

# 1 Scientific Background

The strength of a society depends on the cooperation between its members, and the willingness of people to cooperate depends on them feeling that they are treated fairly. In particular, it is important that all sectors in society feel that they receive a fair share of society's resources and political influence.

A well-known concept that captures the notion of fairness for groups is *proportional representation*. In parliamentary elections, in which each voter votes for a single party, the notion of proportional representation simply says that each party should receive seats in proportion to the number of votes it received.

Recently, there is a growing interest in allowing citizens to vote directly for candidates, for example, using approval ballots. This challenges the existing proportional representation notion.

**Example 1.** Suppose we want to elect a committee of 10 members from a pool of 20 candidates, using approval ballots. There are 1000 citizens, of whom 600 vote for candidates  $1, \dots, 10$  whereas the other 400 vote for candidates  $11, \dots, 20$ .

If committee members are elected by a simple majority vote, as in the common *approval voting* rule, then only candidates  $1, \dots, 10$  will be elected, whereas the group of 400 citizens supporting candidates  $11, \dots, 20$  will have no representation in the committee.

The voting profile in Example 1 is very simple: essentially, the candidates are partitioned into two parties, and each voter votes for exactly one party. In this example, it is clear what proportional representation means: the committee should have 6 candidates from the set  $\{1, \dots, 10\}$  and 4 candidates from the set  $\{11, \dots, 20\}$ .

Recently, researchers have generalized the notion of proportional representation to the committee election settings, in which each voter can vote for an arbitrary set of candidates, not necessarily from the same party. Aziz et al. [4] and Sánchez-Fernández et al. [51] introduced axioms such as *proportional justified representation*, *extended justified representation* and *core stability*. A bit later, Skowron [54] studied a quantitative notion called *proportionality degree*; Peters and Skowron [45] introduced a market-based notion called *priceability*; and Brill and Peters [13] introduced stronger notions called PJR+ and EJR+. See the survey in Faliszewski et al. [23] and chapter 4 in Lackner and Skowron [35]. We refer to all these notions, that emphasize fairness to groups of voters rather than to individuals, by the generic term *fair representation*.

Fair representation notions were extended to many applications beyond committee election:

(1) *Participatory budgeting* (PB): this is a process in which a fixed budget should be used to fund a set of projects, from among a larger set of candidate projects. The set of funded projects is selected democratically by the citizens. PB can be seen as a generalization of committee election, in which each candidate has a cost, and the elected committee should have a bounded total cost (rather than a bounded cardinality). The fair representation notions have been generalized to PB by Aziz et al. [5], Brill et al. [14], Los et al. [38], Maly et al. [40], Peters et al. [47], and others. These notions ensure that different population sectors have control over a fair share of the budget.

(2) *Donor coordination* (DC): this is a process in which several donors to public projects coordinate their donations in order to increase their effectiveness. It is similar to PB, except that the budget is endogenous (comes from the donors) rather than exogenous (from the government). In this context, Brandl et al. [10] introduced a fairness notion called *decomposability* (also called *group fair share*), which intuitively guarantees to each subset of donors that all their donations are spent on projects that they support.

(3) *Perpetual voting* (PV): this is a process in which a set of decisions are to be taken sequentially, over a number of time periods. It is similar to committee election, since the outcome consists of several winning candidates (one winning candidate for each time period), but in PV there may be a different set of candidates in each time period, and the same candidate may be elected multiple times. PV was introduced by Lackner [32], but his analysis focused on fairness notions for individual agents, rather than for groups. In [17], we generalized the notions of fair representation to PV. Our work was later extended by Chandak et al. [20].

(4) *Dynamic ranking of public issues*: this is a new application, which aims to help public servants decide what issues are most important to handle. It allows citizens to vote on issues important to them, and uses the voting data to rank the issues from most to least important. The rank changes dynamically according to the decisions made by the public servant, in order to prioritize citizens whose opinion was not yet taken into account. It can be seen as an extension of committee election in which, instead of constructing an unordered subset of candidates, we construct an order among the candidates. Israel and Brill [27] extended some fair representation notions to dynamic ranking.

Taken together, these applications have a great potential to improve the quality of governance and the citizen participation and cooperation. However, despite the theoretical appeal of fair

representation algorithms, their use for solving practical problems is still quite rare.<sup>1</sup> In recent years, we are trying to implement some of these procedures in practice, and these attempts reveal various theoretical gaps, indicating that they are still not quite ready for solving real-life allocation problems. This research proposal aims to address the following gaps:

(a) *Mixing divisible and indivisible items.* Most of the literature either assumes that all candidates are indivisible (i.e., must be elected as a whole or not elected at all), or that all candidates are divisible (i.e., can be elected at any fraction), but realistic settings contain both indivisible items and divisible items. For example, consider a PB process for determining the budget of a student association. One candidate project is constructing a new student club, which is indivisible (must be constructed as a whole, or not at all); another candidate project is subsidizing the food in the cafeteria, which is divisible (any amount can be used to attain some subsidy). Moreover, there are projects that can be funded in various degrees, for example: the Purim party can be small or large. New algorithms are needed, that can handle both these issues while keeping fair representation.

(b) *Determining the form of agents' utility functions.* Most papers on fair representation algorithms make some assumptions on the agents' utility functions. For example, in the context of donor coordination, Brandl et al. [10] and Brandl et al. [11] assume additive utilities whereas Brandt et al. [12] assume Leontief utilities. The results are substantially different. This raises the need to assess the form of the agents' utility functions in various realistic settings, in order to decide which algorithm to apply.

