

# Participatory Budgeting: Models and Approaches

Haris Aziz and Nisarg Shah

**Abstract** Participatory budgeting is a democratic approach to deciding the funding of public projects, which has been adopted in many cities across the world. We present a survey of research on participatory budgeting emerging from the computational social choice literature, which draws ideas from computer science and microeconomic theory. We present a mathematical model for participatory budgeting, which charts existing models across different axes including whether the projects are treated as “divisible” or “indivisible” and whether there are funding limits on individual projects. We then survey various approaches and methods from the literature, giving special emphasis on issues of preference elicitation, welfare objectives, fairness axioms, and voter incentives. Finally, we discuss several directions in which research on participatory budgeting can be extended in the future.

## 1 Introduction

Participatory budgeting (PB) is a paradigm which empowers residents to directly decide how a portion of the public budget is spent. Specifically, residents deliberate over spending priorities and vote over how the budget should be allocated to different public projects. Projects which receive broad support from the community are then funded through the process.

This process was initiated as a radical democratic project in the city of Porto Alegre, Brazil, led by the leftist Workers’ Party (Cabannes, 2004). Over the next few decades, it quickly spread across the world; it has been implemented by over 1,500 municipalities (Röcke, 2014) and in locations as diverse as Guatemala, Peru, Romania and South Africa (Shah, 2007). The nonprofit organization *Participatory*

---

Haris Aziz  
UNSW Sydney and Data61 CSIRO, Australia, e-mail: [haris.aziz@unsw.edu.au](mailto:haris.aziz@unsw.edu.au)

Nisarg Shah  
University of Toronto, e-mail: [nisarg@cs.toronto.edu](mailto:nisarg@cs.toronto.edu)

*Budgeting Project*<sup>1</sup> alone has helped allocate more than US\$300 million of public budget in 29 cities across the US and Canada. In Europe, the effort is led by Paris and Madrid, each spending at least \$100 million a year on public projects through participatory budgeting (Legendre et al., 2017; Gutiérrez, 2017). In particular, Madrid developed an online open-source platform, *Decide Madrid*, which has been used by more than 30 local governments (Samir, 2018). PB is still spreading in more regions across the world. For example, Toronto recently completed its three-year pilot of PB (Murray, 2019), and the state of New South Wales in Australia recently started PB under the name *My Community Project*.<sup>2</sup>

Participatory budgeting is typically a long process; in many municipalities, one PB cycle takes one full year. While the exact implementation details differ from one PB instance to another, at a high level the process is composed of the following stages, which allow residents to come up with effective project proposals and provide their preferences over budget allocation (Shah, 2007).

- First, the geographical region may be divided into several subregions (e.g. districts), and one PB process may be conducted in each district separately. The goal of this stage is to allow residents to focus on the projects in their own neighborhood and community. The total available budget for each district is also typically decided at this stage.
- Next, residents share and discuss ideas through neighborhood meetings and online tools. This allows them to come up with preliminary project proposals.
- These initial proposals are then developed into feasible proposals by focus groups and vetted by experts. Often, this is the stage where a proposal may be broken into distinct stages of implementation, and a cost estimate may be derived for each implementation stage. Projects may also be classified into categories such as recreation, infrastructure, health, education, transportation, etc.
- This may be followed by several rounds of deliberation through which residents finalize a small number of proposals to be included in the final vote.
- The final stage of the PB process is the voting stage. In this stage, eligible residents vote over how the public budget should be spent across the finalized proposals, and these votes are aggregated into a final budget allocation.

Additional effort is often required to make PB a success. For example, advertisement and promotion through various channels, including social media, can be the key to increasing civic engagement. This can be crucial in encouraging various minorities to participate and ensuring that their preferences are incorporated into the decision-making. See the edited volume by Shah (2007) for a detailed survey of the entire PB process.

In this chapter, we limit our attention to the final *voting* stage of the process. That is, we assume that the project proposals have been detailed and finalized, and projects to be included on the final ballot have been filtered. We also assume that

---

<sup>1</sup> <http://www.participatorybudgeting.org>

<sup>2</sup> <https://www.nsw.gov.au/improving-nsw/projects-and-initiatives/my-community-project/>

project costs have been estimated and the total available budget is known. Research on the voting stage of the PB process focuses on four important considerations:

- (i) Decision space: Is the space of possible outcomes discrete (e.g. because a project can only be “funded” or “not funded”)? Or is it continuous (e.g. because a project can be allocated different amounts of funds to implement it to different degrees of effectiveness)?
- (ii) Preference modeling: How will residents’ preferences over the projects be represented for the purpose of mathematical analysis?
- (iii) Ballot design (aka preference elicitation): It is often desirable for the modeling to allow complex preferences. However, it may be infeasible to ask residents to report such complex preferences. What should the ballot ask from residents which will serve as a proxy for their preferences and allow them to effectively convey their preferences?
- (iv) Vote aggregation: How will the votes cast by residents be aggregated into a final allocation of the available public budget? How should this process fairly incorporate aggregation the preferences of different communities of residents and efficiently allocate the public budget?

The need for systematic modeling and study of these considerations has recently inspired a body of research in the computational social choice literature Brandt et al. (2016); Endriss (2017). This body of research sits at the intersection of computer science, social science, and economics, and aims to use mathematical modeling and algorithmic techniques to design different approaches to the voting stage of the PB process.<sup>3</sup> The goal of this chapter is to present a comprehensive overview of the research on PB in computational social choice. Some of the original contributions of this chapter include providing a unifying theoretical framework which allows viewing the different research works through a single lens, and providing a taxonomy of the different PB models which highlights their unique modeling choices.

*Outline:* The organization of this chapter is as follows. In Section 2, we present a general mathematical formulation of PB, and list several prominent features that distinguish different implementations of PB. Next, we focus on popular PB models that make specific design choices in terms of these features, and present a taxonomy of these models. We pay special attention to the representation and elicitation of preferences, and popular desiderata underlying the vote aggregation methods.

In Section 3, we survey the research on the “integral model” of PB, in which each project can only be fully funded or not funded. In Section 4, we survey the research on the “continuous model” of PB, in which the funding level of each project can lie on a continuum. We highlight the differences between motivations and results

---

<sup>3</sup> In several PB programs, the final decision of budgets is entirely done by deliberation by a small subset of residents. Deliberation has its advantages and disadvantages. While it allows an in-depth discussion of possible outcomes from different perspectives, it also poses the danger of the process being dominated by an unrepresentative subset of the residents, thus thwarting the democratic objective.

under both discrete and continuous models. Finally, in Section 5, we discuss possible extensions and directions for future research.

