Truthful Aggregation of Budget Proposals

RUPERT FREEMAN, Microsoft Research
DAVID M. PENNOCK, Microsoft Research
DOMINIK PETERS, University of Oxford
JENNIFER WORTMAN VAUGHAN, Microsoft Research

We consider a participatory budgeting problem in which each voter submits a proposal for how to divide a single divisible resource (such as money or time) among several possible alternatives (such as public projects or activities) and these proposals must be aggregated into a single aggregate division. Under ℓ_1 preferences—for which a voter's disutility is given by the ℓ_1 distance between the aggregate division and the division he or she most prefers—the social welfare-maximizing mechanism, which minimizes the average ℓ_1 distance between the outcome and each voter's proposal, is incentive compatible [21]. However, it fails to satisfy the natural fairness notion of proportionality, placing too much weight on majority preferences. Leveraging a connection between market prices and the generalized median rules of Moulin [26], we introduce the *independent markets* mechanism, which is both incentive compatible and proportional. We unify the social welfare-maximizing mechanism and the independent markets mechanism by defining a broad class of *moving phantom* mechanisms that includes both. We show that every moving phantom mechanism is incentive compatible. Finally, we characterize the social welfare-maximizing mechanism as the unique Pareto-optimal mechanism in this class, suggesting an inherent tradeoff between Pareto optimality and proportionality.

1 INTRODUCTION

Participatory budgeting gives members of a community power by allowing them to collectively decide how to divide a portion of the community's budget among a set of proposed alternatives [10]. While participatory budgeting was first introduced in Brazil [30], it has now been used in more than 3,000 cities around the world, including New York, Boston, Chicago, San Francisco, and Toronto [27]. It has been used to determine how to allocate the budgets of states and cities as well as housing authorities and schools.

Many participatory budgeting elections are run using a variant of k-approval voting, in which each voter chooses up to k projects to approve, and the projects with the highest number of approvals are funded, subject to budget constraints [21]. Under such a voting scheme, each proposed project is either fully funded or not funded at all. This makes sense for well-delineated projects such as renovating a school or adding an elevator to a public library. For other kinds of projects, funding decisions need not be all-or-nothing. For example, participatory budgeting could be used to decide how to divide a city's tax surplus between its departments of health, education, infrastructure, and parks. A voter might propose a division of the tax surplus among the four departments into the fractions (30%, 40%, 20%, 10%). The city could invite each citizen to submit such a budget proposal, and they could then be aggregated by a suitable mechanism.

A first idea for aggregating the proposals would be to take the mean. But the mean has a serious flaw as an aggregator: it's easily manipulated. A voter preferring a (60%, 40%) division across two alternatives may vote (100%, 0%) instead in order to distort the mean calculation and move the aggregate closer to his or her true preference if the first alternative has little community support.

In this paper, we seek mechanisms that are resistant to such manipulation. In particular, we require that no voter can, by lying, move the aggregate division toward his or her preference on one alternative without moving it away from his or her preference by an equal or greater amount on other alternatives. In other words, we seek budget aggregation mechanisms that are incentive compatible under ℓ_1 preferences, with each voter's disutility for a budget division equal to the ℓ_1 distance between that division and the division she prefers most.

Goel et al. [21] showed that choosing an aggregate budget division that maximizes the welfare of the voters—that is, a division that minimizes the total ℓ_1 distance from each voter's report—is both incentive compatible and Pareto-optimal under this voter utility model. However, this utilitarian aggregate has a tendency to overweight majority preferences, creeping back towards all-or-nothing allocations. For example, imagine that a hundred voters prefer (100%, 0%) while ninety-nine prefer (0%, 100%). The utilitarian aggregate is (100%, 0%) even though the mean is close to (50%, 50%). In many participatory budgeting scenarios, the latter solution is more in the spirit of consensus. For example, imagine that each family votes for all education dollars to go to their own neighborhood school. The utilitarian aggregate would earmark the entire budget to the most populous school district, while we may prefer that funds are split in proportion to the districts' populations. To capture this fairness constraint, we define a notion of *proportionality*, requiring that when voters are single-minded (as in this example), the fraction of the budget assigned to each alternative is equal to the proportion of voters who favor that alternative. Do there exist aggregators that are both incentive compatible and proportional?

