

Knapsack Voting for Participatory Budgeting

ASHISH GOEL, ANILESH K. KRISHNASWAMY, SUKOLSAK SAKSHUWONG, and
TANJA AITAMURTO, Stanford University

We address the question of aggregating the preferences of voters in the context of participatory budgeting. We scrutinize the voting method currently used in practice, underline its drawbacks, and introduce a novel scheme tailored to this setting, which we call “Knapsack Voting.” We study its strategic properties—we show that it is strategy-proof under a natural model of utility (a dis-utility given by the ℓ_1 distance between the outcome and the true preference of the voter) and “partially” strategy-proof under general additive utilities. We extend Knapsack Voting to more general settings with revenues, deficits, or surpluses and prove a similar strategy-proofness result. To further demonstrate the applicability of our scheme, we discuss its implementation on the digital voting platform that we have deployed in partnership with the local government bodies in many cities across the nation. From voting data thus collected, we present empirical evidence that Knapsack Voting works well in practice.

CCS Concepts: • **Theory of computation** → **Algorithmic game theory and mechanism design**;
• **Computing methodologies** → *Artificial intelligence*; • **Applied computing** → Economics;

Additional Key Words and Phrases: Participatory budgeting, social choice, digital voting

ACM Reference format:

Ashish Goel, Anilesh K. Krishnaswamy, Sukolsak Sakshuwong, and Tanja Aitamurto. 2019. Knapsack Voting for Participatory Budgeting. *ACM Trans. Econ. Comput.* 7, 2, Article 8 (July 2019), 27 pages.

<https://doi.org/10.1145/3340230>

1 INTRODUCTION

Direct democracy has gained a lot of significance lately, with many novel attempts at engaging citizens very directly in policy-making (Pateman 2012; Smith 2009). An exciting new development in this space is participatory budgeting (Cabannes 2004), in which a local government body asks residents to vote on project proposals to decide how they should allocate their budgetary spending. The proposals could be, in a particular city/ward, for instance, resurfacing streets, adding street lights, building playgrounds for children, or renovating recreational facilities like parks. Participatory budgeting has had a long history in South America (Schneider and Goldfrank 2002), where it was born out of a need to address inequalities through democratic reform. It is now gaining

This work is supported by the Army Research office (grant # 116388), the Office of Naval Research (grant # 11904718), the National Science Foundation (grant # 1637397), and the Stanford Cyber Initiative.

Authors’ addresses: A. Goel, 475 via Ortega, Room 359, Stanford CA 94035; email: ashishg@stanford.edu; A. K. Krishnaswamy, 15 Coleman Pl, Apt 2, Menlo Park CA 94025; email: anilesh@stanford.edu; S. Sakshuwong, 933 Florence Ln Apt D, Menlo Park CA 94025; email: sukolsak@stanford.edu; T. Aitamurto, Department of Communication, University of Illinois at Chicago, 1845 Sheridan Road, Evanston, Illinois, 60208; email: tanja.aitamurto@gmail.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Association for Computing Machinery.

2167-8375/2019/07-ART8 \$15.00

<https://doi.org/10.1145/3340230>

popularity in the US, with cities like San Francisco, Vallejo, Boston, Chicago, and New York adopting this paradigm (PBP 2016). This development necessitates a detailed look at the voting method used in current ballots and motivates the following question: When there are projects with different costs, and a fixed budget, how can the varied preferences of voters be best aggregated? We address this question by proposing Knapsack Voting and discuss its advantages over existing methods. We show that it has desirable strategic properties and provide empirical evidence that it works well in practice. We have been successful in implementing Knapsack Voting as the official ballot procedure in the Youth Lead the Change PB election in 2016 in Boston (<https://youth.boston.gov/youth-lead-the-change/>). Independently of our work, other implementations of the same method have been used in some elections in Europe (see Section 1.1.3 for a discussion)—adding to the relevance of our work in studying the advantages of Knapsack Voting. Overall, this piece of work is a step in the direction of addressing the need for mechanisms to facilitate complex decision-making processes. In designing the Knapsack Voting method, we incorporate ideas from the classical Knapsack Problem to make the voter choose projects under the budget constraint. In doing so, we align the constraints on the voters' decisions, to those of the decision maker's. We believe that this approach of aligning incentives is useful more broadly in designing mechanisms for complex participatory decision-making processes. Moreover, we want our schemes to be amenable to implementation as convenient digital interfaces to help individuals make informed budgeting decisions.

1.1 Participatory Budgeting in Practice

PB first began in Porto Alegre in Brazil in 1988 and since then has become increasingly popular in Brazil.¹ Over the past couple of decades, PB has spread to Europe and North America and more recently taken root in Asia and Africa (Dias et al. 2014). As of 2015, there have been nearly 1,500 instances of PB around the world (Ganuza and Baiocchi 2012).

The first PB process in the United States was initiated in the 49th ward of Chicago in 2009. Since then PB has spread to many other cities like Vallejo, New York, and Cambridge, to name a few. These elections traditionally used only paper ballots and a voting method known as K-approval that we will describe shortly.

1.1.1 Our Digital Voting Platform. To popularize digital voting, we built a digital voting platform (<https://pbstanford.org/>) for the PB process of Chicago's 49th ward in 2012. We have since customized this tool to be easily adaptable to other elections. Many of the cities/municipalities doing PB have taken to digital voting by adopting our platform, while some cities, like Long Beach and Dieppe, have used our platform for their very first attempt at PB. In all, our platform has been used in over 25 different elections in around a dozen cities/wards over the past few years. One thing to note is that digital voting here means voting in person at a polling booth but via the digital platform. Most of these elections are required to have a paper ballot for those voters that opt out of digital voting. Some elections, for example, PB Cambridge, implemented internet voting, where voters can authenticate themselves and vote from anywhere, via a voting website designed using our platform. These elections had accompanying paper ballots also.

1.1.2 K-approval Voting. The voting method currently used by most participatory budgeting elections is *K-approval* voting. Our interface for this method is shown in Figure 1. This is similar to the standard approval voting method but with a cap on the number of alternatives a voter can choose.

¹https://siteresources.worldbank.org/INTEMPowerment/Resources/14657_Particip-Budg-Brazil-web.pdf.

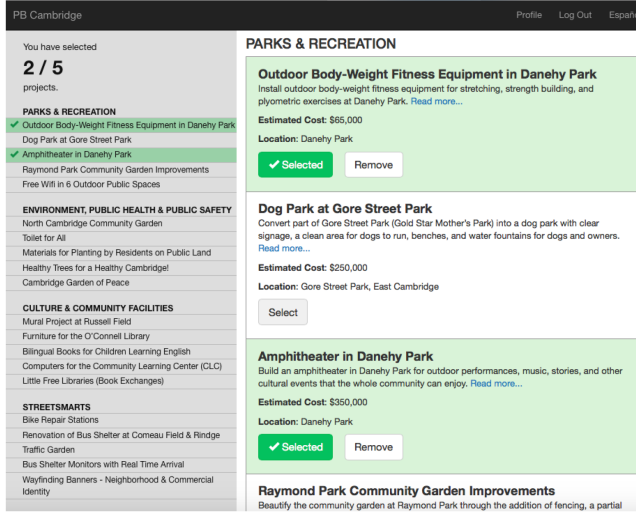


Fig. 1. Our interface for K-approval: Example from PB Cambridge 2015.

Table 1. An Example:
Three Projects and
Additive Utilities

	Cost
Project A	\$300
Project B	\$200
Project C	\$100

Definition 1.1 (K-approval Voting).

- Each voter chooses (“approves”) at most K projects.
- The projects are ranked in descending order according to the total number of approving voters.
- The outcome is decided by picking projects in this order until the budget is exhausted.

While this method is simple and has wide acceptance, it does not make the voters factor in the costs of the projects into their votes. Consider the following example:

Example 1.2. Think of one voter (or a homogeneous pool of voters) doing 1-approval. Let us say the budget is \$300 and there are three projects on the ballot (see Table 1 above): Project A, Project B, and Project C. Their costs are \$300, \$200, \$100, respectively. Imagine that Project A is the single most beneficial project among all, but Project B and Project C are together more valuable in terms of benefit than just Project A.

It is natural to expect that the voter(s) will vote for Project A under 1-approval. Although Projects B and C are together feasible under the budget of \$300, there is no way for the voter to express these preferences. Although this is a very simple example that makes a strong assumption on the voters’ response to the K-approval rule, it is an apt illustration of observed voting behavior (see Section 4.0.1).

A point to note here is that even truthful reporting under K-approval voting (choosing projects with highest utility) does not lead to good outcomes, as seen in Example 1.2 above. It is conceivable

that the outcome might be different if the voters vote strategically, taking aggregation with the budget constraint into account. However, we would like a natural way for the voters to express their preferences with respect to costs and benefits—one that works at least for a homogeneous population under truthful reporting.

This leads us to the following questions:

- What is a good way of eliciting voter preferences?
- What utility models can we use to uncover some insights on the above problem?

We utilize ideas derived from ways of solving the Knapsack Problem: corresponding to our Knapsack Voting and value-for-money schemes, which will be treated in detail in later sections of this article Sections 2 and 3.2. The remainder of this section is organized as follows: We introduce some modeling preliminaries in Section 1.2 and give a short overview of our work in the Section 1.3. We discuss the implications and properties of our model in Section 1.4

1.1.3 Other Implementations Similar to Knapsack Voting. We mention here that independent of our work on Knapsack Voting, other implementations of the same voting method have been in use in some places in Europe—this adds to the relevance of our work in studying the advantages of Knapsack Voting. The most prominent (with a total allocated budget, across many districts, of roughly € 40 million) among these are the “shopping cart votes” use in the PB process in Madrid,² Spain, where it has been in use since 2016. The origins of this method come from the “Open Active Voting” method in the My Neighbourhoods Project in Reykjavik,³ Iceland, where it has been in use since 2012. The same platform has been used in other places like Kòpavogur and Mosò in Iceland⁴ and a pilot implementation in Argyll and Bute,⁵ Scotland.

1.2 Modeling Assumptions: Voting and Utility Models

The participatory budgeting problem addresses the following scenario: The residents of a city, collectively the set of voters $\mathcal{V}$, vote on a set $\mathcal{P}$ of projects that they have identified to be worthwhile, where project $j \in \mathcal{P}$ has a cost c_j and there is a fixed total budget of B dollars.

1.2.1 Voting Models. We will mainly consider a fractional model of voting—each voter can allocate the budget B in any which way among the projects. Our results for this model also carry over (approximately) to an integral setting. For ease of exposition, we defer a discussion of the above to Section A.1 in the appendix.

Definition 1.3 (Fractional Vote). Each voter v chooses an allocation $\{w_p^v\}_{p \in \mathcal{P}}$ such that $0 \leq w_p^v \leq c_p$ for all $p \in \mathcal{P}$ and $\sum_{p \in \mathcal{P}} w_p^v = B$.

In effect, we make the following assumption:

ASSUMPTION 1. *Every project is allowed to be fractionally implemented, and the budget is completely used up.*

Note: It may or may not be realistic to partially fund projects (e.g., renovate only one floor of a library rather than the entire building), but partial funding and the requirement of spending all of the budget are both common in practice.

²<https://decide.madrid.es/presupuestos-participativos-resultados>.

³https://citizens.is/portfolio_page/my-neighbourhood/.

⁴<https://citizens.is/portfolio/>.

