurors have the same evidence. *Synthese http://www.jstor.org/stable/20118506*, 142(2):175–202, 2004.

[25] Shange Fu et al. FTX collapse: A Ponzi story. *arXiv preprint arXiv:2212.09436*, 2022.

[26] Ethereum. The history of ethereum. *https://ethereum.org/en/history/*, 2023.

[27] Ethereum. Ethereum roadmap. *https://ethereum.org/en/roadmap*, 2023.

[28] Louis Tremblay Thibault, Tom Sarry, and Abdelhakim Senhaji Hafid. Blockchain scaling using rollups: A comprehensive survey. *IEEE Access*, 2022.

[29] Gang Wang, Zhijie Jerry Shi, Mark Nixon, and Song Han. SoK: Sharding on blockchain. In *Proceedings of the ACM Conference on Advances in Financial Technologies (AFT)*, pages 41–61, 2019.

[30] Sen Yang, Fan Zhang, Ken Huang, Xi Chen, Youwei Yang, and Feng Zhu. SoK: MEV countermeasures: Theory and practice. *arXiv preprint arXiv:2212.05111*, 2022.

[31] John Kuszmaul. Verkle trees. *Verkle Trees*, 1:1, 2019.

[32] Verkle trees. *https://vitalik.ca/general/2021/06/18/verkle.html*, 2023.

[33] Qin Wang and Shiping Chen. Account abstraction, analysed. *arXiv preprint arXiv:2309.00448*, 2023.

[34] Yulin Liu, Yuxuan Lu, Kartik Nayak, Fan Zhang, Luyao Zhang, and Yinhong Zhao. Empirical analysis of EIP-1559: Transaction fees, waiting times, and consensus security. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 2099–2113, 2022.

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

[36] Peter Gaži, Aggelos Kiayias, and Alexander Russell. Stake-bleeding attacks on proof-of-stake blockchains. In *Crypto Valley conference on Blockchain Technology (CVCBT)*, pages 85–92. IEEE, 2018.

[37] Sarah Azouvi, George Danezis, and Valeria Nikolaenko. Winkle: Foiling long-range attacks in proof-of-stake systems. In *Proceedings of the ACM Conference on Advances in Financial Technologies (AFT)*, pages 189–201, 2020.

[38] Caspar Schwarz-Schilling, Joachim Neu, Barnabé Monnot, Aditya Asgaonkar, Ertem Nusret Tas, and David Tse. Three attacks on proof-of-stake Ethereum. In *International Conference on Financial Cryptography and Data Security (FC)*, pages 560–576. Springer, 2022.

[39] Joachim Neu, Ertem Nusret Tas, and David Tse. Two more attacks on proof-of-stake GHOST/Ethereum. In *Proceedings of the ACM Workshop on Developments in Consensus (ConsensusDay)*, pages 43–52, 2022.

[40] Aggelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol. In *Annual International Cryptology Conference (CRYPTO)*, pages 357–388. Springer, 2017.

[41] Bernardo David, Peter Gaži, Aggelos Kiayias, and Alexander Russell. Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, pages 66–98. Springer, 2018.

[42] Sarah Azouvi and Marko Vukolić. Pikachu: Securing PoS blockchains from long-range attacks by checkpointing into Bitcoin PoW using Taproot. In *Proceedings of the ACM Workshop on Developments in Consensus (ConsensusDay)*, pages 53–65, 2022.

[43] Ertem Nusret Tas, David Tse, Fangyu Gai, Sreeram Kannan, Mohammad Ali Maddah-Ali, and Fisher Yu. Bitcoin-enhanced proof-of-stake security: Possibilities and impossibilities. In *IEEE Symposium on Security and Privacy (SP)*, pages 126–145. IEEE, 2023.

[44] Thomas Kerber, Aggelos Kiayias, Markulf Kohlweiss, and Vassilis Zikas. Ouroboros crypsinous: Privacy-preserving proof-of-stake. In *IEEE Symposium on Security and Privacy (SP)*, pages 157–174. IEEE, 2019.

[45] Shichen Wu, Zhiying Song, Puwen Wei, Peng Tang, and Quan Yuan. Improving privacy of anonymous proof-of-stake protocols. *Cryptology ePrint Archive*, 2023.

[46] Christian Badertscher, Peter Gaži, Aggelos Kiayias, Alexander Russell, and Vassilis Zikas. Ouroboros genesis: Composable proof-of-stake blockchains with dynamic availability. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 913–930, 2018.