For the special case of two alternatives, ℓ_1 preferences are a special-case of *single-peaked* preferences, well studied in the voting literature. The seminal results of Moulin [26] imply that, in this setting, all incentive compatible voting schemes correspond to inserting $n + 1$ "phantom" proposals, where n is the number of voters, and returning the median of the n true proposals and the $n + 1$ phantoms. We show that there exists a way of placing the phantoms that results in a proportional mechanism for two alternatives.

Generalizing Moulin's phantom median mechanisms to allow for more than two alternatives is difficult. Existing proposals for such generalizations take a median in each dimension independently [5, 9, 28], but this strategy is doomed in our application with normalization constraints; unlike the mean, taking a coordinate-wise median will by default fail to normalize. We address this problem by allowing the set of phantoms to continuously shift upwards, increasing the sum of the generalized medians until the aggregate becomes normalized. This idea allows us to define a very general class of *moving phantom* mechanisms. Although one might think that allowing the final phantom locations to depend on voters' reports might give voters an incentive to misreport, we prove that every moving phantom mechanism is incentive compatible under ℓ_1 preferences.

Among this large family of incentive compatible mechanisms, we find one that satisfies our proportionality requirement. This moving phantom mechanism is obtained when phantoms are placed uniformly between 0 and a value $x \geq 0$ which increases until the coordinate-wise medians sum to one. To analyze this mechanism, we prove that the aggregate found by this mechanism can be interpreted as the clearing prices in a market system, and hence call it the *independent markets* mechanism. This reveals an unexpected connection between market prices and generalized medians that may be of broader interest. The independent markets mechanism can also be justified from a game-theoretic perspective as the unique Nash equilibrium of a natural voting game. Thus, this proportional moving phantom mechanism has two alternative interpretations:

(1) Market interpretation: For each alternative, we set up a market in which x units of a divisible good are sold. This amount x is the same across all markets. Each voter has a value for the good in market j that is equal to the fraction of the budget that the voter would prefer be allocated to alternative j in the budget division setting, and has \$1 to spend in each market. Increase x until the point at which the market clearing prices across these independent markets sum to \$1. At this point, the market clearing prices are exactly the aggregate division output by the independent markets moving phantom mechanism.
(2) Voting game: Each agent receives one credit for each alternative and may choose any amount of that credit to spend on the alternative. The outcome of the game is a normalized vector

proportional to the amount spent on each alternative. Agents choose their spending in such a way as to minimize the ℓ_1 distance between their preferred division in the budget division setting and the vector output by the game. The Nash equilibrium of this game is exactly the aggregate division output by the independent markets mechanism.

By analyzing the market and Nash equilibria of these systems, we can show that our mechanism satisfies several important social choice properties.

In contrast, the independent markets mechanism unfortunately fails to satisfy Pareto optimality. We show that this is unavoidable, as no proportional moving phantom mechanism is Pareto-optimal. In fact, we prove that there is a *unique* moving phantom mechanism that is. In this mechanism, all phantoms start at 0 and then, one by one, transition to 1, with no two phantoms moving at the same time. This mechanism turns out to also have a phantom-free interpretation: it is equivalent to selecting the maximum-entropy budget division out of all those that maximize social welfare—the same mechanism studied by Goel et al. [21] up to the choice of tie-breaking rules.

While the motivation of our formal model is participatory budgeting, it applies to the division of other resources such as time. For example, one might imagine using such a mechanism as a way to reach consensus among a team of conference organizers who wish to divide a day between talks, poster sessions, and social activities. Another example would be a government that needs to decide on a target energy mix (that is, how much energy should come from fossil fuels, nuclear, or renewable sources) and wishes to aggregate expert proposals. In all these applications, our class of moving phantom mechanisms can be used to make better decisions.

