Tyranny of the Minority in Social Choice: a Call to Arms

Blue Sky Ideas Track

Reshef Meir*
Technion—Israel Institute of Technology, Haifa, Israel
reshefm@technion.ac.il

ABSTRACT

The vast majority of research in social choice theory, be it axiomatic characterizations, welfare bounds, or equilibrium analysis, makes the implicit assumption that the entire affected population votes. However this ideal description is very far from the situation in practice, and even more so in the age of online voting and various direct democracy initiatives—often, a handful of avid voters effectively make the decisions for the silent majority.

Our starting point is that abstention and partial participation are not a curious exception but the norm. We therefore argue that models of abstention should be given at least as much attention (if not more) as models of strategic behavior, and theoretical properties of voting rules such as welfare and fairness guarantees must be evaluated in light of such models in order to be relevant.

To that end, we call the multiagent systems research community to design more effective tools to better handle both causes and effects of low participation; and highlight promising directions.

KEYWORDS

Voting, Delegation, Abstention

ACM Reference Format:
Reshef Meir. 2025. Tyranny of the Minority in Social Choice: a Call to Arms: Blue Sky Ideas Track. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 5 pages.

> "There are two hundred million people in the country, and with odds like that, I don't think we ought to waste our time wondering about it... ...let's have supper."
> —Norman Muller, *Franchise* by Isaac Asimov

1 INTRODUCTION

Social choice is recognized as one of the core themes in multiagent systems, as evident in the high number of influential papers on the topic published at AAMAS and related venues. The idea is simple—having a number of agents (be it human, machine, organizations, or other) with conflicting opinions or preferences, let them vote and aggregate the ballots.

This process is justified by many theoretical results: From the famous Condorcet Jury Theorem, which states that a majority of independent opinions is likely to find the right answer [14], to various voting rules that are designed to maximize social welfare, guarantee various fairness properties, or hold certain axioms [6]. However, most of these results make the implicit assumption that all the relevant population votes. This observation is far from new, and even more than 20 years ago, Felsenthal and Machover [27] wrote that 'Theoretical discussion of abstention is conspicuous by its almost total absence in the literature.' While they meant specifically the literature on voting power, we argue that this observation is too true in general, even today.

Running example. Consider two alternatives A and B with higher 60% overall support for A. Yet many of A's supporters could be less informed, or face higher voting costs due to geographic or social constraints. It is thus possible that supporters of B turnout in higher numbers, making it win. Yet, the selected alternative would affect the entire population, whether they had voted or not.

Our purpose in this paper is first to convince that partial participation is not some curious anecdote but rather the norm, and that it should be taken seriously by any researcher in the area of social choice and voting. Then, we would like to sketch an initial map of the most relevant research directions for understanding and mitigating the effects of partial participation, focusing on the right type of models we should use and behavior we should incentivize.

Is turnout really low? Voting takes place in diverse circumstances, and turnout depends very much on the context and the voting population. We can perhaps expect full turnout at small high-stake votes like a jury or Papal elections,[1] but this is quite exceptional.

In Table 1 we collected common participation rates in various contexts. This is by no means a systematic study, but it serves to show that: (1) the typical turnout is far from 100% in most contexts; and (2) there is huge variability between and within contexts.

Thus indeed in many situations we may have a minority such as the B-supporters in our running example dictating their preferences over the 'silent majority'.

Moreover, this is only the tip of the iceberg. Recent proliferation of various voting apps and add-ons—from classic doodle.com to fun kahoot.com and polls on Facebook—allow essentially anyone to conduct voting on any issue, small or large, and elicit votes from a designated group. Unfortunately we are unaware of studies that looked at the turnout of such votes,[2] but there is no reason to believe it is higher than in other contexts [58].

The Paradox of Voting. The paradox of voting, which goes back at least to Downs [21] is the observation that even in medium-scale elections with thousands of voters, let alone large ones, a single voter is very unlikely to affect the outcome, and is thus almost

*This project was funded by the Israel Science Foundation (ISF; Grant No. 2539/20).

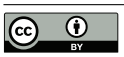

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

[1] Although cardinals as well as Congressmen had been put under arrest to prevent them from voting [36, 37].
[2] It is hard to tell for example how many people *did not* answer a Doodle poll.