(c) *Explaining the outcomes of algorithms.* Although there are sophisticated algorithms that provably satisfy fair representation requirements, their operation may not be transparent to plain humans. As our overarching goal is to encourage citizen participation and cooperation, it is important that they understand why exactly a specific outcome was chosen by the algorithm, and why it is fair towards them.

(d) *Allowing richer ballots.* The most well-studied ballot format, in the context of fair representation, is *approval ballots* (where each vote may approve a subset of the candidates). Richer ballot formats, such as *ranking* (i.e., ordinal ballots) and *cardinal ballots* are available. Another important ballot format is the *trichotomous* one, allowing agents to cast three types of ballots:

---

<sup>1</sup>As an example, participatory budgeting is used in over 7000 municipalities around the world, but only 3 municipalities (as far as we know) use a provably fair allocation procedure — the Method of Equal Shares [47].

supporting a candidate, objecting a candidate, or indifferent. These richer ballots require richer definitions of fair representation notions. These were studied for some applications, but not for others. For example, ordinal and cardinal ballots were not studied for fair-representation perpetual voting, and trichotomous ballots were not studied for participatory budgeting.

(e) *Mitigating the effects of strategic voting.* Algorithms that guarantee fair representation are prone to strategic manipulations. A particularly easy manipulation is *free-riding*: a user can claim to object a candidate that he supports, if he knows that this candidate is going to win anyway, as this will induce the algorithm to “compensate” him by electing other candidates he supports. There are strategyproof algorithms, but they may be very unfair towards minorities. Unfortunately, this trade-off between fairness and strategyproofness is unavoidable: Peters [44] proved that *every* algorithm that guarantees fair representation is manipulable, but we have several ideas on how to mitigate the effects of strategic manipulations.

In theory, each of the gaps (a) to (e) may be relevant for each of the settings (1) to (4). In practice, some gaps are more prevalent in particular settings. In addition, addressing all gaps for all applications may not be achievable in the four years time-frame (though the research team consists of two researchers with graduate students, and additional collaborators from other universities). Still, even bridging some of the gaps for some of the applications is expected to significantly help making the concept of fair representation more applicable to real-world problems. In Section 3, we discuss each of the settings (1) to (4) in detail; we use each setting to exemplify one of the gaps (a) to (e), which we believe is most relevant to this setting.

## 2 Research Objectives and Expected Significance

The main objective is to make the concepts of fair representation applicable to real-world problems, by addressing theoretical gaps in existing fair representation algorithms.

An additional objective is to raise public awareness to the existence of fair representation procedures. To attain this objective, we plan to work with people and organizations that may be interested in such algorithms, such as: student associations, charity organizations, mayors, and other public servants. This cooperation may lead to revealing new gaps that have to be addressed by researchers, and meanwhile, it may increase the public awareness to the potential benefits of the new algorithms.

The expected significance is in having new algorithms that can be used to allocate public resources fairly; using these algorithms may reduce tensions between sectors in society, and increase the social resilience.

### 3 Detailed Description of Proposed Research

This section presents in detail each of the four applications outlined in the introduction.

#### 3.1 Participatory budgeting

*Participatory budgeting* (PB) is one of the most successful democratic innovations in recent years. It is a process in which a municipality dedicates a part of its budget to projects suggested by the public. As the budget usually cannot cover all the suggested projects, the citizens vote by marking the projects that they prefer. Then, an aggregation rule is used to choose which projects to fund [3, 49, 53]. It is implemented in over 7000 municipalities around the world.<sup>2</sup>

The most common aggregation rule (called “Simple Greedy”) simply sorts projects in decreasing order of the number of votes they received, and selects them one by one as long as the total cost of the selected set is below the available budget. This rule displays the same unfairness illustrated in Example 1: if 60% of the population support projects 1, . . . , 10 whereas the other 40% of the population support projects 11. . . , 20, and the budget suffices for only 10 projects, then the Simple Greedy rule would choose only the projects supported by the majority, and no budget would remain for projects supported by the minority.

Following this observation, researchers have extended the notions of fair representation from committee election to PB, and developed various rules that aim to guarantee fair representation [5, 14, 47]. One of these rules — the Method of Equal Shares [47] — has already been applied in three cities in Europe. Computerized simulations showed that this method substantially increased the average number of winning projects that were approved by citizens.<sup>3</sup> As far as we know, no Israeli city has implemented participatory budgeting yet.

In an attempt to bring participatory budgeting into Israel, we have discussed with the student association of our university the possibility of implementing participatory budgeting for their 2024

---

<sup>2</sup>Retrieved from the Participatory Budgeting Project website: <https://www.participatorybudgeting.org/about-pb/#what-is-pb> at 31/10/2023.

<sup>3</sup>Retrieved from their website <https://equalshares.net/> at 30/10/2023.

budget. During the planning, two issues with current participatory budgeting rules came up.

First, most implemented participatory budgeting processes consider *indivisible projects*, that is, projects that are either funded completely, or not funded at all. There are also theoretical works considering the other extreme: all budget items are completely divisible, that is, they can benefit from any amount of money given to them [3]. However, real budget can have both divisible and indivisible items. For example, constructing a new student club is an indivisible project, but subsidizing food in the students' cafeteria is divisible, as any amount of money can be used to decrease the food prices. Theoretically, existing participatory budgeting rules can handle divisible projects by treating each dollar given to the project as an individual item; however, it is not clear that this approach would lead to meaningful fairness guarantees.