⁵<https://www.argyll-bute.gov.uk/abpb/>; <https://citizens.is/portfolio/>.

1.2.2 Utility Models. In what follows, we give concise definitions of the utility models we use in this article. A full discussion of their properties is deferred to Section 1.4.

Let $\{w_p\}_{p \in \mathcal{P}}$ denote an outcome and $\{w_p^v\}_{p \in \mathcal{P}}$ denote voter v 's ideal allocation. We first define the ℓ_1 costs model.

Definition 1.4 (ℓ_1 costs). The ℓ_1 cost of an outcome $\{w_p\}_{p \in \mathcal{P}}$ (where $\sum_{p \in \mathcal{P}} w_p = B$: Assumption 1) for a voter v is given by $\sum_p |w_p - w_p^v|$.

We define another model called “Overlap utility.”

Definition 1.5 (Overlap Utility). The Overlap utility of an outcome $\{w_p\}_{p \in \mathcal{P}}$ (where $\sum_{p \in \mathcal{P}} w_p = B$) for a voter v is given by $\sum_{p \in \mathcal{P}} \min\{w_p, w_p^v\}$.

The above model will be especially useful in our proofs in Section 2. The term corresponding to p in the definition above is equal to the “overlap” between w_p^v and w_p^* , and so we call this expression the Overlap Utility. This means, in other words, that the utility from p for a voter v is equal to w_p^v when $w_p^v \geq w_p^*$, and w_p^* otherwise. For example, if $w_p^v = 5$ and $w_p^* = 10$, then the utility derived by v is 5.

We also define an additive concave utility model, which is a generalization of the Overlap utility model.

Definition 1.6 (Additive concave utility). In the fractional vote model, the additive concave utility of an outcome $\{w_p\}_{p \in \mathcal{P}}$ (where $\sum_{p \in \mathcal{P}} w_p = B$) for a voter v is given by $\sum_{p \in \mathcal{P}} f_v^p(w_p)$, where each $f_v^p(\cdot)$ is a non-decreasing concave function.

We use the definitions from above to summarize our results in the next section. A discussion of the properties of the above utility models is deferred to Section 1.4.

1.3 Our Contributions

Our overall goal in this article will be to present voting mechanisms that are particularly suited to the Participatory Budgeting problem, especially the Knapsack Voting scheme. These mechanisms, by getting voters to either optimize over the total budget or compare projects based on their value-for-money, enable the voter to inherently consider costs and benefits and, moreover, are implementable using interactive digital tools. In building our case for Knapsack Voting, we aim to establish its advantages as follows:

- (1) Knapsack Voting, rather than K -approval voting, is the correct application of the approval voting principle to the participatory budgeting problem, and thereby avoids some issues of strategic voting.
- (2) It is easy to implement in a user-friendly way, and we have been able to deploy it in real elections successfully.
- (3) Empirically it agrees more with binary comparisons of value-for-money: a natural proxy of voter preferences.
- (4) It nudges voters to better consider cost-benefit tradeoffs by making them recognize the budget constraint.

1.3.1 Knapsack Voting. Our primary scheme is called Knapsack Voting. The main idea behind this scheme is that each voter has to adhere to the budget constraint, thus keeping costs in mind while giving her vote—thus internalizing the global constraints, while giving their preferences toward the outcome. Our interface for this method is shown in Figure 2 (integral) and Figure 4 (fractional).




Fig. 2. Our interface for Knapsack Vote: Example from PB NYC District 5 2015.

We primarily look at the Knapsack Vote mechanism in the fractional setting. Most of our results extend to the integral case in an approximate fashion; we defer a discussion of this setting to Section A.1 in the appendix.

We now give an informal definition of Knapsack Voting (for a formal one see Definition 2.1).

Definition 1.7 (Knapsack Vote).

- Each voter $v \in \mathcal{V}$ votes for an allocation $\{w_p^v\}_{p \in \mathcal{P}}$, where $\sum_{p \in \mathcal{P}} w_p^v = B$.
- The outcome is chosen according to approval scores in a per-dollar sense, as in Figure 6 (Section 1.4).

As we can see, while the elicitation method is very different, the aggregation method here is similar to K-approval in that it uses approval scores.

We analyze the strategic properties of Knapsack Voting assuming an ℓ_1 cost model, in which, loosely speaking, the cost to a voter depends on how much the outcome differs from her preferred allocation as determined by the ℓ_1 distance between the two (Definition 1.4). Under this model (and also the related Overlap utility model, via the equivalence Lemma 2.6), we show that Knapsack Voting is strategy proof and welfare maximizing (Theorems 2.3, 2.5).

RESULT 1. *Under the ℓ_1 cost model (and also the Overlap utility model), Knapsack Voting is strategy proof and welfare maximizing.*

This result is analogous to the strategy-proofness results for Approval Voting under dichotomous preferences (Brams and Fishburn 1978).

We also extend Knapsack Voting to more general settings with revenues, surpluses, or deficits (as opposed to one with a fixed budget) where approval voting has no known analog and prove similar results (Section 2.2, Theorem 2.9).

Also, under the more general case of additive concave utilities, we characterize a weaker yet interesting notion of strategy-proofness of a voter's best response, based on the concept of sincerity (Niemi 1984): It is in a voter's best interest to vote for projects that she favors among those that are winning without her vote (Theorem 2.10).

PB Cambridge

This is only a survey. It will not affect your vote.

Which project gives the best value for money, that is, provides the most benefit to the community per dollar spent?

Traffic Garden A Dutch Traffic Garden, called "Safety School" in the U.S., is a permanent park for children to learn transportation etiquette by role playing. Kids take turns biking, walking, and driving a miniature streetscape in a safe environment. Estimated cost: \$200,000	Renovation of Bus Shelter at Comeau Field & Rindge This project entails a major redesign and renovation, with community participation, of the #83 bus shelter. New features will include acrylic panels with cut-out figures and a roof element evocative of nearby recreation and natural habitat. Estimated cost: \$75,000
---	---

Fig. 3. Value-for-money comparison: Example from PB Cambridge 2015.

RESULT 2. *Under additive concave utilities, a voter's best response under Knapsack Voting is partially strategy proof.*

Together, these amount to substantial evidence that Knapsack Voting aligns the incentives of the voters with that of the decision maker. Note that these two results do not hold for K-approval voting.⁶

An interesting, and standard, way of understanding voting rules is viewing them as Maximum Likelihood estimators (MLEs) (Conitzer and Sandholm 2012; Young 1988). Votes are assumed to be drawn from a suitable noisy model parametrized by a "ground-truth" outcome, such that the voting rule is the MLE of the "ground truth" given any realization of votes. For example, the Kemeny Young rule is the MLE of rankings drawn from the Mallows model. The Mallows model defines a probability distribution over all possible rankings, with the probability of each ranking depending on how much it differs from a given ground-truth order. Since we are concerned with the problem of determining an outcome that is a subset of projects that satisfies the budget constraint, we will define a model that takes such a set as the ground truth, and defines a probability for all valid sets, depending on how much they differ from the ground truth. By interpreting Knapsack Voting (Section 2.4) as the MLE of this natural noise model, we reinforce the aggregation method of Knapsack Voting.

1.3.2 Value-for-money Comparisons. In our second scheme, we elicit the voters' preferences based on their perceived value-for-money from projects. By "value-for-money," we mean the utility of a project normalized by its cost: If $v_{i,p}$ is the utility of voter i from project p , then its value-for-money as perceived by i is $\frac{v_{i,p}}{c_p}$. See Figures 3 and 7 for voting interfaces based on this idea, where voters compare pairs of projects, or rank their top projects, based on value-for-money.

In the previous section, we discussed how most elections use paper ballots for those voters that opt out of the digital platform. Using the idea of value-for-money, we designed a paper ballot that accompanies Knapsack Voting.

We will discuss this and other details about value-for-money in Section 3.2.

⁶The utility functions are defined with respect to the voter's favorite budget allocation and the budget outcome chosen and not what the voter expresses. Under such models, the results do not hold for K-approval, e.g., it is not strategy proof under such utility models.




Fig. 4. Experimental interface for fractional Knapsack.

1.3.3 Deployments and Data Analysis. Using our digital voting system that has gained widespread acceptance in many cities/municipalities across the nation, we tested our methods across various participatory budgeting elections. In most of these elections, K-approval was used as the official ballot, and, in addition, we tested the Knapsack, the value-for-money comparisons (Figure 3), or both. We do value-for-money to estimate the aggregate pairwise preferences of voters.

We describe our experimental procedure in greater detail in Section 4. Based on the data from our experiments, we make the following observations:

- (1) Knapsack Voting leads to a more economical consideration of the projects as compared to K-approval (Section 4.0.1),
- (2) The time taken by a voter to do Knapsack Voting is similar to that for K-approval. Value-for-money comparisons take much less time (Section 4.0.2).
- (3) Knapsack Voting does better than K-approval in terms of agreement with pairwise comparisons (Section 4.1).

We present empirical evidence in support of these observations in Section 4.

The above observations suggest a significant qualitative difference between the outcomes of Knapsack Voting and K-approval. In particular, Observation 3 above suggests that the outcome of Knapsack Voting is more in alignment with the aggregate preferences of the voters than that of K-approval.

Observation 1 just reinforces the fact that Knapsack Voting leads to a more economical consideration of projects by the voters. It is natural to expect voters to pay more attention to costs of projects under Knapsack Voting. And Observation 2 suggests that Knapsack Voting does not involve a much larger cognitive load on the voters than K-approval.

However, we stop short of claiming that Knapsack Voting leads to outcomes that are more beneficial to society as a whole. It could be worthwhile to think of ways to compare the outcomes in terms of their long-term societal value.

1.4 A Discussion of Utility Models

As mentioned before, one of our goals will be to prove some good properties of Knapsack Voting, in particular strategy proofness. To do so, we first need to make some modeling assumptions.⁷

There is a wide range of budget problems where the allocation per project is not a fixed cost but can be variable, possibly up to a certain upper limit (see Figure 4 for an experimental interface).

⁷With no specific assumptions, strategy proof mechanisms are impossible, see Section 1.5.

		Projects		
		P1	P2	P3
Voters	A	\$4	\$5	\$1
	B	\$3	\$5	\$2
	C	\$0	\$0	\$10

Budget = \$10

Fig. 5. Preferred allocations of voters.




Fig. 6. Votes per dollar for each project.

This setting is also very relevant to PB.⁸ The fractional setting is also relevant outside of PB. For example, the federal budget consists of a variable allocation of money to various sectors of the economy at large such as education, healthcare, defense, and so on. To accommodate variable allocation, and to avoid combinatorial difficulty, we consider a fully fractional version of Knapsack Voting using a “per-dollar” approach.

Example 1.8. Given a budget of 10, and three voters, A, B, C , let us say that the preferred allocations of the voters for projects $P1, P2, P3$ of maximum costs (5, 5, 10) are (4, 5, 1), (3, 5, 2), (0, 0, 10), respectively (Figure 5). We divide each project into as many dollars as its maximum cost, and given that A allocates 4 to $P1$, we give one vote to each of the first four dollars of $P1$, and so on. Figure 5 shows how many votes each dollar of each project gets. If we select those dollars with the highest number of votes, then we get the allocation (3, 5, 2). Notice that here we did not have to break ties to get a feasible allocation.

As mentioned earlier, we study the above problem under different utility models. We now discuss them below; these models were formally defined earlier in Section 1.2.