[47] Aggelos Kiayias, Saad Quader, and Alexander Russell. Consistency of proof-of-stake blockchains with concurrent honest slot leaders. In *IEEE International Conference on Distributed Computing Systems (ICDCS)*, pages 776–786. IEEE, 2020.

[48] Erica Blum, Aggelos Kiayias, Cristopher Moore, Saad Quader, and Alexander Russell. The combinatorics of the longest-chain rule: Linear consistency for proof-of-stake blockchains. In *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1135–1154. SIAM, 2020.

[49] Francesco D'Amato and Luca Zanolini. A simple single slot finality protocol for ethereum. *arXiv preprint arXiv:2302.12745*, 2023.

[50] Alistair Stewart and Eleftherios Kokoris-Kogia. Grandpa: a Byzantine finality gadget. *arXiv preprint arXiv:2007.01560*, 2020.

[51] Francesco D'Amato and Luca Zanolini. Recent latest message driven GHOST: Balancing dynamic availability with asynchrony resilience. *arXiv preprint arXiv:2302.11326*, 2023.

[52] Yujin Kwon, Jian Liu, Minjeong Kim, Dawn Song, and Yongdae Kim. Impossibility of full decentralization in permissionless blockchains. In *Proceedings of the ACM Conference on Advances in Financial Technologies (AFT)*, pages 110–123, 2019.

[53] B David JH-y Chiang, I Eyal, and T Gong. FairPoS: Input fairness in proof-of-stake with adaptive security. *Proceedings of the ACM Conference on Advances in Financial Technologies (AFT)*, 2023.

[54] Muhammad Saad, Zhan Qin, Kui Ren, DaeHun Nyang, and David Mohaisen. E-PoS: Making proof-of-stake decentralized and fair. *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 32(8):1961–1973, 2021.

[55] @corwintines. Ethereum: Proof-of-stake (PoS). *https://ethereum.org/en/developers/docs/consensus-mechanisms/pos/*, 2023.

[56] Olivier Moindrot and Charles Bournhonesque. Proof of stake made simple with Casper. *ICME, Stanford University*, 2017.

[57] Sheikh Munir Skh Saad and Raja Zahilah Raja Mohd Radzi. Comparative review of the blockchain consensus algorithm between proof of stake (PoS) and delegated proof of stake (DPoS). *International Journal of Innovative Computing*, 2020.

[58] Polkadot. Polkadot: NPoS election algorithms. *https://wiki.polkadot.network/docs/learn-phragmen*, 2023.

[59] Hanaa Abbas, Caprolu Maurantonio, and Di Pietro Roberto. Analysis of Polkadot: Architecture, internals, and contradictions. *IEEE International Conference on Blockchain (Blockchain)*, 2022.

[60] Iddo Bentov, Charles Lee, Alex Mizrahi, and Meni Rosenfeld. Proof of activity: Extending Bitcoin's proof of work via proof of stake [extended abstract]. *ACM SIGMETRICS Performance Evaluation Review (PER)*, 42(3):34–37, 2014.

[61] Binance Academy. A quick guide to BNB staking on BNB smart chain (BSC). *https://academy.binance.com/en/articles/a-quick-guide-to-bnb-staking-on-binance-smart-chain-bsc*, 2023.

[62] Wang Qin, Li Rujia, Wang Qi, Chen Shiping, and Xiang Yang. Exploring unfairness on proof of authority: Order manipulation attacks and remedies. *In Proceedings of the ACM on Asia Conference on Computer and Communications Security (AsiaCCS)*, 2022.

[63] Dominic Grandjean, Lioba Heimbach, and Roger Wattenhofer. Ethereum proof-of-stake consensus layer: Participation and decentralization. *arXiv preprint arXiv:2306.10777*, 2023.

[64] Ping He, Dunzhe Tang, and Jingwen Wang. Staking pool centralization in proof-of-stake blockchain network. *Available at SSRN 3609817*, 2020.

[65] Kose John, Thomas J Rivera, and Fahad Saleh. Equilibrium staking levels in a proof-of-stake blockchain. *Available at SSRN 3965599*, 2021.

[66] Tarun Chitra and Kshitij Kulkarni. Improving proof of stake economic security via MEV redistribution. In *Proceedings of the ACM CCS Workshop on Decentralized Finance and Security (DeFi@CCS)*, pages 1–7, 2022.

[67] Kose John, Thomas J Rivera, and Fahad Saleh. Economic implications of scaling blockchains: Why the consensus protocol matters. *Available at SSRN 3750467*, 2020.

[68] Hans Gersbach, Akaki Mamageishvili, and Manvir Schneider. Staking pools on blockchains. *arXiv preprint arXiv:2203.05838*, 2022.