## 2 Mathematical Formulation

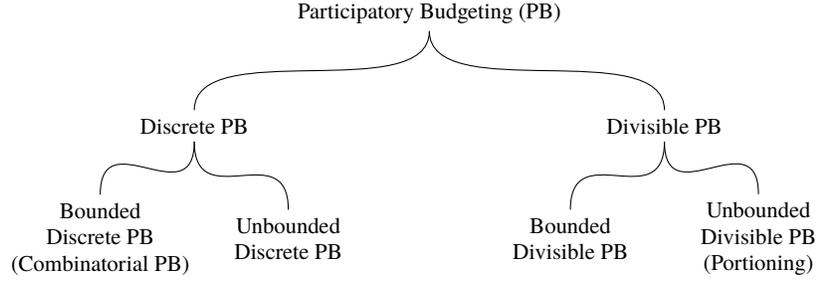
We begin by reviewing various parameters which play a key role in formulating the participatory budgeting problem. Later, we consider specific choices of these parameters, which give rise to popular PB models.

- *Residents*: There is a set of residents (a.k.a. agents or voters)  $N = \{1, \dots, n\}$ . In some applications of PB, residents are divided across different geographical regions (e.g. districts or wards), and each region conducts a separate PB election, in which residents of that region vote.
- *Resources and budgets*: There is a set of  $d$  resources, denoted  $R$ . The available budget is  $\mathbf{B} = (B_1, \dots, B_d)$ , where  $B_r$  denotes the amount of resource  $r$  available. In most applications of PB, money is the only limiting resource. Thus, much of the work on PB in computational social choice has focused on the case of a single resource, in which case the available budget is denoted by the scalar  $B$ .
- *Projects*: There is a set of projects  $P = \{p_1, \dots, p_m\}$ . In some applications, projects are categorized into various domains (e.g. infrastructure, security, education, health & wellness, etc).
- *Degree of completion*: Some projects can be completed to different degrees. For example, completing 65% of a project proposing neighborhood clean-up may mean a part of the neighborhood being cleaned up. Some projects may have discrete milestones, and only the first few milestones may be achieved; e.g., a project proposing the creation of a public park may include construction of the park as the first milestone and construction of a children’s playground within the park as the second milestone. For other projects, the degree of completion may be binary; they must be either fully implemented or not implemented at all (e.g. installing a fountain in a park). For project  $p$ , let  $\mathcal{X}_p$  denote the set of its possible degrees of completion and  $x_p$  denote its actual degree of completion in an outcome. In some models of PB, which we refer to as *bounded models*, there is an upper bound (a.k.a. *cap*) on the degree of completion of each project  $p$ , denoted by  $q_p$ . In *unbounded models*, there are no such caps.<sup>4</sup> We discuss popular choices of  $\mathcal{X}_p$  in Section 2.1. We assume  $0 \in \mathcal{X}_p$  for all projects  $p$ . Let  $\mathcal{X} = \mathcal{X}_{p_1} \times \dots \times \mathcal{X}_{p_m}$ . Various models of PB studied in the literature crucially differ in this key choice; we elaborate on this in Section 2.1.
- *Costs*: Each project  $p \in P$  has an associated cost function  $c_p : \mathcal{X}_p \rightarrow \mathbb{R}^d$ , where the vector  $c_p(x_p)$  gives the amount of each resource required to complete project  $p$  to degree  $x_p$ . We assume that  $c_p(0) = \mathbf{0}$  and  $c_p$  is monotonically non-decreasing.

---

<sup>4</sup> Note that the cost function of a project and the total available budget may still induce an effective upper bound on its degree of completion.

- *Budget Allocations*: An outcome  $\mathbf{x} = (x_p)_{p \in P} \in \mathcal{X}$  is characterized by the degree of completion of each project  $p$ . Note that this also specifies the amounts of different resources that will be devoted to each project. The outcome is feasible if  $\sum_{p \in P} c_p(x_p) \leq \mathbf{B}$ , where addition is component-wise. We refer to feasible outcomes as *(budget) allocations*. Let  $\mathcal{A}$  denote the space of allocations.
- *Resident Preferences*: Each resident  $i \in N$  has preferences over which projects should be implemented and to what degree. These can be represented through an ordinal preference relation  $\succsim_i$  or a cardinal utility function  $u_i$  over the space of allocations  $\mathcal{A}$ . We elaborate on this later in Section 2.2.1.



**Fig. 1** A taxonomy of participatory budgeting models.

## 2.1 Decision Space and Popular PB Models

Figure 1 presents a taxonomy of the popular PB models in the literature, which crucially differ in their modeling of the possible degrees of completion of projects. We review each model in further detail.

*Bounded Discrete PB (Combinatorial PB)*: This is perhaps the most widely studied and applied model of PB. In this model, projects must be either fully implemented or not implemented at all. Hence, the set of possible degrees of completion is  $\mathcal{X}_p = \{0, 1\}$  for each project  $p \in P$ . Note that this model has unit caps ( $q_p = 1$ ). Consequently, the cost function  $c_p$  of project  $p$  is effectively a vector, which indicates the amounts of different resources needed to fully implement project  $p$ . Feasible allocations in  $\mathcal{A}$  are subsets of projects which can be implemented simultaneously subject to budget constraints. In essence, this is a multi-agent variant of a multi-dimensional knapsack problem.

*Discrete PB*: In this model, each project  $p$  has discrete possible degrees of completion. However, unlike combinatorial PB, there may be more than two possible

degrees of completion. These degrees of completion can also mathematically capture funding of multiple of units of the same project. For example installing 10 units of public toilets can be viewed as having a single project with 10 different degrees of completion. Typically, we assume that  $\mathcal{X}_p = \mathbb{N} \cup \{0\}$  for each project  $p$ . Note that this model has no cap on the degree of completion, but limited availability of resources may still place a natural upper bound on the degree of completion of each project in any feasible allocation. This model is applicable when each project is ambitious, and its full implementation can potentially use up all the available resources.

*Divisible PB:* In this model, it is assumed that projects can be implemented to any fractional degree. In the version with caps, we can assume without loss of generality that the cap is  $q_p = 1$ , i.e.,  $\mathcal{X}_p = [0, 1]$  for each project  $p \in P$ ; in this case,  $x_p$  denotes the fraction of project  $p$  that is completed. In practice, projects often cannot be executed to arbitrary fractional degrees. However, in cases of sufficient granularity, assuming divisible PB can help for computational reasons; more details are provided in Section 4.

*Unbounded Divisible PB (Portioning):* In this model of divisible PB, there are no caps on the degrees of completion. Thus,  $\mathcal{X}_p = \mathbb{R}_+$  for each project  $p$ . When there is a single resource type (e.g. money), under mild assumptions<sup>5</sup>, deciding the degrees of completion of different projects is equivalent to deciding how the available budget will be divided among the projects. This setting is known as *portioning* in the literature (Airiau et al., 2019).

### Example

Suppose there is a set of 3000 residents  $N = \{1, \dots, 3000\}$ , a set of four projects  $P = \{A, B, C, D\}$ , and a single resource (money) with a total budget of \$7 million. Suppose the cost functions of the projects are as follows.

- $c_A(x_A) = \$3 \text{ million} \times x_A$
- $c_B(x_B) = \$3 \text{ million} \times x_B$
- $c_C(x_C) = \$2 \text{ million} \times x_C$
- $c_D(x_D) = \$2 \text{ million} \times x_D$

Suppose 2000 residents like projects  $A$  and  $B$  equally, but derive no value from  $C$  or  $D$ . 500 residents derive value only from  $C$ , and the remaining 500 only from  $D$ .