There are several related works in the area of *fair item allocation*. In particular, Bei et al. [8] presented algorithms for fair allocation of goods, where some goods are known in advance to be divisible and others are known to be indivisible. In [52], we have presented algorithms for efficient fair allocation of goods and bads, where all items are divisible, but the number of split items should be as small as possible. In [9], we study the computational complexity of fair allocation among agents with different entitlements, when the number of split items is upper-bounded. However, fair item allocation is substantially different from participatory budgeting in that the allocated items are *private goods* (each item benefits a unique agent), whereas participatory budgeting — by its very nature — involves *public goods* (each item benefits all agents simultaneously).

Recently, Lu et al. [39] studied fair representation in committee election with mixed divisible and indivisible candidates; we aim to extend their work to PB.

**Research Objective 1.** Develop algorithms for participatory budgeting guaranteeing fair representation when: (a) the list of potential projects contains both divisible and indivisible items; (b) all projects can be considered divisible, but the number of divisible projects should be minimized.

A related issue is that even indivisible projects have two or more degrees of funding. For example, a Purim party can be held either in a small and cheap facility, or in a large and expensive facility (or not at all). Very recently, Sreedurga [55] studied this setting, but his study focused on computing welfare-maximizing allocations, and did not consider fairness.

**Research Objective 2.** Develop algorithms for participatory budgeting guaranteeing fair representation when projects are indivisible, but may have two or more degrees of funding.

### 3.2 Donor coordination

*Donor coordination* is a process closely related to participatory budgeting. The main difference is that the budget does not come from outside (e.g. from the municipality), but donated by the agents themselves. Thus, a major design consideration in donor coordination is to encourage *participation* — ensuring the donors that donating more money cannot decrease their utility. Simultaneously, it is desirable to find allocations that are Pareto-efficient.

The axiomatic study of donor coordination was initiated by Brandl et al. [10]. They assumed that the projects are divisible, and that donors have *additive dichotomous* preferences. That is, for each agent  $i$  there is a set  $A_i$  of *approved projects*, and the utility of  $i$  from a budget-allocation is the sum of money allocated to projects in  $A_i$ . The following example illustrates the importance of donor coordination in this case.

**Example 2.** Suppose there are two donors: Alice donates 500 and Bob donates 100. There are three potential projects, Alice approves projects 1 and 2, and Bob approves projects 2 and 3. Without donor coordination, each donor would typically donate half of his/her donation to each of his/her supported projects, resulting in a distribution of 250,300,50 and a utility profile of 550,350. However, this distribution is not efficient: it is Pareto-dominated by the distribution 0,600,0, which yields a utility profile of 600,600.

A simple way to attain an efficient distribution is using the *utilitarian rule*, which maximizes the sum of agents' utilities. In Example 2, this rule would indeed yield the efficient distribution 600,600. However, this rule might display the same unfairness as in Example 1: if 60% of the donors approve project 1 and 40% of them approve project 2, the utilitarian rule would allocate all the donations to project 1. This is unfair towards the supporters of project 2. Over time, such unfair allocations will decrease their willingness to donate.

Here, the main fairness notion is *decomposability*: it means that the budget allocation  $\mathbf{d}$  can be decomposed into individual allocations  $\mathbf{d}_i$ , such that for each agent  $i$ , the sum of elements in  $\mathbf{d}_i$  equals the donation of  $i$ , and  $\sum_i \mathbf{d}_i = \mathbf{d}$ . Effectively, it means that each donor and group of donors spend their donations only on issues that matter to them.

Brandl et al. [10] identified allocation rules that satisfy participation, efficiency and decomposability. On the negative side, they proved that no distribution rule that is efficient and strategyproof can satisfy even a fairness property much weaker than decomposability: guaranteeing a positive utility to each donor.

Brandl et al. [11] extended the model to *additive non-dichotomous preferences*. That is, for each agent  $i$  and project  $x \in A_i$  there is a value  $v_{i,x}$  describing the importance of  $x$  to  $i$ , and the utility of  $i$  from budget-allocation  $\mathbf{d}$  is the sum:  $u_i(\mathbf{d}) = \sum_{x \in A_i} v_{i,x} d_x$ . They proved that the rule maximizing the product of utilities (the *Nash welfare*) satisfies participation, efficiency and decomposability, though (in accordance with the impossibility results proved by Brandl et al. [10]) it is not strategyproof. In Example 2, the Nash rule would indeed find the efficient distribution 0,600,0; in the case of 60%:40% divide, the Nash rule would find the fair allocation, which gives 60% of the donations to project 1 and 40% to project 2.

While considering how to implement the donor coordination idea to charitable donations, we have noticed that the distribution 0, 600, 0 does not make sense when the “projects” are charities, as charitable donations aim to give as much possible to the poorest recipients, and not leave any recipient with 0 money. This lead to the conclusion that the additive utility model is not a good fit for charities. In [12], we have suggested an alternative model, where the utility an agent wants to maximize is the *smallest* amount given to a charity he approves (scaled by the charity value to the agent):  $u_i(\mathbf{d}) = \min_{x \in A_i} d_x / v_{i,x}$ . For this utility model, we identified a rule called *equilibrium distribution rule*, that is efficient, strategyproof and decomposable, and moreover, all donations of each donor group are distributed *optimally* for the group, given the other donors’ distributions. In Example 2, it would yield the distribution 250, 250, 100, with a utility profile (in the minimum-based utility model) of 250, 100. In contrast, the utility profile of 0, 600, 0 in this model is 0, 0, and the utility profile of the uncoordinated distribution 250, 300, 50 is 250, 50. The same rule satisfies various other desirable properties, such as monotonicity and group-strategyproofness. This result stands in stark contrast to the impossibility result of Brandl et al. [10] in the additive utility model.