[69] Lars Brünjes, Aggelos Kiayias, Elias Koutsoupias, and Aikaterini-Panagiota Stouka. Reward sharing schemes for stake pools. In *IEEE european symposium on security and privacy (EuroSP)*, pages 256–275. IEEE, 2020.

[70] Michael Neuder, Daniel J Moroz, Rithvik Rao, and David C Parkes. Defending against malicious reorgs in tezos proof-of-stake. In *Proceedings of the ACM Conference on Advances in Financial Technologies (AFT)*, pages 46–58, 2020.

[71] Michael Neuder, Daniel J Moroz, Rithvik Rao, and David C Parkes. Selfish behavior in the tezos proof-of-stake protocol. 2021.

[72] Christoph Mueller-Bloch, Jonas Valbjørn Andersen, Jason Spasovski, and Jungpil Hahn. Understanding decentralization of decision-making power in proof-of-stake blockchains: an agent-based simulation approach. *European Journal of Information Systems (EJIS)*, pages 1–20, 2022.

[73] Johnnatan Messias, Vabuk Pahari, Balakrishnan Chandrasekaran, Krishna P Gummadi, and Patrick Loiseau. Understanding blockchain governance: Analyzing decentralized voting to amend DeFi smart contracts. *arXiv preprint arXiv:2305.17655*, 2023.

[74] Robin Fritsch, Marino Müller, and Roger Wattenhofer. Analyzing voting power in decentralized governance: Who controls DAOs? *arXiv preprint arXiv:2204.01176*, 2022.

[75] Qin Wang, Guangsheng Yu, Yilin Sai, Caijun Sun, Lam Duc Nguyen, Sherry Xu, and Shiping Chen. An empirical study on snapshot DAOs. *arXiv preprint arXiv:2211.15993*, 2022.

[76] Guangsheng Yu et al. Leveraging architectural approaches in Web3 applications-a DAO perspective focused. In *IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, pages 1–6. IEEE, 2023.

[77] BNB smart chain (BSC). *https://docs.bnbchain.org/docs/learn/intro*, 2023.

[78] Peter Gaži, Aggelos Kiayias, and Alexander Russell. Fait accompli committee selection: Improving the size-security tradeoff of stake-based committees. *Cryptology ePrint Archive*, 2023.

[79] Davide Grossi. Social choice around the block: On the computational social choice of blockchain. *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2022.

[80] Ziyao Liu, Nguyen Cong Luong, Wenbo Wang, Dusit Niyato, Ping Wang, Ying-Chang Liang, and Dong In Kim. A survey on applications of game theory in blockchain. *arXiv preprint arXiv:1902.10865*, 2019.

[81] Aggelos Kiayias, Elias Koutsoupias, Maria Kyropoulou, and Yiannis Tselekounis. Blockchain mining games. In *Proceedings of the ACM Conference on Economics and Computation (EC)*, pages 365–382, 2016.

[82] Yoad Lewenberg, Yoram Bachrach, Yonatan Sompolinsky, Aviv Zohar, and Jeffrey S Rosenschein. Bitcoin mining pools: A cooperative game theoretic analysis. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 919–927. Citeseer, 2015.

[83] Hongyin Chen, Yukun Cheng, Xiaotie Deng, Wenhan Huang, and Linxuan Rong. Absnft: securitization and repurchase scheme for non-fungible tokens based on game theoretical analysis. In *International Conference on Financial Cryptography and Data Security (FC)*, pages 407–425. Springer, 2022.

[84] Matthias Lohr, Kenneth Skiba, Marco Konersmann, Jan Jürjens, and Steffen Staab. Formalizing cost fairness for two-party exchange protocols using game theory and applications to blockchain. In *IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, pages 1–5. IEEE, 2022.

[85] Simon Janin, Kaihua Qin, Akaki Mamageishvili, and Arthur Gervais. Filebounty: Fair data exchange. In *IEEE European Symposium on Security and Privacy Workshops (EuroSP-W)*, pages 357–366. IEEE, 2020.

[86] Bo Qin et al. BDTS: A blockchain-based data trading system with fair exchange. *arXiv preprint arXiv:2211.10001*, 2022.

[87] Ittay Eyal. The miner's dilemma. In *IEEE Symposium on Security and Privacy (SP)*, pages 89–103. IEEE, 2015.

[88] Yujin Kwon, Dohyun Kim, Yunmok Son, Eugene Vasserman, and Yongdae Kim. Be selfish and avoid dilemmas: Fork after withholding (FAW) attacks on Bitcoin. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 195–209, 2017.

