

Temporal Fairness in Multiwinner Voting

Edith Elkind^{1,2}, Svetlana Obraztsova³, Nicholas Teh¹

¹ University of Oxford, United Kingdom

² Alan Turing Institute, United Kingdom

³ Carleton University, Canada

elkind@cs.ox.ac.uk, svetlanaobraztsova@cunet.carleton.ca, nicholas.teh@cs.ox.ac.uk

Abstract

Multiwinner voting captures a wide variety of settings, from parliamentary elections in democratic systems to product placement in online shopping platforms. There is a large body of work dealing with axiomatic characterizations, computational complexity, and algorithmic analysis of multiwinner voting rules. Although many challenges remain, significant progress has been made in showing existence of fair and representative outcomes as well as efficient algorithmic solutions for many commonly studied settings. However, much of this work focuses on single-shot elections, even though in numerous real-world settings elections are held periodically and repeatedly. Hence, it is imperative to extend the study of multiwinner voting to temporal settings. Recently, there have been several efforts to address this challenge. However, these works are difficult to compare, as they model multi-period voting in very different ways. We propose a unified framework for studying temporal fairness in this domain, drawing connections with various existing bodies of work, and consolidating them within a general framework. We also identify gaps in existing literature, outline multiple opportunities for future work, and put forward a vision for the future of multiwinner voting in temporal settings.

Introduction

Elections for the AAAI Executive Council take place annually, with one-third of the positions being up for election each time. Artificial Intelligence research encompasses a broad variety of topics, and it is important that the council has representatives from all major subfields. Moreover, the community is very international, and the council should represent the interests of members around the world, so there is a desire for *proportionate representation* in terms of geographical regions. Diversity in terms of gender and academic seniority is another key desideratum.

There are several multiwinner voting rules that provide fair representation guarantees (Aziz et al. 2017; Peters and Skowron 2020). However, standard notions of fairness for multiwinner voting do not fully capture the complexity of our setting: it may be the case that a voter is not represented by any of the candidates elected in the current round, yet she is satisfied with the overall composition of the council because she is well-represented by some of the council

members elected in the previous rounds. Also, given that the council is renewed regularly, perhaps a group of voters that is too small to be assured permanent representation on the council can be guaranteed a spot every few rounds?

In a somewhat different spirit, consider proof-of-stake blockchain protocols where a primary concern is the “rich gets richer effect” (Fanti et al. 2019; Huang et al. 2021). In such scenarios, lotteries for a single-shot interaction provide a form of fairness, as the probability of winning is proportional to the invested effort (Orda and Rottenstreich 2019; Pan et al. 2022). However, such mechanisms may fail to maintain fairness *over time* when lotteries are held repeatedly over multiple rounds (Fanti et al. 2019; Grossi 2022). Formalizing fairness concepts and successfully designing mechanisms for achieving these in the repeated-interaction setting is therefore highly relevant to blockchain systems.

There are further examples of real-world elections where temporal considerations play an important role, and many authors explored multiwinner voting with temporal elements. However, this body of work does not yet offer a systematic exploration of issues associated with voting over time, with each strand of research considering a slightly different model and a different set of applications. Against this background, we propose a unified framework to facilitate a principled study of temporal fairness in multiwinner voting, highlighting several key challenges, consolidating existing bodies of work, and identifying gaps in the literature. We aim to set the groundwork for a more coherent, targeted approach to tackling conceptual and algorithmic challenges in temporal multiwinner voting.

Several elements of our framework originate from a larger body of work within the social choice literature and beyond. This includes, in particular, works on *perpetual voting* (Lackner 2020; Lackner and Maly 2023)—and, more broadly, perpetual participatory budgeting (Lackner, Maly, and Rey 2021)—which consider multiple rounds of single-winner voting. Lackner (2020) and Lackner and Maly (2023) conduct an axiomatic analysis of perpetual voting rules (temporal extensions of traditional voting rules) with respect to fairness across time. Relatedly, Harrenstein, Lackner, and Lackner (2022) argue that several common sequential election mechanisms may be detrimental to welfare. Another relevant strand of work is that on conference scheduling (Patro et al. 2022), slot assignment (Elkind, Kraczy,