1.4.1 ℓ_1 Costs and Overlap Utility. In Section 2, we look at a model where the disutility of a voter for a given allocation of the budget is given by how much her preferred allocation differs from it in terms of the ℓ_1 distance between the two allocations. Minimizing ℓ_1 distances to reach a compromise decision has some precedence in public policy literature (André et al. 2010). Assume that the ideal preference of a voter v corresponds to allocating an (integral) amount w_p^v to project $p \in \mathcal{P}$, such that $\sum_{p \in \mathcal{P}} w_p^v = B$, where B is the total budget. Then we know from Definition 1.4 that the cost for a voter v for an outcome $\{w_p\}_{p \in \mathcal{P}}$ is given by $\sum_{p \in \mathcal{P}} |w_p^v - w_p|$.

We must note here that applying the ℓ_1 cost model to the case where outcomes are allowed to be of size smaller than B (i.e., violating Assumption 1, e.g., the purely integral setting) does not lead to anything meaningful. Indeed doing so could violate the free disposal of utilities.⁹

Example 1.9. Consider a total budget of 2, with three projects a, b , and c . Take a voter that has a preferred allocation of 1 to a . Then an allocation to b and c of 1 each has a cost of -3 for the voter, whereas just an allocation of 1 to project b has a cost of -2 .

As mentioned in Definition 1.5, we define another model called Overlap Utility, which we will see is the analog of dichotomous preferences as applied to approval voting (Brams and Fishburn 1978): For a voter’s preferred allocation, we look at each dollar subproject (see Section 2 for a formal

⁸See <https://pbstanford.org/boston16internal/knapsack> for a PB election where most of the projects allowed for flexible allocations.

⁹We thank one of the reviewers for pointing this out via the described example.

definition of subprojects) in it and count it as 1 toward the utility if the same dollar subproject is in the outcome allocation and 0 if not.

However, it is easy to see that the Overlap utility model satisfies¹⁰

- (1) convexity of preferences, and
- (2) free disposal of utilities.

To see why it satisfies free disposal, let us look at Example 1.9 again, where a voter has a preferred allocation of 1 to a . An allocation to b and c of 1 has a utility of 0, since there is no overlap. If the voter agreed completely with this allocation, then she would have a utility of 2. And the allocation differs from her ideal point with respect to allocating 2 units of the budget, thereby giving a net utility of 0. If we look at an allocation of 1 to just b , then the utility is still 0. If the voter agreed completely with this allocation, then she would have a utility of 1. And the allocation differs from her ideal point with respect to 1 unit of the budget, again giving her a net utility of 0.

In the fractional setting, the ℓ_1 model is equivalent to the Overlap utility. This equivalence (see Lemma 2.6) is useful in the proof of Result 1.

1.4.2 Additive Concave Utilities. The most general model we consider is one where voters have utilities that are additive over the projects in $\mathcal{P}$. Moreover, the utility from each project is concave in the amount of allocation given to that project. It is easy to see that this model (see Definition 1.6) is a generalization of the Overlap utility model. The Overlap utility is obtained as a special case as follows: Given a voter's ideal allocation $\{w_p^v\}_{p \in \mathcal{P}}$, $f_p^v(w_p)$ is defined to be $\min\{w_p, w_p^v\}$. We use this model in Section 2.3 to prove our partial strategy-proofness result (Result 2 in the previous section).

1.5 Related Work

A wide variety of voting procedures, both ranked and non-ranked, have been studied in social choice literature (Brams and Fishburn 2002). Plurality and Approval Voting among non-ranked procedures, and Borda, Copeland, and Kemeny-Young (Levin and Nalebuff 1995) among the ranked procedures, are perhaps the most well known. In approval voting, each voter can vote for, or “approve of,” any number of candidates on the ballot. Each candidate gets a vote from every voter that voted for her, and the candidate with the most votes wins. Approval voting has been widely studied as an alternative to plurality voting, with some compelling advantages (Brams 1993). Various modifications to approval voting, such as added constraints (Brams 1990), and cumulative voting (Bhagat and Brickley 1984), which helps improve the representation of minorities or under-represented groups in committees, have been examined. In this article, we devise the right adaptation of the approval paradigm to participatory budgeting elections.

Another parallel stream of literature in social choice theory has delved deeply into the question of manipulability of voting rules. It has been shown that the only reasonable social choice functions with single/multiple winners that are not susceptible to strategic manipulation by voters are degenerate forms of dictatorships (Gibbard 1973; Satterthwaite 1975). Such impossibility results (Duggan and Schwartz 2000) rule out the existence of strategy-proof mechanisms for our setting in general.

However, strategy-proof schemes were shown to be possible in restricted domains like single-peaked preferences (Moulin 1980; Barberà et al. 1993; Nehring and Puppe 2002). Inspired by Black's Median Voter Theorem (Black 1948), Moulin (1980) characterized Generalized Median Voter Schemes as the only strategy-proof rules under single-dimensional single-peaked

¹⁰Even in the approximately integral model mentioned in Appendix A.1.

preferences. This was generalized to multiple dimensions for preferences on the Cartesian box (Border and Jordan 1983; Barberà et al. 1993) and general subsets of the Cartesian box (Barberà et al. 1997). The results of (Barberà et al. 1997) imply that, under single-peaked preferences, the only voting rules that could be both strategy-proof and non-dictatorial for our problem are Generalized Median Voter Schemes (GMVS). Unfortunately, GMVSs do not respect the budget constraint for three or more projects.

We propose an aggregation rule that finds the geometric median of the votes on the simplex that satisfies the budget constraint and show that strategy-proofness holds under specific utility models like the ℓ_1 cost model (Definition 1.4), an adaptation of dichotomous preferences to our setting. It is also interesting to compare our setting with the axioms of Arrow's Impossibility Theorem (Arrow 2012), especially the Independence of Irrelevant Alternatives (IIA). With a weaker form of IIA (Campbell and Kelly 2000), Knapsack Voting, akin to Approval Voting with dichotomous preferences, satisfies all three of Arrow's axioms. In recent work,¹¹ Freeman et al. characterize a broad class of strategy-proof "moving phantom" mechanisms under the ℓ_1 cost model (Freeman et al. 2019). Interestingly, their social-welfare maximizing method (which is equivalent to Knapsack Voting upto tie-breaking) is the unique Pareto-optimal mechanism in the above-mentioned class.

From the more recent literature in Computational Social Choice, our work is similar in spirit to some recent work on identifying the tradeoffs between the utilities from various possible societal activities (Conitzer et al. 2015). Our work also involves capturing tradeoffs, but only so far as to obtain the best allocation of the budget, and we propose voting methods that use ideas from the Knapsack Problem to do so. There has also been some work on selecting committees under weight or cost constraints given complete rankings from voters (Klamler et al. 2012). However, our work is different in that we do not get complete rankings from voters and design ways of eliciting preferences of voters. Under the utilitarian distortion framework, there has been some recent work on designing a mechanisms for PB (Benade et al. 2017; Dani 2017). However, the outcomes here could be a factor $\sqrt{m} \log m$ off of the optimal and, in general, do not extend to the fractional case. Another drawback of such mechanisms is that they are not transparent enough to be used in formal ballots.

There has also been some work on developing visualization tools for voters to make sense of different items in complex budget problems, e.g., the federal budget (Kim et al. 2016). In our setting, the items on the ballot are simple and well defined, for instance, building a park or improving streets, and can be described concisely on our digital voting platform for voters to make an informed decision.

The problem of choosing the best subset of candidates given global preference information has received some attention lately (Klamler et al. 2012; Lu and Boutilier 2011), and various methods have been proposed. For the problem of selecting optimal sets of weighted candidates given a sum weight constraint, there is some work on algorithms that, given the knowledge of a global preference relation among the candidates, compute optimal sets for various objective functions (Klamler et al. 2012). Another treatment of this problem involves a recommendation that pairs each voter to one of these candidates, thereby assuming that each voter derives satisfaction from one of the candidates in the chosen subset (Lu and Boutilier 2011). In our context, eliciting full rankings on the candidates is not feasible and, moreover, we motivate our methods in terms of social welfare, with each voter benefiting from all the chosen candidates. There is also some work (Fain et al. 2016) with an overlapping set of co-authors that studies the fairness properties of mechanisms for participatory budgeting, largely using the notion of core.

¹¹Which appeared while our manuscript was under review.

2 KNAPSACK VOTING: IMPOSING BUDGET CONSTRAINTS

In this section, we show that Knapsack Voting is strategy proof (Theorem 2.3) and welfare maximizing (Theorem 2.5) under the ℓ_1 costs model. By Lemma 2.6, the same results hold for the Overlap utility model, too. We then extend Knapsack Voting to more complex settings with revenues, deficits, or surpluses (Equation (1)) for which there is no known analog of Approval voting. We prove that Knapsack Voting is strategy proof in this setting (Theorem 2.9). We also prove a partial strategy-proofness result in the case of additive concave utilities (Theorem 2.10).

In what follows, we adopt the Fractional vote model, as in Definition 1.3. More formally, since each project is fractionally implementable, we split each project $p \in \mathcal{P}$ into c_p different “per-dollar sub-projects” $D_1^p, D_2^p, \dots, D_{c_p}^p$, and collect all the sub-projects into a set $\mathcal{P}'$. The problem then reduces to choosing B of the C candidates in $\mathcal{P}'$. Given a vote from a voter v using our Knapsack interface (Figure 2), it is guaranteed to be *consistent* after the per-dollar conversion, i.e., a set $S_v \subset \mathcal{P}'$ such that for all $p \in \mathcal{P}$, if $D_t^p \in S_v$, then $D_{t'}^p \in S_i$ for all $1 \leq t' < t$.

For any $j \in \mathcal{P}'$, define its score as $\text{score}(j) \triangleq |\{v \in \mathcal{V} : j \in S_i\}|$. We will use a *consistent deterministic tie-breaking order*: a strict ordering $<$ on $\mathcal{P}'$ such that for all $p \in \mathcal{P}$, $1 \leq t' < t \leq c_p \iff D_{t'}^p < D_t^p$ (if $j < k$, then j gets priority over k when $\text{score}(j) = \text{score}(k)$).

Definition 2.1 (Knapsack Vote).

- Each voter $v \in \mathcal{V}$ submits a consistent subset $S_v \subseteq \mathcal{P}'$, such that it satisfies a budget constraint $|S_v| = B$.
- The winning set is given by $\arg \max_{S: |S|=B} \sum_{j \in S} \text{score}(j)$, using a consistent deterministic tie-breaking order $<$.

Note that given consistent votes, and a consistent tie-breaking order, the Knapsack Vote gives a “consistent” outcome.

We also redefine Overlap utility (Definition 1.5) in terms of the per-dollar notation above—this will be of use shortly.

Overlap Utility. Let $S_v \in \mathcal{P}'$ and $S^* \in \mathcal{P}'$ be the allocations in the per-dollar sense corresponding to $\{w_p^v\}_{p \in \mathcal{P}}$ and $\{w_p^*\}_{p \in \mathcal{P}}$. Restating Definition 1.5 in terms of subprojects:

Definition 2.2. The utility of the voter v is given by $|S_v \cap S^*|$ or, equivalently, $\sum_{p \in \mathcal{P}} \min\{w_p^v, w_p^*\}$.