[89] Kevin Alarcón Negy, Peter R Rizun, and Emin Gün Sirer. Selfish mining re-examined. In *International Conference on Financial Cryptography and Data Security (FC)*, pages 61–78. Springer, 2020.

[90] Ayelet Sapirshtein, Yonatan Sompolinsky, and Aviv Zohar. Optimal selfish mining strategies in Bitcoin. In *International Conference on Financial Cryptography and Data Security (FC)*, pages 515–532. Springer, 2016.

[91] Qianlan Bai, Yuedong Xu, Nianyi Liu, and Xin Wang. Blockchain mining with multiple selfish miners. *IEEE Transactions on Information Forensics and Security (TIFS)*, 2023.

[92] Dimitris Karakostas, Aggelos Kiayias, and Thomas Zacharias. Blockchain nash dynamics and the pursuit of compliance. *arXiv preprint arXiv:2201.00858*, 2022.

[93] Yue Wang, Changbing Tang, Feilong Lin, Zhonglong Zheng, and Zhongyu Chen. Pool strategies selection in PoW-based blockchain networks: Game-theoretic analysis. *IEEE Access*, 7:8427–8436, 2019.

[94] Wenbai Li, Mengwen Cao, Yue Wang, Changbing Tang, and Feilong Lin. Mining pool game model and Nash equilibrium analysis for PoW-based blockchain networks. *IEEE Access*, 8:101049–101060, 2020.

[95] Yujin Kwon, Hyoungshick Kim, Jinwoo Shin, and Yongdae Kim. Bitcoin vs. Bitcoin cash: Coexistence or downfall of Bitcoin cash? In *IEEE Symposium on Security and Privacy (SP)*, pages 935–951. IEEE, 2019.

A ADDITIONAL EXPERIMENT RESULTS

We show more details corresponding to the description in Sec.5.2. Fig.3 provides the results of running Alg.1 more rounds than results shown in Fig.2. Fig.4 shows the average ratio of users' delegation token numbers to their budgets. Fig.5 shows users' trust on a validator by varying the number of homogeneous delegators, in different settings of validator integrity and user accuracy and error.

B STAKING SERVICES

Following Sec.2.3, we further present a summary of custodial staking services offered by various platforms (majorly CEXes, Tab.3). These platforms vary in terms of user entry requirements, supported proofs, unstaking periods, and maximum staking rates.

Table 3: Staking services (*custodial*)

	Mini.	Proofs	Unst.	M. Shr	Valid.	C. Fee	Sprt.	Max.
Coinbase	Any	cbETH	T+2	14%	121k	25%	9	6.12%
Binance	Any	BETH	T+2	4.3%	37k	10%	200+	157.81%
Kraken	Any	“staked"	Secs	3%	26k	15%	15+	26%
KuCoin	Any	ksETH	T+5	0.1%	2k	-	40+	180.15%
Cake DeFi	Any	csETH	-	0.1%	2k	-	4	12.4%
Crypto.com	$1e^{-8}$	“staked"	T+5	-	-	-	20+	12%
Nexo	US$10	NETH	-	-	-	0.20%	30+	24%
	Ethereum 2.0						Other PoS	

Then, we also present a comparison of two staking types across various key parameters (Tab.4). Custodial staking relies on a single third party, typically requiring less prior knowledge from users and often not imposing a minimum deposit. Unstaking periods are short due to the providers maintaining a flexible pool of unfrozen tokens for liquidity. However, custodial staking comes with the drawback that users do not have control over their private keys, exposing them to higher delegation risks (e.g., FTX collapse [25]).

Conversely, non-custodial staking presents contrasting attributes in several dimensions. It involves multiple validators, sets a higher entry threshold for users, and enforces extended unbonding (equiv. unstaking/unfreezing/redemption) periods. Additionally, it comes with increased operational expenses for running a full node. On the plus side, users retain control of their accounts through private keys and benefit from reduced delegation risks.

Table 4: Comparisons for staking types

	Custodial	**Non-custodial**
Reliance	one single third-party	multiple validators
User perception	less required	technical knowledge required
Mini. deposit	mostly no required	often required
Unstaking period	flexible (seconds to days)	long (due to the network)
Maintainance	fully delegated no further costs	need to run a full node
Account	do not hold private keys (*pk*)	owned by self with *pk*s
Delegation cost	high (service fees, low APR, etc.)	low (commission fee only)

Additionally, we provide an overview (Tab.6) of the relationships between non-custodial staking providers and PoS blockchains.

Staking types. Lastly, we explain four staking options.

⋄ *Solo home staking* is seen as the most impactful, offering full control and rewards to users who are ready to commit at least 32 ETH. This method enhances network decentralization but demands technical know-how and a dedicated setup.