In the divisible PB model with unit caps (i.e.  $\mathcal{X}_p = [0, 1]$  for each project  $p$ ), we have numerous choices. For instance, we could implement  $7/10$  fraction of each project. Or we could implement  $1/2$  fraction of  $A$  and  $B$  each, and fully implement  $C$  and  $D$ .

In the combinatorial PB model (i.e.  $\mathcal{X}_p = \{0, 1\}$  for each project  $p$ ), given the budget of \$7 million, we can implement  $A$  and  $B$ , which would make 2000 residents

<sup>5</sup> E.g., if the cost function of each project is strictly increasing, then there is a one-to-one correspondence between the amount of money allocated to the project and its degree of completion.

very happy but 1000 residents very unhappy, or we can implement one of  $A$  and  $B$  together with both  $C$  and  $D$ , which would make 2000 residents partially happy and 1000 residents very happy.

---

How do we quantify how happy the residents are? How do we make tradeoffs between such decisions? To understand this better, we need to understand modeling of resident preferences and goals of the PB process, which we review below.

## 2.2 Preference Modeling and Ballot Design

Two important decisions in the PB process are the modeling of residents' preferences (which helps in the mathematical analysis of the effectiveness of the process) and the format in which the residents cast their votes.

### 2.2.1 Preference Modeling

Recall that  $\mathcal{A}$  is the set of feasible allocations. In an expressive model, each resident  $i$  has a cardinal utility function  $u_i : \mathcal{A} \rightarrow \mathbb{R}$ .<sup>6</sup> If only comparisons among allocations are required, one may work instead with an ordinal preference relation  $\succsim_i$  over  $\mathcal{A}$ .<sup>7</sup>

While such an expressive modeling allows capturing all possible preferences, it does not utilize the structure that is often present in the preferences. Below, we discuss three common approaches to modeling such structure.

The first approach is to directly impose a structural assumption on the utility function or the preference relation. For example, in combinatorial PB, where allocations are effectively feasible subsets of projects and a utility function can be represented as  $u_i : 2^P \rightarrow \mathbb{R}$ , one may impose that  $u_i$  satisfies subadditivity or submodularity (when projects are substitutes of each other), or superadditivity or supermodularity (when projects are complements of each other).

The second approach is to use spatial models, where allocations are embedded in a metric space, each resident has a preferred allocation, and her utility for another allocation is a (typically non-increasing) function of its distance to her preferred allocation. For example, Garg et al. (2019) study  $\ell_p$  normed utilities, under which the utility of resident  $i$  (with preferred allocation  $\mathbf{x}_i^*$ ) for allocation  $\mathbf{x}$  is  $u_i(\mathbf{x}) = -\|\mathbf{x} - \mathbf{x}_i^*\|_p$ , where  $\|\cdot\|_p$  is the  $\ell_p$  norm.

---

<sup>6</sup> We note that in this case, much of the existing research on PB in computational social choice, inspired from classical research on voting, implicitly assumes that  $u_i(\mathbf{0}) = 0$  and  $u_i$  satisfies monotonicity (i.e.  $u_i(\mathbf{x}) \geq u_i(\mathbf{x}')$  for  $\mathbf{x} \geq \mathbf{x}'$ ). In words, implementing more projects or implementing projects to a greater degree cannot bring less happiness. In the PB context, implementing projects requires spending costly resources, which makes this assumption questionable. See Section 5 for further discussion.

<sup>7</sup> The ordinal preference relation is typically assumed to be transitive and complete.

The third — and perhaps the most popular — approach is to model residents' preferences over individual projects and use a natural rule to extend them to preferences over allocations. Below, we review common ways to achieve this for both cardinal utility functions and ordinal preference relations.

*Cardinal extensions:* Fain et al. (2016a) study the class of *scalar separable* utility functions, whereby resident  $i$  derives utility  $u_{i,p} \cdot f_p(x_p)$  from each project  $p$ ,<sup>8</sup> and her utility for an allocation  $\mathbf{x}$  is simply additive across projects, i.e.,  $u_i(\mathbf{x}) = \sum_{p \in P} u_{i,p} \cdot f_p(x_p)$ . This is a reasonable assumption when the different project proposals are completely independent of each other. In case of combinatorial PB, where  $x_p \in \{0, 1\}$  for each project  $p$ , this effectively reduces to  $u_i(S) = \sum_{p \in S} u_{i,p}$  for each  $S \subseteq P$ ; this is commonly known as an *additive* utility function (Benade et al., 2017; Fain et al., 2018; Bhaskar et al., 2018). A further restriction of  $u_{i,p} \in \{0, 1\}$  for each resident  $i$  and project  $p$  gives rise to *dichotomous* preferences (Aziz et al., 2018; Faliszewski and Talmon, 2019), under which each resident approves or disapproves each project and only cares about the number of her approved projects that are implemented. In combinatorial PB, another common extension is the *max set extension* (Caragiannis et al., 2017), whereby the utility of a resident for an allocation is her utility for her most favorite project that is funded:  $u_i(S) = \max_{p \in S} u_{i,p}$  for each  $S \subseteq P$ . This is applicable when the projects are substitutes of each other, and the resident will derive value from only one of the implemented projects.

*Ordinal extensions:* When residents are assumed to have ordinal preferences over projects, we use the following notation. Each resident  $i$  has a weak order preference relation  $\succsim_i$  over  $P$ . We denote by  $\succ_i$  the strict part and by  $\sim_i$  the indifference part of the relation  $\succsim_i$ . We denote by  $E_i^1, \dots, E_i^{k_i}$  the equivalence classes of the relation  $\succsim_i$  in decreasing order of preferences; thus, each set  $E_i^j$  is a maximal equivalence class of objects among which agent  $i$  is indifferent, and  $k_i$  is the number of equivalence classes. Given an ordinal preference relation  $\succsim_i$ , one can induce resident  $i$ 's preference relation over the set of allocations  $\mathcal{A}$  in several natural ways.

The (first order) stochastic dominance (SD) extension is defined as follows (see, e.g., Brandl et al. (2016)). For allocations  $\mathbf{x}, \mathbf{y} \in \mathcal{A}$ , we say that  $\mathbf{x} \succsim_i^{SD} \mathbf{y}$  iff  $\sum_{j=1}^l \sum_{p \in E_i^j} x_p \geq \sum_{j=1}^l \sum_{p \in E_i^j} y_p$  for all  $l \in \{1, \dots, k_i\}$ . Since this extension requires adding the degrees of completion of different projects, it makes more sense in models where the degrees of completion of different projects are comparable (e.g. in combinatorial PB or divisible PB with unit caps). The desirable aspect of the SD extension is that  $\mathbf{x} \succsim_i^{SD} \mathbf{y}$  is equivalent to saying that  $\mathbf{x}$  yields at least as much utility to resident  $i$  as  $\mathbf{y}$  does, for all additive utilities over projects that are consistent with preference relation  $\succsim_i$ . However, the SD relation is not complete (i.e. it does not compare every pair of allocations).