Context	turnout	source
US presidential elections	52%-62%	[(Wikipedia)]
US company boards	73%	[17]
MKs in Israeli Parliament	53%	oknesset.org
Swiss referenda	30%-80%	[42]
Participatory budgeting	1%-14%	[34]

Table 1: Typical turnout in various voting contexts. Some of these refer to a specific instance or time period, see refs.

indifferent between voting and not voting at all. If we add even a small cost to the act of voting (say, the effort required to get to the ballot or to form an informed opinion) then voting essentially has negative expected utility. Therefore we should not be surprised that turnout is far from 100%, but rather that it is far from 0!

1.1 What's wrong with partial participation?

Distortion. First and foremost, abstention **distorts the outcome** even for simple Majority, as the running example above shows. A Borda winner[3] of the active votes may not be the Borda winner (and may even be the Borda loser) of the entire population; a participatory budgeting outcome that is guaranteed to be proportional may violate proportionality once inactive voters are also considered; a facility placed in the optimal location to serve the active voters may turn out to provide poor service in general; and so on.

In practice, it is estimated that more than 3% of board decisions would flip if the entire board had voted [17], and this is the domain with the *highest* turnout in Table 1. All of these problems become more severe if the tendency to participate is somehow correlated with the preferences themselves. E.g. it was shown that in participatory budgeting some age/race/socioeconomic groups vote in substantially higher numbers [34].

Engagement. A second motivation comes from the fields of political and organizational science, which often see voting not just as a tool to find a 'good outcome' but as a form of *civic participation*, and means of boosting **engagement** within one's community [1, 7]. In that respect, low turnout, even if perfectly representative, may indicate and even exacerbate low engagement and distrust. For example, Leininger and Heyne [42] claim that a higher turnout in Swiss referenda would probably not change their outcome, but increase their public legitimacy.

In the next three sections, we describe three sub-areas of AAMAS in which researchers can contribute in mitigating both the *distortion* and the *engagement* problems mentioned above. In each area, we briefly review some existing research directions, emphasizing possible limitations, and suggest alternative ideas.

2 MEASUREMENT AND MODELING

The Social Choice community has developed three types of tools that could be valuable in analyzing partial participation:

(1) Axioms and quantitative notions of social welfare, fairness, and proportionality, in order to compare different voting rules.
(2) Measures of distortion due to imperfect input (e.g. ordinal instead of cardinal).
(3) Models of strategic voting.

These tools are already being used to some extent in order to understand the impact of abstention.

Worst-case and Best-case abstention. Suppose that we know or estimate that in our running example the abstention rate is at most 8%. Since the gap between candidates is larger, we know that the better candidate A is still guaranteed to win, regardless of *who* fails to vote. That is, we have a *worst-case* guarantee on the outcome of the Majority rule, that depends on the margin.

Conversely, the *best-case* is when the active voters are sampled uniformly from the population. In this case both the margin and the *size of the population* matter. E.g. even if half of a population of a million would vote, the probability that B wins is negligible.

When considering other voting rules, both worst-case and best-case effects of abstention can be far less trivial, and may apply to any measure of welfare, fairness, etc. Yet to the best of our knowledge, these questions are rarely studied, except in specific settings. Some notable exceptions are [48], where authors bound both worst-case and best-case effect of abstention on protection from sybils; and [25], which empirically studies robustness of various Participatory budgeting rules to random abstentions. In contrast, there is ample literature on the (far more complicated!) issue of *sample complexity*, where voters are restricted to communicate only part of their preferences, see [45] and references therein. We believe that almost any result on positive properties of voting rules should be complemented with some analysis of their robustness to abstention.

Strategic abstention. Curiously, while game-theoretic considerations require more complex analysis than best- or worst-case models, they were studied from the outset. Calculus of Voting was suggested already back in the 1960's as an economic approach to explain the 'Paradox of Voting' [55], and was later extended to analyze strategic voting with multiple candidates [49, 52, 54]. A survey of economic and game theoretic approaches to voter turnout is in [20]. There are also newer and simpler models of strategic abstention such as 'lazy bias' [19] where a voter abstains whenever she is not pivotal.