and Teh 2022), and line-up elections (Boehmer et al. 2020), where the goal is to output a sequence of non-repeating winners, whilst maintaining fairness to voters across the entire time horizon. Sequential committee election models consider settings where an entire committee is elected in each round, and impose constraints on the extent a committee can change, whilst ensuring candidates continue to have support from the electorate (Bredereck, Kaczmarczyk, and Niedermeier 2020; Bredereck, Fluschnik, and Kaczmarczyk 2022; Deltl, Fluschnik, and Bredereck 2023). Other models in the social choice literature that include temporal elements include sequential decision-making (Kahana and Hazon 2023; Chandak, Goel, and Peters 2024), public decision-making (Brandt et al. 2016; Conitzer, Freeman, and Shah 2017; Fain, Munagala, and Shah 2018; Skowron and Górecki 2022; Lackner, Maly, and Nardi 2023), repeated matching (Golapudi, Kollias, and Plaut 2020; Trabelsi et al. 2023; Caragiannis and Narang 2023), resource allocation over time (Bampis, Escoffier, and Mladenovic 2018; Igarashi et al. 2024), online committee selection (Do et al. 2022), dynamic social choice (Parkes and Procaccia 2013; Freeman, Zahedi, and Conitzer 2017), and temporal liquid democracy (Markakis and Papisotiropoulos 2023).

As illustrated by the many examples mentioned above, maintaining welfare and fairness guarantees over time is relevant to many domains. Research into temporal multiwinner elections would thus allow us to capture many practical scenarios, discover notions of representation that are appropriate for temporal settings, and offer insights into the possibility of achieving desirable goals by society in a computationally efficient manner.

The remainder of the paper is structured as follows. First, we introduce the framework. We then proceed to look at each element of the model, fleshing out various options available, consolidating existing work and positioning it with regard to these options, and pointing out gaps or ideas for expansion. Finally, to set the foundations for further study, we formulate several research challenges and potential directions to push this field forward, and highlight the immense potential of research in this area.

The Temporal Multiwinner Voting Framework

The basic components of our framework are a set $N = \{1, \dots, n\}$ of n agents, a set $P = \{p_1, \dots, p_m\}$ of m distinct projects, or candidates, and a set $[\ell] = \{1, \dots, \ell\}$ of ℓ timesteps. For each timestep $r \in [\ell]$, each agent $i \in N$ has a preference $s_{i,r}$. This can come in various forms (see the Preferences section for a discussion). We write $\mathbf{s}_i = (s_{i,1}, s_{i,2}, \dots, s_{i,\ell})$ and refer to \mathbf{s}_i as i 's temporal preference. The goal is to select an outcome $\mathbf{o} = (o_1, \dots, o_\ell)$, which is a sequence of ℓ sets of candidates such that for every $r \in [\ell]$, the set of candidates $o_r \subseteq P$ is chosen at timestep r .

Next, we discuss various elements of the framework, contrasting several options within each element, and positioning the existing work in the computational social choice literature with respect to these options. We also outline several directions for future work.

Outcomes

The first parameter we explore is what is considered a permissible outcome at each timestep. In particular, we distinguish between selecting one candidate at a time and choosing multiple candidates at each timestep. We also consider constraints that can be imposed on the outcome vector.

Structure

In a single-shot multiwinner election, the outcome is a single set of candidates. There are two possible generalizations to the temporal setting. The first (and one that is more common in the literature) is having a single winning candidate chosen at each timestep (i.e., $o_r \in P$ for each $r \in [\ell]$). While one can view this variant of the model as a temporal extension of single-winner elections, the multiwinner interpretation is justified, too, as one can treat the (multi-)set $O = \{o_r : r \in [\ell]\}$ as the winning committee and apply fairness concepts that originate in multiwinner voting literature to the entire set O ; e.g., Bulteau et al. (2021) and Page, Shapiro, and Talmon (2022) reason about justified representation provided by O . This model is considered in numerous existing works, including scheduling problems (Elkind, Kraczy, and Teh 2022; Patro et al. 2022), perpetual voting (Lackner 2020; Lackner and Maly 2023), and public decision-making (Conitzer, Freeman, and Shah 2017; Fain, Munagala, and Shah 2018). The second model assumes that the goal is to select a fixed-size set of winning candidates (i.e., an entire committee) at each timestep. This approach is taken, e.g., by Bredereck, Kaczmarczyk, and Niedermeier (2020), Bredereck, Fluschnik, and Kaczmarczyk (2022), and Deltl, Fluschnik, and Bredereck (2023) in their study of sequential committee elections.