In this quantity, the term corresponding to p is equal to the “overlap” between w_p^v and w_p^* , and so we call this expression the Overlap Utility.

2.1 Strategy Proofness under the ℓ_1 Cost Model

Given the set of projects $\mathcal{P}$, let us say that the true preference S_v of each voter v corresponds to allocating an (integral) amount w_p^v to project $p \in \mathcal{P}$, such that $\sum_{p \in \mathcal{P}} w_p^v = B$. Any outcome that deviates from this allocation would result in some dis-utility for the voter. Let $\{w_p^*\}_{p \in \mathcal{P}}$ denote the final outcome. By Definition 1.4, the dis-utility that each voter gets is equal to the ℓ_1 distance between $\{w_p^v\}_{p \in \mathcal{P}}$ and $\{w_p^*\}_{p \in \mathcal{P}}$, i.e., $\sum_{p \in \mathcal{P}} |w_p^v - w_p^*|$.

In the above, the budget allocation to each project p is flexible and can take any integral value in $[0, c_p]$, with the votes are restricted to the simplex determined by $\sum_{p \in \mathcal{P}} w_p^v = B$. The outcome of Knapsack Voting can be defined as an ℓ_1 -median restricted to the above-mentioned simplex (follows from Lemma 2.6 and Theorem 2.5), i.e., an allocation that minimizes the sum of the ℓ_1 distances to the votes.¹² Under the Overlap utility model, by defining the outcome as the geometric ℓ_1 median, we get the following result:

¹²This works in continuous space too, with a suitably defined tie-breaking rule.

THEOREM 2.3. *For Knapsack Voting, under the ℓ_1 cost model, voting for her preferred allocation S_v is a weakly dominant strategy for voter v .*

PROOF. Let S_{-v} and $\text{score}_{-v}(\cdot)$ be the outcome determined by the votes of everyone except v using the Knapsack Voting rule (Definition 2.1). As mentioned earlier in this section, for any $i \in \mathcal{P}'$, $\text{score}_{-v}(i) = |\{u \in \mathcal{V} \setminus \{v\} : i \in S_u\}|$, where S_u denotes the vote of any voter $u \in \mathcal{V} \setminus \{v\}$ (since we are looking for a dominant strategy for voter v , S_u here can be any valid vote, not necessarily the truthful report of voter u). Assume that $T_v \neq S_v$ is a best response of i . Let the outcome after incorporating T_v be $\text{score}'(\cdot)$ and S' .

Let $j \in T_v \setminus S_v$ such that if $j = D_t^p$ for some $p \in \mathcal{P}$ (recall the definition of per-dollar sub-projects from earlier in this section), then $D_{t'}^p \notin T_v$ for all $t' > t$. Choose some $k \in S_v \setminus T_v$ such that if $k = D_z^q$ for some $q \in \mathcal{P}$, then $D_{z'}^q \in T_v$ for all $z' < z$. Such a k exists because T_v and S_v are consistent and of the same size B . Let $T'_v \triangleq T_v \cup \{k\} \setminus \{j\}$, and the outcome here be $\text{score}'(\cdot)$ and S' . We will show that T'_v is also a best response for v . If $S = S'$, then we have nothing to prove. In what follows, we will assume $S \neq S'$.

We have that $\text{score}'(j) = \text{score}(j) - 1$, $\text{score}'(k) = \text{score}(k) + 1$, and, for all $l \in \mathcal{P}' \setminus \{j, k\}$, $\text{score}'(l) = \text{score}(l)$. Note that the only change here is that the score of j decreases, and the score of k increases. As a result, for any given consistent tie-breaking order (as discussed at the head of this section), $S' \setminus S$ must be singleton. Further, we must have either $j \in S \setminus S'$ or $k \in S' \setminus S$, i.e., any change in outcome must involve either j moving from within the winning set to without, or k from without to within (or both).

If $j \in S \setminus S'$, then it must be that $S \setminus \{j\} \subset S'$, since for any $l \in S \setminus \{j\}$, $\text{score}'(l) \geq \text{score}(l)$. And since $S' \setminus S$ is a singleton set, say, it contains $m \in \mathcal{P}'$ (possibly $m = k$), the change in utility of voter v is $\mathbb{1}(m \in S_v) - \mathbb{1}(j \in S_v) = \mathbb{1}(m \in S_v) \geq 0$ (here $\mathbb{1}(\cdot)$ is the indicator function that takes the value 1 if the argument is true and 0 if not). In other words, the change in utility from removing j is 0 and that from adding any m in its place cannot be negative.

Similarly, if $k \in S' \setminus S$, then $S \setminus S' = \{m'\}$ for some $m' \in \mathcal{P}'$. And the change in utility is $\mathbb{1}(k \in S_v) - \mathbb{1}(m' \in S_v) = 1 - \mathbb{1}(m' \in S_v) \geq 0$.

By repeating this process until we only have elements in S_v , we have not decreased the utility of voter v . Therefore, the utility obtained by voting for S_v cannot be strictly dominated by that for any $T_v \neq S_v$. $\square$

By a simple modification, the above proof extends to the Overlap utility model as well.

COROLLARY 2.4. *For Knapsack Voting, under the Overlap utility model, voting for her preferred allocation S_v is a weakly dominant strategy for voter v .*

Under truthful voting, the Knapsack Voting outcome maximizes the social welfare.

THEOREM 2.5. *The truthful dominant strategy equilibrium for Knapsack Voting is welfare maximizing under ℓ_1 costs.*

Before proving the above theorem, we will show an equivalence between the ℓ_1 cost model and the Overlap utility model.

2.1.1 Overlap Utility. Let S_v and S^* be the allocations in the per-dollar sense corresponding to $\{w_p^v\}_{p \in \mathcal{P}}$ and $\{w_p^*\}_{p \in \mathcal{P}}$. In the budgeted case being treated in this section (cf. Section 2.2), the utility given by the ℓ_1 model is equal to the Overlap utility. This equivalence is of independent interest and is also useful in the proof of Theorem 2.3. We now prove this equivalence formally.

LEMMA 2.6. $|S_v \cap S^*| = B - \frac{1}{2} \sum_{p \in \mathcal{P}} |w_p^v - w_p^*|$.

PROOF. As discussed above, we know that $|S_v \cap S^*| = \sum_{p \in \mathcal{P}} \min\{w_p^v, w_p^*\}$. Partition $\mathcal{P}$ into two sets $\mathcal{P}_L$ and $\mathcal{P}_H$, where $\mathcal{P}_L \triangleq \{p \in \mathcal{P} : w_p^v > w_p^*\}$ is the set of projects for which the amount allocated in the outcome is less than that in the voter v 's preferred allocation, and $\mathcal{P}_H \triangleq \{p \in \mathcal{P} : w_p^v \leq w_p^*\}$ is the set of projects where the amount allocated in the outcome is at least as much as that in the voter v 's preferred allocation. Then we have $|S_v \cap S^*| = \sum_{p \in \mathcal{P}_L} w_p^* + \sum_{p \in \mathcal{P}_H} w_p^v$. First we note that $\sum_{p \in \mathcal{P}_L} w_p^* + \sum_{p \in \mathcal{P}_H} w_p^v = \sum_{p \in \mathcal{P}} w_p^* - \sum_{p \in \mathcal{P}_H} (w_p^* - w_p^v) = B - \sum_{p \in \mathcal{P}_H} (w_p^* - w_p^v)$. Similarly, $\sum_{p \in \mathcal{P}_L} w_p^* + \sum_{p \in \mathcal{P}_H} w_p^v = -\sum_{p \in \mathcal{P}_L} (w_p^v - w_p^*) + \sum_{p \in \mathcal{P}} w_p^v = -\sum_{p \in \mathcal{P}_L} (w_p^v - w_p^*) + B$. Therefore $|S_v \cap S^*| = \frac{1}{2}(2B - \sum_{p \in \mathcal{P}_L} (w_p^v - w_p^*) - \sum_{p \in \mathcal{P}_H} (w_p^* - w_p^v)) = B - \frac{1}{2} \sum_{p \in \mathcal{P}} |w_p^v - w_p^*|$. $\square$

PROOF OF THEOREM 2.5. The Knapsack Voting rule outputs a set S^* that maximizes $\sum_{j \in S} \text{score}(j)$ among all consistent sets S that satisfy the budget constraint.

If voters vote truthfully, then

$$\sum_{j \in S^*} \text{score}(j) = \sum_{j \in S^*} \sum_{v \in V} \mathbb{1}(j \in S_v) = \sum_{v \in V} \sum_{j \in S^*} \mathbb{1}(j \in S_v) = \sum_{v \in V} |S_v \cap S^*|,$$

and by Lemma 2.6, the proof follows. $\square$

In effect, by way of Lemma 2.6, the equivalence of ℓ_1 costs and Overlap utilities model, we have established the following:

COROLLARY 2.7. *Knapsack Voting is strategy-proof and welfare-maximizing under either ℓ_1 costs or Overlap utilities.*

We take a moment here to discuss Knapsack Voting in terms of the conditions in Arrow's Theorem. Choosing a set of projects that fit the budget can be thought of as a partial ordering $<$ between those within, and those without. In our case, if the preferences of voters between two projects, say, a and b , with respect to this partial ordering remains the same, then the final score of a and b remains the same. If the outcome is $a < b$, then changing the other preferences cannot result in b being chosen over a , i.e., change from $a < b$ to $b < a$. A similar property has been characterized as a weaker form of the Independence of Irrelevant Alternatives in the context of Approval Voting (Campbell and Kelly 2000).

We must also mention here that under the ℓ_1 cost model, Knapsack Voting is *not group strategy proof*. For example, consider the following:

Example 2.8. There are five projects $\{a, b, c, d, e\}$, four voters $\{1, 2, 3, 4\}$, and a budget of \$2. The true preferences are as follows: Voters 1 and 2 allocate \$2 to a , voter 3 allocates \$1 to b and \$1 to c , and voter 4 allocates \$1 to d and \$1 to e . If they voted truthfully, then the outcome is \$2 to a , giving a utility of 0 to both voters 3 and 4. Assuming ties are broken in the order b, d, c, e, a . Then if voters 3 and 4 both voted for \$1 to b , and \$1 to d , then the outcome changes to b and d , giving both voters a utility of 1.

2.2 An Extension to Scenarios with Revenues, Deficits, or Surpluses

A similar result can be obtained in a case where there is no hard budget and there are both expenditure terms and revenue terms by extending the framework of approval voting under dichotomous preferences. Such considerations are common in real settings that are a little beyond PB, e.g., designing mechanisms for determining the federal budget. There has been some work as a follow-up to ours (Garg et al. 2017), with an overlapping set of co-authors, which uses models with revenue and spending as categories for the design and testing of new adaptive voting mechanisms.

For our purposes, a voter proposes both how to generate revenue from among various avenues in $\mathcal{R}$ and how to spend it on various projects in $\mathcal{P}$. We will now discuss a case where the budget

is balanced. Extensions to cases where there are multiple revenue items, or where the budget is unbalanced, are easy to see.