Yet another utility model is studied by Buterin et al. [18]: they assume that donors’ utilities are quasilinear, and suggest a mechanism based on *quadratic voting*.

The difference in results between different utility models naturally raises two challenges. The first challenge is to find out what are the actual utility functions of donors in various settings. It is reasonable to assume that utility functions for charitable donations differ from utility functions for other purposes, such as funding public projects. In order to decide what mechanism to use in each setting, we need a way to assess the format of donors’ utility functions.

There is a vast literature on empirical measurements of utility functions; see e.g. Kirby [31] and Ramos-Goñi et al. [48] and the many references therein. However, most of these studies focus

on private goods, such as income or health status. In contrast, donations create public goods. Recently, Rosenfeld and Talmon [50] conducted a related experiment in indivisible participatory budgeting. One of their aims was to assess the relation between voters' approval ballots, and their utilities from particular budget-allocations (e.g. do users prefer that a large *number* of their approved projects be funded, or that a large amount of *money* be allocated to their approved projects?). However, their assessment was indirect: they used each potential utility function as an input to an optimization algorithm, and asked the subjects to assess the outcomes of the optimization algorithms from the point of view of an impartial manager. We want to assess donors' utility functions directly, in the context of *divisible* donor coordination. We also want to assess the effects of the scenario on the utility function, e.g. compare utility functions for donations to charity organizations vs. donations to other purposes. Our plan is to adapt the procedures in the book of Keeney and Raiffa [30] for multi-objective decision-making.

**Research Objective 3.** Conduct experiments for assessing the form of donors' utility functions in various donor coordination settings.

This objective requires IRB approval; as the planned experiment is a relatively straightforward opinion-poll, with no risk to subjects, we believe we will receive the IRB approval.

The second challenge is to develop good distribution rules for realistic utility functions. In preliminary work we found out that, for *dichotomous* utilities, our equilibrium distribution rule [12] can be applied whenever each agent  $i$  has a separably-additive, strictly-concave utility function, i.e.,  $u_i(\mathbf{d}) = \sum_{x \in A_i} g_i(d_x)$ , where  $g_i$  is strictly-concave, strictly-increasing and continuously differentiable. However, so far we could not extend these results to general (non-dichotomous) utilities. We also do not have a complete classification of distribution rules and their properties for arbitrary utility functions.

**Research Objective 4.** Develop distribution rules that satisfy desirable properties for utility functions of real donors, as found in Research Objective 3.

### 3.3 Perpetual voting

In a standard voting setting, a single collective decision has to be made in a single point in time. In this setting, choosing the candidate with the highest number of votes (i.e., approval voting rule) is a sensible choice. There are, however, many situations in which a set of collective decisions are to

be taken sequentially, over a number of time periods. We refer to these settings as *perpetual voting* (PV). The most common examples are presidential elections in democratic countries, which occur every fixed number of years. Using approval voting in each time-period might result in the same unfairness illustrated in Example 1. Suppose that there are 10 time periods and two parties. In each period, each party nominates one candidate. 60% of the population support the candidates nominated by the first party in all periods, and the other 40% support the candidates nominated by the second party in all periods. Approval voting would choose the candidate nominated by the first party in all periods, leaving the minority group with no influence on the outcome. Over time, this might lead to frustration and alienation among the minority group.

The study of fairness in PV setting was initiated by Lackner [32], who suggested several perpetual voting rule, and analyzed their fairness via three axiomatic properties. Indeed, Lackner concentrated on fairness towards each individual voter, and analyzed only simple proportionality, which applies only to simple instances (where the set of candidates and the votes remain the same in each time period, and each voter supports just a single candidate). In [17], we generalized fair representation notions to PV. Recently, Chandak et al. [20] further extended our work, by adapting stronger notions of fair representation to PV, and by showing three PV rules that satisfy some fair representation notions. Lackner and Maly [33] suggested additional fairness notions towards groups of voters in PV (“perpetual lower quota for closed groups”, “perpetual upper quota for closed groups”, and “perpetual priceability”).

As a first step towards using PV for practical problems, we conducted a large scale human study [17]. Our study was aimed at identifying what potential voters (i.e., ordinary people of different backgrounds) deem as desirable outcomes in PV settings. To our surprise, our results suggested that using approval voting was generally preferred to any other examined method by our study participants, regardless of their background. That is, even though the other methods satisfy one of the fair representation notions, they were inferior to approval voting, which violates all the fair representation notions. Therefore, our first challenge is to make PV rules, which satisfy fair representation notions, more favorable to human subjects. We plan to do so using explanations.

Explanations are useful, since it is important that users are not only treated fairly, but also *see* that they are treated fairly. Several works in computational social choice discuss explainability of voting rules [19, 46]. In [42], we developed an automatic method that generates intuitive explanations for a Shapley-based payoff allocation. A very recent empirical work [58] studies the

effects of various explanation schemes on the human perception of various rules for participatory budgeting. We are not aware of similar studies on explainability of PV rules<sup>4</sup>.