Feasibility Constraints

In the standard multiwinner voting setting, it is typically assumed that there is a parameter k such that every subset of P of size k is a feasible outcome (we note, however, that there is also work on single-shot multiwinner voting with constraints on feasible committees (Yang and Wang 2018)). However, in temporal settings it may be the case that at each timestep r only a subset of candidates $P_r \subseteq P$ is available, and the winning candidate(s) have to be chosen from this set. The simplest variant of this model is where the set P_r is given in advance and is independent of committees selected at steps r' with $r' \neq r$. This model is relevant if changes in P_r are caused by external constraints on candidate availability (i.e., perhaps candidate X does not want to run for the AAAI executive council this year due to heavy administrative load in their department, but will be happy to serve in the future). In the context of public decision-making (Conitzer, Freeman, and Shah 2017; Fain, Munagala, and Shah 2018; Skowron and Górecki 2022), such constraints can be influenced by the suitability of a project being implemented at a particular timestep due to manpower or geographical constraints (Lodi et al. 2022).

However, it may also be the case that a decision at timestep r constrains the options available at timestep r' : e.g., in fair scheduling problems it is common to assume

that each project in P can be implemented at most once, so if we select $p \in P$ at timestep r , it is no longer available at $r' \neq r$ (Elkind, Kraicz, and Teh 2022; Patro et al. 2022). More generally, for each project $p \in P$ there may be a bound $\alpha_p \in \mathbb{N}$, indicating the maximum number of times that p can be selected. One can also consider more sophisticated constraints, where at most (at least) a certain fraction of the winning committee needs to be replaced at each timestep; this approach is taken by Bredereck, Kaczmarczyk, and Niedermeier (2020) and Bredereck, Fluschnik, and Kaczmarczyk (2022). There are other constraints of this form that can be found in practice, but, to the best of our knowledge, have not been modeled in the literature: e.g., an AAAI executive council member, once elected, remains a member of the winning committee for three timesteps, but is then not eligible to participate in the next election (but can run again later on).

Preferences

Another important component of our framework is agents' preferences. The aspects that need to be considered include ballot types (e.g., approval, ranked, or cardinal), and if preferences can evolve over time (static or dynamic).

Ballot Types

The computational social choice literature typically focuses on three ballot types: approval ballots (Aziz et al. 2015; Lackner and Skowron 2023), ranked ballots (Faliszewski et al. 2017; Elkind et al. 2017), and cardinal ballots (Conitzer, Freeman, and Shah 2017; Freeman, Zahedi, and Conitzer 2017; Fain, Munagala, and Shah 2018).

Approval ballots Much of the recent work on proportionality in computational social choice focuses on *approval* ballots (Lackner and Skowron 2023): each agent reports which candidates they like and dislike. Approval preferences are relatively easy to elicit and reason about (Kilgour 1983; Brams and Fishburn 2005; Aragones, Gilboa, and Weiss 2011), yet they can capture a wide variety of scenarios from city budget planning to elections for the board of trustees. In temporal settings, approval ballots have been considered in the context of sequential committee elections (Bredereck, Kaczmarczyk, and Niedermeier 2020; Bredereck, Fluschnik, and Kaczmarczyk 2022; Deltl, Fluschnik, and Bredereck 2023) and scheduling (Bulteau et al. 2021; Elkind, Kraicz, and Teh 2022).

Ranked ballots Under ranked ballots, each agent reports a *ranking* over the candidates at each timestep; voting with ranked ballots has been extensively studied in single-shot multiwinner elections (Faliszewski et al. 2017), with applications ranging from parliamentary elections to movie selection (Elkind et al. 2017).

Cardinal ballots Cardinal ballots—where each voter explicitly specifies the *utility* they obtain from each candidate—offer a lot of expressivity, albeit at a higher cost of elicitation. Such ballots have been studied in settings such as public decision-making (Conitzer, Freeman, and Shah 2017; Freeman, Zahedi, and Conitzer 2017; Fain, Munagala, and Shah 2018), portioning (Freeman et al. 2021;

Elkind, Suksompong, and Teh 2023), and line-up elections (Boehmer et al. 2020).

Temporal Evolution

Next, we consider the temporal evolution of preferences. Preferences are said to be *static* if they do not change over time; otherwise, they are said to be *dynamic*.

Static preferences In this setting, each voter's preferences remain the same across the entire time horizon, and temporal considerations arise because of candidate availability issues or constraints on possible outcomes (see the Outcomes section for a discussion). For instance, Bulteau et al. (2021) consider static preferences in perpetual voting, and obtain positive algorithmic results for achieving notions of proportionality in this setting.