However, in almost all of these models we get that from a large population, only a tiny fraction will vote in equilibrium (or sometimes a single voter [23]). Notably, this is true even for simple voting rules that hold the *participation axiom* (i.e. a voter is never better-off by abstaining). Thus proving existence of this axiom or its relaxations [2], would be at most a partial solution to our concerns.

Experiments. The literature is abundant with experiments of strategic abstention. However these experiments (much like experiments of strategic voting in general) typically focus on *aggregate voting behavior*, i.e., overall turnout, and in particular on reproducing or refuting certain notions of equilibria in specific voting games [4, 41]. Studies that track individual behavior (like [5]) are more helpful in that respect, but less common.

2.1 Future focus

We argue that theoretical and experimental work should focus on:

- Testing behavior on the *individual level*, in order to account for heterogeneity and different types of behaviors;

[3]Borda is a voting method where each voter ranks candidates, and higher-ranked candidates get proportionally higher score.

- Trying to separate what voters *believe* (in particular about chances of being pivotal) from their actual decisions. This separation allows us to maintain the self-interest assumption and apply game-theoretic analysis, while still considering a broad set of behaviors due to variability in perception.

These directions are already used in the study of strategic compromise in voting, both in theory [46, 51, 53] and in experiments [26, 56], and can be extended to abstention.

For example, in a preliminary work with Ganesh Ghalme [31], we showed that such biased perception of pivotality may lead to substantial turnout in equilibrium, and characterized conditions under which Condorcet Jury Theorem is maintained in a population with heterogeneous voting costs.

Similar questions can be studied for other properties and under different abstention models. This is crucial, as change of a few percents in voter turnout is likely to have more impact than strategic behavior that happens in the fringes.

3 MECHANISM DESIGN

If our focus is on distortion due to non-representative set of voters, then social choice literature offers two very useful tools, although they are often discussed in different contexts.

Sampling. The premise of most abstention models mentioned in the previous section, is that voters abstain because they are not *pivotal enough* to justify the effort. Then, as in our running example, those who happen to be more motivated or face lower costs, may support the less popular alternatives. However, suppose we only sample 3 voters and ask them to vote. Since each of them has a substantial influence on the outcome, it is highly unlikely they will choose to abstain. The chances of the wrong candidate B getting elected are about 35%, but drop sharply as we increase the size of the sample. This idea is highly related to the ancient practice of *sortition*, i.e. randomly selecting representatives from a population [10, 35].

Indeed, several recent works showed that sortition with a reasonable sample size can guarantee various notions of welfare, proportionality, and representation [9, 22, 47], and some argue for its wide adoption [12]. In another preliminary work with Ghalme that is built on the bounded-rational models mentioned above [30], we suggest to use a *second round* in case the first round is nearly tied. The idea is that voters will conclude they are more pivotal than previously thought and then vote in higher numbers in the second round. In case the margin in the first round is sufficiently large, we can call the results as they are very unlikely to favor the minority candidate.

Delegation. Direct delegation of voting rights has been suggested and analyzed explicitly as a solution to low participation [13, 33], and direct or transitive delegation can be shown in some circumstances to even beat full participation on various criteria [3, 32, 44]. This is true especially when voters avoid active participation not due to high cost, but since they are unsure about the right decision, and can now delegate to others who are more informed.

3.1 The problem with sampling and delegation

While both tools have been shown to be effective in *improving the outcome*, thereby tackling our *distortion* problem, they are inherently problematic when considering our second motivation of *engagement*. This is particularly obvious with sampling, as evident from the *Franchise* quote in the opening: if only few people are ultimately active in the voting process (regardless of how they are being selected), then the rest of us don't have to wrap our head around health and economics, transportation and social security dilemmas. We might as well go and "have supper".

With delegation, the engagement problem may seem less acute. After all, voters can still vote, if they want to. So ideally, only some of those who would otherwise abstain would delegate. But there is some evidence that reality is worse than that.

Theoretically, some delegation schemes might create a strict dis-incentive to accept delegated votes from others [39]. Further, experiments show that delegation results in poor outcomes, not because of power concentration, but simply because too many voters choose to delegate. This is true even compared to another experiment where voters could have abstained [8].