**Research Objective 5.** (a) Develop algorithms for generating automatic explanations for PV rules. (b) Test the generated explanations on human subjects.

A second issue that can make PV rules more applicable, is to allow the users richer options to express their preferences. That is, all of the existing PV rules, which satisfy notions of fair representation, allow the users to vote using approval ballots. Indeed, Conitzer et al. [22] studied fair public decision making, which is a variant of PV in which the voters can use cardinal ballots, and all decisions are made offline. Freeman et al. [24] study the online version of public decision making, in which the voting rule has to choose the winner in each time period based only on the preferences of the voters up to that time period. In [28], we studied both offline and online fair public decision making. All of these works consider fairness towards individual agents, and not towards group of agents.

In committee election and PB, there are several works that study notions of fair representation with ordinal or cardinal ballots, for example, [1, 2, 13, 47]. To the best of our knowledge, there is no work that studies PV with ordinal or cardinal ballots.

**Research Objective 6.** Adapt notions of fair representation to PV with ordinal or cardinal ballots, and develop PV rules that satisfy them.

Another issue with PV is related to its potential application to presidential elections. These elections are usually anonymous; therefore, in each period  $t$ , we do not know which candidate each citizen voted for in periods  $1, \dots, t - 1$ ; we only know the total number of votes each candidate received. However, fair representation rules need to know the entire ballot of each voter. Without this information, we must make some simplifying assumptions. For example, if the candidates in time  $t$  are the same as in time  $t - 1$ , and the distribution of votes is the same, it is reasonable to assume that the same voters voted for the same candidates in each period, and apply fair representation rules accordingly. However, it is not clear how to extend this idea to settings in which the candidates or the vote distributions are different.

Another potential approach is to relax the anonymity requirement: instead of keeping all votes of all candidates in each period, we “compile” the voting data and keep a small amount of

---

<sup>4</sup>Lackner and Maly [33] commented that win-based and loss-based weighted approval methods can be easily explained to voters, but they are too restrictive to achieve proportional outcomes.

information on each period  $1, \dots, t - 1$ , that allows us to produce the correct outcome at period  $t$ . The *compilation complexity* of a voting rule measures the smallest amount of information that has to be kept regarding a partial set of votes, in order to compute the voting outcome given the entire set of votes. Compilation complexity has been studied for common single-winner voting rules [21, 57] and multi-winner voting rules [29]. Here, we are interested in a slightly different concept: the amount of information that has to be kept in order to compute an outcome that guarantees fair representation — not necessarily the outcome of a specific rule. This concept can be called “compilation complexity of fair representation”.

**Research Objective 7.** (a) Develop PV algorithms satisfying approximate fair representation properties when the ballots in each time-period are entirely anonymous (no compiled information is saved, except the public vote counts). (b) Compute the amount of information that has to be saved in each period, in order to guarantee fair representation in the following periods.


### 3.3.1 Direct election of ministers in the government

Page et al. [43] have introduced a novel application, closely related to PV: allowing citizens to directly vote for the ministers. Such direct elections may potentially increase the citizens’ trust in the government and their willingness to cooperate with it. It is similar to PV in that multiple issues (= ministries) have to be decided by the same group of voters, and therefore it is prone to the same unfairness problem illustrated in Example 1. The differences are technical: the sets of candidates for each issues are pairwise-disjoint (in the model of Page et al. [43]), each agent must vote for exactly one candidate in each office, and all decisions are made simultaneously. Page et al. [43] extended the fair representation notions to this setting, presented a novel algorithm called GreedyPAV, and showed via simulations that it guarantees fair representation to minorities.


In an undergraduate student project [37], we set out to check how the GreedyPAV algorithm [43] fares with real Israeli voters. We conducted an opinion poll with 600 subjects, who are a representative sample of Israeli voters.<sup>5</sup> We focused on the 12 offices with the highest budget in the Israeli government, as of that time. We chose 4 potential candidates for each office, such as past and present ministers, and other politicians whose activity is related to the office. We then asked each participant to vote for a single candidate for each office. We also explained the new method to the participants, and asked them whether they prefer the existing or the new method. The new

---

<sup>5</sup>The poll was conducted at 14.8.2022 by Panel4All — an established Israeli polling company.



(a) Answers to “which method for electing cabinet ministers do you prefer?” (%).



(b) Answers to “are you satisfied with the government elected by GreedyPAV?” (%).

Figure 1

method won a clear majority (see Figure 1a). In a follow-up poll among the same participants, we showed them the hypothetical government that would be formed by the GreedyPAV algorithm applied to the first poll data, and asked them whether they are satisfied with this hypothetical government. Surprisingly, more people were dissatisfied than satisfied (see Figure 1b)!<sup>6</sup>

To better understand the reasons for the dissatisfaction, we examined the verbal arguments provided by the subjects to their dissatisfaction. Prevalent arguments were “I do not want person X in the government at all”, or “person Y must not be in office Z”. We checked why these “unwanted” candidates were elected, and found out that they represented substantial minority groups. In other words, the GreedyPAV algorithm did exactly what it was designed to do — it computed a government that guarantees fair representation! The advantage of the algorithm is also its disadvantage — guaranteeing fair representation to minorities might alienate the majority, who do not like these particular representatives.