Dynamic preferences Dynamic preferences capture the idea that agents' preferences may evolve with time. Just as in case of outcomes, we distinguish between *non-adaptive* preferences, which evolve due to external considerations (e.g., when expressing preferences over restaurants for each day of the week, an agent may prefer a non-meat option on Fridays), and *adaptive* preferences, which change based on decisions that have been made so far (e.g., an agent may be unwilling to eat at the same restaurant twice in a row). For instance, the literature on the building of public projects (Conitzer, Freeman, and Shah 2017; Fain, Munagala, and Shah 2018; Skowron and Górecki 2022) and the work of Bulteau et al. (2021) on perpetual voting consider dynamic non-adaptive preferences, whereas Parkes and Procaccia (2013) consider dynamic adaptive preferences.

Intertemporal constraints An important special case that is not captured by the static/dynamic dichotomy is intertemporal restrictions on voters' ballots: e.g., in the context of scheduling, an agent may be asked to report her ideal schedule (Elkind, Kraicz, and Teh 2022). This means that agents approve exactly one project per timestep, with no repetition of projects—this is equivalent to a constraint imposed across timesteps.

We also note that, at least for approval ballots, simple candidate availability constraints can be incorporated into voters' preferences: e.g., if candidate p is not available at timestep r , we can simply remove p from all voters' ballots at r . By doing so, one can simplify the description of the input instance. However, this approach is not well-suited to more complex feasibility constraints or ranked ballots.

Observability

The next element we consider is the observability of future timesteps. The setting is said to be *online* if the preferences of agents at future timesteps are not known in advance, and *offline* otherwise.

Online settings Online scenarios are most prevalent in the dynamic social choice literature (Freeman, Zahedi, and Conitzer 2017; Freeman et al. 2018), where candidates are

selected sequentially without information on agents' preferences at future timesteps. This is reminiscent of the secretary problem, or, more specifically, its k -winner variant (Albers and Ladewig 2021); however, the key difference is that in the standard secretary problem the decision is made by a single stakeholder. Do et al. (2022) consider a multiple-stakeholder variant of the k -secretary problem that specifically addresses the question of whether committees selected online can proportionally represent voters. A similar model is explored by Israel and Brill (2021). Another line of work considers settings where voters rather than candidates are the ones that appear online (Oren and Lucier 2014; Dey, Talmon, and van Handel 2017). Other related online models in social choice include those in fair division (see, e.g., the survey by Aleksandrov and Walsh (2020)).

Offline settings In offline settings, agents' preferences are fully known in advance, prior to the computation of the outcome at each timestep. Several works on fair scheduling (Elkind, Kraicz, and Teh 2022; Patro et al. 2022) consider an offline model where preferences are fully available at the onset, and the goal is to study the computational problems associated with finding a desirable outcome or to conduct an axiomatic analysis of the mechanisms for obtaining an outcome. Offline models usually admit better solutions than online models; this is the case in, e.g., temporal fair allocation settings (Bampis, Escoffier, and Mladenovic 2018; Igarashi et al. 2024) and in the context of proportional representation in perpetual voting (Chandak, Goel, and Peters 2024).

Sequentiality

Another element of our framework is sequentiality of the problem instance. This emphasizes the importance of *order* over the timesteps. The order matters when agents have adaptive preferences (refer to the discussion of dynamic preferences in the Temporal Evolution subsection under Preferences), or where the contiguity of timesteps is important, as well as in online settings. For instance, in the context of AAAI executive council elections, a candidate agrees to serve on the council for three consecutive timesteps (and cannot participate immediately after completing their term); contiguity also matters for sequential committee elections (Bredereck, Kaczmarczyk, and Niedermeier 2020; Bredereck, Fluschnik, and Kaczmarczyk 2022; Deltl, Fluschnik, and Bredereck 2023). In contrast, in the scheduling model of Elkind, Kraicz, and Teh (2022) or Bulteau et al. (2021), timeslots can be rearranged arbitrarily, i.e., this model is fundamentally non-sequential. While non-sequential settings have interesting computational problems in their own right (Boehmer and Niedermeier 2021), sequential settings pose additional challenges.

Solution Concepts

The final element of our framework is the definition of suitable (temporal) solution concepts. For instance, one could consider the temporal extensions of popular fairness notions such as proportionality (Freeman et al. 2018), equitability (Elkind, Kraicz, and Teh 2022), or justified representation (Bulteau et al. 2021; Chandak, Goel, and Peters 2024). One

may also wish to use time as a tool for allowing agents to be *represented* eventually (global), consistently (local), or periodically (frequency-based). We classify these approaches into three (non-exhaustive) broad categories.