As before, in the per-dollar sense, let $\mathcal{P}'$ be the set of per-dollar sub-projects for expenditure, and $\mathcal{R}'$ be the set of per-dollar sub-projects for revenue. We use a similar ℓ_1 cost metric for the project terms as for the budgeted case, with revenue included. Formally, we will assume that if $\{x_r^*\}_{r \in \mathcal{R}}$ is the outcome revenue level, and the voter v prefers $\{x_r^v\}_{r \in \mathcal{R}}$, then her dis-utility from the revenue term will be $\sum_{r \in \mathcal{R}} |x_r^* - x_r^v|$ and her total dis-utility will be $\sum_{p \in \mathcal{P}} |w_p^v - w_p^*| + \sum_{r \in \mathcal{R}} |x_r^* - x_r^v|$. Because we are considering *balanced budgets*, we will assume $\sum_{r \in \mathcal{R}} x_r^v = \sum_{p \in \mathcal{P}} w_p^v$ for all voters v .

Let $R_v \subseteq \mathcal{R}'$ and $S_v \subseteq \mathcal{P}'$ denote the vote of voter v satisfying $|R_v| = |S_v|$. For $j \in \mathcal{R}'$ we define $\text{score}(j) \triangleq -|\{v \in \mathcal{V} : j \notin R_v\}|$. The Knapsack Vote outcome $R^* \subseteq \mathcal{R}'$ and $S^* \subseteq \mathcal{P}'$ is defined as follows:

$$(R^*, S^*) = \arg \max_{(R, S): |S|=|R|} \left(\sum_{i \in S} \text{score}(i) + \sum_{j \in R} \text{score}(j) \right) \quad (1)$$

In this setting, revenue corresponds to voters paying the government in fees/taxes, and hence we make the score corresponding to revenue terms negative. We will need a consistent tie-breaking rule here as well. But since the size of the outcome is not fixed, we can just break ties in favor of the bigger (or smaller) sets. We state a similar result as in Theorem 2.3 for this case:

THEOREM 2.9. *With a balanced budget, Knapsack Voting is strategy-proof under the ℓ_1 cost model.*

We defer the proof to the appendix (A.2), as it is similar to the proof of Theorem 2.3 and not crucial to the exposition here.

We can also extend this to *unbalanced* settings by including the budget deficit (expenditure - revenue) as part of the voters' preferences. If $\Delta_v = |S_v| - |R_v|$ is the preferred deficit of voter v , and the deficit in the outcome is Δ^* , then the voter incurs an additional dis-utility of $|\Delta^* - \Delta_v|$, i.e., how much the deficit in the outcome differs from her preferred level. As such, it is equivalent to adding an additional revenue term.

2.3 “Partial” Strategy-proofness Under Additive Concave Utilities

In the case with additive concave utilities, we illustrate a weaker, yet interesting, property that we call “partial” strategy-proofness. Consider a focal voter i responding to the votes of all others. Assume that she has full knowledge about the how the others voted in aggregate. If S_{-i} denotes the cumulative votes of all voters except i , then she knows $\mathcal{W}_{-i} \subseteq \mathcal{P}'$, the winning set as determined by S_{-i} . Let $\mathcal{W}(S_i, S_{-i})$ denote the set of winners if her vote S_i is added. A *best response* for i is a consistent set

$$S_i^* \in \arg \max_{S_i \subseteq \mathcal{P}': |S_i|=B} \sum_{j \in \mathcal{W}(S_i, S_{-i})} v'_{i,j},$$

where $v'_{i,j}$ is the utility of voter i for sub-project $j \in \mathcal{P}$. As mentioned in Definition 1.6, the utility from per-dollar sub-projects is monotone non-increasing (as the utility from a project is concave in the amount allocated), i.e., if $x < y$ and $a = D_x^p$ ¹³ and $b = D_y^p$ for some $p \in \mathcal{P}$, then $v'_{i,a} \geq v'_{i,b}$. With respect to voter i , we say a candidate k *dominates* j if and only if

$$k \in \mathcal{W}_{-i} \text{ and } v'_{i,k} \geq v'_{i,j}.$$

A consequence of the budget constraint is that it allows for *partial strategy-proofness* in the best response of a focal voter responding to all other votes.

¹³ D_x^p is the x th dollar subproject in project p .

THEOREM 2.10 (PARTIAL STRATEGY-PROOFNESS). *Under Knapsack Voting, there exists a best response S_i^* such that if k dominates j , and $j \in S_i^*$, then $k \in S_i^*$.*

The proof is not critical to the exposition here, and we defer it to the appendix.

We can think of $\mathcal{W}_{-i}$ as representing the candidates that i thinks are popular. And projects with a higher “value-for-money” are preferred from her perspective. The theorem states that a simple way for a voter to act in her best interest is to vote for projects that to her are both *popular* and *favorite*. This notion is similar in spirit to the ideal of *sincerity* in approval voting (Brams and Fishburn 1978; Niemi 1984).

The reason partial strategy-proofness holds under Knapsack Voting is that the voters face the same constraints as the outcome they are collectively deciding. It is interesting to note that this property does not hold under K-approval voting, as there is a mismatch between the constraints on the voters (choosing a fixed number of projects k) and those on the outcome (a fixed budget B). We illustrate this with a small example.

Example 2.11. Consider five projects a, b, c, d, e of costs \$200, \$100, \$100, \$100, and \$200 and using the 2-approval rule. The budget is \$400 and they poll 100, 50, 50, 50, 20 votes, respectively, without counting in i ’s vote. The tie is broken in the order e, d, c, b, a , and in this case d wins over b, c . Let us say that i ’s utilities for these are 500, 100, 150, 200, 500, respectively. Based on the above, i ’s best response is to not vote for a but for c, d .

2.4 Maximum Likelihood Interpretation

One way of looking at voting rules is as follows: There exists a “ground-truth” outcome, and each voter has a noisy perception of it. And the the voting rule is the Maximum likelihood estimator of the “ground truth” given any realization of votes. The Knapsack Voting rule can be interpreted in this way since it belongs to a family of voting rules known as scoring rules (Conitzer and Sandholm 2012). Selecting subsets of winners based on the MLE approach has received some attention lately (Procaccia et al. 2012) wherein voting rules to select subsets based on their performance on various metrics with respect to noisy comparisons or rankings drawn from a Mallows model.

In this section, we will use the per-dollar approach and explicitly construct a natural noise model for the votes for which it is the Maximum Likelihood estimator. The noise model we construct is similar to the Mallows model but defines a distribution over subsets directly as opposed to rankings.

Definition 2.12 (Noisy Knapsack Vote Model). There is a “ground-truth” set $S^* \in \mathcal{P}'$, which satisfies $|S^*| = B$. Each vote S_i is drawn i.i.d. according to a distribution that is given by

$$\begin{aligned} \Pr\{S_i|S^*\} &\propto \exp(|S^* \cap S_i|), & \text{if } |S_i| \leq B \\ \Pr\{S_i|S^*\} &= 0, & \text{otherwise} \end{aligned}$$

The quantity $|S^* \cap S_i|$ is equal to the number of dollars in the allocation given by S^* that agrees to that given by voter i in her vote S_i and by Lemma 2.6 is related to the Overlap/ ℓ_1 utility we discussed previously.

By taking the logarithm of the probabilities $\Pr(S_i|S^*)$, it is easy to see that the Maximum Likelihood estimate of the “ground-truth” S^* given all the votes is the set $S \subseteq \mathcal{P}'$ satisfying $|S| = B$ that maximizes

$$\frac{1}{|V|} \sum_{i \in V} |S \cap S_i|. \quad (2)$$

THEOREM 2.13. *The Knapsack rule returns the maximum likelihood estimate under the Noisy Knapsack Vote Model.*

PROOF. Let us rewrite the quantity in Equation (2) as

$$\sum_{i \in \mathcal{V}} |S \cap S_i| = \sum_{i \in \mathcal{V}} \sum_{j \in S} \mathbb{1}(j \in S_i) = \sum_{j \in S} \sum_{i \in \mathcal{V}} \mathbb{1}(j \in S_i) = \sum_{j \in S} \text{score}(j),$$

where $\mathbb{1}(\cdot)$ is a $\{0, 1\}$ variable that takes on a value 1 when the statement in the argument is true and 0 otherwise. This quantity is maximized by picking B candidates from $\mathcal{P}'$ that have the highest score (see Definition 2.1), which is essentially what the Knapsack rule chooses. $\square$

3 VOTING BASED ON VALUE-FOR-MONEY

We have seen that Knapsack Voting has many advantages with respect to its strategic properties and implementation in Participatory Budgeting elections. We now turn to another way of eliciting voters' preferences in this setting: *value-for-money* comparisons. Given a single agent's Knapsack Problem, one way of computing the optimal solution is to order the items according to their value-to-size ratio and pick the higher-ranked ones until the knapsack capacity is used up. This order can be ascertained by comparing pairs of projects according to their value-to-size ratio. Asking voters to compare/rank projects based on their *value-for-money* is a natural analog (see Figure 3) of this idea in a setting where multiple agents together have to decide an outcome.

Unlike Knapsack Voting, aggregation schemes based on value-for-money comparisons cannot be guaranteed to have good strategic properties. Further, these schemes are not as transparent as Knapsack Voting, since different information is elicited from each voter, and the aggregation method is not as straightforward as that of Knapsack Voting. Value-for-money schemes also do not extend naturally to settings with revenue, deficits, and surpluses.

Despite these difficulties, they are useful in practice, because they can be used to design paper ballots with Knapsack Voting and also elicit the aggregate preferences of voters to make empirical observations from data. We will discuss these presently.

In addition, value-for-money schemes have the following potential advantages:

- a smaller cognitive load on voters: especially with large ballots (see Section 4.0.2)
- aggregation in cases where the budget is not known or fixed *a priori* (perhaps using Kemeny-Young like ranking rules)

We will delve into these in more detail in the next section.

3.1 Value-for-money Rankings and Paper Ballots

As mentioned before, some elections that employ digital voting also require a corresponding paper ballot. This was the case in the Youth Lead The Change 2016 PB election held in Boston, a set of projects was put to vote to youth between the ages 12 and 25. We implemented Knapsack Voting as the official ballot process in this election. Since it is tedious to do Knapsack Voting on paper, we designed a paper ballot that asked voters to rank their top four projects according to value-for-money (Figure 7). From a total of around 4,000 voters, about 10% used the digital Knapsack interface, and the rest used the paper ballot with ranking.

Here we ask each voter to rank the projects on the ballot taking both costs and benefits into account.¹⁴ Voters take into account costs and benefits into account, and ranking reinforces this consideration. Interpreting these votes as Knapsack Votes, we observe that the outcome from the rankings was exactly the same as that of Knapsack Voting. This gives us a way of using rankings to solve the budgeting problem in practice.

¹⁴There is an implicit assumption here that taking both costs and benefits into account is the same as considering value-for-money.

Step 1: Million Dollar Vote

Rank your top projects in order at the bottom of this page.
Take both costs and benefit into account.
The most highly ranked projects will be funded.

