crypto.com

A Discussion on Efficiency and Security for Committee-Based Consensus Protocols

May 2022

Research and Insights

Research Manager
Kevin Wang

Head of Research and Insights
Henry Hon PhD, CFA, CAIA

Published on 16 May 2022

RESEARCH DISCLAIMER

The information in this report is provided as general commentary by Crypto.com and its affiliates, and does not constitute any financial, investment, legal, tax, or any other advice. This report is not intended to offer or recommend any access to products and/or services. The views expressed herein are based solely on information available publicly, internal data, or information from other reliable sources believed to be true.

While we endeavour to publish and maintain accurate information, we do not guarantee the accuracy, completeness, or usefulness of any information in this report nor do we adopt nor endorse, nor are we responsible for, the accuracy or reliability of any information submitted by other parties. This report includes projections, forecasts, and other predictive statements that represent Crypto.com's assumptions and expectations in light of currently available information. Such projections and forecasts are made based on industry trends, circumstances, and factors involving risks, variables, and uncertainties. Opinions expressed herein are our current opinions as of the date appearing in this report only.

No representations or warranties have been made to the recipients as to the accuracy or completeness of the information, statements, opinions, or matters (express or implied) arising out of, contained in, or derived from this report or any omission from this document. All liability for any loss or damage of whatsoever kind (whether foreseeable or not) that may arise from any person acting on any information and opinions contained in this report or any information made available in connection with any further enquiries, notwithstanding any negligence, default, or lack of care, is disclaimed.

This report is not meant for public distribution. Reproduction or dissemination, directly or indirectly, of research data and reports of Crypto.com in any form is prohibited except with the written permission of Crypto.com. This report is not directed or intended for distribution to, or use by, any person or entity who is a citizen or resident of, or located in a jurisdiction, where such distribution or use would be contrary to applicable law or that would subject Crypto.com and/or its affiliates to any registration or licensing requirement.

The brands and the logos appearing on this report are registered trademarks of their respective owners.

Contents

Executive Summary **5**

1. Introduction **6**

1.1 Proof of Stake (PoS) 6

1.2 Committee-Based Consensus & Delegated Proof of Stake (DPoS) 6

1.3 Approval Voting 7

1.4 Motivation 7

2. Model **9**

2.1 Voting with Limited Information 10

2.2 Class of Voting Strategies 11

3. Analysis **12**

3.1 The Probability of Electing an Honest Committee 12

3.2 Complexity of the Optimal Voting Strategy 12

3.3 Special Cases: Single-Voter & Signal Pooling 14

3.4 General Case: Multiple Voters 15

3.5 Asymptotic Optimality 16

3.6 Efficiency Gains over Alternative PoS Protocols 17

4. Limitations **19**

5. Conclusion **19**

References **20**

Published on 16 May 2022

Executive Summary

- We introduce the key points of "**Scaling Blockchains: Can Elected Committees Help?**" (referred to as the 'paper' in this report) by Alon Benhaim, Brett Hemenway Falk, and Gerry Tsoukalas.
- In 'committee-based consensus', voters delegate to a small committee the rights to produce and certify blocks.
    - Often associated with Delegated Proof of Stake (DPoS) systems, stake-weighted voters elect a small group of block producers as a committee responsible for producing and validating blocks.
    - Most DPoS systems adopt approval voting, where each stake-holding voter "approves" preferred candidates to join the committee.
- The paper lays out a model where a blockchain's token holders vote to elect a committee of block producers according to an **up-to-t**, **k-winner** approval voting system; **t** is the maximum number of candidates each voter can vote for, and **k** indicates the candidates with the highest scores.
    - The goal of the model is to seek an optimal strategy to maximise the probability of electing an honest committee.
    - Data collected from EOS is used as empirical observations to analyse the real-world voting strategies.
    - There are two particular voting strategies illustrated: threshold voting and cardinal voting.
- In the model:
    - The authors derive the formula to calculate the objective function of the optimisation, which is the probability that an elected committee is honest.
    - The voters' behaviour dramatically depends on the low vs. high number of voters.
    - The probability of success rapidly and exponentially converges to 100% when the signals are deemed not entirely uninformative.
    - The minimum committee size required to achieve a failure probability is much smaller for the DPoS-like system compared with the randomly chosen committee (as in Algorand).
- The paper also illustrates the problems of approval voting:
    - One drawback of electing committees, compared to selecting random committees, is that elections seem to lead to stagnation.
    - A small, static set of block producers reduces decentralisation — the core tenet of blockchains and cryptocurrencies.