Related Work. Several recent papers study voting rules for participatory budgeting, considering both axiomatics and computational complexity, but under the assumption that indivisible projects can either be fully funded or not funded at all [3, 7, 21, 24]. The setting in which partial funding of alternatives is permitted has also been studied, but generally under a different utility model in which voters assign utility scores to the alternatives rather than having an ideal distribution [2, 8, 16]. This body of work includes positive results for weak versions of incentive compatibility, but impossibilities for obtaining full incentive compatibility. Garg et al. [19] perform a Mechanical Turk study exploring preference structure in a high-dimensional continuous setting similar to ours.

Closest to our work is that of Goel et al. [21]. The primary focus of their paper is on *knapsack voting*, in which each voter submits her preferred set of projects to fully fund. However, they also consider the use of *fractional* knapsack voting in a setting in which partial funding of alternatives is permitted and voters have ℓ_1 preferences. This corresponds exactly to our setting. They show that the mechanism that maximizes social welfare (with some fixed tie-breaking) is incentive compatible. We replicate this result by showing that the welfare-maximizing mechanism (with an arguably more natural way to break ties) is a member of the large class of moving phantom mechanisms, all of which are incentive compatible. Goel et al. do not consider other mechanisms for this setting.

The truthful aggregation of preferences over numerical values (such as the temperature for an office) has been extensively studied. A famous result of Moulin [26] characterizes the set of incentive compatible voting rules under the assumption that voters have single-peaked preferences over values in $[0, 1]$. These voting rules are generalized median schemes. The best-known example is the standard median, in which each voter reports her ideal point in $[0, 1]$ and the median report is selected. Other voting rules in this class insert "phantom voters" who report a fixed top choice. Barberà et al. [4] obtained a multi-dimensional analogue of this result for $[0, 1]^m$, and there are further generalizations that characterize truthful rules if other constraints are imposed on the feasible set [6]. Crucially, the constraints allowed by Barberà et al. [6] do not include the normalization constraint that is fundamental to our setting. Several other papers [5, 9, 28] introduced multidimensional models in which one can achieve truthfulness by taking a generalized median in each coordinate, but such a strategy does not work with normalization constraints. We are not aware of results (prior to this work) that extend generalized medians to multiple dimensions without using a mechanism that decomposes into one-dimensional mechanisms.

In the computer science literature, the above-mentioned generalized median schemes have also been studied in the context of truthful facility location [1, 29]. In this context, the aim is to approximate social welfare subject to incentive compatibility.

One could apply our results to the aggregation of probabilistic beliefs. There is a large literature on *probabilistic opinion pooling* [12, 17, 20, 23] which studies aggregators in this context. The main focus of that literature is to preserve stochastic and epistemic properties. To the best of our knowledge, strategic aspects have not been considered.

Finally, recently proposed rules for crowdsourcing societal tradeoffs [13, 14] can also be used to aggregate budget divisions (with full support) after converting them into pairwise ratios of funding amounts, but this setting has also not been analyzed from a strategic viewpoint.

2 PRELIMINARIES

Let $N = \{1, \ldots, n\}$ be a set of voters and $M = \{1, \ldots, m\}$ be a set of possible alternatives. Voters have structured preferences over budget *divisions* $\mathbf{p} \in [0, 1]^m$, with $\sum_{j\in[m]} p_j = 1$, where p_j is the fraction of a public resource (such as money or time) allocated to alternative j. Each voter i has a most preferred division $\mathbf{p_i} = (p_{i,1}, \ldots, p_{i,m})$, with their preference over other divisions induced by ℓ_1 distance from $\mathbf{p_i}$. Specifically, each voter i has a disutility for division $\mathbf{q}$ equal to $d(\mathbf{p_i}, \mathbf{q})$, where $d(\mathbf{x}, \mathbf{y}) = \sum_{j=1}^{m} |x_i - y_i|$ denotes the ℓ_1 distance between $\mathbf{x}$ and $\mathbf{y}$. Note that a voter's complete preference over all possible divisions can be deduced from their most preferred division $\mathbf{p_i}$.