Moreover, it has been argued (in other contexts) that delegation allows for reduced accountability [28]. This does not mean that delegation is a-priori bad, but that it *might* exacerbate the very problem it is trying to solve, and should be carefully applied.

3.2 Future focus

In light of the above discussion, we again argue in favor of two research directions:

A Principal-Agent approach. We would like to incentivize voters to *invest effort* in getting more information to guide their vote, rather than abstaining or guessing.

Mechanisms for incentivizing effort have been studied in other areas of mechanism design [29, 43], but these techniques may not be suitable for voting (e.g. paying the players).

Focus on the candidates. Instead of trying to incentivize voters directly, we can focus on those who already have a strong incentive to recruit voters. That is, the candidates or lobbyists.

Delegation schemes can play an important role in both directions, and we next outline one such proposal: The general idea is to allow (direct) delegation, but where more followers have *decreasing marginal impact* on the proxy's weight. That is, an active voter with 25 delegated votes has less power than, say, two proxies with 12 votes each. In particular, under the *Penrose Proxy Voting* (PPV) rule,[4] the weight of each proxy is set to the square root of its number of followers. One benefit of this idea is to keep the power of proxies in check, without a formal cap.[5] However in our context, our goal is to allow voters to choose their *level of participation*: Delegating to a large and popular proxy would typically require very little effort, but also carries little benefit to the voter. If she wants to influence more, she could 'do some research' and find a smaller proxy, becoming more informed in the process. Then of course, our voter can decide she is sufficiently informed on a topic to vote

[4] A reference to Lyonel Penrose' square-root law for proportional representation [38].
[5] Indeed, a similar idea has been very recently proposed with this goal in mind [57].

Figure 1: The cover art of *Franchise* shows Norman Muller connecting to Multivac. The ending sentence says: "In this imperfect world, the sovereign citizens of the first and greatest Electronic Democracy had, through Norman Muller exercised once again its free, untrammeled franchise."

directly for maximal influence, or maybe even to try and attract votes from others.

We can also consider this process from the perspective of the proxies/lobbyists, who need to make quadratic effort in order to linearly increase their voting weight on a particular binary vote. This may sound familiar as this is exactly the logic behind *quadratic voting* [40] (where our lobbyists are the strategic voters). However the dynamics would be different as there is no fixed set of players and no clear distinction between proxies and followers.

4 ARTIFICIAL AGENTS

Learning and representation of preferences has received much attention, including at AAMAS, and even has dedicated workshops such as M-PREF. In the context of prevalent abstention, learning preferences may look like a silver bullet. Who needs real voters? once our laptop/phone/smart watch learn our preferences, they could vote on our behalf!

If this sounds far-fetched, recall that we already have review assignment systems that do not require bidding, and instead infer the reviewer's preferences from her list of publications [11].

The extreme potential of such systems was demonstrated in Asimov's *Franchise* from 1955, where the all-powerful computer Multivac uses its vast understanding of human preferences and desires to select a single representative voter on behalf of the entire US population.

The recent and sudden introduction of LLMs to our lives, brought us some steps closer to this reality. In [15], Conitzer et al. critically discuss various ways LLMs may learn from human preferences. Some of the models they review, like *Reinforcement Learning from Human Feedback* deal with using human preferences to guide the decisions of AI systems, which is the opposite direction of our intention. Other methods such as *Constitutional AI* and the authors' novel proposal of *Simulated Collective Decisions* suggest generating 'virtual votes' on various issues by a system that has already learned collective preferences to a sufficient degree.

In the half-sarcastic closing paragraph of *Franchise* (see Fig. 1 caption), Asimov emphasizes both the utopian and dystopian aspects of the end point of such progress. This tension mirrors how our two motivations for solving the abstention problem are conflicting: delegating our votes to a computer may, in some circumstances, improve the outcome and mitigate the distortion that results from low and biased participation. But this may come at a cost of lowering voters' engagement and accountability even further. Indeed, this is a very similar dilemma to the one we highlighted in Section 3, regarding delegating votes to other people.

Delegating-to-the-machine might be suitable or even necessary for tuning the parameters of a recommendation system, or a smart vehicle safety module [50]. But our conclusion is that we would (and should!) be wary of using it to decide about budget allocation, educational policies, presidential elections, and any scenario where we see value in participation itself.

