Towards a More Burkean Approach to Computational Social Choice

Omer Lev

Ben-Gurion University of the Negev
omerlev@bgu.ac.il

Abstract

In the last few years, a lot of the activity of the computational social choice community has focused on novel mechanisms for reaching decisions by large groups of people. While this research makes meaningful scientific contributions, many of these mechanisms are not quite useful in realistic decision-making settings. Moreover, their radicalism ignores the centuries-old experience we have with large-scale human decision-making, and what it teaches us about what works. We believe it is important the community engage with mechanisms which are widely-used in the real world, as they may hold a key to a deeper understanding of how people reach decisions and the way that helps them do that productively. Moreover, letting the community bring its analysis and understanding to these will allow for algorithmic suggestions that have some chance of being implemented (and, thus, can contribute to the public debate on these topics). In particular, we highlight the relatively less-investigated role of parties and grouping of voters and candidates, and the role of executive capacity in analyzing decision-making structures.

> *By preserving the method of nature in the conduct of the state, in what we improve we are never wholly new; in what we retain we are never wholly obsolete.*
>
> Edmund Burke, *Reflections on the Revolution in France*, 1790

In the past several years, prominent computational social choice researchers working on voting mechanisms began exploring several novel topics, expanding the COMSOC realm far and wide. Some of these novel topics have to do with trying to understand existing mechanisms, which are widely used throughout the world and in various political systems (e.g., multi-winner elections (Lackner and Skowron 2023; Aziz, Lang, and Monnot 2016; Aziz et al. 2017; Elkind et al. 2017; Faliszewski et al. 2017)). Other directions investigate less common systems, some of which seek to supplant existing decision-making systems. Among the most commonly researched are:

Liquid democracy / proxy voting A method in which each voter can transfer their vote to another voter, i.e., let the other voter vote on their behalf. This voter can, in turn, transfer their own vote (along with the votes transferred to them by others) onwards to other voters[1]. This results in an election in which only a subset of voters participate, each with a different weight (depending on the number of voters who have – directly or through others – given each agent their votes). There have been several papers discussing various aspects and variations of the model (Brill 2019; Brill and Talmon 2018; Zhang and Grossi 2021), and while Kahng, Mackenzie, and Procaccia (2021) (later strengthened by Caragiannis and Micha (2019)) raised the issue of too high concentration of power, several responses have been authored to tackle these issues, among other topics (Gölz et al. 2021; Bloembergen, Grossi, and Lackner 2019; Abramowitz and Mattei 2019; Colley, Grandi, and Novaro 2021; Becker et al. 2021).

Liquid democracy came to the fore when it was adopted by the Pirate Party in Germany, and later, in various variants, in a few other organizations.

Blockchain-related As part of a general vision (Shapiro and Talmon 2022), some researchers attempt to move decision-making to a blockchain-based mechanism (Poupko, Shapiro, and Talmon 2022), exploring various coalitional and consensual mechanisms to make it work (Abramowitz et al. 2021; Elkind et al. 2021; Meir et al. 2022).

Blockchain-like mechanisms were discussed for several decades, but their current meaning began with Bitcoin in 2008. Smart contracts (contracts that – as long as they involve assets kept on the blockchain – can be enforced through the blockchain) in their current form were popularized by Ethereum since 2015. For decision-making purposes, the idea of DAOs (Decentralized Autonomous Organization) has been suggested, but so far there has been very little practical use.

Sortition The process of reaching a decision by choosing a group of people by lottery, and letting them decide. In the last few years there has been research – including very high-profile (Flanigan et al. 2021) – on how to select agents while representing various predetermined groups appropriately (Flanigan et al. 2020; Flanigan, Kehne,

[1] In proxy voting (Cohensius et al. 2017), the transfer is a step only, so longer chains cannot exist.

and Procaccia 2021), and also noting representability issues (Ebadian et al. 2022).

Sortition of sorts was used in the past for some judicial selections (such as juries) and for some arcane election processes (the best known of which is, probably, part of the process of appointing the Doge of Venice (Mowbray and Gollmann 2007)). A comprehensive overview of its use can be found in a recent OECD report (OECD 2020). Its most famous successes were the processes leading to the Irish referendums on abortion policy and gay marriage, though several attempts are made throughout Europe to establish such forums as "citizens' assemblies". Some of the mainstream organizations pushing for sortition (e.g., the Sortition Foundation) explicitly refer to existing institutions as not real democracy and disenfranchising.