1. Introduction

The Bitcoin White Paper introduced a new consensus protocol (known as Nakamoto's consensus) that allows participants to reach an agreement on a sequence of events, even without their stable identities. This paved the way to a new form of decentralised currency and more. But permissionless blockchains suffer a challenging problem: How can trustless, decentralised nodes all consent to, for instance, state updates? The key behind Nakamoto's consensus is that participants can continuously establish trust by consuming verifiable computational power via proof of work (PoW).

Despite the proven security of PoW-based consensus mechanisms for more than a decade, it is a wasteful process that has been unable to generate the required throughput to handle massive global transactions. Finding a better alternative may contribute to the global economy from the innovation of blockchain technology, and also determine the winners and losers of the ongoing cryptocurrency arms race in the long run.

Presently, the proof of stake (PoS) system has become one of the promising consensus protocols widely used in blockchains, as it is more scalable than PoW.

1.1 Proof of Stake (PoS)

In PoS, a stake-weighted lottery elects block producers in proportion to their token stake on the blockchain instead of their computational power, like in PoW. But this selection process is considered problematic, as malicious producers who try to fork the chain can also be selected. Thus, PoS systems require additional methods to avoid forks.

1.2 Committee-Based Consensus & Delegated Proof of Stake (DPoS)

In PoS, block producers (validators) can earn considerable returns, including transaction fees and block rewards. However, being an efficient block producer usually requires powerful computing resources plus a significant amount of tokens; thus, It is difficult for regular token holders to take on this role. Most PoS systems support some kind of delegation mechanism to solve this problem, allowing token holders to delegate their stake to professional block producers (usually in exchange for profit-sharing).

Committee-based consensus maximises this separation between token holders and block producers. As the name suggests, **in 'committee-based consensus', the rights to produce and certify blocks are delegated to a small committee.** Some PoS systems (e.g., Cardano) adopted this idea to elect leaders randomly with probability equal to their proportional token stake.

Another mechanism involves users voting for block producers proportional to their staked tokens. Validators with the highest number of votes become producers for some fixed time slot. **Often associated with Delegated Proof of Stake (DPoS) systems, this type of stake-weighted vote elects a small committee of block producers who are responsible for producing and validating blocks.** This 'committee-based consensus' is the main topic discussed in this article.

Several prominent blockchains, including Cosmos, EOS, TRON, and Algorand use this committee-based approach, though they differ in how to select the committee members: random selections (e.g., Algorand), single-vote election (e.g., Cosmos), or approval voting in DPoS consensus (e.g., EOS and TRON).

1.3 Approval Voting

Committee-based consensus protocols have some variants, but the key difference is how committees are selected. Most DPoS systems adopt **approval voting, where each stake-holding voter "approves" preferred candidates; in this single-winner electoral system, the candidate approved by the largest number of voters is the winner.** Approval voting is used in multi-winner systems in the DPoS, in which multiple candidates can be elected to join the committee. This is different in nature from traditional voting schemes, where voting for two candidates could split the vote. In approval voting, when voters choose multiple candidates, each candidate receives the same 'approval' as the other chosen ones. Adopted into the blockchain space via DPoS, approval voting is generally considered a more efficient and democratic version than the standard PoS mechanism.

1.4 Motivation

The core of committee-based consensus is straightforward: Participants continuously vote to elect their preferred nodes to the committee. Since the committee size is usually small, the efficiency of the blockchain can be improved, including increasing throughput, decreasing latency, and allowing for member specialisation. However, the security of the entire blockchain can be undermined since malicious nodes can also be elected as committee members. Hence, there is

a concern between performance and robustness — a small committee is efficient but may compromise security.