Various authors have studied extensions of approval voting to *trichotomous preferences*, which allow the voters to cast negative votes for candidates they do not like, in addition to casting positive votes for candidates they like. See Baumeister and Dennisen [7], Zhou et al. [59] and section 6.2 in Lackner and Skowron [35] for a recent survey. Recently, Freeman et al. [25], Talmon and Page [56] and Masařík et al. [41] suggested several extensions of the fair representation axioms to trichotomous preferences. Still, “many questions regarding the trichotomous model remain unanswered” [35]. In particular, fair representation in direct election of ministers has not been studied with trichotomous ballots.

<sup>6</sup>In this preliminary poll, we made a mistake: we did not ask the subjects whether they prefer the hypothetical government over the real government. It is possible that people are dissatisfied with the hypothetical government, but even more dissatisfied with the real government. Thus, we do not view the poll’s result as an argument against the GreedyPAV algorithm; we only view it as an indication for a potential direction for improvement.

**Research Objective 8.** (a) Develop an algorithm for direct election of ministers, that ensures fair representation, while taking into account negative votes.<sup>7</sup> (b) Perform extensive experiments with real voters, to assess the voter satisfaction from the resulting hypothetical governments.

### 3.4 Dynamic ranking of public issues

An opinion poll from 2009 showed that 63% of the Israeli citizens think that politicians do not care about the citizens' opinions; only 18% think that they have some influence on the government actions.<sup>8</sup> In order to improve trust and participation, it is important to enable direct communication channels between politicians and citizens. There already exist various online applications,<sup>9</sup> that allow meeting organizers to engage their audience through questions and answers: typically, people can ask questions and vote for questions asked by others; the meeting organizer sees a ranking of the questions by the number of votes, and can choose which question to answer. It is easy to adapt these platforms for politicians who want to engage their citizens by letting them ask questions or suggest issues that require action.

A main problem with these systems is the lack of fair representation. As in Example 1, suppose 60% of the population are interested in issues 1, ..., 10 whereas 40% are interested in issues 11, ..., 20. The existing systems will probably show the 10 issues supported by the majority at the top, while the 10 issues supported by the minority might be pushed to the bottom and ignored.

Recently, Israel and Brill [27] extended the notions of fair representation from committee election to ranking. They focused particularly on the dynamic nature of the ranking in such Q&A platforms: when the meeting organizer (the politician) chooses an issue to reply to, the ranking may change due to the fair representation requirements. They present four adaptations of existing committee election rules to the dynamic ranking setting. Their algorithms provide a promising route for improving communication and cooperation between politicians and citizens. However, making these algorithms practically useful poses several challenges.

One of these challenges is *strategic manipulation*. It is well-known that all rules for committee election that guarantee fair representation are susceptible to manipulation by strategic voters. In fact, Peters [44] proved the following strong impossibility result: every committee election rule

---

<sup>7</sup>Another potential issue with direct election of government ministers is that some ministers do not get along with each other. This imposes some constraints on the allocation, e.g. "Person X cannot be elected if person Y is elected". Such constraints were recently studied by Brill et al. [16].

<sup>8</sup>See <https://www.idi.org.il/parliaments/3467/8205>; retrieved 31/10/23.

<sup>9</sup>For example: [slido.com](https://www.slido.com), [mentimeter.com](https://www.mentimeter.com), [pigeonholelive.com](https://www.pigeonholelive.com).

that guarantees a very weak variant of fair representation and a very weak variant of efficiency is susceptible to *free-riding manipulation*. That is, voters who approve a set of candidates have an incentive to remove from their approval-set some candidates who are going to win anyway, in order to pretend to be worse-off than they actually are. This induces the rule to “compensate” the manipulator by electing more of their approved candidates.

The theoretic results on strategic manipulation usually assume that the manipulator has complete information on the actions of all other players. In most settings, this assumption is not realistic. Moreover, without complete information, the manipulation might make the manipulator worse-off. Therefore, in practice, the problem of strategic voting might not be as harmful as theory predicts. However, in our envisioned Q& A system, strategic voting might be particularly prevalent, as users see the voting status of all issues in real-time. Thus, it is important to find practical solutions that could mitigate the free-riding problem.

A potential solution is suggested by the empirical works of Barrot et al. [6], Lackner and Skowron [34] and Lackner et al. [36]. Each work studies a family of committee election rules (based on ordered-weighted averaging or Thiele’s rules). The rules in their family are parameterized by a real parameter  $p$ . They show, using simulation experiments, that a larger  $p$  makes the rule fairer to minorities, whereas a smaller  $p$  makes the rule less prone to strategic manipulations (a smaller fraction of randomly-generated profiles can be manipulated). Thus, by adjusting the value of  $p$ , one can control the trade-off between fairness and strategyproofness. However, all of these works do not present a formal proof of the fairness-strategyproofness trade-off. Moreover, they only consider specific sub-families of rules, which do not contain important rules such as Phragmén’s rule [15] or the Method of Equal Shares. Another potential direction is to control the amount of information shown to the users: instead of showing the accurate number of votes for each issue, the platform can show only approximate figures (e.g. rounded to the nearest hundred). This increases the risk faced by a strategic voter, and may reduce the incentive to manipulate.

**Research Objective 9.** (a) Design parameterized families of voting-rules based on Phragmén and the Method of Equal Shares, as well as ordered-weighted averaging and Thiele, with a *provable* trade-off between fairness and strategyproofness.