It should be noted that in some cases, the work on these topics is not done simply as a cold analysis of theoretical constructs, but some of the researchers involved are proponents of these systems, seeking to expand their use by people around the world (Hennig 2017; Shapiro and Talmon 2022; Procaccia 2022, 2019; Degani 2021), and claiming they should be implemented to improve democracy. Of course, there is nothing problematic about researchers having personal views and pursuing research into topics related to them. However, as a wider vision for the community, the interest in these particular mechanisms as part of computational social choice (rather than abstract voting settings, which were a catalyst for much COMSOC research), over other endeavors, indicates an ideological choice: opting for analyzing mechanism designs which are relevant for quite narrow circumstances, or wishing to put aside millennia of human experience in large-scale decision-making systems for a more revolutionary approach (Fukuyama 2011, 2014).

We wish to strongly emphasize that this is not a criticism of researchers exploring these topics, nor of their work. The research itself is amazingly good and deserving, and the results are interesting and meaningful. However, it seems that many of the community "hot topics" seem to be around very extreme changes to how people reach decisions, and we wish to propose different directions, which focus on understanding key components of common large-scale decision systems which have received relatively little attention.

A Brief Critique

> *Very plausible schemes, with very pleasing commencements, have often shameful and lamentable conclusions... The means taught by experience may be better suited to political ends than those contrived in the original project.*
>
> Edmund Burke, *Reflections on the Revolution in France*, 1790

It is not the goal of this paper to focus on the flaws of the above-mentioned systems, but to advocate for a different direction. But we would be remiss not to highlight, briefly, few reasons why we believe these directions – despite impressive results that display the brilliance and theoretical insights of the researchers involved – are, perhaps, not the right direction for the computational social choice community as a whole. Since the suggested systems are mainly relevant for large human groups (and not artificial agents), we focus here on issues relating to people's interaction with these mechanisms.

Principled objections to novel, untried systems have been expressed by many thinkers, from Edmund Burke (Burke 1790) to Francis Fukuyama (Fukuyama 1992, 2014). These have included fear of unintended negative consequences as well as the understanding that existing mechanisms have developed for a reason, and their existing shape has to do with the various requirements they serve and properties they have, which may not all be explicitly stated and understood. In addition to this, some of the issues posed by these systems are more specific to them:

Liquid democracy and blockchain Any direct democracy model (including liquid democracy and blockchain-based) needs to deal, first and foremost, with the inability of the public to skillfully analyze legislation or to access expertise (due to work, other interests, etc.). Liquid democracy deals with this by having voters give their vote to other, better informed, agents (possibly different agents for different issues). Beyond the intuitive understanding that this will concentrate power with a rather small group of activists (and in the blockchain variants – a tech-oriented elite[2]), we know that more knowledgable people tend to be more opinionated, and – even worse – more partisan (Hannon 2022). This means that it is very easy to formulate a theorem showing that even if voters give their vote to marginally more knowledgable voters than themselves, instead of relying on "wisdom of the crowd", this will lead, ultimately, to a wider societal divide than exists in practice. Moreover, one of the key benefits of representative democracy is the ability to deliberate and to establish give-and-take (one side gives up on one issue, to get its way on another) and to reconcile impossible requirements (low taxes with high expenditures, for example, both of which are highly popular on their own), or judgement aggregation issues (Blum and Zuber 2016). These would become quite impossible to reach without any representatives, though even if there were few enough non-delegating voters to make deliberations practical, not only would their more extreme partisan position make this far harder, a minority of voters – those who feel more strongly on the issue "given up" that the one "taken" – can scupper any such deal by changing their delegation.

Sortition While sortition was widely used in the past to avoid corruption (one cannot bribe in advance a decision-maker if they are selected at random), its most successful recent use cases – in Ireland – were in allowing divisive issues with dividing lines orthogonal to the usual political system's fault-lines to be settled. This is not the space to discuss all issues with sortition (for some more of these, see Lee,

[2]In general, we will not delve here into the issues with various variants of the blockchain-infused "techno-utopia" suggested in various talks and seminars. The issues detail here are mostly shared with liquid democracy, and often they are tied together.