PROJECT A

\$80,000

More Trash Cans and Recycling Bins

Install newer trash cans and recycling bins to make our neighborhoods cleaner - Mattapan, Dorchester, Roxbury

PROJECT B

\$475,000

Gym Renovation at K-8

Improve the gym facility for students in elementary schools - students in Roxbury - Higginson/Lewis K-8

PROJECT C

\$135,000

Job and Resource Finder App

Create an app that lets teens find jobs, programs, free events, and scholarships - Citywide

PROJECT D

\$500,000

Parks are for Everyone

Up to 10 new park renovations will have inclusive play elements for people with disabilities - Dorchester, South End, Roxbury, Jamaica Plain, Rosindale

PROJECT E

\$300,000

Performance Space in Chinatown

Replace seats and renovate the auditorium to support the theatre performances and church that use the space throughout the week - Josiah Quincy School

PROJECT F

\$500,000

Multi-use space for Chinatown

Fix the gym divider, scoreboard, and bleachers in order to make the space more usable for students and the community - Josiah Quincy School

PROJECT G

\$150,000

Wicked Free Wifi 3.0

Bring internet and phone charging to side-walks and bus stops in key locations across Boston - Citywide

PROJECT H

\$280,000

Turn up on Community Centers

Install a dance floor, coding, a DJ booth, and sound system in order to throw parties - ICYF Curtis Hall Community Center

PROJECT I

\$100,000

Link up: Digital Billboards in Schools

Install TV screens in schools that display information about jobs, programs, and other resources for students in BPS High Schools

PROJECT J

\$25,000

Sunshine and Books

Develop an app that make it easy to find outdoor study spaces in parks across Boston - Citywide

YOUR TOP FOUR PROJECTS:

#1 CHOICE

#2 CHOICE

#3 CHOICE

#4 CHOICE

Fig. 7. Value-for-money ranking paper ballot for Knapsack Voting–PB Boston 2016.

[illegible]

Fig. 8. Cambridge 2014: Comparison table.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
1 Afterschool Sports Program - Soccer & Basketball		58	54	5	73	68	74	59	77	66	8	5	64	75	76	9	
2 Safer Streets - Bike Lanes & Markings	42	61	61	48	44	62	51	49	47	64	63	65	62	6	77	84	9
3 U Help Hello Kids Join - Afterschool Tech Program	6	46	39	46	43	58	61	54	62	59	72	65	7	8	63	68	9
4 Pedestrian Safety and Traffic Calming	5	52	48	38	53	63	53	63	5	55	75	64	65	7	82	93	
5 Enhanced Disaster Response	27	56	57	52	51	66	48	52	65	43	55	71	78	73	83	96	
6 Waterfront History Arts Park	32	38	42	67	49	49	48	54	55	58	48	57	81	7	82	92	97
7 Vallejo Firefighter Youth Academy	39	49	39	37	34	51	51	57	56	68	49	54	68	68	7	66	9
8 Beautify Downtown and Georgia Park	21	46	37	37	52	43	5	5	58	46	5	55	68	75	6	93	
9 Enhance Neighborhood Watch - Outreach Van for IPV	41	33	38	4	4	4	4	4	4	4	4	4	4	4	4	4	4
10 Path 2 Valley's Future! Waterfront Promenade Ramp	23	36	41	4	35	44	32	43	52	67	57	69	58	61	73	79	
11 School Roads Repair	32	38	28	45	57	41	51	54	51	33	43	35	69	67	67	7	
12 Mobile Nutritional Cooking Carts for Kids	2	35	35	25	45	52	46	5	65	43	57	59	48	42	55	66	9
13 Safety Through School Fences	2	38	3	36	29	43	52	45	49	31	45	41	5	45	51	56	81
14 Improve Valley's Image by Upgrading Public Events	36	4	2	35	23	28	33	28	33	15	32	31	52	55		6	5
15 Commercial/Event Kitchen at OVRD Community Center	25	23	37	3	19	3	3	25	29	39	33	38	4	4	9		
16 Waterfront Flags and Welcome Signs	24	16	24	18	18	18	34	22	3	27	33	5	44	35	61	63	
17 Fountains of Native Flora	11	68	14	67	68	03	03	1	31	26	21	3	32	19	3	47	37

Fig. 9. Vallejo 2015: Comparison table.

3.2 Value-for-money Comparisons

We can elicit fine-grained information about the preferences of voters between pairs of projects by doing value-for-money comparisons between randomly chosen projects.

Definition 3.1 (Value-for-money comparison). For each pair of projects $\{j, k\}$ from $\mathcal{P}$ shown to her, voter i chooses a winner $w_i(\{j, k\}) = \arg \max_{t \in \{j, k\}} \frac{v_{i,t}}{c_t}$.

Doing these comparisons in practice, leads us to an interesting empirical observation of the structure in the aggregate preferences of voters. We observe that compiling this data from our experiments reveals a (nearly) transitive majority relation among the projects (Figure 8). By this we mean that there exists a Condorcet winner, and upon removal of that candidate, there again exists a Condorcet winner, and so on.

In Figures 8 and 9, we show the aggregate strength of each comparison across the two elections Cambridge 2015 and Vallejo 2015. More details about the two elections are in the next section. The number in the cell denotes the fraction of comparisons where the row project beats the column

project. Green represents a fraction greater than 0.5, and red represents a fraction smaller than 0.5. The darker the green (red), the closer it is to 1 (0).

Given this structure, there is a clear order among projects based on the aggregate strength of the comparisons. In fact, any Condorcet rule leads to the same outcomes on such data. If this structure does not hold, then the aggregation of these comparisons becomes more complicated.

3.2.1 Comparing Outcomes of Knapsack and K-approval. We will define a quantitative measure for any outcome that represents the level of its agreement with the pairwise comparisons. To do this, we generalize the Borda rule and define a Borda score for outcomes that are sets of projects.

In our experiments, we compare one project against another as shown in Figure 3, and to each voter we present randomly chosen pairs so as to ascertain the aggregate preferences of voters. Given the results of these comparisons, let $n(j, k)$ denote the number of voters that chose j over k .

We now build on the standard for a standard definition of Borda's rule (Young 1974), and its classical interpretation as an MLE (Young 1988). Generalizing the Borda score for a single winner, we define

Definition 3.2 (Set-Borda score). For any set of projects $S \subseteq \mathcal{P}$, the Set-Borda score of S is given by $\frac{1}{C(M-C)} \sum_{j \in S} \sum_{k \notin S} c_j c_k (n(j, k) - n(k, j))$, where $C = \sum_{p \in S} c_p$ and $M = \sum_{p \in \mathcal{P}} c_p$.

We can interpret a pairwise value-for-money comparison between two projects j and k , as a pairwise relation between a dollar sub-project of j and a dollar sub-project of k . The Set Borda score of a set of S corresponds to the average number of such dollar versus dollar comparisons that agree, minus those that disagree, with the partition induced by it. We can use this score, as a measure of social welfare that embodies the essence of the Borda rule, to empirically compare the outcome of Knapsack Voting with that of K-approval.

4 OUR DIGITAL PLATFORM AND RESULTS FROM OUR EXPERIMENTS

Our work with Participatory Budgeting began with a partnership with Chicago's 49th Ward to develop a digital voting system for their election. Since then we have worked with almost a dozen cities/districts over the past couple of years (see <http://pbstanford.org>), but we will primarily look at data collected from the following elections:

- Boston 2016—Youth Lead the Change
- Boston 2015—Youth Lead the Change
- Cambridge 2015
- Cambridge 2014
- Vallejo 2015
- New York City District 5 2015 (NYC5)
- New York City District 8 2015 (NYC8)

The total number of voters in each of these elections were 4,176, 2,600, 3,273, 2,194, 1,834, 704, and 271, respectively.

In the Boston 2016 election, we implemented Knapsack Vote as the official ballot mechanism. In the other elections, besides implementing an interactive interface for the official K-approval vote (Figure 1), our goal has been to experiment with Knapsack Voting and value-for-money comparisons and use the data so collected to complement a theoretical understanding of these methods. After the K-approval vote used for the formal election, we ask the voters to participate in our experiments on either Knapsack Voting or value-for-money comparisons or both.

- In Cambridge 2015 and Boston 2015, we showed each voter either Knapsack and value-for-money experiments with a 50% chance.





Fig. 10. Cumulative fraction of total votes versus cost of projects: Cambridge 2015.

Fig. 11. Cumulative fraction of total votes versus cost of projects: NYC District 8.

Table 2. Average Cost of Winning Projects, as a Fraction of the Budget

	<i>K-approval</i>	<i>knapsack</i>
NYC District 5	0.18	0.14
NYC District 8	0.20	0.12
Boston 2015	0.27	0.18
Cambridge 2015	0.15	0.10

- In NYC5 and NYC8, we did only the Knapsack Vote.
- In Cambridge 2014 and Vallejo 2015 we did only pairwise comparisons.

The Knapsack interface (Figure 2 in Section 1.3) has a live budget bar that shows how much of the budget has used up with the current selection. For the value-for-money comparisons, we present each voter with a fixed number of pairs of projects chosen uniformly at random without replacement and ask them the following question (in the spirit of Definition 3.1): “Which of these projects gives a higher benefit to the community per dollar spent?” (see Figure 3). We must mention that only a subset of the entire set of voters actually chose to take part in our experiments, and we report this percentage in Tables 5 and 4. This procedure was approved by Stanford University’s Institutional Review Board.

4.0.1 Cost Consideration under Knapsack. Our data suggest that there is a bias toward projects of larger costs in the K-approval method, as compared to the Knapsack method. We empirically verify this effect in two ways. First, we present data from Cambridge and NYC District 8, where a budget of \$600,000 and 6-approval and \$1,000,000 and 5-approval were used, respectively. In Figures 10 and 11, we lay out the projects in descending order of cost and plot the cumulative fraction of votes for projects above every cost threshold. We then compare K-approval and Knapsack against the uniform distribution. We see that this function for K-approval dominates that for Knapsack, which means that costlier projects are over-represented, thereby supporting our hypothesis. While this observation is not direct evidence that Knapsack Voting is better than K-approval, it does suggest a significant qualitative difference between the outcomes of the two methods.

Second, we look at the average cost of the winning projects under each method. Table 2 shows the average cost of the winning projects (normalized by the total budget) in each of those elections. On average, across the three places, there is a reduction of about 30% in the average cost of the

Table 3. Timing Data for K-approval
N = number of projects on ballot;
T = median completion time

City	K	N	T
Boston 2015	4	10	61s
Cambridge 2015	6	23	213s

Table 4. Timing data for value-for-money
comparisons n = number of comparisons;
T = median completion time

City	n	T	% of voters
Boston 2015	4	14s	10
Cambridge 2015	4	53s	40

Table 5. Timing Data for Knapsack Voting N = number
of projects; T = median completion time

City	Budget	N	T	% of voters
Boston 2015	\$1,000,000	10	70s	7
Cambridge 2015	\$600,000	23	115s	30
Boston 2016	\$1,000,000	10	86s	100

winning projects. The above two observations clearly suggest that Knapsack Voting leads to voters' being more frugal while choosing which projects to vote for.

4.0.2 Timing Data. The data from Tables 3 and 5 suggest that the Knapsack interface is not much more time-consuming than the K-approval interface. Of course, since the Knapsack interface follows the official K-approval interface, and so the voters were familiar with the projects when they attempted the Knapsack Vote. Even if the time taken by Knapsack were the sum of times for both in our experiments, we can see that it is very reasonable. This is borne out in the Boston 2016 election (see Table 5), where the number of projects was the same as in Boston 2015 and Knapsack was presented first as the official ballot.

Here are some more interesting observations on the ease of each voting method based on Tables 3, 4, and 5:

- we presented to each voter either Knapsack or value-for-money comparisons with a 50% chance in both Boston 2015 and Cambridge 2015, and the percentage of voters that completed the comparisons was at least a third greater than that of Knapsack;
- in Boston 2015, the median time taken by the voters was 61s for K-approval vote (K = 4 of 10 projects on the ballot), 70s for Knapsack, and 14s for value-for-money comparisons of four pairs. A similar trend is seen in the Cambridge experiment.