Additionally, there are several questions related to the committee-based consensus:

1. How should nodes vote for their preferred candidates with limited information?
2. What is the minimum size of a committee to ensure security?
3. How efficient is committee-based consensus with approval voting compared to other PoS protocols?

To answer the questions above, the authors of the paper developed a 'simple' voting model using the DPoS protocol of **EOS** as a carrier. Since block producers can behave either honestly or dishonestly, the vote is considered successful if a **2/3 majority** of the elected committee is honest, according to the Byzantine Fault Tolerance.

2. Model

This section briefs the core parts of the model (original paper here), covering the real-world voting data collected from EOS from 20 August 2021 to 30 August 2021. Using the up-to-t, k-winner approval voting system, block producers on EOS were elected by token holders according to an up-to-**30**-vote, **21**-winner model. There were up to **30** candidates the voters could vote for and **21** candidates to form the committee.

Definition 1 (k-winner approval voting)

A set of voters V votes on a set of candidates C. Let n be defined as the number of voters in V, and m be defined as the number of candidates in C. Voter v chooses a subset of candidates $C_v \subseteq C$. For each candidate $c \in C$, the score of candidate c is defined as the number of votes the candidate receives:

$$score(c) = \left|\{v | c \in C_v\}\right|$$

The elected committee is determined to be the k candidates with the highest scores.

In DPoS protocols, token holders vote for a group of 'block-producers', changing the k-winner approval voting system to a limited number of candidates, as defined below:

Definition 2 (up-to-t-vote, k-winner approval voting)

With notation as in Definition 1, we limit the maximum number of candidates t each voter can vote for, so that voter v chooses a subset of candidates $C_v \subseteq C$ restricted to $|C_v| \leq t \leq m$. As in Definition 1, the elected committee is determined to be the k candidates with the highest scores.

The model first characterised voters' optimal voting strategies and the obtainable pure-strategy Bayesian Nash equilibrium. In game theory, Bayesian Nash equilibrium is defined as a strategy profile that maximises the expected result for each player given their beliefs and the strategies played by other players. Bayesian Nash equilibrium is derived from the Bayesian games used to model the games with incomplete information. (For example, in committee-based consensus, voters have limited information about the candidates.) Next, two simple voting strategies were analysed by their basic empirical observations: threshold voting (where participants vote for all candidates whose probability of being honest is above a

certain threshold) and cardinal voting (where voters approve their top k candidates).

2.1 Voting with Limited Information

The paper lays out a model where the blockchain's token holders vote according to an *up-to-t*, *k-winner* approval voting system (see Definition 2) **in order to elect a committee of block producers.**

There is a pool of **m** producers to choose from, $c_1, \ldots, c_m$, and **n** voters on the platform, $v_1, \ldots, v_n$. Every producer has an unknown type, either 'honest' (H) or 'malicious' (M). **The goal of each voter is to maximise the probability that a 2/3 majority of the elected committee is honest.**

Definition 3 (honest committee)

Suppose the k producers with highest number of votes are elected to be on the block-producer committee, T. If there are less than k candidates with non-zero score, then the committee is filled adversarially (i.e., in a worst-case fashion), and if there are ties between the candidates such that there are more than k producers with the highest score, then they are broken adversarially (between the ones with least score). Since most Byzantine Agreement protocols require at least $\frac{2k}{3}$ honest members, we say the committee T is honest, $T = H$, if at least $\frac{2k}{3}$ of the elected block producers are honest.

With this definition and the knowledge of probability, the authors modelled the probability of honest and malicious producers when they sent a raw private signal to voters. Additionally, **a bijective function was built to compute the posterior probability that producer j was honest conditioned on the raw signals.**

The core model assumes all voters are strategic (entirely rational) and seeks to characterise the pure-strategy Bayesian Nash equilibrium of the game using the platform voting system described in Definition 2. Specifically, each voter v_i maximises the **success probability —** that which the elected committee, T, is honest conditioned on the private signal vector they received.

Definition 4 (voting strategy)

A voting strategy is an algorithm used by all voters that inputs the parameters accessible to the voter and outputs a subset of candidates in which the voter wishes to vote.