---

<sup>8</sup> Note that  $f_p : \mathcal{X}_p \rightarrow \mathbb{R}$  is a resident-independent function. Typically, it is assumed to be non-decreasing, and when  $\mathcal{X}_p$  is continuous, smooth and concave as well.

One popular complete extension is the lexicographic extension  $\succsim_i^{lex}$ , under which resident  $i$  is assumed to care significantly more about project  $p$  than about project  $p'$  whenever  $p \succ_i p'$ . Formally,  $\mathbf{x} \succ_i^{lex} \mathbf{y}$  iff for the smallest (if any)  $l$  with  $\sum_{p \in E_i^l} x_p \neq \sum_{p \in E_i^l} y_p$ , we have  $\sum_{p \in E_i^l} x_p > \sum_{p \in E_i^l} y_p$ . Once again, since this requires adding the degrees of completion of different projects, it makes more sense when the degrees of completion of different projects are comparable. However, if each resident submits a strict ordering over the projects, the equivalence classes become singletons and the definition makes sense for other models of PB too. For combinatorial PB, under the lexicographic extension, resident  $i$  compares two allocations by comparing her most favorite project that implemented in each allocation, breaking ties by her next most favorite project, and so on. This is similar to the max set extension mentioned above for cardinal utility functions, and is a reasonable assumption when projects are substitutes of each other.<sup>9</sup> Klamler et al. (2012) study a broader class of preference extensions that includes the lexicographic extension.

Another complete extension is derived by converting ordinal preferences to cardinal preferences using *scoring rules*. A scoring vector is denoted by  $\mathbf{s} = (s_1, \dots, s_m)$ , where  $s_1 \geq \dots \geq s_m \geq 0$ . Given a ranking  $\succ_i$  over projects, it is assumed that resident  $i$  has utility  $u_{i,p} = s_k$  for project  $p$ , where  $k$  is the rank of project  $p$  under  $\succ_i$ . Then, any of the cardinal extensions mentioned above can be used to induce the resident's preferences over allocations.

### 2.2.2 Ballot Design

Even in the simplest case of combinatorial PB, there can be exponentially many allocations. This makes it infeasible to ask residents to communicate their full preferences, thus motivating the use of more restrictive preference elicitation techniques.

For example, even when residents' preferences over allocations are induced from their ranked preferences over individual projects, asking residents to rank as many as 20 projects<sup>10</sup> can be tiresome, and the cognitive burden can lead to fewer residents voting or residents making poor choices (Iyengar and Lepper, 2000). Hence, most real-world applications of PB use easier ballot formats such as  $k$ -approval voting (where residents approve the  $k$  projects they like the most), approval voting (where residents approve all projects that they like), range voting (where residents rate projects), or knapsack voting (where residents provide the ideal allocation according to their preferences) (Goel et al., 2019).

<sup>9</sup> For instance, the projects may propose to build public parks in different locations, but a resident may only be interested in using a single park that is built closest to her home.

<sup>10</sup> The 2018 PB in Cambridge, USA involved 20 projects: <https://pb.cambridgema.gov/pbcycle5>

## 2.3 Vote Aggregation

Once the representation of residents' preferences and the format in which they cast their votes are decided, the next challenge is to decide how to aggregate their votes into a collective outcome, which can be a single allocation or a distribution over allocations (if randomization is permitted).

This is perhaps the most well-studied aspect of the entire PB process. We review several prominent approaches to vote aggregation in computational social choice.

### 2.3.1 Welfare Maximization

This approach is applicable when individual preferences are represented as cardinal utility functions. It uses a *social welfare function*, which combines individual utility functions of residents into a societal utility function, and finds an allocation maximizing this societal utility function. There are three popular social welfare functions:

- The *utilitarian welfare* of an allocation is the sum of utilities it gives to residents:  $UW(\mathbf{x}) = \sum_{i \in N} u_i(\mathbf{x})$  for  $\mathbf{x} \in \mathcal{A}$ .
- The *egalitarian welfare* of an allocation is the minimum utility it gives to any resident:  $EW(\mathbf{x}) = \min_{i \in N} u_i(\mathbf{x})$  for  $\mathbf{x} \in \mathcal{A}$ . Maximizing this welfare function is seen as an extreme form of one interpretation of fairness.
- The *Nash welfare* of an allocation is the product of utilities it gives to residents:  $NW(\mathbf{x}) = \prod_{i \in N} u_i(\mathbf{x})$  for  $\mathbf{x} \in \mathcal{A}$ . Maximizing this welfare function is seen as a compromise between utilitarianism and egalitarianism.

### 2.3.2 The Axiomatic Approach

This approach entails specifying intuitive normative axioms that the vote aggregation rule must satisfy and searching for rules that satisfy as many of the axioms as possible. Many compelling axioms have been proposed for PB (see, e.g., Aziz et al. (2018); Faliszewski and Talmon (2019)); the two examples given below are selected because they apply to all the models of PB we described above.

**Definition 1 (Exhaustiveness (Aziz et al., 2018))** A feasible outcome (i.e. allocation)  $\mathbf{x}$  is called exhaustive if an outcome  $\mathbf{y}$  is not feasible whenever  $y_p \geq x_p$  for all projects  $p$  and a strict inequality holds for at least one project. In words, it should not be possible to implement any more projects or any projects to a greater degree with the leftover budget. A vote aggregation rule is exhaustive if it always outputs an exhaustive allocation. A similar axiom is termed budget monotonicity by Faliszewski and Talmon (2019).

**Definition 2 (Discount Monotonicity (Faliszewski and Talmon, 2019))** Suppose a vote aggregation rule outputs allocation  $\mathbf{x}$ . Suppose project  $p$  receives a revised cost

function  $c'_p$  such that  $c'_p(x_p) \leq c_p(x_p)$  for all  $x_p \in X_p$ , and, all else being equal, the vote aggregation rule now outputs allocation  $\mathbf{y}$ . Then,  $y_p \geq x_p$ . In words, if a project becomes more affordable, it should not be implemented to a lesser degree.

Faliszewski and Talmon (2019) study a number of additional axioms that are specifically applicable to settings where residents have dichotomous preferences.

### 2.3.3 Fairness

Finally, an important consideration in democratic decision making is to fairly take into account the preferences of all residents when making the collective decision. We review one fairness axiom that is applicable in our general PB framework. For other axioms that are applicable in more specific environments (e.g. with dichotomous preferences), we refer interested readers to the work of Aziz et al. (2018).

**Definition 3 (Core (Fain et al., 2016a, 2018))** An allocation  $\mathbf{x}$  is said to be in the core if no subset of residents  $S \subseteq N$  can find an outcome  $\mathbf{x}'$  that is feasible given a budget of  $|S|/n \cdot \mathbf{B}$  such that  $u_i(\mathbf{x}') \geq u_i(\mathbf{x})$  for every resident  $i \in S$  and a strict inequality holds for at least one  $i \in S$ .