4.1 Future focus

Our main message regarding artificial agents is not for a particular technique or result, but rather about the way they should be applied.

We mentioned above systems like TMPS that *enable* conferences to skip explicit bidding. However, some conferences use such systems to *facilitate* bidding. Perhaps this is done only due to lack of sufficient input, or to improve the matching quality, but it outlines the correct way (in our opinion) to use learned preferences, at least in sensitive situations.

There is already a large economy of 'Voting Advice Applications' (VAA) that started from simple booklets handed to school pupils [18, 24]. The diverse expertise in the AAMAS community could be put to use to take VAAs to the next level.

The success of such VAAs tightly depends on the issues and tools we mentioned in Section 3. For example, we could obtain the benefits of delegation without the reduced accountability, by explicitly exposing voters to the preferences of their peers.

One possible source of inspiration for the right information flow could be the ways in which heterogeneous flocks of animals reach consensus in the work of Lain Couzin [16].

5 CONCLUSION

Low turnout is the unavoidable shadow of the blessed ability to vote anywhere and on anything. It means that even when taking democratic decisions, in practice it may be a small minority that dictates the outcome for everyone.

Much like strategic behavior, the amount and identity of those who participate may have a huge effect on the outcome, making any naïve analysis of voting rule guarantees utterly irrelevant.

But despite the fact that abstention has far greater impact than strategic voting, and should in principle be easier to model and to deal with, it is far less studied both generally and within the AAMAS research community. Moreover, in this article I warn that some of the techniques that might seem promising in curbing the tyranny of the minority—in particular delegating our voting rights to other voters and/or computers—may be a double edge sword.

I therefore hope that the community will take this 'call to arms' seriously, in making sure the tools we develop help voters understand their own preferences better, inform them and empower them in making their own decisions, rather than voting on their behalf. Indeed, this may leave us with somewhat less time for personal activities (like supper), but with more accountability for decisions that affect us all.

REFERENCES

[1] Ben Abramowitz and Nicholas Mattei. 2024. Flexible representative democracy: An introduction with binary issues. *Social Choice and Welfare* (2024), 1–46.

[2] Stéphane Airiau, Haris Aziz, Ioannis Caragiannis, Justin Kruger, Jérôme Lang, and Dominik Peters. 2019. Portioning Using Ordinal Preferences: Fairness and Efficiency. In *IJCAI'19)*. 11–17.

[3] Georgios Amanatidis, Aris Filos-Ratsikas, Philip Lazos, Evangelos Markakis, and Georgios Papasotiropoulos. 2024. On the Potential and Limitations of Proxy Voting: Delegation with Incomplete Votes. In *AAMAS'24*. 49–57.

[4] Marco Battaglini, Rebecca B Morton, and Thomas R Palfrey. 2008. Information aggregation and strategic abstention in large laboratory elections. *American Economic Review* 98, 2 (2008), 194–200.

[5] André Blais, Jean-Benoit Pilet, Karine Van der Straeten, Jean-François Laslier, and Maxime Héroux-Legault. 2014. To vote or to abstain? An experimental test of rational calculus in first past the post and PR elections. *Electoral studies* 36 (2014), 39–50.

[6] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. 2016. *Handbook of computational social choice*. Cambridge University Press.

[7] Alon Brav, Matthew Cain, and Jonathon Zytnick. 2022. Retail shareholder participation in the proxy process: Monitoring, engagement, and voting. *Journal of Financial Economics* 144, 2 (2022), 492–522.

[8] Joseph Campbell, Alessandra Casella, Lucas de Lara, Victoria R Mooers, and Dilip Ravindran. 2022. *Liquid democracy. Two experiments on delegation in voting*. Technical Report. National Bureau of Economic Research.

[9] Ioannis Caragiannis, Evi Micha, and Jannik Peters. 2024. Can a Few Decide for Many? The Metric Distortion of Sortition. In *ICML'24*.

[10] Lyn Carson, Professorial Fellow Lyn Carson, and Brian Martin. 1999. *Random selection in politics*. Greenwood Publishing Group.

[11] Laurent Charlin and Richard Zemel. 2013. The Toronto paper matching system: an automated paper-reviewer assignment system. (2013).

