Artificial Delegates Resolve Fairness Issues in Perpetual Voting with Partial Turnout

Apurva Shah[1,2,3][0009−0003−6838−7133], Axel Abels[1,2,3][0000−0003−2784−8653], Ann Nowé[2,3][0000−0001−6346−4564], and Tom Lenaerts[1,2,3,4][0000−0003−3645−1455]

[1] Machine Learning Group, Université Libre de Bruxelles, Brussels, Belgium
[2] AI Lab, Vrije Universiteit Brussels, Brussel, Belgium
[3] FARI Institute, Université Libre de Bruxelles - Vrije Universiteit Brussel,Brussels, Belgium
[4] Center for Human-Compatible AI, UC Berkeley, Berkeley, CA 94720, USA

Perpetual voting has been proposed as a mechanism for ensuring long-term fairness in sequential collective decision making (CDM), where sequences of decisions are made over time [5]. A key challenge in this setting is to avoid "Tyranny of the Majority", i.e., ensuring that voters who are repeatedly in the minority receive adequate representation. Traditional aggregation methods, such as majority voting or Borda count, fail as they do not guarantee long-term fairness or cumulative influence. In response, perpetual voting rules, for example, Perpetual Phragmén, Perpetual Consensus, and Perpetual Quota, have been proposed to balance representation across rounds and prevent persistent marginalisation of minority groups [6]. However, these voting rules rely on an assumption of full voter participation, which is rarely satisfied in real-world settings where absenteeism is common. Our paper [7] addresses the open challenge of maintaining fairness in perpetual voting under partial turnout.

While approaches such as Multiple Imputation [1] can statistically estimate missing preferences, they typically require large datasets and are less reliable in smaller electorates. Another popular approach is liquid democracy [2], where voters can delegate their vote, often transitively. This can lead to power concentration [3, 4] and, in practice, makes it difficult to identify appropriate delegates. To address these challenges, we introduce Artificial Delegates, preference-learning agents that act as proxies for absent voters. Each delegate is personalised and represents one voter. Using logistic preference-learning models trained on approval data, these agents learn to approximate voters' preferences from historical participation and cast votes on their behalf in rounds where the voter is absent.

Our study investigates two main research questions: (1) How does absenteeism affect fairness and representativeness in perpetual voting systems? and (2) To what extent can Artificial Delegates mitigate or exacerbate these effects? To evaluate this, we conduct simulations of sequential elections under varying parameters, analysing outcomes across four voting rules (Approval Voting, Perpetual Phragmén, Perpetual Consensus, and Perpetual Quota). Fairness is quantified using established metrics: quota compliance (representation proportionality), longest dry spell (periods of prolonged dissatisfaction), and the Gini influence coefficient (inequality in decision-making influence).

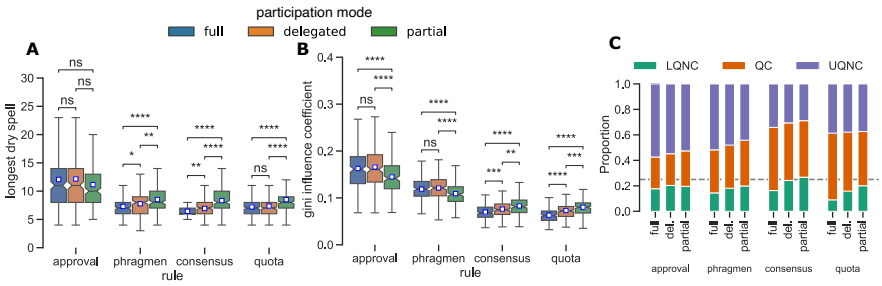

Fig. 1. Comparison of voting rules across participation modes — *full*, *partial*, and *delegated* — for different evaluation metrics. **A,B** Each boxplot represents the distribution for a specific rule over a given voter set type. Notches indicate CI around the median, and squares denote the mean. Statistical differences across participation modes for each rule are tested using Kruskal–Wallis tests, followed by post-hoc Dunn tests. Significance levels: ns ($p > 0.05$), * ($p \leq 0.05$), ** ($p \leq 0.01$), *** ($p \leq 0.001$), **** ($p \leq 0.0001$). **C** Distribution of LQNC, QC, and UQNC per rule and voter set types.

Our results, as illustrated in Figure 1, show that absenteeism systematically undermines fairness, particularly harming Perpetual Quota, which directly relies on satisfaction and quota alignment. Absentees are treated as having empty preferences, leading to underestimation of their entitlements and distortion of load-balancing mechanisms. In contrast, Perpetual Phragmén and Perpetual Consensus exhibit some resilience, as their balancing procedures implicitly account for underrepresented voters. Strikingly, approval voting may sometimes appear to benefit from absenteeism, as reduced turnout occasionally enables minority wins.

Artificial Delegates consistently restore fairness levels close to those under full participation. They reduce dry spells, equalise voter influence, and improve quota compliance across scenarios, while preserving outcome fidelity: delegated outcomes largely overlap with those of either full or partial turnout, ensuring that delegation does not distort decisions. Robustness analyses further demonstrate that the benefits of delegation persist across variations in absenteeism rates, electorate size, number of alternatives, cluster density (polarisation), and voter noise, though performance declines when absenteeism is extreme, as insufficient data limits learning accuracy. This work contributes three main advances: (i) systematic analysis of absenteeism in perpetual voting; (ii) a novel framework for integrating Artificial Delegates into sequential decision-making; and (iii) empirical evidence that delegation mitigates fairness loss due to inconsistent participation. By bridging computational social choice, preference learning, and voting theory, we show that Artificial Delegates make perpetual voting more robust and representative in realistic settings where full turnout cannot be guaranteed.

While this study provides an initial exploration of absenteeism and the benefits of Artificial Delegates, future work should address more practical aspects. In real-world settings, voters are influenced by complex behavioural dynamics such as deliberation effects and cognitive limitations when faced with many alternatives. Additionally, beyond approval-based approximations, learned utility models may provide richer insights for enhancing fairness in dynamic CDM.

References

1. Bernhagen, P., Marsh, M.: The partisan effects of low turnout: Analyzing vote abstention as a missing data problem. Electoral studies **26**(3), 548–560 (2007)
2. Blum, C., Zuber, C.I.: Liquid democracy: Potentials, problems, and perspectives. Journal of political philosophy **24**(2), 162–182 (2016)
3. Gölz, P., Kahng, A., Mackenzie, S., Procaccia, A.D.: The fluid mechanics of liquid democracy. ACM Transactions on Economics and Computation **9**(4), 1–39 (2021)
4. Kahng, A., Mackenzie, S., Procaccia, A.: Liquid democracy: An algorithmic perspective. Journal of Artificial Intelligence Research **70**, 1223–1252 (2021)
5. Lackner, M.: Perpetual voting: Fairness in long-term decision making. In: Proceedings of the AAAI conference on artificial intelligence. vol. 34, pp. 2103–2110 (2020)
6. Lackner, M., Maly, J.: Proportional decisions in perpetual voting. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 5722–5729 (2023)
7. Shah, A., Abels, A., Nowé, A., Lenaerts, T.: Artificial delegates resolve fairness issues in perpetual voting with partial turnout. In: Proceedings of the ACM Collective Intelligence Conference. pp. 71–82 (2025)