The notion of the core is based on the idea that a group of residents  $S$  collectively deserve at least  $|S|/n$  fraction of the budget spent on their needs. This is formalized by requiring that they not be able to allocate this fraction of the budget in a way that they prefer more.

## 3 Discrete Participatory Budgeting

We are now ready to review how these approaches have been applied to different models of PB. Recall that in the discrete model, the decision space is discrete. This model is applicable when projects can only be executed up to discrete levels. For example, a project  $p$  that proposes to build public toilets in the community can have degree of completion  $x_p \in \mathbb{N} \cup \{0\}$ , where  $x_p$  may denote the number of public toilets that are actually built; the assumption here is that each individual toilet is either built fully or not built at all. This discrete aspect affects the modeling of residents' preferences as well as the design of ballots and vote aggregation procedures.

This model is a natural generalization of *multi-winner voting* (alternatively known as *committee selection*), which has been widely studied in social choice theory (Aziz et al., 2017; Faliszewski et al., 2017). In multi-winner voting, there are  $m$  candidates, and  $k$  of them are to be selected based on voters' preferences. This is a special case of discrete PB (more specifically, of combinatorial PB) where each candidate  $p$  is a project with unit cap ( $q_p = 1$ ) and unit cost ( $c_p = 1$ ), and the budget limit is  $k$ . We first review prior literature on multi-winner voting and other similar models of decision-making, and then provide an overview of various approaches to discrete PB.

### 3.1 Review of the Literature on Settings Related to Discrete PB

The simplest special case of combinatorial PB is multi-winner voting. As explained above, this is where each project has unit cost. A natural goal in this setting is to maximize the utilitarian welfare. In this case, each voter approves a subset of candidates, and the goal is to select  $k$  candidates with the highest number of total approvals. This can be accomplished efficiently by a greedy algorithm which selects candidates one-by-one in the decreasing order of the number of approvals they receive. However, when we replace the selection of  $k$  candidates by a knapsack constraint to extend this to combinatorial PB with dichotomous preferences, the problem becomes NP-hard even for a single agent (Karp, 1972). The greedy algorithm is still well-defined in this setting; it selects projects one-by-one in the decreasing order of the number of approvals they receive, skipping any project if its inclusion exceeds the budget limit. This aggregation rule is often used in practice.<sup>11</sup> The literature also focuses on several other objectives in multi-winner voting; an excellent overview is given by Faliszewski et al. (2017).

Klamler et al. (2012) also focus on committee selection under knapsack constraints. However, in their model, there is a single agent with a ranking over individual candidates (projects). They explore different ways to extend this to a ranking over sets of candidates, and then select a committee of  $k$  candidates subject to a knapsack constraint. A similar model has been considered by Delort et al. (2011).

Another model related to PB is *combinatorial public projects* (CPP) (Papadimitriou et al., 2008). In this model, a set of public projects are to be selected subject to a resource constraint, similarly to PB. However, research on CPP focuses on designing truthful mechanisms by charging payments to agents (Dughmi et al., 2016), which is unrealistic in the PB setting.

Conitzer et al. (2017) propose the *public decision making* model, in which there are a finite number of *issues*, and for each issue, there is a set of alternatives which can be implemented. A feasible outcome consists of choosing a single alternative corresponding to each issue. One can view this as a special case of discrete PB, where each issue corresponds to a unique type of resource of which one unit is available and implementing any alternative of an issue requires the full one unit of the corresponding resource. One of the fairness definitions they propose is *proportionality up to one issue*, which is a relaxation of the core defined in Section 2.3.3.

Lu and Boutilier (2011) introduce *budgeted social choice*. In this framework, each alternative has a cost and the goal is to select a set of alternatives subject to budget constraint. This framework generalizes the PB model we discuss as it allows the cost of an alternative to be dependent on the number of voters who derive utility from that alternative. However, they work with a limited modeling of preferences wherein there is a common *positional scoring rule* which maps an alternative's rank in a voter's preference ranking to the voter's utility for the alternative. This is also a common approach in the resource allocation literature (Bouveret and Lang, 2011).

---

<sup>11</sup> See, for example, the PB process in Toronto (Murray, 2019) and in Cambridge (Mundt, 2017).

### 3.2 Approaches to Discrete PB

We now provide an overview of various approaches to PB in the discrete model.

*Welfare maximization:* This is a natural approach when residents have cardinal utilities. Maximizing the utilitarian welfare subject to the budget constraint is the classic knapsack problem. While the problem is NP-hard (Karp, 1972), there exists a pseudo-polynomial time dynamic programming algorithm and a fully polynomial time approximation scheme (FPTAS) (Vazirani, 2013).

Fluschnik et al. (2017) consider the combinatorial PB model and study the computational complexity of maximizing (1) the utilitarian social welfare with respect to additive utilities, (2) the utilitarian social welfare with respect to the max extension, and (3) the Nash social welfare with respect to additive utilities. All of these problems are NP-hard except under severe restrictions on residents' utility functions.

*Elicitation:* Eliciting cardinal utilities in the PB context can be challenging and may impose high cognitive burden on the voters, potentially even deterring them from voting in the first place (Garg et al., 2019). This has inspired a line of research that aims to propose simpler ballot formats. Goel et al. (2019) introduce *knapsack voting*, under which residents simply report their favorite allocation. They suggest aggregating knapsack votes using a greedy approach. In combinatorial PB, the number of "approvals" that a project receives is the number of residents who include the project in their favorite allocation, and projects are selected in the decreasing order of their number of approvals subject to the budget constraint. In discrete PB, the algorithm starts by setting the degree of completion of each project to zero. Then, at each step, it considers increasing the degree of completion of each project to its next possible value, and among the feasible improvements, selects the one which is part of the favorite allocation of most voters. They show that this approach has compelling welfare and incentive guarantees under restrictive models of resident preferences.

Benade et al. (2017) compare knapsack voting to three other ballot formats: ranking projects by value, ranking projects by value-for-money, and a new format that they call *threshold approval voting*. They use the *implicit utilitarian voting* framework (Procaccia and Rosenschein, 2006; Boutilier et al., 2015), in which residents are assumed to have underlying cardinal utilities and the information they provide on the ballot is treated as a proxy for their underlying utilities. They show that in the worst case, knapsack voting leads to an exponentially bad approximation of utilitarian welfare, ranking projects by value or value-for-money leads to a polynomially bad approximation, and threshold approval voting leads to a logarithmic approximation.

*Incentives:* Another line of research studies residents' incentive to misreport their preferences in order to induce an outcome that is better than what would be implemented had they reported their honest preferences. While all reasonable deterministic

aggregation rules are susceptible to such manipulations due to the classic impossibility result by Gibbard (1973) and Satterthwaite (1975), Bhaskar et al. (2018) show that when residents rank projects by value, there exist randomized aggregation rules that are *strategyproof* (i.e. provide no incentive to residents to misreport) and achieve nearly optimal approximation of utilitarian welfare among all (even non-strategyproof) aggregation rules, implying that, at least when randomization is allowed, strategyproofness does not impose a significant burden in this setting.