Global solution concepts Global solution concepts evaluate agents' welfare across the *entire* time horizon. This means emphasizing the eventual outcome with respect to the welfare goal, even if it means that within certain timesteps the treatment of agents can be highly unequal. An example of this approach is taking a global utilitarian or egalitarian view (Elkind, Kraicz, and Teh 2022), achieving a pre-defined notion of fairness and welfare in offline models of scheduling (Patro et al. 2022), or providing justified representation with respect to the choices made over all timesteps (Bulteau et al. 2021; Chandak, Goel, and Peters 2024; Elkind, Obraztsova, and Teh 2024).

Local solution concepts Local solution concepts evaluate agents' welfare at specified time intervals. This means emphasizing the welfare of agents either at each timestep, or for each pre-specified number of timesteps. For instance, in the perpetual voting model (Lackner 2020), decisions are made at each round, and welfare properties are considered for each agent up to and including the current round. In other models, such as sequential committee elections (Bredereck, Kaczmarczyk, and Niedermeier 2020; Bredereck, Fluschnik, and Kaczmarczyk 2022; Deltl, Fluschnik, and Bredereck 2023), the *quality* of a committee may be defined and maintained as a goal at each timestep.

Frequency-based solution concepts Another possibility is to consider frequency-based properties, which can be formulated in terms of bounds on the *number* of timesteps (be it consecutive or not) that have elapsed since the welfare of an agent was last addressed. For instance, one could mandate that there should not be κ timesteps that have elapsed such that an agent has not received a utility of at least γ . This idea was briefly mentioned by Boehmer and Niedermeier (2021) in a different, more abstract social choice scenario. Several other problems arise as well: e.g., one could ask what combinations of γ and κ can be accomplished by a specific voting rule (in the worst case, or on a specific instance), or by all rules satisfying a given set of axioms; one could also ask what is the minimum attainable γ for some fixed κ .

Research Directions/Challenges

We have suggested several dimensions according to which temporal multiwinner voting scenarios can be classified. Some of the "points" in the resulting multidimensional space have been considered already, but others remain unexplored and present opportunities for future work. On top of that, we will now highlight several interesting and challenging, yet fundamental directions for this field moving forward.

Formalism of solution concepts and goals The study of fairness across time opens up opportunities for formulating novel solution concepts, either ones that are specific to the temporal setting, or generalized (temporal) variants of concepts developed for the traditional multiwinner election model. For instance, one can extend popular concepts

of *representation* formulated for the setting of multiwinner elections with approval ballots, i.e., justified representation and its variants (Aziz et al. 2017; Sánchez-Fernández et al. 2017; Peters, Pierczyński, and Skowron 2021), to temporal scenarios, and study the associated computational problems; this direction has been explored by Bulteau et al. (2021), Chandak, Goel, and Peters (2024), and Elkind, Obraztsova, and Teh (2024), but many open problems remain. Other welfare goals in multiwinner elections such as diversity (Bredereck et al. 2018; Celis, Huang, and Vishnoi 2018; Relia 2022) can similarly be explored. Another direction that one could pursue would be considering a generalization of existing multiwinner voting rules to the temporal setting (see, e.g., the work of Lackner (2020), Lackner and Maly (2023) and Chandak, Goel, and Peters (2024)), and investigating whether these rules satisfy the corresponding generalized temporal axioms (e.g., representation, as discussed above).

The Temporal Dimension In the multiwinner voting literature, it is common to consider structural constraints on voters’ preferences, as a means to circumvent impossibility and computational hardness results (Betzler, Slinko, and Uhlmann 2013; Elkind and Lackner 2015). The temporal setting may benefit from this approach, too. In particular, it may be of interest to investigate the impact of structural constraints that are specific to the temporal setting. For instance, when considering evolving preferences, one can place additional restrictions on how agents’ preferences may change over time: e.g., perhaps agents’ approval sets can only expand (as they learn about benefits of projects they were previously not aware of), or, alternatively, the preferences cannot change too drastically between two consecutive time periods; similarly, the set of available candidates may evolve in a predictable manner, i.e., perhaps each candidate is only present for a number of consecutive steps. Such constraints on candidates’ availabilities and agents’ temporal preferences can be viewed as a novel structured preference domain (Elkind, Lackner, and Peters 2022), and can therefore offer a pathway towards positive algorithmic and axiomatic results.

Furthermore, we can consider settings where agents assign different importance to different timesteps (on top of preferences over candidates): e.g., in case of public projects, an agent may plan to be overseas during a specific time period and is therefore less interested in projects implemented during that period. Allowing more expressive ballots will allow solutions that provide better welfare guarantees, but is likely to introduce a whole new set of computational challenges that must be dealt with.