McQuarrie, and Walker (2015)), but we mention two that are relevant if – as its proponents advocate – it is used as a permanent fixture of democratic decision-making, beyond "community advisory" roles. First is the strong influence of the sortition panel organizer, in a classic control problem issue. This is not manifested in the selection process itself (researchers have shown very impressive algorithms for that), but in the role of providing the information to the panel on the discussed issues, thus being able to tilt the discussions' direction (the organizer's role in providing information is clear both in instructions (OECD 2020) and practice, as the lottery selects people with, most probably, very little familiarity or expertise in the issue at hand). In addition, the organizer decides which population properties are relevant to maintain in the randomized selection (e.g., gender, ethnicity, age, etc.). The assumption that the planner knows the relevant population fault lines is not trivial. Perhaps on some issues the relevant divisions are about technology proficiency; public/private employment; religious observance; or television watching habits.

Second, and perhaps more important for actual democratic systems (as recently suggested in the New-York Times (Grant 2023)), graft and corruption are much easier and more common when the sanction of not being elected again is removed (Tsur 2022; Ferraz and Finan 2011; Laurent 2021; Raveh and Tsur 2021). Thus, such panels allow a far wider opening for graft than elected democracy, particularly if the panel selected by sortition serve for an extended period of time, as is being implemented in some locations[3]. The lack of any need for reelection, nor any process of electoral accountability, means each participant may do whatever they want to further their own personal desires.

The (Almost) Undiscovered Country

> *Old establishments... are not often constructed after any theory; theories are rather drawn from them.*
>
> Edmund Burke, *Reflections on the Revolution in France*, 1790

When discussing voting mechanisms that can be useful in large human societies, it makes intuitive sense to examine how people have been reaching decisions in the past several centuries in which large-scale human societies have been doing so. In particular, focusing on what seem to be very common features of human voting systems has particular appeal, as those features might be crucial in ways we may not yet fully understand. By understanding them better, researchers can suggest optimizations and tweaks to the existing system, and attempt to nudge it towards what seems to be a more useful direction. While it is, possibly, less ground-breaking than a wholesale replacement of the decision making system, it may have better chance of being implemented, or at least attempted, on a very large scale, with much fewer unintended outcomes.

There are many such directions, from the need for voting rule *simplicity* (explaining, perhaps, the prevalence of plurality around the world, despite its many flaws) and some form of communication complexity-like definition for aggregation of votes (as used in opposition tally of Zimbabwe 2008 votes (IHT 2008)) to robustness from interference or external influence. In order to be concise, we will focus here on one particular such feature, which has been relatively little explored in the computational social choice literature – parties and sub-groups of voters and candidates. We will examine the model, what has been investigated so far, and suggest further research paths to expand the existing research into this topic. We stress that there are many further such topics, and many of them (like parties) apply not only to nation-sized elections, but to much smaller groups that can benefit from the insights that the well-known existing systems provide.

Parties: The Work So Far

> *In all questions, whether concerning foreign or domestic affairs, the whole generally turns more upon some party-matter than upon the nature of the thing itself.*
>
> Edmund Burke, *A Vindication of Natural Society*, 1756

Ad-hoc parties, and those based on a particular leader or a shared ancestor, are known since antiquity. But *ideological* parties form much later, and begin to take their modern shape with the "Tory" and "Whig" parties in the aftermath of the Glorious Revolution in 1689 Britain (Harris 2006), and then, more clearly, following the American and the French revolutions in the late 18th century. Despite the explicit desire of the American founding fathers to avoid political parties (Hamilton, Madison, and Jay 1787; Wilentz 2005), they were unable to stop this development, and some of them were deeply involved in the creation of the modern party framework. Today, of course, many countries incorporate political parties as a fundamental building block of their political systems (e.g., any country with an electoral threshold, such as Turkey, Germany, or Israel).

In computer science, and specifically in computational social choice, dealing with parties – basically, sub-groups of voters and candidates – has been quite limited, but it is already clear that their existence means various existing results for non-partisan settings need to be examined. Parties affect two main settings:

Voter aggregation Parties let the same "candidate", in a sense, compete in several elections, by having several different elections with a candidate affiliated with the party. Thus, the separate elections (each its own "district") can be aggregated together, so the winning party is, for example, the one that won a plurality of districts. This is termed *district-elections*. This district structure changes the way votes are aggregated, so that, for example, if each district uses the plurality rule, if there are m parties, some party A can win the most votes overall, and in particular, $\theta(m^2)$ more votes than party B, and yet party B will still win the plurality of districts (Bachrach et al. 2016).