[12] David Chaum. 2016. Random-sample voting. *White Paper* (2016).

[13] Gal Cohensius, Shie Mannor, Reshef Meir, Eli Meirom, and Ariel Orda. 2017. Proxy voting for better outcomes. In *AAMAS'17*. 858–866.

[14] Marie J Condorcet et al. 1785. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Vol. 252. American Mathematical Soc.

[15] Vincent Conitzer, Rachel Freedman, Jobst Heitzig, Wesley H Holliday, Bob M Jacobs, Nathan Lambert, Milan Mossé, Eric Pacuit, Stuart Russell, Hailey Schoelkopf, et al. 2024. Position: Social Choice Should Guide AI Alignment in Dealing with Diverse Human Feedback. In *Forty-first International Conference on Machine Learning*.

[16] Iain D Couzin, Christos C Ioannou, Güven Demirel, Thilo Gross, Colin J Torney, Andrew Hartnett, Larissa Conradt, Simon A Levin, and Naomi E Leonard. 2011. Uninformed individuals promote democratic consensus in animal groups. *science* 334, 6062 (2011), 1578–1580.

[17] Dragana Cvijanovic, Moqi Groen-Xu, and Konstantinos E Zachariadis. 2020. Free-riders and underdogs: Participation in corporate voting. *ECGI-Finance Working Paper (649/2020)* (2020).

[18] Jochum De Graaf. 2010. The irresistible rise of Stemwijzer. *Voting advice applications in Europe: the state of the art* (2010), 35–46.

[19] Yvo Desmedt and Edith Elkind. 2010. Equilibria of plurality voting with abstentions. In *EC'10*. ACM, 347–356.

[20] A Dhillon and Susana Peralta. 2002. Economic theories of voter turnout. *The Economic Journal* 112, 480 (2002), F332.

[21] Anthony Downs. 1957. An economic theory of political action in a democracy. *Journal of political economy* 65, 2 (1957), 135–150.

[22] Soroush Ebadian, Gregory Kehne, Evi Micha, Ariel D Procaccia, and Nisarg Shah. 2022. Is sortition both representative and fair? *Advances in Neural Information Processing Systems* 35 (2022), 3431–3443.

[23] Edith Elkind, Evangelos Markakis, Svetlana Obraztsova, and Piotr Skowron. 2015. Equilibria of plurality voting: Lazy and truth-biased voters. In *SAGT'15*. Springer, 110–122.

[24] Zsolt Enyedi. 2016. The influence of voting advice applications on preferences, loyalties and turnout: An experimental study. *Political Studies* 64, 4 (2016), 1000–1015.

[25] Roy Fairstein, Gerdus Benadè, and Kobi Gal. 2023. Participatory budgeting designs for the real world. In *AAAI'23*, Vol. 37. 5633–5640.

[26] Roy Fairstein, Adam Lauz, Reshef Meir, and Kobi Gal. 2019. Modeling People's Voting Behavior with Poll Information. In *AAMAS'19*. 1422–1430.

[27] Dan S Felsenthal and Moshé Machover. 2001. Models and reality: the curious case of the absent abstention. In *Power indices and coalition formation*. Springer, 87–103.

[28] Justin Fox and Stuart V Jordan. 2011. Delegation and accountability. *The Journal of Politics* 73, 3 (2011), 831–844.

[29] Peter Frazier, David Kempe, Jon Kleinberg, and Robert Kleinberg. 2014. Incentivizing exploration. In *EC'14*. 5–22.

[30] Ganesh Ghalme and Reshef Meir. 2024. Boosting Voter Participation: A Case for Two-Round Voting. *In preparation* (2024).

[31] Ganesh Ghalme and Reshef Meir. 2024. Condorcet's Jury Theorem with Abstention. *arXiv preprint arXiv:2408.00317* (2024).

[32] Paul Gölz, Anson Kahng, Simon Mackenzie, and Ariel D Procaccia. 2021. The fluid mechanics of liquid democracy. *ACM Transactions on Economics and Computation* 9, 4 (2021), 1–39.

[33] James Green-Armytage. 2015. Direct voting and proxy voting. *Constitutional Political Economy* 26 (2015), 190–220.