An optimal strategy is to maximise the success probability — the probability of electing an honest committee. Since computing the objective function is challenging, the types of acceptable voting strategies for voters are defined below.

2.2 Class of Voting Strategies

There are two particular voting strategies illustrated simply: **threshold voting** and **cardinal voting.**

Definition 5 (threshold voting)

Voter v_i is said to follow the threshold voting strategy if (prior to seeing the realisation his or her signals), voter v_i chooses a threshold $z_i \in [0, 1]$ and voter v_i votes for all producers c_i with probability of being honest higher than the threshold.

Define p_{ij} as the probability that voter i chooses producer j. Summing up all n voters, the number of votes received by producer j is distributed as the sum of n Bernoulli random variables with parameters $p_{1j}, \ldots, p_{nj}$, suggesting that the number of votes received by producer j is:

1) a binomial random variable, if $p_{1j} = \cdots = p_{nj}$
2) a Poisson binomial random variable, if the p_{ij} numbers are distinct.

Definition 6 (cardinal voting)

Voter v_i is said to follow the cardinal voting strategy if (prior to seeing the realisation his or her signals), voter v_i creates a strategy $z_i \in \{1, \ldots, t\}$, then voter v_i orders producers according to their probability of being honest and votes for the top z_i producers in the list.

3. Analysis

3.1 The Probability of Electing an Honest Committee

The first step in the analysis is to define the **objective function** of the optimisation — **the probability that an elected committee is honest.**

This section covers the authors' proposal of two theorems and two propositions to formulate and calculate the probability of electing an honest committee.

Theorem 1 (success probability) derives the formula that at least **2k/3** honest producers are in a committee (the size of committee is k). **Theorem 2 (distribution of votes)** gives the corresponding probability distribution of the number of votes received for honest and dishonest producers. The combination of Theorems 1 and 2 gives a closed-form expression for the **success probability. Propositions 1 (threshold voting) and Propositions 2 (cardinal voting)** derive the calculation of those probabilities for threshold voting and cardinal voting.

For those interested in the mathematics behind the formula, please see the original paper.

3.2 Complexity of the Optimal Voting Strategy

The combination of **Theorems 1 and 2** with **Propositions 1 or 2** results in highly complex objective functions for the probability that an elected committee is honest. In order to understand the origin of this complexity, the paper focusses on the case where voters follow a simple threshold voting strategy. Additionally, the objective functions are visualised correspondingly.

For the first step, the paper gives the success probability as a function of the threshold chosen. The figure below shows the success probability under threshold voting with small numbers of voters (n = 1 to 4). Although systems have more voters practically, these graphs highlight the complex dynamics of approval voting.

crypto.com

Success Probability as a Function of Threshold Chosen (Small No. of Voters)

The number of local optima increases with n (number of voters)