The district structure opens up a variety of research prob-

[3] https://congress.crowd.law/files/Belgian_Sortition_Models.pdf

lems, some of which have been explored in the social-sciences and legal circles, but much less in computer science. What was done in a computational angle includes adapting existing models to districts (e.g., iterative voting and variants thereof (Lev and Lewenberg 2019; Bervoets and Merlin 2012; van Bevern et al. 2015)), but a more popular direction, research wise, deals with the division of districts when each district is a geographical unit, as in political systems in which each elected candidate represents a geographical unit. Dividing districts according to specific desiderata is an NP-hard problem (Dyer and Frieze 1985; Yang 2014; Apollonio et al. 2009; Lewenberg, Lev, and Rosenschein 2017; Cohen-Zemach, Lewenberg, and Rosenschein 2018), but the interesting game-theoretical questions have to do with dividing it specifically with partisan (or fairness) requirements, termed *gerrymandering* when done for partisan gain. Several papers – mostly empirical (though not all (Pegden, Procaccia, and Yu 2017)) – delved into this setting, showing how the geographical spread of partisan preferences affects the possibility of gerrymandering, and highlighting, in particular, urban and rural differences (Borodin et al. 2018, 2022).

Candidate elimination Another function of parties has to do with their use in selecting candidates to run. That is, instead of multiple candidates with similar ideology, some process happens, which leads to the selection of only a few candidates to run in the general election. The most famous such intra-party election happens in the US' *primaries* process, in which the main US parties select their presidential nominee by running elections in each state, aggregating those state results, and reaching an outcome. Borodin et al. (2019) showed that adding such an election stage can, at worse, increase the *distortion*[4] of an election by $\mathcal{O}(1)$, but it can reduce it by $\mathcal{O}(n)$ (n being the number of voters), when parties are separable in an ideological space (i.e., each party occupies some non-overlapping part of the ideological space). A few papers have begun approaching the strategic aspects induced by this model (Karpov, Lev, and Obraztsova 2022; Harrenstein et al. 2021).

It is important to note that while research on parties is inspired by the political world, it has applicability throughout the whole range of computational social choice settings. It seems that whenever decisions need to be reached in a very large group, sub-groups arise organically or are created according to some criterion (bureaucratically, geographically, etc.). In a university, many departments (which are bureaucratically created sub-groups) interview potential candidates for new positions, passing the selected candidates to the faculty level, and then the university level, much like in a party system. Similarly, corporate decisions are often made by input of various company divisions, each with its own decision process. But we are not limited to human decision mechanisms alone. One can consider a sensor array, where each sensor induces a ranking of what phenomenon is it observing, and each sensor type is its own "party", aggregating the sensors of a particular type and reporting the outcome to the array itself, which incrporates various sensor types, and uses all of them to reach an ultimate ranking of what is being observed.

[4]A quality metric for voting rules, introduced in Procaccia and Rosenschein (2006), and refined by Anshelevich et al. (Anshelevich and Postl 2017; Anshelevich, Bhardwaj, and Postl 2015).

Parties: The Next Generation

At once to preserve and to reform...

Edmund Burke, *Reflections on the Revolution in France*, 1790

The existence of sub-groups within an election changes it profoundly in ways we have only started to understand. The strategic concerns have only begun to be explored – even in the most basic model, of a two-stage process which includes an inter-party election followed by a general election, we still do not understand how voters will vote under various informational settings. We also have very little inkling of how candidates behave if they cannot alter their ideological position freely, but are limited in their movement (e.g., can change their location in the ideological metric space from the intra-party election to the general election by ε at most). We believe understanding this is crucial to comprehending this very fundamental component of human (and other) decision-making. Furthermore, in district-elections we do not yet understand the full impact of the districts on various different election rules, and various district aggregation methods beyond plurality have barely been investigated at all. Thus, our understanding of the effects of having parties is still quite limited.

However, looking ahead, we wish to discuss a few exciting – yet unexplored – research directions, which will significantly expand our understanding of decision-making processes, allowing us to improve them without requiring a wholesale remake of the voting process.

Axiomatic Approach & Mechanism Design In both district-elections and intra-party elections, we lack a normative framework that will allow us to discuss and compare various mechanisms. There are a few metrics (e.g., distortion; or comparing how close is the district-election result, in a particular setting, to an outcome that is close to the proportional outcome if the whole country was a single district), and cooperative game theory has some contributions (such as power indices), but generally, we do not yet have a rich enough way to discuss relevant partisan mechanisms. What is needed is to establish desiderata – a set of properties which may be desired in a district-voting mechanism or an intra-party voting system (for example, some ε distance from proportionality, or compactness constraints[5]). Once this set of axioms, or properties, is established, research can examine whether there are mechanisms that fulfill some subsets of these properties, and, more interestingly, which properties conflict with one another.