⋄ *Staking as a Service* suits those who want to stake 32 ETH but prefer a simpler approach. Users delegate the validation process, though trust in the provider is necessary.

⋄ *Pooled staking* is an alternative for users with any amount of ETH. It introduces liquidity tokens, making staking more flexible and accessible. Users keep custody of their assets but must be aware that these solutions are third-party creations.

⋄ *CEXes* are the least involved option, offering minimal oversight and effort for stakers uncomfortable with self-custody. However, they consolidate large ETH pools, posing centralization risks.

C RELATED WORK

Ethereum evolution. Ethereum has been developed for years. The 1.0 version (*Frontier*) was launched in 2015 and introduced the concept of smart contracts to enable decentralized applications (DApps). Subsequently, a series of upgrades [26] was added to Ethereum to improve its scalability, security, and functionality. Notable milestones include *Homestead* (block index #1, 150, 000), *Byzantium* (#4, 370, 000), *Constantinople* (#7, 280, 000), *Istanbul* (#9, 069, 000). The 2.0 version (*Serenity*, the focus of this paper) is a massive upgrade to the Ethereum blockchain that will bring about many transitions including PoW to PoS, EVM to eWASM, and rollups integration. The transformation commenced with the establishment of the parallel Beacon Chain (Dec 1, 2020), and has since witnessed a sequence of significant updates, including *Berlin* (#12, 965, 000), *London* (#12, 965, 000), *Paris* (a.k.a., *Merge*, #15, 537, 394), and *Shanghai* (#17, 034, 870). Furthermore, Ethereum's roadmap outlined a series of near-future updates [27], including *Merge* (PoW to PoS), *Surge* (rollups [28], data sharding [29]), *Scourge* (PBS [12], MEV protection [30]), *Verge* (Verkle tree [31][32]), *Purge* (protocol simplification) and *Splurge* (account abstraction [33], EIP-1559 [34]).

PoS consensus. The first instance of PoS adoption in blockchain occurred with PPcoin [35], pioneering efforts to replace PoW with PoS within the Bitcoin ecosystem. This motivates a series of new constructions in different aspects, including investigations into potential attacks [36–39], advancements in provable secure constructions [40, 41], and improvements in properties covering security [42, 43], privacy [44, 45], availability [14, 46], consistency [47, 48], finality [49, 50], dynamicity [51], decentralization [52] and fairness [53, 54]. Recognizing the wide-reaching potential of PoS, Ethereum also initiated its incorporation of PoS into its core development by proposing a series of guiding principles and protocols [13, 15, 55, 56] (details in Sec.2).

PoS variants. PoS is adaptable and has given rise to various variants. Delegated Proof of Stake (DPoS) [57], adopted by EOS, TRON and Steem, enables users to stake tokens without becoming validators. Validators have the flexibility to adjust the rewards they share with their delegators as an incentive. Nominated Proof of Stake (NPoS) is a consensus model developed by Polkadot [58, 59] that shares many similarities with DPoS. One significant difference is that if a nominator (delegator) stakes behind a malicious validator, they may also risk losing their stake. Proof of Activity [60] is a hybrid consensus protocol that integrates elements of both PoW and PoS. Participants have the capability to engage in both mining and staking activities for block validation. Proof of Staked Authority (PoSA), implemented by BNB Smart Chain [61], combines Proof of Authority (PoA) [62] (e.g., OpenEthereum, Substrate) and PoS. Validators in this model take turns to forge blocks. A group of 21 active validators is eligible to participate, selected based on the amount of BNB they stake or have delegated behind them.

Staking analyses. Grandjean [63] conducted an analysis of the decentralization of staking power within Ethereum's beacon chain, highlighting the centralized distribution of validators' influence, primarily held by a small number of large entities. A similar study by He et al. [64] also identified stake concentration and established the existence of a stable equilibrium where no staking pool has an incentive to deviate. John et al. [65] dived into the relationship between block rewards and the equilibrium level of staking, demonstrating that staking levels do not consistently rise with increasing block rewards. Chitra et al. [66] explored the impact of MEV on validators, revealing that rational validators tend to remain active by sharing a portion of the MEV revenue through redistribution, thereby maintaining economic security. John et al. [67] investigated an economic model of PoS, exploring how adoption decisions are influenced by security risks and network congestion. Gersbach [68] conducted analyses on the existence and uniqueness of equilibria within staking pools, considering the presence of malicious agents, and identified potential risks. Brunjes et al. [69] proposed reward-sharing schemes that encourage the fair formation of stake pools involving a large number of stakeholders. Furthermore, several studies have also focused on other PoS-based blockchain platforms, such as Cardano [69] and Tezos [70, 71].