Moreover, it is interesting to explore extensions of normative analysis into the temporal realm. One could consider fairness criteria evaluated according to the worst-case, average-case or best-case with respect to timesteps, possibly with a discount factor. More broadly, when defining temporal extensions of notions of representation, rather than considering a simple additive variant across timesteps, one could define novel concepts that take into account the timesteps themselves in the definition.

Does time hurt or heal? Another important question is to understand the effect of time in comparison to traditional single-shot multiwinner election models. Apart from the impact (positive or negative) that it can have on agents’ welfare (as discussed in previous sections), one can approach this question from a computational perspective: does time make computing certain solution concepts easier (i.e., more algorithmically efficient), or does it add a computational hurdle?

For instance, in single-shot elections, outcomes that provide proportional justified representation (Sánchez-Fernández et al. 2017) can only be obtained by fairly sophisticated algorithms (Brill et al. 2017), whereas in perpetual voting with static preferences a simple greedy algorithm produces outcomes with this property (Bulteau et al. 2021). It would be interesting to explore whether similar results can be obtained for other proportionality and fairness concepts.

Beyond multiwinner voting: Participatory budgeting

Throughout the paper, we focused on the setting where the number of candidates to be selected is specified as part of the input. One can also consider a more general setting of *participatory budgeting*, where candidates (projects) may have distinct costs, and there is a budget, so that the total cost of the selected projects must remain within the budget (Aziz and Shah 2021; Rey and Maly 2023). While some of the work we discussed considers this more expressive setting (Lackner, Maly, and Rey 2021), most of the literature focuses exclusively on the basic multiwinner scenario; it would be interesting to see to what extent the existing positive results carry over to the richer participatory budgeting domain.

Conclusion

The multiwinner voting literature has burgeoned in recent years, and significant progress has been made towards designing robust mechanisms for a fairer and more just society. The inclusion of temporal considerations introduces a plethora of novel research directions, which should be explored systematically and in-depth, with a focus on real-life applications.

We note that, while we aimed to be comprehensive in our analysis and systematization of the prior work, and proposed some new research directions to be investigated, there may be further dimensions worth exploring. Overall, voting over time is a topic that is both practically important and theoretically challenging, as we hope that our paper will stimulate further research in this domain.

Acknowledgments

This work was supported by the AI Programme of The Alan Turing Institute.

References

- Albers, S.; and Ladewig, L. 2021. New Results for the k -Secretary Problem. *Theoretical Computer Science*, 863: 102–119.
- Aleksandrov, M.; and Walsh, T. 2020. Online Fair Division: A Survey. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 13557–13562.