[34] Carolin Hagelskamp, Chloe Rinehart, Rebecca Silliman, and David Schleifer. 2016. Public Spending, by the People, Participatory Budgeting in the United States and Canada in 2014–15. Public Agenda report.

[35] James Wycliffe Headlam. 1891. *Election by Lot at Athens*. The University Press.

[36] Frank Key Howard. 1863. *Fourteen Months in American Bastiles*. Kelly, Hedian & Piet.

[37] Thomas Laurence Kington-Oliphant. 1862. *History of Frederick the Second, Emperor of the Romans*. Macmillan.

[38] Werner Kirsch. 2013. On Penrose's square-root law and beyond. *Power, Voting, and Voting Power: 30 Years After* (2013), 365–387.

[39] Grammateia Kotsialou and Luke Riley. 2020. Incentivising Participation in Liquid Democracy with Breadth-First Delegation. In *AAMAS'20*. 638–644.

[40] Steven P Lalley and E Glen Weyl. 2018. Quadratic voting: How mechanism design can radicalize democracy. In *AEA Papers and Proceedings*, Vol. 108. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 33–37.

[41] David Dreyer Lassen. 2005. The effect of information on voter turnout: Evidence from a natural experiment. *American Journal of political science* 49, 1 (2005), 103–118.

[42] Arndt Leininger and Lea Heyne. 2017. How representative are referendums? Evidence from 20 years of Swiss referendums. *Electoral Studies* 48 (2017), 84–97.

[43] Yang Liu and Yiling Chen. 2016. Learning to incentivize: eliciting effort via output agreement. In *IJCAI'16*. 3782–3788.

[44] Malik Magdon-Ismail and Lirong Xia. 2018. A mathematical model for optimal decisions in a representative democracy. *Advances in Neural Information Processing Systems* 31 (2018).

[45] Debmalya Mandal, Ariel D Procaccia, Nisarg Shah, and David Woodruff. 2019. Efficient and thrifty voting by any means necessary. *Advances in Neural Information Processing Systems* 32 (2019).

[46] Reshef Meir, Omer Lev, and Jeffrey S Rosenschein. 2014. A local-dominance theory of voting equilibria. In *EC'14*. 313–330.

[47] Reshef Meir, Fedor Sandomirskiy, and Moshe Tennenholtz. 2021. Representative committees of peers. *Journal of Artificial Intelligence Research* 71 (2021), 401–429.

[48] Reshef Meir, Gal Shahaf, Ehud Shapiro, and Nimrod Talmon. 2022. Sybil-resilient social choice with partial participation. In *EUMAS'22*.

[49] Roger B Myerson and Robert J Weber. 1993. A theory of voting equilibria. *American Political Science Review* 87, 1 (1993), 102–114.

[50] Ritesh Noothigattu, Snehalkumar Gaikwad, Edmond Awad, Sohan Dsouza, Iyad Rahwan, Pradeep Ravikumar, and Ariel Procaccia. 2018. A voting-based system for ethical decision making. In *AAAI'18*.

[51] Martin J Osborne and Ariel Rubinstein. 2003. Sampling equilibrium, with an application to strategic voting. *Games and Economic Behavior* 45, 2 (2003), 434–441.

[52] Thomas R Palfrey. 1988. A mathematical proof of Duverger's law. (1988).

[53] Annemieke Reijngoud and Ulle Endriss. 2012. Voter response to iterated poll information. In *AAMAS'12*. 635–644.

[54] William H Riker. 1982. The two-party system and Duverger's law: an essay on the history of political science. *American political science review* 76, 4 (1982), 753–766.

[55] William H Riker and Peter C Ordeshook. 1968. A Theory of the Calculus of Voting. *American political science review* 62, 1 (1968), 25–42.

[56] Maor Tal, Reshef Meir, and Ya'akov Kobi Gal. 2015. A study of human behavior in online voting. In *AAMAS'15*. 665–673.

[57] Wenpin Tang and David D Yao. 2024. Polynomial voting rules. *Mathematics of Operations Research* (2024).

[58] Kristjan Vassil and Till Weber. 2011. A bottleneck model of e-voting: Why technology fails to boost turnout. *New media & society* 13, 8 (2011), 1336–1354.