Voting power analyses. Mueller et al. [72] employed an agent-based simulation approach to investigate the decentralization of validators' decision-making power. Messias [73] investigated the

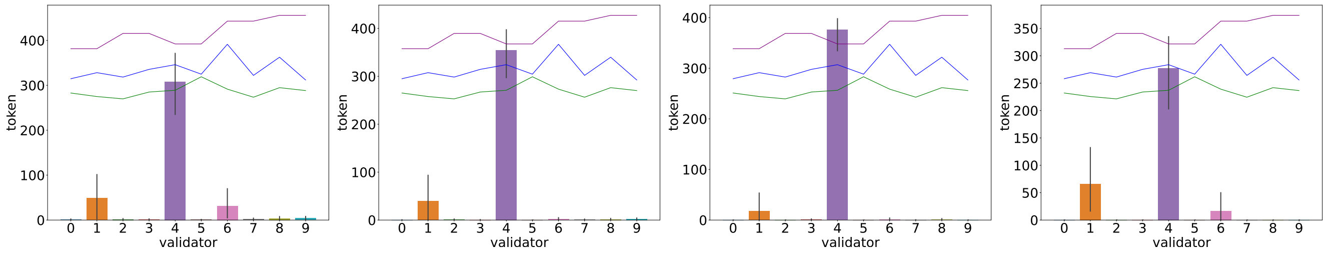

(a) ϵ = 0.7 **with** 3 **round.**

(b) ϵ = 0.8 **with** 3 **round.**

(c) ϵ = 0.9 **with** 3 **round.**

(d) ϵ = 1 **with** 3 **round.**

Figure 3: *Bars*: the average values of each validator's received delegation tokens in 20 **simulation instances, varying** ϵ **and fixing** $RL = 3$ **(i.e., the round limit).** ***Lines*: — validator integrities** p_j**, — validator evidences** $\Pr(b_j = 1)$**, — validator commission rates** c_j**.**