- Aragones, E.; Gilboa, I.; and Weiss, A. 2011. Making Statements and Approval Voting. *Theory and Decision*, 71: 461–472.
- 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: 461–485.
- Aziz, H.; Gaspers, S.; Gudmundsson, J.; Mackenzie, S.; Mattei, N.; and Walsh, T. 2015. Computational Aspects of Multi-Winner Approval Voting. In *Proceedings of the 14th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 107–115.
- Aziz, H.; and Shah, N. 2021. Participatory budgeting: Models and approaches. In *Pathways Between Social Science and Computational Social Science: Theories, Methods, and Interpretations*, 215–236.
- Bampis, E.; Escoffier, B.; and Mladenovic, S. 2018. Fair Resource Allocation Over Time. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 766–773.
- Betzler, N.; Slinko, A.; and Uhlmann, J. 2013. On the Computation of Fully Proportional Representation. *Journal of Artificial Intelligence Research*, 47: 475–519.
- Boehmer, N.; Brederbeck, R.; Faliszewski, P.; Kaczmarczyk, A.; and Niedermeier, R. 2020. Line-Up Elections: Parallel Voting with Shared Candidate Pool. In *Proceedings of the 13th International Symposium on Algorithmic Game Theory (SAGT)*, 275–290.
- Boehmer, N.; and Niedermeier, R. 2021. Broadening the Research Agenda for Computational Social Choice: Multiple Preference Profiles and Multiple Solutions. In *Proceedings of the 20th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 1–5.
- Brams, S.; and Fishburn, P. 2005. Going From Theory to Practice: The Mixed Success of Approval Voting. *Social Choice and Welfare*, 25: 457–474.
- Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. 2016. *Handbook of Computational Social Choice*. Cambridge University Press.
- Brederbeck, R.; Faliszewski, P.; Igarashi, A.; Lackner, M.; and Skowron, P. 2018. Multiwinner Elections with Diversity Constraints. In *Proceedings of the 32th AAAI Conference on Artificial Intelligence (AAAI)*, 933–940.
- Brederbeck, R.; Fluschnik, T.; and Kaczmarczyk, A. 2022. When Votes Change and Committees Should (Not). In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, 144–150.
- Brederbeck, R.; Kaczmarczyk, A.; and Niedermeier, R. 2020. Electing Successive Committees: Complexity and Algorithms. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 1846–1853.
- Brill, M.; Freeman, R.; Janson, S.; and Lackner, M. 2017. Phragmén’s voting methods and justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 406–413.
- Bulteau, L.; Hazon, N.; Page, R.; Rosenfeld, A.; and Talmon, N. 2021. Justified Representation for Perpetual Voting. *IEEE Access*, 9: 96598–96612.
- Caragiannis, I.; and Narang, S. 2023. Repeatedly Matching Items to Agents Fairly and Efficiently. In *Proceedings of the 16th International Symposium on Algorithmic Game Theory (SAGT)*, 347–364.
- Celis, L. E.; Huang, L.; and Vishnoi, N. K. 2018. Multi-winner Voting with Fairness Constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 144–151.
- Chandak, N.; Goel, S.; and Peters, D. 2024. Proportional Aggregation of Preferences for Sequential Decision Making. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*.
- Conitzer, V.; Freeman, R.; and Shah, N. 2017. Fair Public Decision Making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, 629–646.
- Deltl, E. M.; Fluschnik, T.; and Brederbeck, R. 2023. Algorithmics of Egalitarian versus Equitable Sequences of Committees. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, 2651–2658.
- Dey, P.; Talmon, N.; and van Handel, O. 2017. Proportional Representation in Vote Streams. In *Proceedings of the 16th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 15–23.
- Do, V.; Hervouin, M.; Lang, J.; and Skowron, P. 2022. Online Approval Committee Elections. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, 251–257.
- Elkind, E.; Faliszewski, P.; Skowron, P.; and Slinko, A. 2017. Properties of Multiwinner Voting Rules. *Social Choice and Welfare*, 48: 599–632.
- Elkind, E.; Kraiczy, S.; and Teh, N. 2022. Fairness in Temporal Slot Assignment. In *Proceedings of the 15th International Symposium on Algorithmic Game Theory (SAGT)*, 490–507.
- Elkind, E.; and Lackner, M. 2015. Structure in Dichotomous Preferences. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI)*, 2019–2025.
- Elkind, E.; Lackner, M.; and Peters, D. 2022. Preference Restrictions in Computational Social Choice: A Survey. *arXiv preprint arXiv:2205.09092*.
- Elkind, E.; Obraztsova, S.; and Teh, N. 2024. Verifying Proportionality in Temporal Multiwinner Voting. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*.
- Elkind, E.; Suksompong, W.; and Teh, N. 2023. Settling the Score: Portioning with Cardinal Preferences. In *Proceedings of the 26th European Conference on Artificial Intelligence (ECAI)*, 621–628.
- Fain, B.; Munagala, K.; and Shah, N. 2018. Fair Allocation of Indivisible Public Goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, 575–592.
- Faliszewski, P.; Skowron, P.; Slinko, A.; and Talmon, N. 2017. Multiwinner Voting: A New Challenge for Social