*Axiomatic desiderata:* Aziz et al. (2018) begin the formal study of defining axioms for PB that encode principles of proportional representation. They design algorithms for combinatorial PB with dichotomous preferences which satisfy such axioms. Their central result is an algorithm called the *generalized Phragmen's sequential rule*, which computes an allocation satisfying a reasonable proportional representation property. An alternative simple method, which is motivated by proportional representation concerns, is discussed by Aziz (2019). Faliszewski and Talmon (2018), in the same model, focus on monotonicity-style axioms of a range of aggregation rules. The common approach in this line of work is to focus on axioms and algorithms for multi-winner voting (Aziz et al., 2017; Faliszewski et al., 2017) and extend them to combinatorial PB.

Since combinatorial PB is more general than multi-winner voting, the negative existential and computational results from multi-winner voting carry over directly. For example, strategyproofness and any weak notion of proportional representation are inherently incompatible (Peters, 2018), and finding Pareto optimal allocations is typically NP-hard (Aziz et al., 2016).

Shapiro and Talmon (2017) generalize Condorcet's principle (also referred to as popularity in allocation settings) to PB and devise an algorithm to compute allocations satisfying this principle. Since Condorcet's principle is based on majoritarian comparisons, the approach is not well-suited for proportional representation of minorities, which is a prime concern in real-world applications of PB (World Bank, 2008; Ganuza and Francés, 2012).

*Fairness:* Fain et al. (2018) investigate fairness in a setting that generalizes discrete PB. In particular, they show that allocations in the core may not exist in discrete PB with additive utilities, but allocations that provide a logarithmic approximation of the core are guaranteed to exist and can be computed efficiently. For the special case of discrete PB with binary additive utilities (where each resident has utility 1 or 0 for each project, and utilities are additive across projects), it is an open question whether allocations in the core always exist. Other work on proportional representation for the combinatorial PB model (see e.g. Aziz et al. (2018)) can also be viewed as focussing on fairness.

*Other approaches:* Rios and Insua (2008) study the combinatorial PB model with additive utilities. They propose to use a generalization of the Kalai-Smorodinsky (1975)

cooperative game solution to find a desirable allocation. A similar approach is taken by Rios et al. (2005).

## 4 Divisible Participatory Budgeting

The divisible model of PB applies to scenarios in which funding for a project can lie on a continuum. For instance, recall the example from Section 3, which involved a project proposing to build five public toilets in the neighborhood. We argued that this falls under the discrete model of PB, as one can only provide funding to build an integral number of toilets. In contrast, a project on maintaining the cleanliness of the toilets may be funded partially without any integral restrictions, and thus may fall under the divisible model of PB.

While a divisible model can often be approximated by a discrete model, the divisible model is interesting in its own right as it often leads to better existential and computational results (see, e.g., the work of Fain et al. (2018)).

### 4.1 Review of the Literature on Settings Related to Divisible PB

A strand of the literature focuses on a multi-dimensional continuous space, where the dimensions correspond to categories of projects such as defense, health and education (Garg et al., 2019; Freeman et al., 2019). The model is especially relevant when deciding the composition of funding across different categories is more important than deciding particular projects within a category. This literature typically works with the spatial model of resident preferences, in which each resident has an ideal point in the space. This is closely related to spatial voting models studied in political science (Enelow and Hinich, 1984; Arrow, 1990)

Recall that in bounded divisible PB, each project has a cap of  $q_p = 1$  (without loss of generality). One might consider a special case of this setting, where the cost function of each project  $p$  is given by  $c_p(x_p) = x_p$ , and the total budget is  $B = 1$ . Thus, the set of feasible allocations is given by  $\{\mathbf{x} : \sum_{p \in P} x_p \leq 1\}$ . This can be thought of as portioning (also called fair mixing) (Aziz et al., 2019), where  $x_p$  can be thought of as the fraction of budget devoted to project  $p$  (after appropriate renormalization of cost functions). As noted in Section 2.1, this is equivalent to the unbounded divisible PB model. Alternatively, one can also think of this model as representing probabilistic voting (Gibbard, 1977), where each project  $p$  is a candidate and  $x_p$  denotes the probability of choosing candidate  $p$  as the winner of the election.

## 4.2 Approaches to Divisible PB

*Welfare maximization:* Goel et al. (2019) study the divisible model of PB with unit caps. They consider cardinal additive utilities as well as a spatial model of utilities. They point out that a simple greedy algorithm finds an allocation maximizing the utilitarian welfare. This algorithm sorts the projects in decreasing order of their value-for-money ( $\sum_{i \in N} u_i(p)/c_p$ ), and fully funds projects in this order until no more projects can be fully funded. The remaining budget is allocated to partially fund the next project in the list.

Garg et al. (2019) consider a model of decision-making in continuous spaces that is more general than PB. In their model, each resident  $i$  has a favorite point  $\mathbf{x}_i$  in the space, and her utility for an allocation  $\mathbf{x}$  is based on the  $\ell_p$  norm, specifically,  $u_i(\mathbf{x}) = -\|\mathbf{x} - \mathbf{x}_i\|_p$ . They study a class of algorithms called iterative local voting (ILV), in which residents are iteratively asked to modify the current allocation by moving it towards their favorite allocation within a local neighborhood until convergence. This approach adopts the classic stochastic gradient descent (SGD) method from optimization to multi-agent decision-making.

*Fairness:* Fain et al. (2016b) study proportional representation in the divisible model of PB with scalar separable utility functions. They argue that a variant of the classic game-theoretic notion of core captures fairness in the setting. Recall that an allocation  $\mathbf{x}$  is in the core if no subset of voters  $S$  can use  $|S|/n$  of the budget to find an allocation which is no less appealing than  $\mathbf{x}$  to any member of  $S$  and strictly more appealing than  $\mathbf{x}$  to some member of  $S$ . They show that in the divisible model, there always exists an allocation in the core, and present a polynomial-time algorithm for computing one.

Recall that in portioning (i.e. unbounded divisible PB), there are no caps on the degrees of completion. Hence, as discussed in Section 2.1, one can normalize the model such that feasible allocations satisfy  $\sum_{p \in P} x_p \leq 1$ , and  $x_p$  can be thought of as the fraction of budget devoted to project  $p$  (Airiau et al., 2019). Aziz et al. (2019) consider portioning — they refer to it as fair mixing — with dichotomous preferences, which was originally studied by Bogomolnaia et al. (2005). They consider the relative merits of several rules. They show that maximizing Nash welfare satisfies several strong fairness properties including the fact that it finds an allocation in the core. It also satisfies the *average fair share* (AFS) guarantee, which requires that for any set of residents  $S$ , the average utility of residents in  $S$  is at least  $|S|/n$ . Another rule with several desirable properties that emerges from their study is the *conditional utilitarian rule* (CUT). In CUT, each resident is assumed to have a ‘personal budget’ that is  $1/n$  fraction of the total budget.<sup>12</sup> Each resident spreads her personal budget among the projects that she likes that are liked by the most number of residents. Unlike the maximum Nash welfare solution, CUT does not satisfy Pareto optimality, AFS, or core fairness.