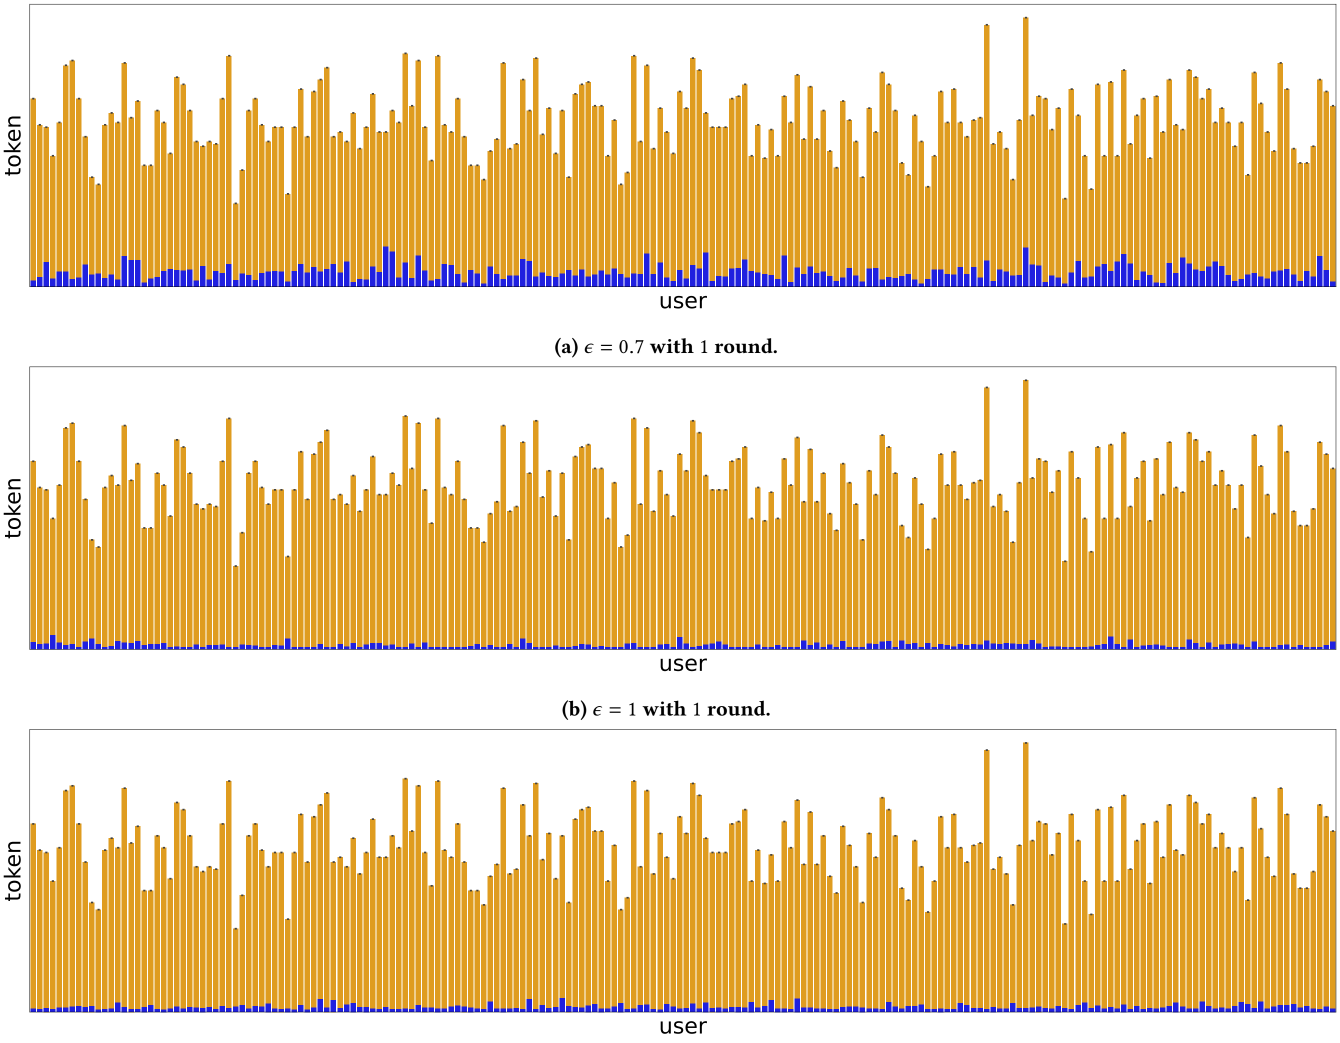

(a) ϵ = 0.7 **with** 1 **round.**

(b) ϵ = 1 **with** 1 **round.**

(c) ϵ = 0.7 **with** 3 **round.**

Figure 4: Users' token usage. 🟧: users' budgets, 🟦: users' average number of delegating tokens.

voting distribution and its impact on several leading DeFi protocols. Fritsch et al.[74] and Wang et al.[75] conducted research on the distribution of voting power among governance delegates and its impact on governance decisions within blockchain DAOs [76]. Li

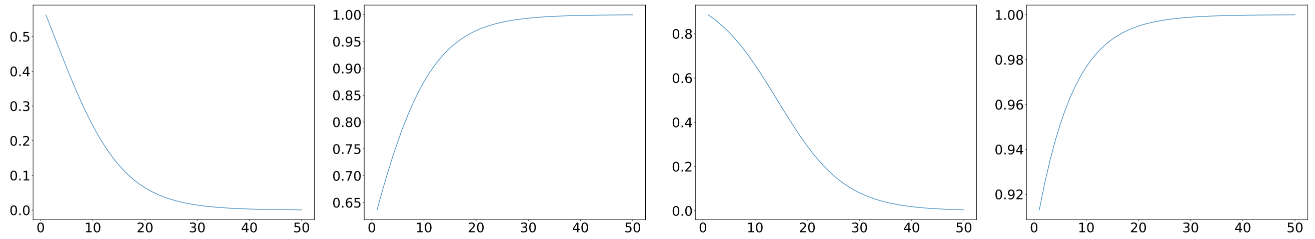

(a) $p_j = 0.6$, $(q, \bar{q}) = (0.6, 0.7)$. **(b)** $p_j = 0.6$, $(q, \bar{q}) = (0.7, 0.6)$. **(c)** $p_j = 0.9$, $(q, \bar{q}) = (0.6, 0.7)$. **(d)** $p_j = 0.9$, $(q, \bar{q}) = (0.7, 0.6)$.

Figure 5: Users' trust on a validator v_j, by varying integrity p_j and users' accuracy and error $(q, \bar{q})$. The x-axis denotes the number of users delegate to v_j.

et al. [4] analyzed the security issues of token-based voting governance in DPoS blockchains.