[5]Some states in the US encourage compact districts, though there are multiple different mathematical definitions for that, and no legal definition.

There are some obvious impossibilities – it is easy to see that dividing a geographic area into single-member districts, each using plurality, cannot ensure an outcome that is proportional to the population vote. But it remains an open problem what can we guarantee? Are there bounds that can be reached? In intra-party votes, we know even less. We hypothesize that a rich vein of research can be reached by dividing each party's voters into *pragmatists* (that is, voters which put a high premium on selecting a candidate that will win in the general election) vs. voters which are *ideologues* (i.e., voters which prefer a candidate close to their ideological position over other concerns). In many political settings we have such a simple division of each party's electorate (e.g. in the British party system, Members of Parliament are generally interested in candidates that are popular in the wider population, which helps their own reelection campaigns. These may not be the same candidates preferred by the party's members, which are more ideological).

The ability to propose a range of mechanisms, each with its own properties (e.g., to reach a particular influence of pragmatic/ideological party sides; or particular proportional influence on district outcomes), to a mechanism designer, is the sign of ultimate success. This will require understanding some of the basic trade-offs inherent in the system, of which we still have very limited understanding today.

Executive Capacity The notion of executive capacity has to do with a winner's ability to actually implement their agenda, and, in a sense, it is a form of winner robustness. While cooperative game theory (Chalkiadakis, Elkind, and Wooldridge 2011), hedonic games (Aziz and Savani 2016), and multi-winner elections all have contributions in this domain, we believe discussing the ability of winning parties to command enough support to implement their policies can be better formulated in a party-based framework. For example, considering voters in an ideological metric space (Schofield 2008), with parties forming separable subsets of the space, can allow for considering what location of candidate can give a party the implicit support of many voters, even if they vote for other candidates.

There are several potential ways to define such a robustness (e.g., using repeated games), but one possibility is the following:

Definition 1 *Let V be the set of voters, C the set of candidates, (M, d) a metric space with the metric d, and $\rho : V \cup C \to M$ the embedding function such that a voter $v \in V$ ranks a candidate $a \in C$ higher than $b \in C$ iff $d(\rho(v), \rho(a)) < d(\rho(v), \rho(b))$. Let f be a decision-making mechanism (possibly involving coalition formation and structure) such that the winner is $c \in C$. We say candidate c is ε-robust if for any $\bar{\rho} : V \cup C \to M$ such that $d(\bar{\rho}(x), \rho(x)) < \varepsilon$ for $x \in V \cup C \setminus \{c\}$, candidate c still wins.*

f's executive capacity $XC \in \mathbb{R}$ is such that f is ε-robust for any $\varepsilon < XC$ and f is not ε-robust for $\varepsilon > XC$.

This definition means that even if the voters slightly "move" in the ideological space, the outcome does not change. There are, potentially, many other possible definitions, but the key issue is to have a winner that is "stable" in its victory, and can therefore do various actions without fear of being toppled by someone else. This definition can easily encompass multiple agents and power relationships *within* each party, as well as candidates shifting to more popular positions between primary and general elections, making it a potentially useful one in this context as well.

Once such a measure is established, one can begin looking for mechanisms that provide better executive capacity, or – more interestingly – examine if there is a trade-off between representability (i.e., a parliamentary system that provides representatives more-or-less in the same proportion as their supporters in the electorate) and executive capacity. Naturally, this analysis may involve, as noted above, cooperative game theory or hedonic games as well, as they may be needed to establish a winner in a coalition-like setting.

Conclusion

> *The whole, at one time, is never old or middle-aged or young, but, in a condition of unchangeable constancy, moves on through the varied tenor of perpetual decay, fall, renovation, and progression.*
>
> Edmund Burke, *Reflections on the Revolution in France*, 1790

Edmund Burke, in the late 18th century, presented an intricate viewpoint, based on the historical events of his time (first the American revolution, and then, more than a decade later, the French revolution). At its core, this view considers radical revolutionary steps as misguided and prone to unexpected consequences, regardless of the original intentions. On the other hand, Burke, despite having a great deal of respect towards the established ways of governance of his time, believed that constant change is needed, and it should be done incrementally and carefully, as to ensure that unintended consequences are dealt with, and are never allowed to grow so bad that they present a fundamental challenge to the system of government itself.