All this suggests that Knapsack Vote and K-approval have comparable times. The *value-for-money* comparisons involve a smaller cognitive load than the other voting methods, especially with a large number of projects as seen in Cambridge 2015.

4.1 Comparison of Knapsack and K-approval against Value-for-money Comparisons

Using the data from random pairwise comparisons in the value-for-money experiment, and based on the Set Borda score (Definition 3.2), we calculate the average number of comparisons that agree and disagree with the winning sets as determined by the K-approval and Knapsack methods. If we picked a dollar allocated in the winning outcome, and a dollar not in it, both at random, then the numbers in Table 6 are a measure of the average fraction of votes that agree with the outcome minus those that disagree. We see that on this score, *Knapsack* does better than the *K-approval*, thereby indicating a higher level of agreement with the voters' preferences. We

Table 6. Comparison of Knapsack and K-approval Based on Agreement with Pairwise Comparisons: Knapsack Has a Greater Level of Agreement

	K-approval		Knapsack	
	Agreement	Standard error	Agreement	Standard error
Boston 2015	0.20	0.047	0.26	0.045
Cambridge 2015	0.30	0.022	0.40	0.021

also note the standard error in the sample mean calculated above, assuming the sample mean is normally distributed. We see that the difference between the sample means of Knapsack and K-approval is greater than the standard error.

5 CONCLUSIONS AND ONGOING WORK

Knapsack Voting and value-for-money comparisons are intuitive ways of eliciting voters' preferences for budgetary decisions. The Knapsack Vote admits interesting strategic properties: In particular, it is strategy-proof under a natural utility model that depends on the overlap (and equivalently the ℓ_1 distance) between the outcome and the voters' true preferred allocations. It can also be extended to more complicated settings with revenues, deficits, and surpluses. However, Knapsack Voting is not very practical when there are a large number of projects. While *value-for-money* comparisons have some drawbacks, they provide a way of eliciting the voters' preferences with a small cognitive load, especially in the case of large ballots. The fact that our schemes do better on many different measures finds support in the data we collected from participatory budgeting elections in various cities/municipalities. All our schemes are amenable to implementation using interactive digital tools, thereby enhancing the ability of voters to make more informed decisions in *participatory budgeting*. We have been able to make some initial progress along implementing Knapsack Voting as the official ballot, and we hope that this article makes a strong case for its wider adoption in practice.

APPENDIX

A.1 Extending Results (Approximately) to an Integral Model

The integral version of Knapsack Voting is defined as follows:

Definition A.1 (Integral Knapsack Vote).

- (1) Each voter $v \in V$ votes for a (integral) subset $S_v \in P$, such that it satisfies the budget constraint $\sum_{p \in S_v} c_p \leq B$.
- (2) The projects are arranged in decreasing order of number of their approval scores, which for any project p is given by $\#\{v \in V : p \in S_v\}$.
- (3) The projects are chosen in this order till there is not enough budget to include the next project.

Our results, for example, Theorem 2.3, do not hold in the above integral model. The following example (similar ones can be given for Theorems 2.5 and 2.10) shows us why:

Example A.2. We have a budget of 5 and three projects a , b , and c of cost 2, 2, and 3, respectively. Let us say that ties are broken in favor of projects with larger cost whenever possible. For a voter v , say that her favorite outcome is b and c . Without v 's vote being tallied, assume a and b are winning—with scores 15, 11, and 10 for a , b , and c . Now if she votes for b and c , then the outcome

is a and b . But if she switches to a and c , the outcome becomes a and c , leading to higher utility (since she likes c more than b).

However, our results do hold in an *approximately integral* model, i.e., with some assumptions.

Definition A.3 (Approximately Integral Knapsack Vote).

- (1) Each voter $v \in V$ votes for a (integral) subset $S_v \in P$, such that it satisfies the budget constraint $\sum_{p \in S_v} c_p \leq B$.
- (2) The projects are arranged in decreasing order of number of their approval scores, which for any project p is given by $\#\{v \in V : p \in S_v\}$.
- (3) The projects are chosen in this order until there is not enough budget to include the next project. And the next one in line is fractionally implemented.

In other words, although the votes are constrained to be integral, *approximately integral outcomes* are allowed—the first project that gets the most votes and cannot be fully funded is allowed to be fractionally implemented, so that the entire budget is used up. When a project j in S_v (voter v 's favorite set) is funded fully, the utility derived is equal to c_p . If it is fractionally funded, then the utility is pro-rated appropriately. It may or may not be realistic to partially fund projects (e.g., renovate only one floor of a library rather than the entire building), but partial funding and the requirement of spending all of the budget are both common in practice.

Note: As mentioned in Section 1.4, the ℓ_1 costs model does not make sense in a fully integral model (on account of violating free disposal). However, with approximately integral outcomes, both ℓ_1 costs and Overlap utilities satisfy free disposal.

For Lemma 2.6, it is easy to see that we have the following analog in the approximately integral case.

$$\text{COROLLARY A.4. } |S_v \cap S^*| = \frac{B + |S_v|}{2} - \frac{1}{2} \sum_{p \in P} |w_p^v - w_p^*|.$$

Using the above, and by a simple modification of the proofs of Theorems 2.3, 2.5, and 2.10, we can see that they extend to the approximately integral model.

OBSERVATION 1. *The results of Theorems 2.3, 2.5, and 2.10 hold under the approximately integral model.*

The above observation means that our results hold approximately in the strictly integral model under Overlap utilities (see note above). First, let us consider Theorem 2.3. For example, let us say we have a voter whose favorite (integral) set of projects is S_v . Then under the approximately integral model, voting for S_v is a weakly dominant strategy for voter v . This implies that the corresponding approximately integral outcome, composed of some integral set S^* and a part of project p^* , yields at least as much utility as any other approximately integral outcome that v can obtain by voting some other integral set S'_v . Consequently, we have under the integral model that the set S^* yields at least as much utility as any other integral set (that v can obtain by changing her vote) minus the value of the project p^* . A similar observation can be made for Theorems 2.5 and 2.10.

OBSERVATION 2. *The results of Theorems 2.3, 2.5, and 2.10 hold approximately, up to the value of one project, under the integral model.*

We are drawing on ideas such as ϵ -strategyproofness (strategy-proof upto a utility value of ϵ), and relaxations in combinatorial fair division such as envy-freeness to one good (Caragiannis et al. 2016; Budish 2011; Lipton et al. 2004).

A.2 Proof of Theorem 2.9

Let R_{-v}, S_{-v} and $\text{score}_{-v}(\cdot)$ be the outcome determined by the votes of everyone except v using the Knapsack Voting rule (Equation (1)). Assume that $Q_v \subseteq \mathcal{R}'$ and $T_v \subseteq \mathcal{P}'$ be a best response for voter v such that $Q_v \neq R_v$ and $T_v \neq S_v$. Let the outcome after incorporating Q_v, T_v be $\text{score}(\cdot)$ and S .

We will discuss the case when $|Q_v| = |T_v| < |S_v| = |R_v|$. The other case of $|Q_v| = |T_v| \geq |S_v| = |R_v|$ follows analogously (equality is similar to the proof of Theorem 2.3).

Choose some $j \in S_v \setminus T_v$ such that if $k = D_t^p$ for some $p \in \mathcal{P}$, then $D_{t'}^p \in T_v$ for all $t' < t$. Such a k exists because of consistency and the fact that $|T_v| < |S_v|$. Let $T'_v \triangleq T_v \cup \{j\}$. Similarly, choose the $k \in R_v \setminus Q_v$ such that if $k = D_z^q$ for some $q \in \mathcal{R}'$, then $D_{z'}^q \in T_v$ for all $z' < z$ and define $Q'_v \triangleq Q_v \cup \{k\}$. Let the outcome here be $\text{score}'(\cdot)$ and R', S' . We will show that Q'_v, T'_v is also a best response for v .

If $R = R'$ and $S = S'$, then the utility is unchanged and we have nothing to prove.

We have that $\text{score}'(j) = \text{score}(j) + 1$, $\text{score}'(k) = \text{score}(k) + 1$, and for all $l \in \mathcal{P}' \cup \mathcal{R}' \setminus \{j, k\}$, $\text{score}'(l) = \text{score}(l)$. Note that the only change from $\text{score}(\cdot)$ to $\text{score}'(\cdot)$ is that the score of j and k increases. So the outcomes must satisfy $R' \supseteq R$ and $S' \supseteq S$ and $|R'| = |S'| \leq |S| + 2 = |R| + 2$. For any given tie-breaking rule, if $S' \neq S$, then we must have $j \in S' \setminus S$, $k \in R' \setminus R$, or both, i.e., any change in outcome must involve either j or k moving from outside the winning set to within.

If $j \in S' \setminus S$, then there is a corresponding sub-project that is added to R to maintain the budget balance, say, $m \in R' \setminus R$. The change in utility of voter v is $\mathbb{1}(j \in S_v) - \mathbb{1}(m \notin R_v) = 1 - \mathbb{1}(m \notin R_v) \geq 0$. In other words, the change in utility from adding j to S is 0, and any m to R to maintain budget balance cannot be negative.

Similarly, if $k \in R' \setminus R$, then there is a corresponding $m' \in S' \setminus S$. And the change in utility is $-\mathbb{1}(k \notin R_v) + \mathbb{1}(m' \in S_v) = 0 + \mathbb{1}(m' \in S_v) \geq 0$.

By repeating this process, we reach a point at which we have a best response Q_v, T_v equal in size to R_v, S_v , respectively. From here we do a procedure similar to the proof of Theorem 2.3 until we only have elements in S_v . In this entire process, we do not decrease the utility.

To prove that the outcome is welfare maximizing, note that it is given by (R^*, S^*) , where R^* and S^* are both consistent and $|R^*| = |S^*|$, which maximizes the following:

$$\begin{aligned} \sum_{i \in S} \text{score}(i) + \sum_{j \in R} \text{score}(j) &= \sum_{i \in S} \sum_{v \in \mathcal{V}} \mathbb{1}(i \in S_v) - \sum_{j \in S} \sum_{v \in \mathcal{V}} \mathbb{1}(j \notin R_v) \\ &= \sum_{v \in \mathcal{V}} \sum_{i \in S} \mathbb{1}(i \in S_v) - \sum_{v \in \mathcal{V}} \sum_{j \in R} \mathbb{1}(j \notin R_v) = \sum_{v \in \mathcal{V}} |S_v \cap S| - |R \setminus R_v|, \end{aligned}$$

and the last quantity in the above is the welfare according to the Overlap Utility Model.

A.3 Proof of Theorem 2.10

We will use the per-dollar approach (see definition in Section 2), i.e., each voter $i \in \mathcal{V}$ submits a vote $S_i \subseteq \mathcal{P}'$ such that S_i is consistent and $|S_i| = B$. Let us first reinterpret the notation from Section 2.3 per-dollar, and restate Theorem 2.10 in technical terms.

Consider a focal voter i responding to the votes of all others. Assume that she has full knowledge about the how the others voted in aggregate. If S_{-i} denotes the cumulative votes of all voters except i , then she knows $\mathcal{W}_{-i}$, the set of winners as determined by S_{-i} . Let $\mathcal{W}(S_i, S_{-i})$ denote the set of winners if her vote S_i is added.