Airiau et al. (2019) consider the portioning problem where voters express preferences over individual projects and use the stochastic dominance (SD) relation to

<sup>12</sup> Note that CUT easily generalizes to the case where agents have unequal personal budgets.

reason about their preferences over probability distributions over projects. They find that using a scoring rule as a proxy for utilities consistent with the ordinal preferences and then maximizing the Nash social welfare with respect to these proxy utilities achieves desirable notions of fairness. Another notable example of algorithms for portioning is the egalitarian simultaneous reservation algorithm (Aziz and Stursberg, 2014) for agents with weak and transitive ordinal preferences over projects.

*Incentives:* Aziz et al. (2019) prove that for the portioning problem with dichotomous preferences, both CUT and the Nash welfare maximizing rule provide strict incentives for residents to participate in the voting process. Additionally, CUT is strategyproof (i.e. provides voters no incentive to misreport), while maximizing Nash welfare is not.

Freeman et al. (2019) consider the spatial voting model and focus on the setting in which a resident's disutility for an allocation is its  $\ell_1$  distance from the resident's ideal allocation (a.k.a. bliss point). They present the independent markets mechanism, which is both strategyproof and satisfies a basic notion of proportional representation.

## 5 Extensions and Future Directions

We presented a survey of research on participatory budgeting. Specifically, we presented a mathematical model which classifies existing research across different axes, and surveyed computational and axiomatic approaches proposed in the literature.

Research on participatory budgeting within computational social choice is still in its infancy. As such, only limited models of resident preferences and decision-making processes have been explored. Further research is required to design better theoretical models and approaches for participatory budgeting, and bridge the gap between theory and practice. To that end, it is crucial to leverage insights from various disciplines such as computer science, social science, microeconomics, and public policy. In this section, we review several directions in which these models can (and should) be extended to bring them closer to reality and for them to inform real-world implementations of PB.

- *Multi-dimensional constraints:* Much of the literature focuses on a single, knapsack-style constraint which stems from a limit on the available money. While current real-world implementations of PB also only deal with money, this is one dimension where further research on practical approaches to conducting PB with multi-dimensional constraints that capture other types of costs (e.g. costs to the environment) can lead the frontier of novel PB implementations.
- *Voter or voter group entitlements:* In the classic PB setting, giving equal consideration (or weight) to all voters is a normative desideratum. However, one can imagine a more general setting in which different voters or groups of voters bring different resources to the system, and their entitlements need to reflect

their contributions. In other words, different voters or groups of voters may have a claim over different parts of the universal budget.

- *Distributional constraints*: Some applications of PB impose lower and upper bounds on the amount of funding which can be allocated to each project or each category of projects (e.g. education or healthcare). Studying the effects of such constraints on the welfare, fairness, and incentives is an interesting direction.
- *Hybrid models*: While our taxonomy (and the literature) partitions PB models into discrete and divisible, there can be hybrid models allowing some projects to be funded only at discrete levels while others on a continuous scale.
- *Complex resident preferences*: Most of the positive axiomatic and algorithmic results in the literature rely on stylized modeling of resident preferences. In practice, residents have complex preferences which stem from intricate synergies between different projects. An important research direction is to extend the existing results to more general classes of utility functions. For example, most utility functions considered in the literature treat projects as independent or substitutes; tackling complementarity in projects remains an interesting challenge.
- *Initial endowments*: In reality, some projects may already have some funds allocated to them (e.g. through previous iterations of PB). The choice between extending funding to existing projects versus funding new projects can give rise to novel challenges.
- *Pledges*: In some cases, private organizations may be willing to pledge financial support to a project conditional on certain projects receiving at least a minimum level of funding. This can affect the choice of vote aggregation methods.
- *Market-based approaches*: In recent years, market-based approaches have been investigated for public decision making (Garg et al., 2018). Extending this line of approach to PB can lead to intuitive mechanisms for PB.
- *Strategic agent models*: The work on incentives in PB focuses on strategyproofness, which aims to prevent manipulations by rational agents. However, people often do not act like rational agents. An interesting direction is to use insights from behavioral game theory, develop models of realistic manipulations that residents may use in PB, and design algorithms to prevent such manipulations.
- *The role of information*: It requires significant effort to inform residents of the costs, benefits, and complexities of different projects. The manner in which this information is communicated can have significant effect on the preferences of the residents; this is a complex issue which requires a detailed study.
- *An end-to-end model*: Finally, as we mentioned at the beginning of this chapter, we only focus on the final stage of voting within the entire PB pipeline. However, the initial stages have a direct impact on the final outcome. For example, the agenda-setting phase where projects are proposed by the residents themselves crucially affects the latter stages. Formally modeling the entire PB process and designing end-to-end solutions is a complex challenge of paramount importance.

To conclude, participatory budgeting is an important grassroots approach to democracy. Research on models, methods, and axioms for PB will provide insights that will be valuable to both the theory and practice of PB.

## 6 Acknowledgements

Haris Aziz is supported by a UNSW Scientia Fellowship. Nisarg Shah is supported by an NSERC Discovery grant. The authors thank Aditya Ganguly, Barton Lee and Dominik Peters for very helpful feedback.

## References

- S. Airiau, H. Aziz, I. Caragiannis, J. Kruger, J. Lang, and D. Peters. Portioning using ordinal preferences: Fairness and efficiency. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 2019.
- K. Arrow. *Advances in the spatial theory of voting*. Cambridge University Press, 1990.
- H. Aziz. Participatory budgeting: Are we really giving a voice to everyone?, 2019. URL <https://medium.com>.
- H. Aziz and P. Stursberg. A generalization of probabilistic serial to randomized social choice. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pages 559–565. AAAI Press, 2014.
- H. Aziz, J. Lang, and J. Monnot. Computing Pareto optimal committees. In *Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016*, pages 60–66, 2016.
- H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2):461–485, 2017.
- H. Aziz, B. E. Lee, and N. Talmon. Proportionally representative participatory budgeting: Axioms and algorithms. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 23–31. IFAAMAS, 2018.
- H. Aziz, A. Bogomolnaia, and H. Moulin. Fair mixing: the case of dichotomous preferences. In *Proceedings of the 20th*, pages 753–781, 2019.
- G. Benade, S. Nath, A. D. Procaccia, and N. Shah. Preference elicitation for participatory budgeting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 376–382. AAAI Press, 2017.
- U. Bhaskar, V. Dani, and A. Ghosh. Truthful and near-optimal mechanisms for welfare maximization in multi-winner elections. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, pages 925–932, 2018.
- A. Bogomolnaia, H. Moulin, and R. Stong. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2):165–184, 2005.
- C. Boutilier, I. Caragiannis, S. Haber, T. Lu, A. D. Procaccia, and O. Sheffet. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227:190–213, 2015.