In recent years, the computational social choice community has been exploring various novel ideas to replace "regular" human democracy. Indeed, the new ideas undermine our understanding of the operation of representative democracy. While we may have specific, particular issues with the suggested novel ideas, we also believe that being as radical as they are, they stand little chance of being adopted widely and significantly replace or meaningfully augment the existing systems. Thus, we believe, following Burke, that *"a politic caution, a guarded circumspection, a moral rather than a complexional timidity"* is a necessary approach to constructively engage both in the public debate around us, as well as to productively suggest better mechanisms that have fewer unexpected negative consequences.

Therefore, we believe that it is more constructive to understand why existing mechanisms are structured the way they are, and seek out ways to improve them in a gradual way. Eventually, as we more fully understand the interaction between the different parts of the mechanism, we can improve

different parts more and more, until, as in Theseus' ship, the whole has been made better.

A possible first step on this road to understanding existing mechanisms is to look at one significant component of almost any human-related decision-making system: parties, or, more generally, sub-groups. The ubiquity of these sub-groups indicates they may serve a meaningful role, and therefore, they are a great example of this Burkean approach. While some research has begun in this field, we believe a far larger effort is needed to understand them, as our current comprehension is quite limited. This means not only to explore the game-theoretic properties of party systems, but also to examine how to improve them using computational tools – from establishing better metrics, to utilizing the axiomatic approach and suggesting better mechanisms. But this is but one example of such topics that, we believe, could use this approach. Rather than a radical break with existing mechanisms, we believe exploring what makes successful mechanisms work, and expanding on them is an approach that can yield interesting research topics, and focus more research intelligence and conversation on issues which, lately, have not been central to the research community.

Acknowledgements

This work was supported in part by Israel Science Fund (ISF) grants #1965/20 and #3152/20, as well as the National Science Fund-Bi-national Science Foundation (NSF-BSF) grant #2021659.

References

Abramowitz, B.; Elkind, E.; Grossi, D.; Shapiro, E.; and Talmon, N. 2021. Democratic Forking: Choosing Sides with Social Choice. In *Proceedings of the 7th International Conference on Algorithmic Decision Theory (ADT)*, 341–356. Toulouse, France.

Abramowitz, B.; and Mattei, N. 2019. Flexible Representative Democracy: An Introduction with Binary Issues. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 3–10. Macau.

Anshelevich, E.; Bhardwaj, O.; and Postl, J. 2015. Approximating Optimal Social Choice under Metric Preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, 777–783. Austin, Texas.

Anshelevich, E.; and Postl, J. 2017. Randomized Social Choice Functions Under Metric Preferences. *Journal of Artificial Intelligence Research*, 58(1): 797–827.

Apollonio, N.; Becker, R.; Lari, I.; Ricca, F.; and Simeone, B. 2009. Bicolored graph partitioning, or: gerrymandering at its worst. *Discrete Applied Mathematics*, 157: 3601–3614.

Aziz, H.; Brill, M.; Conitzer, V.; Elkind, E.; Freeman, R.; and Walsh, T. 2017. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2): 461–485.

Aziz, H.; Lang, J.; and Monnot, J. 2016. Computing Pareto Optimal Committees. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 60–66. New York City, New York.

Aziz, H.; and Savani, R. 2016. *Handbook of Computational Social Choice*, chapter 15: Hedonic Games. Cambridge University Press.

Bachrach, Y.; Lev, O.; Lewenberg, Y.; and Zick, Y. 2016. Misrepresentation in District Voting. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 81–87. New York City, New York.

Becker, R.; D'Angelo, G.; Delfaraz, E.; and Gilbert, H. 2021. Unveiling the Truth in Liquid Democracy with Misinformed Voters. In *Proceedings of the 7th International Conference on Algorithmic Decision Theory (ADT)*, 132–146. Toulouse, France.

Bervoets, S.; and Merlin, V. 2012. Gerrymander-proof representative democracies. *International Journal of Game Theory*, 41(3): 473–488.

Bloembergen, D.; Grossi, D.; and Lackner, M. 2019. On Rational Delegations in Liquid Democracy. In *Proceedings of the 33rd Conference on Artificial Intelligence (AAAI)*, 1796–1803. Honolulu, Hawaii.

Blum, C.; and Zuber, C. I. 2016. Liquid Democracy: Potentials, Problems, and Perspectives. *TheJournalofPoliticalPhilosophy*, 24(2): 162–182.

Borodin, A.; Lev, O.; Shah, N.; and Strangway, T. 2018. Big City vs. the Great Outdoors: Voter Distribution and How it Affects Gerrymandering. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 98–104. Stockholm, Sweden.