(b) Evaluate these families by conducting user-studies with users of an online Q&A platform.<sup>10</sup>

---

<sup>10</sup>Recently, Halpern et al. [26] studied another important challenge in Q&A platforms: most users vote only for a small number of issues, so platform knows only a small subset of the users’ preferences.

## References

- [1] Haris Aziz and Barton E Lee. The expanding approvals rule: improving proportional representation and monotonicity. *Social Choice and Welfare*, 54(1):1–45, 2020.
- [2] Haris Aziz and Barton E Lee. Proportionally representative participatory budgeting with ordinal preferences. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI-2021)*, pages 5110–5118, 2021.
- [3] Haris Aziz and Nisarg Shah. Participatory budgeting: Models and approaches. *Pathways Between Social Science and Computational Social Science: Theories, Methods, and Interpretations*, pages 215–236, 2021.
- [4] Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48: 461–485, 2017.
- [5] Haris Aziz, Barton E. Lee, and Nimrod Talmon. Proportionally representative participatory budgeting: Axioms and algorithms. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2018)*, pages 23–31, 2017.
- [6] Nathanaël Barrot, Jérôme Lang, and Makoto Yokoo. Manipulation of hamming-based approval voting for multiple referenda and committee elections. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS-2017)*, pages 597–605, 2017.
- [7] Dorothea Baumeister and Sophie Dennisen. Voter dissatisfaction in committee elections. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2015)*, pages 1707–1708, 2015.
- [8] Xiaohui Bei, Zihao Li, Jinyan Liu, Shengxin Liu, and Xinhang Lu. Fair division of mixed divisible and indivisible goods. *Artificial Intelligence*, 293:103436, 2021.
- [9] Samuel Bismuth, Vladislav Makarov, Erel Segal-Halevi, and Dana Shapira. Number partitioning with splitting, 2022. arXiv preprint 2204.11753.
- [10] Florian Brandl, Felix Brandt, Dominik Peters, and Christian Stricker. Distribution rules under dichotomous preferences: two out of three ain’t bad. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC-2021)*, pages 158–179, 2021.
- [11] Florian Brandl, Felix Brandt, Matthias Greger, Dominik Peters, Christian Stricker, and Warut Suksompong. Funding public projects: A case for the nash product rule. *Journal of Mathematical Economics*, 99:102585, 2022.
- [12] Felix Brandt, Matthias Greger, Erel Segal-Halevi, and Warut Suksompong. Balanced donor coordination. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC-2023)*, 2023.
- [13] Markus Brill and Jannik Peters. Robust and verifiable proportionality axioms for multiwinner voting. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC-2023)*, 2023.

- 
- [14] Markus Brill, Stefan Forster, Martin Lackner, Jan Maly, and Jannik Peters. Proportionality in approval-based participatory budgeting. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-2023)*, pages 5524–5531, 2023.
- [15] Markus Brill, Rupert Freeman, Svante Janson, and Martin Lackner. Phragmén’s voting methods and justified representation. *Mathematical Programming*, pages 1–30, 2023.
- [16] Markus Brill, Evangelos Markakis, Georgios Papatotiropoulos, and Jannik Peters. Proportionality guarantees in elections with interdependent issues. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI-2023)*, pages 2537–2545, 2023.
- [17] Laurent Bulteau, Noam Hazon, Rutvik Page, Ariel Rosenfeld, and Nimrod Talmon. Justified representation for perpetual voting. *IEEE Access*, 9:96598–96612, 2021.
- [18] Vitalik Buterin, Zoë Hitzig, and E. Glen Weyl. A flexible design for funding public goods. *Management Science*, 65:5171–5187, 2019.
- [19] Olivier Cailloux and Ulle Endriss. Arguing about voting rules. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2016)*, pages 287–295, 2016.
- [20] Nikhil Chandak, Shashwat Goel, and Dominik Peters. Proportional aggregation of preferences for sequential decision making, 2023. arXiv preprint arXiv:2306.14858.
- [21] Yann Chevaleyre, Jérôme Lang, Nicolas Maudet, and Guillaume Ravilly-Abadie. Compiling the votes of a subelectorate. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI-2009)*, pages 97–102, 2009.
- [22] Vincent Conitzer, Rupert Freeman, and Nisarg Shah. Fair public decision making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC-2017)*, pages 629–646, 2017.
- [23] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. Multiwinner voting: A new challenge for social choice theory. *Trends in computational social choice*, 74(2017): 27–47, 2017.
- [24] Rupert Freeman, Seyed Majid Zahedi, and Vincent Conitzer. Fair social choice in dynamic settings. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI-2017)*, pages 4580–4587, 2017.
- [25] Rupert Freeman, Anson Kahng, and David M Pennock. Proportionality in approval-based elections with a variable number of winners. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI-2020)*, pages 132–138, 2020.
- [26] Daniel Halpern, Gregory Kehne, Ariel D Procaccia, Jamie Tucker-Foltz, and Manuel Wüthrich. Representation with incomplete votes. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-2023)*, pages 5657–5664, 2023.
- [27] Jonas Israel and Markus Brill. Dynamic proportional rankings. In *Proceedings of the the 30th International Joint Conference on Artificial Intelligence (IJCAI-2021)*, pages 261–267, 2021.