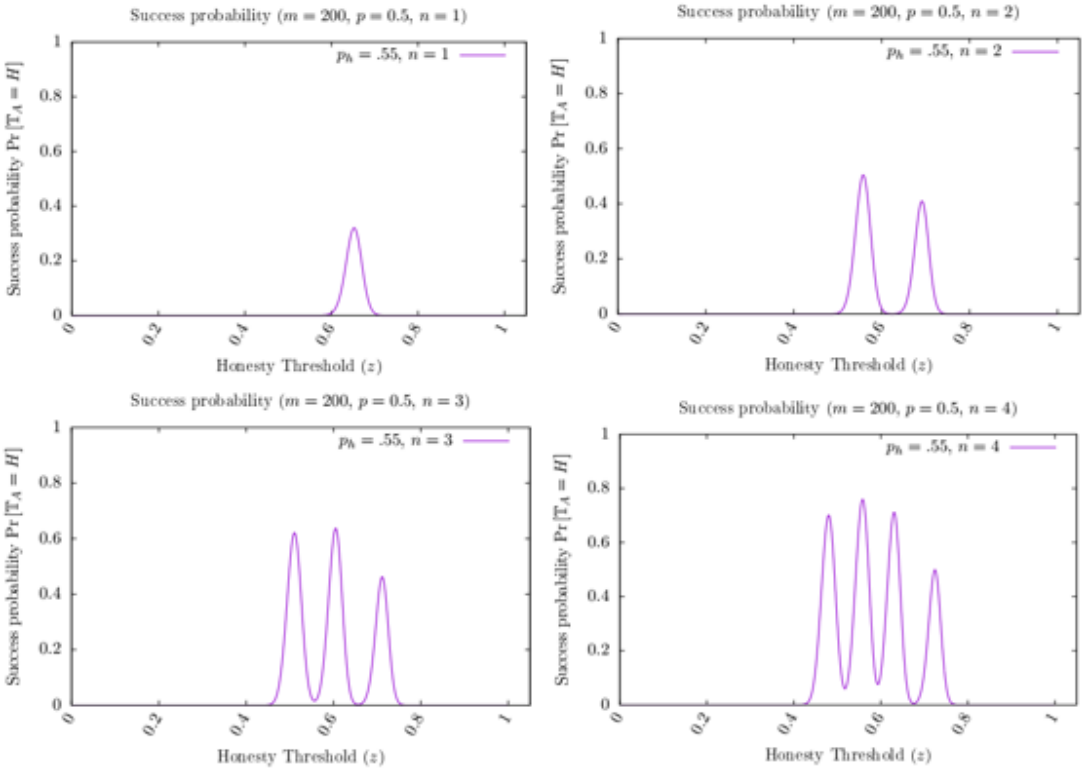

Note:
1) m: number of block producers (candidate); 2) n: number of voters; 3) p: a priori probability that block producer is honest; 4) z: threshold probability of honest producer; 5) p_h: the base signal for an honest candidate producer.

As of 2 Oct 2021
Source: Benhaim, Alon & Hemenway Falk, Brett & Tsoukalas, Gerry, Scaling Blockchains: Can Elected Committees Help? (7 October 2021).

From the graph above, the optimal thresholds are around 0.5–0.7, suggesting that voters should vote for any candidate j, whose posterior probability is above this threshold (producer j is honest conditioned on the raw signals). The thinness of the peaks indicates that even a small bias from the optimal strategy can drastically reduce the success probability. Meanwhile, **the number of local optima increases with the number of voters (n).**

Another experiment examines the situation for a large number of voters, n = 100. It shows that **for large n, the success probability increases to 100% across a wide range of thresholds, meaning that a wide range of voting strategies yields nearly optimal results.**

crypto.com

Success Probability as a Function of Threshold Chosen (Large No. of Voters)

For large n, the probability of success increases to 100% across a wide range of thresholds, and thus a wide range of voting strategies yields nearly optimal results.

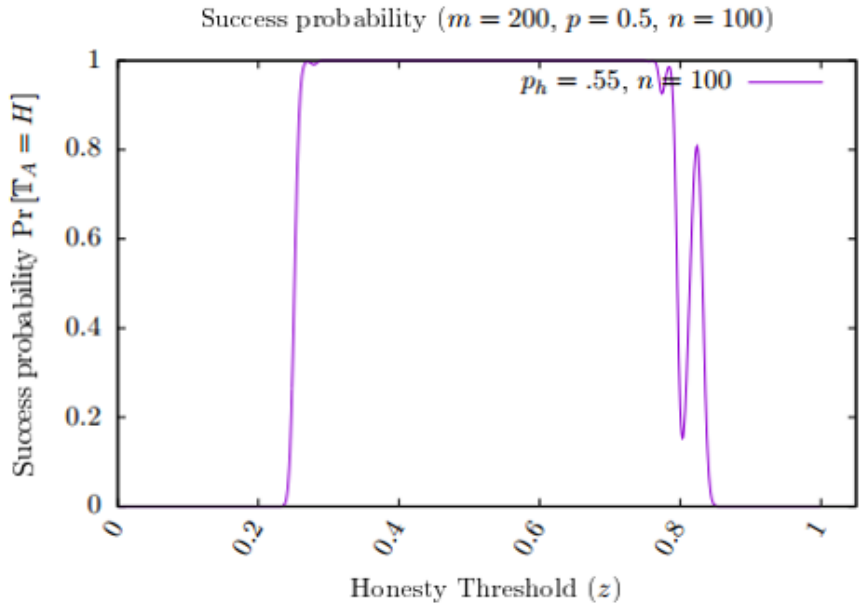

Note:
1) m: number of block producers (candidate); 2) n: number of voters; 3) p: a priori probability that block producer is honest; 4) z: threshold probability of honest producer; 5) p_h: the base signal for an honest candidate producer.