Borodin, A.; Lev, O.; Shah, N.; and Strangway, T. 2019. Primarily About Primaries. In *Proceedings of the 33rd Conference on Artificial Intelligence (AAAI)*, 1804–1811. Honolulu, Hawaii.

Borodin, A.; Lev, O.; Shah, N.; and Strangway, T. 2022. Little House (Seat) on the Prairie: Compactness, Gerrymandering, and Population Distribution. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 154–162. Auckland, New Zealand.

Brill, M. 2019. *Interactive Democracy: New Challenges for Social Choice Theory*, 59–66. Springer Cham.

Brill, M.; and Talmon, N. 2018. Pairwise Liquid Democracy. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 137–143. Stockholm, Sweden.

Burke, E. 1790. *Reflections on the Revolution in France*. James Dodsley.

Caragiannis, I.; and Micha, E. 2019. A Contribution to the Critique of Liquid Democracy. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 116–122. Macau.

Chalkiadakis, G.; Elkind, E.; and Wooldridge, M. 2011. *Computational Aspects of Cooperative Game Theory*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan and Claypool.

Cohen-Zemach, A.; Lewenberg, Y.; and Rosenschein, J. S. 2018. Gerrymandering Over Graphs. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 274–282. Stockholm, Sweden.

Cohensius, G.; Mannor, S.; Meir, R.; Meirom, E.; and Orda, A. 2017. Proxy Voting for Better Outcomes. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 858–866. São-Paulo, Brazil.

Colley, R.; Grandi, U.; and Novaro, A. 2021. Unravelling multiagent ranked delegations. *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)*, 36(9).

Degani, C. 2021. This Israeli Professor Wants to Make Cryptocurrencies More Democratic. Ha'Aretz.

Dyer, M.; and Frieze, A. 1985. On the Complexity of Partitioning Graphs into Connected Subgraphs. *Discrete Applied Mathematics*, 10: 139–153.

Ebadian, S.; Kehne, G.; Micha, E.; Procaccia, A. D.; and Shah, N. 2022. Is Sortition Both Representative and Fair? In *Proceedings of the 36th Conference on Neural Information Processing Systems (NeurIPS)*, volume 35. New Orleans, Louisiana.

Elkind, E.; Faliszewski, P.; Skowron, P.; and Slinko, A. 2017. Properties of multiwinner voting rules. *Social Choice and Welfare*, 48(3): 599–632.

Elkind, E.; Grossi, D.; Shapiro, E.; and Talmon, N. 2021. United for Change: Deliberative Coalition Formation to Change the Status Quo. In *Proceedings of the 35th Conference on Artificial Intelligence (AAAI)*, 5339–5346.

Faliszewski, P.; Skowron, P.; Slinko, A.; and Talmon, N. 2017. Multiwinner Rules on Paths From k-Borda to Chamberlin-Courant. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 192–198. Melbourne, Australia.

Ferraz, C.; and Finan, F. 2011. Electoral Accountability and Corruption: Evidence from the Audits of Local Governments. *American Economic Review*, 101(4): 1274–1311.

Flanigan, B.; Gölz, P.; Gupta, A.; Hennig, B.; and Procaccia, A. D. 2021. Fair algorithms for selecting citizens' assemblies. *Nature*, 596: 548–552.

Flanigan, B.; Gölz, P.; Gupta, A.; and Procaccia, A. D. 2020. Neutralizing Self-Selection Bias in Sampling for Sortition. In *Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS)*, volume 33, 6528–6539.

Flanigan, B.; Kehne, G.; and Procaccia, A. D. 2021. Fair Sortition Made Transparent. In *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS)*, volume 34, 25720–25731.

Fukuyama, F. 1992. *The End of History and the Last Man*. Free Press.

Fukuyama, F. 2011. *The Origins of Political Order*. Farrar, Straus and Giroux.

Fukuyama, F. 2014. *Political Order and Political Decay*. Farrar, Straus and Giroux.

Gölz, P.; Kahng, A.; Mackenzie, S.; and Procaccia, A. D. 2021. The Fluid Mechanics of Liquid Democracy. *ACM Transactions on Economics and Computation (TEAC)*, 9(4): 1–39.

Grant, A. 2023. The Worst People Run for Office. It's Time for a Better Way. *New York Times*, A22.

Hamilton, A.; Madison, J.; and Jay, J. 1787. *The Federalist Papers*. The Independent Journal.

Hannon, M. 2022. Are knowledgeable voters better voters? *Politics, Philosophy & Economics*, 21(1): 29–54.