- S. Bouveret and J. Lang. A general elicitation-free protocol for allocating indivisible goods. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 73–78. AAAI Press, 2011.
- F. Brandl, F. Brandt, and W. Suksompong. The impossibility of extending random dictatorship to weak preferences. *Economics Letters*, 141:44–47, 2016.
- F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- Y. Cabannes. Participatory budgeting: a significant contribution to participatory democracy. *Environment and Urbanization*, 16(1):27–46, 2004.
- I. Caragiannis, S. Nath, A. D. Procaccia, and N. Shah. Subset selection via implicit utilitarian voting. *Journal of Artificial Intelligence Research*, 58:123–152, 2017.
- V. Conitzer, R. Freeman, and N. Shah. Fair public decision making. In *Proceedings of the 18th ACM Conference on Electronic Commerce (EC '17)*, 2017.
- C. Delort, O. Spanjaard, and P. Weng. Committee selection with a weight constraint based on a pairwise dominance relation. In *Proceedings of the Second International Conference on Algorithmic Decision Theory (ADT)*, pages 28–41, 2011.
- S. Dughmi, T. Roughgarden, and Q. Yan. Optimal mechanisms for combinatorial auctions and combinatorial public projects via convex rounding. *Journal of the ACM (JACM)*, 63(4):30, 2016.
- U. Endriss, editor. *Trends in Computational Social Choice*. AI Access, 2017.
- J. M. Enelow and M. J. Hinich. *The spatial theory of voting: An introduction*. CUP Archive, 1984.
- B. Fain, A. Goel, and K. Munagala. The core of the participatory budgeting problem. In *Proceedings of the 12th International Workshop on Internet and Network Economics (WINE)*, Lecture Notes in Computer Science (LNCS), pages 384–399, 2016a.
- B. Fain, A. Goel, and K. Munagala. The core of the participatory budgeting problem. In *Proceedings of the 12th International Conference on Web and Internet Economics (WINE '16)*, pages 384–399, 2016b.
- B. Fain, K. Munagala, and N. Shah. Fair allocation of indivisible public goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (ACM-EC)*, pages 575–592, 2018.
- P. Faliszewski and N. Talmon. A framework for approval-based budgeting methods. *CoRR*, abs/1809.04382, 2018.
- P. Faliszewski and N. Talmon. A framework for approval-based budgeting methods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 2019.
- P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Multiwinner voting: A new challenge for social choice theory. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 2. 2017.
- T. Fluschnik, P. Skowron, M. Triphaus, and K. Wilker. Fair knapsack. *CoRR*, abs/1711.04520, 2017.

- R. Freeman, D. M. Pennock, D. Peters, and J. W. Vaughan. Truthful aggregation of budget proposals. In *Proceedings of the 20th ACM Conference on Electronic Commerce (EC '19)*, 2019.
- E. Ganuza and F. Francés. The deliberative turn in participation: the problem of inclusion and deliberative opportunities in participatory budgeting. *European Political Science Review*, 4(2):283–302, 2012.
- N. Garg, A. Goel, and B. Plaut. Markets for public decision-making. arXiv:1807.10836, 2018.
- N. Garg, V. Kamble, A. Goel, D. Marn, and K. Munagala. Iterative local voting for collective decision-making in continuous spaces. *Journal of Artificial Intelligence Research (JAIR)*, 2019.
- A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973.
- A. Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3):665–681, 1977.
- A. Goel, Anilesh K. K., S. Sakshuwong, and T. Aitamurto. Knapsack voting for participatory budgeting. *ACM Transactions on Economics and Computation (TEAC)*, 7(2):8:1–8:27, 2019. ISSN 2167-8375.
- B. Gutiérrez. Madrid as a democracy lab, 2017. URL <https://www.opendemocracy.net/democraciaabierta/bernardo-guti-rrez/madrid-as-democracy-lab>. Accessed: 2018-05-02.
- S. S. Iyengar and M. R. Lepper. When choice is demotivating: Can one desire too much of a good thing? *Journal of personality and social psychology*, 79(6):995, 2000.
- E. Kalai and M. Smorodinsky. Other solutions to nash’s bargaining problem. *Econometrica*, 43(3):513–518, 1975.
- R. M. Karp. Reducibility among combinatorial problems. In *Complexity of computer computations*, pages 85–103. Springer, 1972.
- C. Klamler, U. Pferschy, and S. Ruzika. Committee selection under weight constraints. *Mathematical Social Sciences*, 64(1):48–56, 2012.
- J. Legendre, H. Madénian, and P. L. Scully. Participatory budgeting in Paris, France, 2017. URL <http://participedia.net/en/cases/participatory-budgeting-paris-france>. Accessed: 2018-05-02.
- T. Lu and C. Boutilier. Budgeted social choice: From consensus to personalized decision making. In *Proceedings of the 22nd Joint Conference on Artificial Intelligence (IJCAI '11)*, pages 280–286, 2011.
- M. Mundt. Participatory budgeting evaluation report. [https://pb.cambridgema.gov/read\\_the\\_pb3\\_evaluation](https://pb.cambridgema.gov/read_the_pb3_evaluation), 2017. Last accessed: Aug 5, 2019.
- C. Murray. Toronto’s participatory budgeting pilot evaluation. <https://www.toronto.ca/legdocs/mmis/2019/bu/bgrd/backgroundfile-124370.pdf>, 2019. Last accessed: Aug 5, 2019.
- C. Papadimitriou, M. Schapira, and Y. Singer. On the hardness of being truthful. In *Proceedings of the 49th Symposium on Foundations of Computer Science (FOCS)*, pages 250–259. IEEE, 2008.

- D. Peters. Proportionality and strategyproofness in multiwinner elections. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, volume 1549–1557, 2018.
- A. D. Procaccia and J. S. Rosenschein. The distortion of cardinal preferences in voting. In *Proceedings of the 10th*, pages 317–331, 2006.
- J. Rios and D. Rios Insua. A framework for participatory budget elaboration support. *Journal of the Operational Research Society*, 59(2):203–212, 2008.
- J. Rios, D. R. Insua, E. Fernandez, and J. A. Rivero. Participatory budget formation through the web. In *E-Government: Towards Electronic Democracy, International Conference, TCGOV 2005, Bolzano, Italy, March 2-4, 2005, Proceedings*, pages 268–276, 2005.
- A. Röcke. *Framing citizen participation: participatory budgeting in France, Germany and the United Kingdom*. Springer, 2014.
- R. Samir. Madrid, part 1: e-democracy hub of europe, 2018. URL <http://netdem.nl/en/articles/madrid-e-democracy-hub-europe/>. Accessed: 2019-04-28.
- M. A. Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.
- A. Shah. *Participatory Budgeting*. Public sector governance and accountability series. The World Bank, 2007.
- E. Shapiro and N. Talmon. A condorcet-optimal participatory budgeting algorithm. *arXiv preprint arXiv:1709.05839*, 2017.
- V. V. Vazirani. *Approximation algorithms*. Springer Science & Business Media, 2013.
- World Bank. *Brazil: Toward a more inclusive and effective participatory budget in Porto Alegre*. World Bank, 2008.