A best response for i is then defined as a vote S_i^* that satisfies

$$S_i^* = \arg \max_{S_i \in \mathcal{A}} \sum_{j \in \mathcal{W}(S_i, S_{-i})} v'_{i,j}, \quad (3)$$

where $A \triangleq \{S : S \text{ is consistent, } |S| = B\}$, and $v'_{i,j}$ is the utility of voter i from sub-project $j \in \mathcal{P}'$.

With respect to voter i , we say a candidate $q \in \mathcal{P}'$ *dominates* $p \in \mathcal{P}'$ if and only if

$$\begin{aligned} & -q \in \mathcal{W}_{-i}, \text{ and} \\ & -v'_{i,q} > v'_{i,p}. \end{aligned}$$

Denote collection of candidates that dominate p with respect to voter i as $\Lambda_{i,p}$, i.e., $\Lambda_{i,p} \triangleq \{j \in W_{-i} : v'_{i,j} > v'_{i,p}\}$.

Let us say that $p \in S_i^*$ and $\Lambda_{i,p} \not\subseteq S_i^*$. We will first claim that $\Lambda_{i,p} \subseteq \mathcal{W}(S_i^*, S_{-i})$.

Let $j \in \Lambda_{i,p}$. We have the following two possible cases:

- (1) If $j \in S_i^*$: Since $\Lambda_{i,p} \subseteq W_{-i}$, then we have $j \in S_i^* \cap \mathcal{W}_{-i}$, and, consequently, $j \in \mathcal{W}(S_i^*, S_{-i})$.
- (2) Else, if $j \notin S_i^*$: Assume $j \notin \mathcal{W}(S_i^*, S_{-i})$. Then S_i^* cannot be a best response, because by switching her vote from S_i^* to $(S_i^* \setminus \{p\}) \cup \{j\}$, i can make j win instead of p , and this strictly increases her total utility (since $v_{i,j} > v_{i,p}$).

In either case, $j \in \mathcal{W}(S_i^*, S_{-i})$.

Now, if $S_1 \triangleq S_i^* \setminus \Lambda_{i,p}$, then $|S_1| = B - |S_i^* \cap \Lambda_{i,p}|$. And since we have proved $\Lambda_{i,p} \subseteq \mathcal{W}(S_i^*, S_{-i})$, it follows that $|S_1 \cap \mathcal{W}(S_i^*, S_{-i})| \leq B - |\Lambda_{i,p}|$. These two facts together imply the following:

$$\begin{aligned} |S_1 \setminus \mathcal{W}(S_i^*, S_{-i})| &= |S_1| - |S_1 \cap \mathcal{W}(S_i^*, S_{-i})| \\ &\geq (B - |S_i^* \cap \Lambda_{i,p}|) - (B - |\Lambda_{i,p}|) \\ &\geq |\Lambda_{i,p}| - |S_i^* \cap \Lambda_{i,p}| \\ &= |\Lambda_{i,p} \setminus S_i^*|. \end{aligned}$$

Hence, there exists $S_2 \subseteq S_1 \setminus \mathcal{W}(S_i^*, S_{-i})$ such that $|S_2| = |\Lambda_{i,p} \setminus S_i^*|$.

Let $S_i^{**} = (S_i^* \setminus S_2) \cup \Lambda_{i,p}$. Clearly, $|S_i^{**}| = B$ (since $S_2 \cap \Lambda_{i,p} = \emptyset$), and so it is a valid vote for voter i .

Also $\mathcal{W}(S_i^{**}, S_{-i}) = \mathcal{W}(S_i^*, S_{-i})$, since we have replaced $S_2 \subseteq \mathcal{W}(S_i^*, S_{-i})^c$ with $\Lambda_{i,p} \setminus S_i^* \subseteq \mathcal{W}(S_i^*, S_{-i})$. Because S_i^* is a best response, so is S_i^{**} .

Note that

$$S_2 = S_i^* \setminus S_i^{**} \subseteq \mathcal{W}(S_i^*, S_{-i})^c. \quad (4)$$

Now if there is a $p' \in S_i^{**}$ such that $\Lambda_{i,p'} \not\subseteq S_i^{**}$, we can do a similar replacement procedure to define another best response $S_i^{***} = (S_i^{**} \setminus S'_2) \cup \Lambda_{i,p'}$, such that $S'_2 \subseteq \mathcal{W}(S_i^{**}, S_{-i})^c$ (from Equation (4)). Since $\Lambda_{i,p} \subseteq \mathcal{W}(S_i^{**}, S_{-i})$, this implies that $S'_2 \cap \Lambda_{i,p} = \emptyset$ and so S_i^{***} includes both $\Lambda_{i,p}$ and $\Lambda_{i,p'}$. By a series of replacements, we have the best response as required by the theorem. $\square$

ACKNOWLEDGMENT

We thank Tim Roughgarden for useful discussions.

REFERENCES

- Francisco J. Andr , M. Alejandro Cardenete, and Carlos Romero. 2010. *Designing Public Policies: An Approach Based on Multi-criteria Analysis and Computable General Equilibrium Modeling*. Vol. 642. Springer Science & Business Media.
- Kenneth J. Arrow. 2012. *Social Choice and Individual Values*. Vol. 12. Yale University Press.
- Salvador Barber , Faruk Gul, and Ennio Stacchetti. 1993. Generalized median voter schemes and committees. *J. Econ. Theory* 61, 2 (1993), 262–289.
- Salvador Barber , Jordi Mass , and Alejandro Neme. 1997. Voting under constraints. *J. Econ. Theory* 76, 2 (1997), 298–321.

- Gerdus Benade, Swaprava Nath, Ariel D. Procaccia, and Nisarg Shah. 2017. Preference elicitation for participatory budgeting. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI'17)*. 376–382.
- Sanjai Bhagat and James A. Brickley. 1984. Cumulative voting: The value of minority shareholder voting rights. *J. Law Econ.* 27, 2 (1984), 339–365.
- Duncan Black. 1948. The decisions of a committee using a special majority. *Econometrica* 16, 3 (1948), 245–261.
- Kim C. Border and James S. Jordan. 1983. Straightforward elections, unanimity and phantom voters. *Rev. Econ. Stud.* 50, 1 (1983), 153–170.
- Steven J. Brams. 1990. Constrained approval voting: A voting system to elect a governing board. *Interfaces* 20, 5 (1990), 67–80.
- Steven J. Brams. 1993. Approval voting and the good society. *PEGS Newslett.* 3, 1 (1993), 10–14.
- Steven J. Brams and Peter C. Fishburn. 1978. *Approval Voting*. Vol. 72. Cambridge University Press. 831–847 pages.
- Steven J. Brams and Peter C. Fishburn. 2002. Voting procedures, Ch. 4 in the *Handbook of Social Choice and Welfare*, Vol. 1 (2002).
- Eric Budish. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *J. Pol. Econ.* 119, 6 (2011), 1061–1103.
- Yves Cabannes. 2004. Participatory budgeting: A significant contribution to participatory democracy. *Environ. Urban.* 16, 1 (2004), 27–46.
- Donald E. Campbell and Jerry S. Kelly. 2000. Weak independence and veto power. *Econ. Lett.* 66, 2 (2000), 183–189.
- Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. 2016. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. ACM, 305–322.
- Vince Conitzer, Markus Brill, and Rupert Freeman. 2015. Crowdsourcing societal tradeoffs. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 1213–1217.
- Vincent Conitzer and Tuomas Sandholm. 2012. Common voting rules as maximum likelihood estimators. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence (UAI'05)*. 145–152.
- Varsha Dani. 2017. Truthful and near-optimal mechanisms for welfare maximization in participatory budgeting. (unpublished).
- Nelson Dias. 2014. Hope for democracy: 25 years of participatory budgeting worldwide. www.in-loco.pt.
- John Duggan and Thomas Schwartz. 2000. Strategic manipulability without resoluteness or shared beliefs: Gibbard-Satterthwaite generalized. *Soc. Choice Welfare* 17, 1 (2000), 85–93.
- Brandon Fain, Ashish Goel, and Kamesh Munagala. 2016. The core of the participatory budgeting problem. In *Proceedings of the International Conference on Web and Internet Economics*. Springer, 384–399.
- Rupert Freeman, David M Pennock, Dominik Peters, and Jennifer Wortman Vaughan. 2019. Truthful aggregation of budget proposals. *arXiv preprint arXiv:1905.00457* (2019).
- Ernesto Ganuza and Gianpaolo Baiocchi. 2012. The power of ambiguity: How participatory budgeting travels the globe. *J. Publ. Delib.* 8, 2 (2012), 8.
- Nikhil Garg, Vijay Kamble, Ashish Goel, David Marn, and Kamesh Munagala. 2017. Collaborative optimization for collective decision-making in continuous spaces. In *Proceedings of the 26th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 617–626.
- Allan Gibbard. 1973. Manipulation of voting schemes: A general result. *Econometrica* 41, 4 (1973), 587–601.
- Nam Wook Kim, Jonghyuk Jung, Eun-Young Ko, Songyi Han, Chang Won Lee, Juho Kim, and Jihee Kim. 2016. Budgetmap: Engaging taxpayers in the issue-driven classification of a government budget. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. ACM, 1028–1039.
- Christian Klamler, Ulrich Pferschy, and Stefan Ruzika. 2012. Committee selection under weight constraints. *Math. Soc. Sci.* 64, 1 (2012), 48–56.
- Jonathan Levin and Barry Nalebuff. 1995. An introduction to vote-counting schemes. *J. Econ. Perspect.* 9, 1 (1995), 3–26.
- Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*. ACM, 125–131.
- Tyler Lu and Craig Boutilier. 2011. Budgeted social choice: From consensus to personalized decision making. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'11)*. 280–286.
- Hervé Moulin. 1980. On strategy-proofness and single peakedness. *Public Choice* 35, 4 (1980), 437–455.
- K. Nehring and Clemens Puppe. 2002. Strategy-proof social choice on single-peaked domains: Possibility, impossibility and the space between. *Unpublished, Department of Economics, University of California at Davis*.
- Richard G. Niemi. 1984. The problem of strategic behavior under approval voting. *Am. Pol. Sci. Rev.* 78, 4 (1984), 952–958.
- Carole Pateman. 2012. Participatory democracy revisited. *Perspect. Pol.* 10, 1 (2012), 7–19.

- Participatory Budgeting Project (PBP). 2016. Where has it worked? Retrieved from <http://www.participatorybudgeting.org/about-participatory-budgeting/where-has-it-worked/>.
- Ariel D. Procaccia, Sashank J. Reddi, and Nisarg Shah. 2012. A maximum likelihood approach for selecting sets of alternatives. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence (UAI'12)*. 695–704.
- Mark Allen Satterthwaite. 1975. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *J. Econ. Theory* 10, 2 (1975), 187–217.
- Aaron Schneider and Ben Goldfrank. 2002. Budgets and ballots in Brazil: Participatory budgeting from the city to the state. *Working Paper Series* 149 (2002). Institute of Development Studies, Brighton, UK.
- Graham Smith. 2009. *Democratic Innovations: Designing Institutions for Citizen Participation*. Cambridge University Press, Cambridge.
- H. Peyton Young. 1974. An axiomatization of Borda's rule. *J. Econ. Theory* 9, 1 (1974), 43–52.
- H. Peyton Young. 1988. Condorcet's theory of voting. *Am. Pol. Sci. Rev.* 82, 4 (1988), 1231–1244.

Received April 2017; revised November 2018; accepted May 2019