Harrenstein, P.; Lisowski, G.; Sridharan, R.; and Turrini, P. 2021. A Hotelling-Downs Framework for Party Nominees. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 593–601.

Harris, T. 2006. *Revolution: The Great Crisis of the British Monarchy, 1685-1720*. Allen Lane.

Hennig, B. 2017. *The End of Politicians: Time for a Real Democracy*. Unbound Digital.

IHT. 2008. Opposition claims win in Zimbabwe on unofficial tally. International Herald Tribune.

Kahng, A.; Mackenzie, S.; and Procaccia, A. D. 2021. Liquid Democracy: An Algorithmic Perspective. *Journal of Artificial Intelligence Research (JAIR)*, 70: 1223–1252.

Karpov, A.; Lev, O.; and Obraztsova, S. 2022. Hotelling-Downs Equilibria: Moving Beyond Plurality Variants. In *Proceedings of the 22nd International Conference on Group Decision and Negotiation (GDN)*, 3–16.

Lackner, M.; and Skowron, P. 2023. *Multi-Winner Voting with Approval Preferences*. SpringerBriefs in Intelligent Systems. Springer Cham.

Laurent, H. 2021. Corruption and politicians' horizon. *Economics of Governance*, 22(1): 65–91.

Lee, C. W.; McQuarrie, M.; and Walker, E. T., eds. 2015. *Democratizing Inequalities: Dilemmas of the New Public Participation*. New York University Press.

Lev, O.; and Lewenberg, Y. 2019. "Reverse Gerrymandering": a Decentralized Model for Multi-Group Decision Making. In *Proceedings of the 33rd Conference on Artificial Intelligence (AAAI)*, 2069–2076. Honolulu, Hawaii.

Lewenberg, Y.; Lev, O.; and Rosenschein, J. S. 2017. Divide and Conquer: Using Geographic Manipulation to Win District-Based Elections. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 624–632. São-Paulo, Brazil.

Meir, R.; Talmon, N.; Shahaf, G.; and Shapiro, E. 2022. Sybil-Resilient Social Choice with Partial Participation. In *Proceedings of the 19th European Conference on Multi-Agent Systems (EUMAS)*. Düsseldorf, Germany.

Mowbray, M.; and Gollmann, D. 2007. Electing the Doge of Venice: Analysis of a 13th century protocol. In *Proceedings of the IEEE Symposium on Computer Security Foundations*, 295–310.

OECD. 2020. *Innovative Citizen Participation and New Democratic Institutions*. Paris, France: OECD Publishing.

Pegden, W.; Procaccia, A. D.; and Yu, D. 2017. A partisan districting protocol with provably nonpartisan outcomes. ArXiv:1710.08781.

Poupko, O.; Shapiro, E.; and Talmon, N. 2022. Self-Sovereign Digital Agents for a Grassroots Digital Society. In *Proceedings of the IEEE 42nd International Conference on Distributed Computing Systems (ICDCS)*, 202–212. Bologna, Italy.

Procaccia, A. 2019. Lotteries Instead of Elections? Not So Arbitrary. *Bloomberg Service*.

Procaccia, A. 2022. A More Perfect Algorithm. *Scientific American*, 327: 52–59.

Procaccia, A. D.; and Rosenschein, J. S. 2006. The Distortion of Cardinal Preferences in Voting. In *Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA)*, 317–331. Edinburgh, Great Britain.

Raveh, O.; and Tsur, Y. 2021. Can Resource Windfalls Reduce Corruption? The Role of Term Limits. SSRN.

Schofield, N. 2008. *The Spatial Model of Politics*. Number 95 in Routledge Frontiers of Political Economy. Routledge.

Shapiro, E.; and Talmon, N. 2022. Foundations for Grassroots Democratic Metaverse. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 1814–1818. Auckland, New Zealand.

Tsur, Y. 2022. Political tenure, term limits and corruption. *European Journal of Political Economy*, 74: 102166.

van Bevern, R.; Bredereck, R.; Chen, J.; Froese, V.; Niedermeier, R.; and Woeginger, G. J. 2015. Network-Based Vertex Dissolution. *SIAM Journal on Discrete Mathematics*, 29(2): 888–914.

Wilentz, S. 2005. *The Rise of American Democracy*. Norton.

Yang, J. 2014. Some NP-complete edge packing and partitioning problems in planar graphs. ArXiv:1409.2426.

Zhang, Y.; and Grossi, D. 2021. Power in Liquid Democracy. In *Proceedings of the 35th Conference on Artificial Intelligence (AAAI)*, 5822–5830.