- [28] Ido Kahana and Noam Hazon. The leximin approach for a sequence of collective decisions. In *Proceedings of the 26th European Conference on Artificial Intelligence (ECAI-2023)*, pages 1198–1206, 2023.
- [29] Neel Karia and Jérôme Lang. Compilation complexity of multi-winner voting rules (student abstract). In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI-2021)*, pages 15809–15810, 2021.
- [30] Ralph L Keeney and Howard Raiffa. *Decisions with multiple objectives: preferences and value trade-offs*. Cambridge university press, 1993.
- [31] Kris N Kirby. An empirical assessment of the form of utility functions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37(2):461, 2011.
- [32] Martin Lackner. Perpetual voting: Fairness in long-term decision making. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI-2020)*, pages 2103–2110, 2020.
- [33] Martin Lackner and Jan Maly. Proportional decisions in perpetual voting. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-2023)*, pages 5722–5729, 2023.
- [34] Martin Lackner and Piotr Skowron. Approval-based multi-winner rules and strategic voting. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI-2018)*, pages 340–346, 2018.
- [35] Martin Lackner and Piotr Skowron. *Multi-winner voting with approval preferences*. Springer, 2023.
- [36] Martin Lackner, Jan Maly, and Oliviero Nardi. Free-riding in multi-issue decisions. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2023)*, pages 2040–2048, 2023.
- [37] Itai Lashover, Liav Weiss, Amichai Kafka, and Shoshana Levin. Election of government ministers, 2022. arXiv preprint 2210.08985.
- [38] Maaïke Los, Zoé Christoff, and Davide Grossi. Proportional budget allocations: Towards a systematization. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI-2022)*, pages 398–404, 2022.
- [39] Xinhang Lu, Jannik Peters, Haris Aziz, Xiaohui Bei, and Warut Suksompong. Approval-based voting with mixed goods. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-2023)*, pages 5781–5788, 2023.
- [40] Jan Maly, Simon Rey, Ulle Endriss, and Martin Lackner. Fairness in participatory budgeting via equality of resources. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2023)*, pages 2031–2039, 2023.
- [41] Tomáš Masařík, Grzegorz Pierczyński, and Piotr Skowron. A generalised theory of proportionality in collective decision making, 2023. arXiv preprint 2307.06077.
- [42] Meir Nizri, Amos Azaria, and Noam Hazon. Improving the perception of fairness in shapley-based allocations. In *Proceedings of the 44th Annual Conference of the Cognitive Science Society (COGSCI-23)*, pages 2285–2291, 2022.

- 
- [43] Rutvik Page, Ehud Shapiro, and Nimrod Talmon. Electing the executive branch, 2020. arXiv preprint 2009.09734.
- [44] Dominik Peters. Proportionality and strategyproofness in multiwinner elections. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS-2018)*, pages 1549–1557, 2018.
- [45] Dominik Peters and Piotr Skowron. Proportionality and the limits of welfarism. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC-2020)*, pages 793–794, 2020.
- [46] Dominik Peters, Ariel D Procaccia, Alexandros Psomas, and Zixin Zhou. Explainable voting. In *Advances in Neural Information Processing Systems (NeurIPS-2020)*, pages 1525–1534, 2020.
- [47] Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron. Proportional participatory budgeting with additive utilities. In *Advances in Neural Information Processing Systems (NeurIPS-2021)*, pages 12726–12737, 2021.
- [48] Juan M Ramos-Goñi, Mark Oppe, Elly Stolk, Koonal Shah, Simone Kreimeier, Oliver Rivero-Arias, and Nancy Devlin. International valuation protocol for the EQ-5D-Y-3L. *Pharmacoeconomics*, 38:653–663, 2020.
- [49] Simon Rey and Jan Maly. The (computational) social choice take on indivisible participatory budgeting, 2023. arXiv preprint 2303.00621.
- [50] Ariel Rosenfeld and Nimrod Talmon. What should we optimize in participatory budgeting? an experimental study. *arXiv preprint 2111.07308*, 2021.
- [51] Luis Sánchez-Fernández, Edith Elkind, Martin Lackner, Norberto Fernández, Jesús Fisteus, Pablo Basanta Val, and Piotr Skowron. Proportional justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI-2017)*, pages 670–676, 2017.
- [52] Fedor Sandomirskiy and Erel Segal-Halevi. Efficient fair division with minimal sharing. *Operations Research*, 70(3):1762–1782, 2022.
- [53] Anwar Shah. *Participatory budgeting*. World Bank Publications, 2007.
- [54] Piotr Skowron. Proportionality degree of multiwinner rules. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC-2021)*, pages 820–840, 2021.
- [55] Gogulapati Sreedurga. Participatory budgeting with multiple degrees of projects and ranged approval votes. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI-2023)*, pages 2870–2877, 2023.
- [56] Nimrod Talmon and Rutvik Page. Proportionality in committee selection with negative feelings, 2021. arXiv preprint 2101.01435.
- [57] Lirong Xia and Vincent Conitzer. Compilation complexity of common voting rules. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence (AAAI-2010)*, pages 915–920, 2010.

- [58] Joshua C Yang, Carina I Hausladen, Dominik Peters, Evangelos Pournaras, Regula Haenggli Fricker, and Dirk Helbing. Designing digital voting systems for citizens: Achieving fairness and legitimacy in digital participatory budgeting. *arXiv preprint arXiv:2310.03501*, 2023.
- [59] Aizhong Zhou, Yongjie Yang, and Jiong Guo. Parameterized complexity of committee elections with dichotomous and trichotomous votes. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS-2019)*, pages 503–510, 2019.