As of 2 Oct 2021
Source: Benhaim, Alon & Hemenway Falk, Brett & Tsoukalas, Gerry, Scaling Blockchains: Can Elected Committees Help? (7 October 2021).

Combining the information from the above two figures, the paper concludes that **the voters' behaviour dramatically depends on the number of voters.** It also supposes that asymptotic analysis may offer more insights.

3.3 Special Cases: Single-Voter & Signal Pooling

This section covers the authors considering the special case of a single voter (n = 1). It can also be treated as collapsing the general n > 1 case to n = 1 when voters can share their signals credibly (and costlessly). Another proposition is given:

Proposition 3 (optimality of cardinal voting when n = 1)

Consider a k-winner approval voting system with n = 1 voter and m ≥ k candidates, then the globally optimal strategy is the cardinal strategy with z = k.

This proposition indicates that voters should vote for the top k candidates (based on their posterior probability of being honest). This strategy is optimal across all possible strategies, not just cardinal or threshold voting.

Proposition 4 (suboptimality of threshold voting when n = 1)

Consider a k-winner approval voting system with n = 1 voter and m ≥ k candidates, then any threshold strategy is not optimal.

This proposition shows that **for n = 1, the threshold strategy gives a strictly lower success probability than the cardinal strategy.**

When voters can share their private signals credibly without cost, they effectively act as a single voter.

Proposition 5 (optimality of cardinal voting with shared signals)

Consider a k-winner approval voting system with n > 0 voters, m ≥ k candidates, and such that voters' private signals are credibly shared. Then the globally optimal strategy is the cardinal strategy with z = k, where each voter v_i is ranked based on the shared signal instead of their private signal.

Proposition 5 shows that **when voters share their signal, the optimal strategy is to follow cardinal voting with threshold z = k (voting for the top z producers in the list).**

3.4 General Case: Multiple Voters

Proposition 6 (suboptimality of cardinal voting when n > 1)

Consider a k-winner approval voting system with n > 1 voter and m ≥ k candidates, then the cardinal strategy can be suboptimal.

This result occurs because a vote for one candidate can actually bump other candidates out of the committee.

3.5 Asymptotic Optimality

This section covers the authors introducing ***Theorem 3*** *(exponential convergence),* which shows that success probability follows an exponential form — as the number of voters (n) increases, the higher the probability of electing an honest committee in most reasonable strategies.

The visualisation of the exponential convergence result displayed below assumes each voter followed a generally suboptimal threshold voting strategy where z = p. The figure shows that **as long as the signals are not completely uninformative** ($p_h \neq 0.5$), **the probability of success rapidly converges to 100%.**

crypto.com

Rate of Convergence to Optimality

The success probability quickly goes to one, as n increases.