- Choice Theory. In *Trends in Computational Social Choice*, chapter 2, 27–47.
- Fanti, G.; Kogan, L.; Oh, S.; Ruan, K.; Viswanath, P.; and Wang, G. 2019. Compounding of Wealth in Proof-of-Stake Cryptocurrencies. In *Proceedings of the 23rd International Conference on Financial Cryptography and Data Security (FC)*, 42–61.
- Freeman, R.; Pennock, D.; Peters, D.; and Wortman Vaughan, J. 2021. Truthful Aggregation of Budget Proposals. *Journal of Economic Theory*, 193: 105234.
- Freeman, R.; Zahedi, S. M.; and Conitzer, V. 2017. Fair and Efficient Social Choice in Dynamic Settings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 4580–4587.
- Freeman, R.; Zahedi, S. M.; Conitzer, V.; and Lee, B. 2018. Dynamic Proportional Sharing: A Game-Theoretic Approach. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 2(1): 1–36.
- Gollapudi, S.; Kollias, K.; and Plaut, B. 2020. Almost Envy-Free Repeated Matching in Two-Sided Markets. In *Proceedings of the 16th International Conference on Web and Internet Economics (WINE)*, 3–16.
- Grossi, D. 2022. Social Choice Around the Block: On the Computational Social Choice of Blockchain. In *Proceedings of the 21st International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 1788–1793.
- Harrenstein, P.; Lackner, M.-L.; and Lackner, M. 2022. A Mathematical Analysis of an Election System Proposed by Gottlob Frege. *Erkenntnis*, 87: 2609–2644.
- Huang, Y.; Tang, J.; Cong, Q.; Lim, A.; and Xu, J. 2021. Do the Rich Get Richer? Fairness Analysis for Blockchain Incentives. *arXiv preprint arXiv:2103.14713*.
- Igarashi, A.; Lackner, M.; Nardi, O.; and Novaro, A. 2024. Repeated Fair Allocation of Indivisible Items. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*.
- Israel, J.; and Brill, M. 2021. Dynamic Proportional Rankings. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 261–267.
- Kahana, I.; and Hazon, N. 2023. The Leximin Approach for a Sequence of Collective Decisions. *arXiv preprint arXiv:2305.18024*.
- Kilgour, D. M. 1983. Approval Voting. *Canadian Journal of Political Science*, 16(4): 839–840.
- Lackner, M. 2020. Perpetual Voting: Fairness in Long-Term Decision Making. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2103–2110.
- Lackner, M.; and Maly, J. 2023. Proportional Decisions in Perpetual Voting. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 5722–5729.
- Lackner, M.; Maly, J.; and Nardi, O. 2023. Free-Riding in Multi-Issue Decisions. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2040–2048.
- Lackner, M.; Maly, J.; and Rey, S. 2021. Fairness in Long-Term Participatory Budgeting. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 299–305.
- Lackner, M.; and Skowron, P. 2023. *Multi-Winner Voting with Approval Preferences*. Springer Nature.
- Lodi, A.; Olivier, P.; Pesant, G.; and Sankaranarayanan, S. 2022. Fairness Over Time in Dynamic Resource Allocation with an Application in Healthcare. *Mathematical Programming*.
- Markakis, E.; and Papatotiroopoulos, G. 2023. As Time Goes By: Adding a Temporal Dimension Towards Resolving Delegations in Liquid Democracy. *arXiv preprint arXiv:2307.12898*.
- Orda, A.; and Rottenstreich, O. 2019. Enforcing Fairness in Blockchain Transaction Ordering. In *Proceedings of the 2022 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 368–375.
- Oren, J.; and Lucier, B. 2014. Online (Budgeted) Social Choice. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, 1456–1462.
- Page, R.; Shapiro, E.; and Talmon, N. 2022. Electing the Executive Branch. *arXiv preprint arXiv:2009.09734*.
- Pan, Y.; Zhao, Y.; Liu, X.; Wang, G.; and Su, M. 2022. FPLotto: A Fair Blockchain-Based Lottery Scheme for Privacy Protection. In *Proceedings of the 2022 IEEE International Conference on Blockchain*, 21–28.
- Parkes, D.; and Procaccia, A. 2013. Dynamic Social Choice with Evolving Preferences. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI)*, 767–773.
- Patro, G.; Jana, P.; Chakraborty, A.; Gummadi, K.; and Ganguly, N. 2022. Scheduling Virtual Conferences Fairly: Achieving Equitable Participant and Speaker Satisfaction. In *Proceedings of the ACM Web Conference (WWW) 2022*, 2646–2656.
- Peters, D.; Pierczyński, G.; and Skowron, P. 2021. Proportional Participatory Budgeting with Additive Utilities. In *Proceedings of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS)*, 12726–12737.
- Peters, D.; and Skowron, P. 2020. Proportionality and the Limits of Welfarism. In *Proceedings of the 21st ACM Conference on Economics and Computation*, 793–794.
- Relia, K. 2022. DiRe Committee: Diversity and Representation Constraints in Multiwinner Elections. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, 5143–5149.
- Rey, S.; and Maly, J. 2023. The (Computational) Social Choice Take on Indivisible Participatory Budgeting. *arXiv preprint arXiv:2303.00621*.
- Sánchez-Fernández, L.; Elkind, E.; Lackner, M.; Fernández, N.; Fisteus, J.; Val, P. B.; and Skowron, P. 2017. Proportional Justified Representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 670–676.
- Skowron, P.; and Górecki, A. 2022. Proportional Public Decisions. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, 5191–5198.

Trabelsi, Y.; Adiga, A.; Kraus, S.; Ravi, S. S.; and Rosenkrantz, D. J. 2023. Resource Sharing Through Multi-Round Matchings. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 11681–11690.

Yang, Y.; and Wang, J. 2018. Multiwinner Voting with Restricted Admissible Sets: Complexity and Strategyproofness. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 576–582.