Validator selection (equiv. committee selection/formation). As discussed in Sec.1, most studies rely on three straightforward methods, *randomized block selection* (a combination of the lowest hash value and the highest stake) [2, 77], *coin age selection* (the duration of the tokens have been staked) [40, 44], and *node's wealth* (holding shares) [35]. Unfortunately, we encountered a lack of sufficient formal studies that analyze committee formation in the context of PoS. Nonetheless, we refer to several relevant materials. Gavzi et al. [78] explored the trade-off between committee size and the probability of selecting a committee that may contain corrupted validators, potentially resulting in system failure.

Delegation in blockchain. Grossi [79] engaged in a qualitative discussion of rational delegations within liquid democracy, particularly in the context of voting for delegators. Li et al. [5] conducted an empirical analysis of participation and delegation behaviors within DPoS-based blockchain systems. Notably, based on existing investigations, most blockchain delegation analyses have predominantly focused on *permissioned* blockchains, rather than exploring the nuances of *permissionless* protocols like PoS. This distinction underscores the motivation behind the present work.

Game theory in blockchain. Game theory has been widely employed as a tool for analyses in the blockchain field due to the profit-driven nature of its participants [80]. The first aspects typically involve defining constraints, which encompass both static and dynamic aspects, or involve two-party and multi-party interactions. Once these constraints are established, a proper game theory model can be applied, such as utilizing stochastic games [81], cooperative games [82], evolutionary games [8], and Stackelberg games [83].

The second aspect entails selecting an appropriate model that aligns with the assumed conditions and applying it to real-world cases. Lohr et al. [84] and Janin et al. [85] delve into two-party exchange protocols. Analyzing deviations from the Nash equilibrium provides insights into protocol designs. Qin et al. [86] develop a data trading platform and conduct an analysis of the interactions among involved parties via subgame perfect Nash equilibrium.

Additionally, the game analyses can be beneficial for constructing fairness and security in mining procedures. Existing research independently examines miner strategies in contexts such as selfish mining [87][88][89][90], multiple miners strategy [91], compliance strategies in PoW/PoS [92], pooled mining strategies [93][94], and fickle mining behaviors across different chains [95].

Table 5: Notations

Name	Symbol	Function
Participants	P	All participants in the market: $A \cup V$.
User	A (or a_i)	The group we are focused on in this paper.
Validator	V (or v_j)	The group who are eligible for producing blocks.
Delegation	$\mathbf{d}$ (or d_i)	A user i's choice of delegation.
Delegation	$\mathbf{t}$ (or t_i)	A user i's choice of token number.
Delegation set	$d(v_j)$	The group who delegated to the validator v_j.
Commission fee	c	The charge of staking services, $c \in [0, 1]$.
Event	$\theta_j = 1/0$	A validator will stay in the market/game (or not).
Delegation event	$\theta_{ij} = 1/0$	User a_i delegates to validator v_j.
Integrity	p_j	A prior that v_j stays in the he market/game.
Evidence quality	z_j	The probability that evidence e_j perfectly reveal p_j.
Accuracy and error	$q_{ij}, \bar{q}_{ij}$	The probability that a_i delegates to v_j conditioned on $e_j = 1$.
Trust	T	A statistic process for stimulating user's trust.
Budget	b_i	All tokens a_i can use.
Utility	u_i	The utility function of a_i.

Table 6: Mappings between *providers* and *PoS blockchains*

	Lido	Ankr	Dokia Capital	Marinade Finance	P2P Validator	Rocket Pool	StakeWise	StakeWithUs	Staked	Stakin	Staking Facilities	Stake.fish
BNB Chain		✓										
Fantom		✓										
Avalanche		✓										
Serum			✓									
Nucypher					✓							
Regen					✓							
DAObet					✓							
Marlin					✓							
LSD	✓	✓					✓					
Flow					✓			✓				
Polygon	✓	✓			✓				✓	✓		
Ethereum	✓	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓
Polkadot		✓			✓				✓	✓	✓	✓
Kusama		✓	✓		✓				✓	✓		✓
Kava			✓		✓			✓	✓	✓		✓
Cosmos			✓		✓			✓	✓	✓	✓	✓
IRISnet			✓		✓					✓		✓
Near			✓		✓				✓	✓		✓
Solana	✓		✓	✓	✓			✓	✓	✓	✓	✓
Tezos					✓				✓	✓	✓	✓
TheGraph					✓			✓	✓		✓	
Cardano					✓				✓			✓
Oasis					✓							✓
Mina					✓				✓	✓		
SUI					✓			✓	✓	✓	✓	✓
APTOS					✓				✓	✓	✓	
Persistence					✓			✓				
Osmosis					✓							✓
JUNO								✓				✓
Band Protocol								✓				✓
Loom								✓				✓
SEI									✓	✓		
Algorand										✓		
ICON										✓		
Celo										✓		
Casper Net.												✓