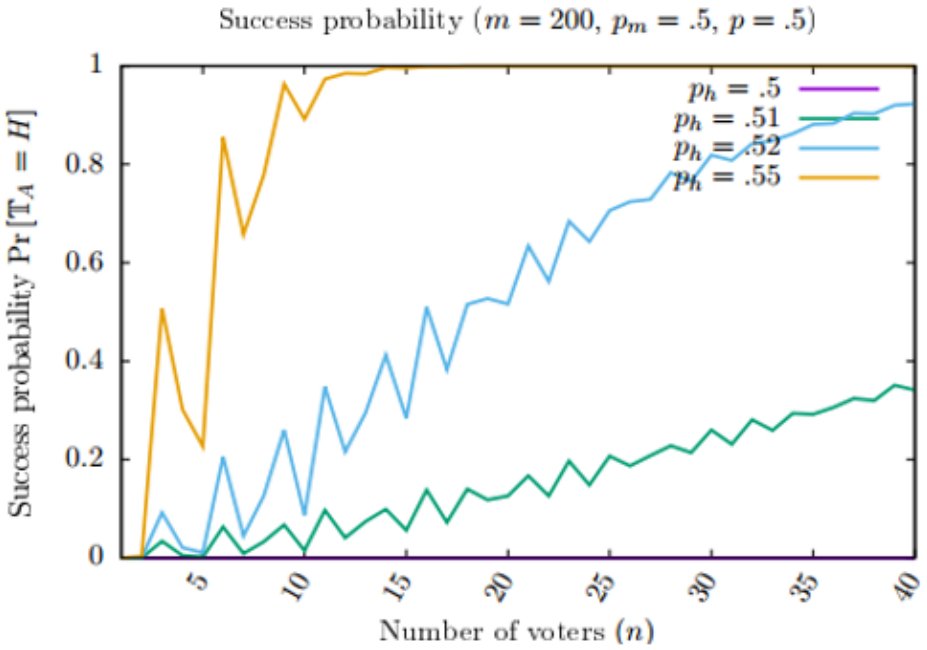

Note:
1) m: number of block producers (candidate); 2) n: number of voters; 3) p: a priori probability that block producer is honest; 4) p_m: the base signal for a malicious candidate producer; 5) p_h: the base signal for an honest candidate producer.

As of 2 Oct 2021
Source: Benhaim, Alon & Hemenway Falk, Brett & Tsoukalas, Gerry, Scaling Blockchains: Can Elected Committees Help? (7 October 2021).

3.6 Efficiency Gains over Alternative PoS Protocols

The paper analyses the minimum committee size necessary in order to achieve a desired failure probability between the two cases: the randomly chosen committee (as in Algorand) versus the elected committee according to an approval vote (as in DPoS). Even when the voters have only minimal information, allowing users to vote for candidates drastically reduces the size of the committee necessary to achieve a specific failure bound. Since the committee executes a Byzantine Agreement protocol with communication cost that is quadratic in the committee size **k**, minimising the committee size is critical for performance.

The figure below clearly shows that **the minimum committee size required to achieve a failure probability is noticeably smaller for the committee elected by voters (as in DPoS) compared to the randomly chosen committee (as in Algorand).**

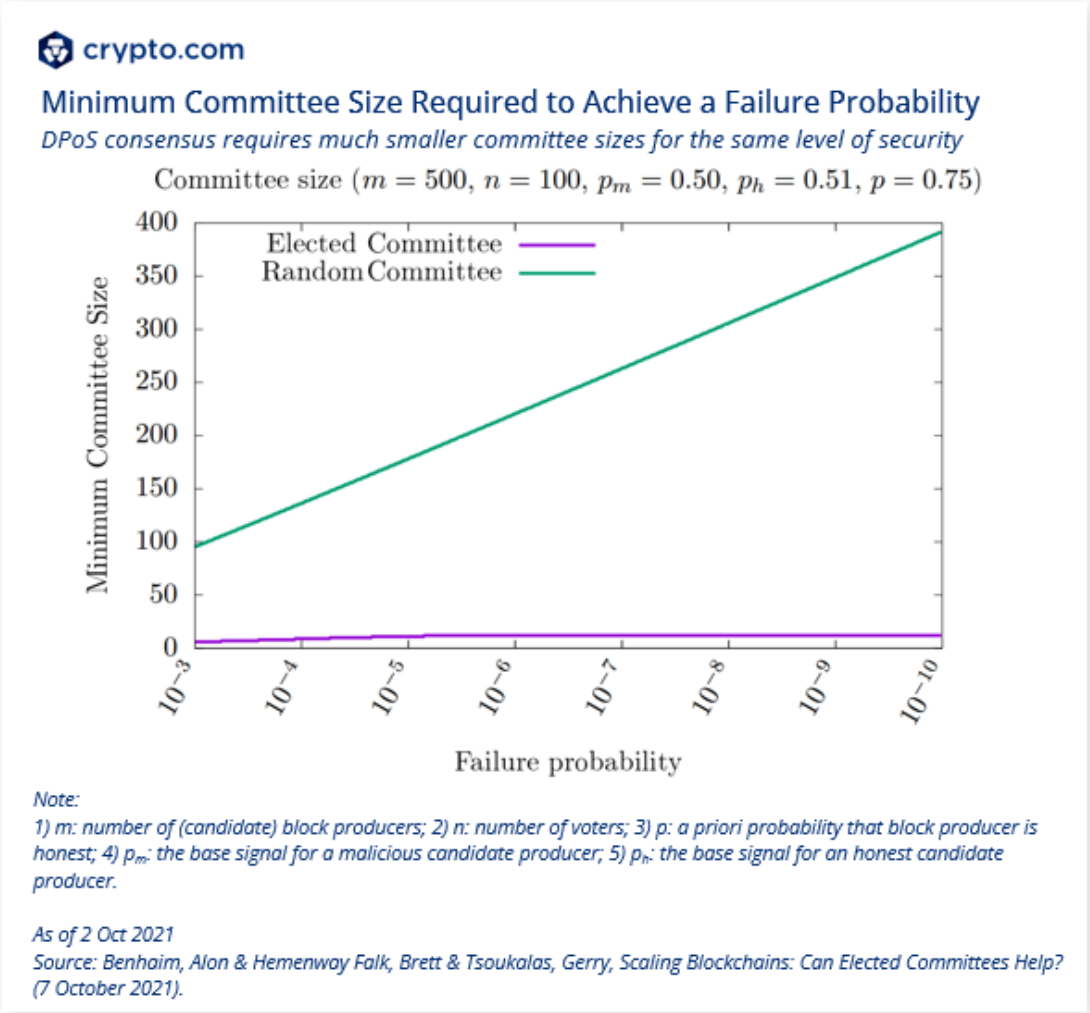

4. Limitations

The results suggest that in a committee-based consensus, there seems to be no rule for voters to follow and they could vote intuitively. Nonetheless, these systems are asymptotically robust and efficient from an election perspective, but the authors address some defects. The model does not, however, consider other factors that voters may care about, except the optimisation to reduce failure rates.

One drawback of electing committees (as in Cosmos, EOS, and TRON) compared with selecting random committees (as in Algorand) is that elections seem to lead to stagnation, especially in the early stages of the blockchain life cycle. Research shows that in EOS there are only 63 distinct producers who mined the first 89 million EOS blocks; whereas, the first 655,000 Bitcoin blocks were mined by more than 275,000 distinct addresses.

Moreover, a small, static set of block producers reduces decentralisation — the core tenet of blockchains and cryptocurrencies. The diversity of block producers is key to the open and democratic blockchain ecosystem. Moreover, the turnover in the set of block producers is deemed chain quality, a measure of fairness.

5. Conclusion

The paper proposes mathematical methods to evaluate the robustness of committee-based consensus protocols, focusing on the approval voting mechanism. The results show that as long as the private signals from committee candidates to voters are not entirely uninformative, the probability of success rapidly converges to 100%. Moreover, the authors also conclude that DPoS consensus requires much smaller committee sizes for the same level of security than those that adopt a randomly chosen committee (as in Algorand).

Using private information and strategic agents, the paper is the first to analyse the efficiency of committee elections in committee-based consensus protocols. The authors also point out that the chain quality (the turnover in block producers) could be used as another metric to measure the election mechanism in future research.

References

Benhaim, Alon, et al. “Scaling Blockchains: Can Elected Committees Help?” 7 October 2021, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3914471.

Nakamoto, Satoshi. “A Peer-to-Peer Electronic Cash System.” *Bitcoin.org*, https://bitcoin.org/bitcoin.pdf.

Zheng, Weilin, et al. “XBlock-EOS: Extracting and Exploring Blockchain Data From EOSIO.” *arXiv*, 26 March 2020, https://arxiv.org/abs/2003.11967.

“Approval voting.” *Wikipedia*, https://en.wikipedia.org/wiki/Approval_voting.

“Bayesian game.” *Wikipedia*, https://en.wikipedia.org/wiki/Bayesian_game#Bayesian_Nash_equilibrium.

e. contact@crypto.com

Published on 16 May 2022