A *preference profile* $\mathbf{P} = (\hat{\mathbf{p}}_1, \hat{\mathbf{p}}_2, \ldots, \hat{\mathbf{p}}_\mathbf{n})$ consists of a reported division $\hat{\mathbf{p}}_\mathbf{i}$ for each voter i. We use $\mathbf{P_{-i}}$ to denote the reports of all voters other than i. A *budget aggregation mechanism* $\mathcal{A}$ takes as input a preference profile $\mathbf{P}$, and outputs an aggregate division $\mathcal{A}(\mathbf{P})$. A mechanism is *continuous* if it is continuous when considered as a function $\mathcal{A} : (\mathbb{R}^m)^n \to \mathbb{R}^m$. We say that a mechanism is *anonymous* if its output is fixed under permutations of the voters, and *neutral* if a permutation of the alternatives in voters' inputs permutes the output in the same way.

We are interested in mechanisms that satisfy *incentive compatibility*. Voters should not be able to change the aggregate division in their favor by misrepresenting their preference.

Definition 2.1. A budget aggregation mechanism $\mathcal{A}$ satisfies incentive compatibility if, for all preference profiles $\mathbf{P}$, voters i, and divisions $\mathbf{p_i}$ and $\hat{\mathbf{p}}_\mathbf{i}$, $d(\mathcal{A}(\mathbf{P}_{-i}, \hat{\mathbf{p}}_\mathbf{i}), \mathbf{p_i}) \geq d(\mathcal{A}(\mathbf{P_{-i}}, \mathbf{p_i}), \mathbf{p_i})$.

We are also interested in the basic efficiency notion of *Pareto optimality*. It should not be possible to change the aggregate so that some voter is strictly better off but no other voter is worse off.

Definition 2.2. A budget aggregation mechanism $\mathcal{A}$ satisfies Pareto optimality if, for all preference profiles $\mathbf{P}$, and all divisions $\mathbf{q}$, if $d(\mathcal{A}(\mathbf{P}), \hat{\mathbf{p}}_\mathbf{i}) > d(\mathbf{q}, \hat{\mathbf{p}}_\mathbf{i})$ for some voter i, then there exists a voter j for which $d(\mathcal{A}(\mathbf{P}), \hat{\mathbf{p}}_\mathbf{j}) < d(\mathbf{q}, \hat{\mathbf{p}}_\mathbf{j})$.

Observe that the definitions of incentive compatibility and Pareto optimality depend only on the voters' preference relations, not the exact utility model. Results pertaining to these properties therefore hold for any utility function that induces the same ordinal preferences as linear utilities.

We also consider a fairness property that we call proportionality: Suppose each voter is single-minded, in that they prefer a division in which the entire resource goes to a single alternative. Then it is natural to split the resource in proportion to the number of voters supporting each alternative. For example, if 6 voters report $(1, 0, 0)$, 3 voters report $(0, 1, 0)$, and 1 voter reports $(0, 0, 1)$, then the aggregate should be $(0.6, 0.3, 0.1)$. We call this property *proportionality*.

Definition 2.3. A voter is *single-minded* if their preferred division is a unit vector. A budget aggregation mechanism $\mathcal{A}$ is *proportional* if, for every preference profile $\mathbf{P}$ consisting of only single-minded voters, and every alternative j, $\mathcal{A}(\mathbf{P})_j = n_j/n$, where n_j is the number of voters that support alternative j.

We note that proportionality is a fairly weak definition, only applying to a small subset of possible profiles. However, as we will see later, it is already strong enough to be incompatible with Pareto optimality within the class of moving phantom mechanisms that we introduce in this paper.

3 TWO ALTERNATIVES

To build intuition, we begin by considering the case in which $m = 2$. Due to the normalization of inputs and of the output, and with ℓ_1 preferences, the problem is perfectly one-dimensional in this case. This allows us to directly import Moulin's [26] famous characterization of *generalized median* rules as the only incentive compatible mechanisms for voters with single-peaked preferences over a single-dimensional quantity.[1]

Theorem 3.1 ([26]). *For $m = 2$, an anonymous and continuous budget aggregation mechanism $\mathcal{A}$ is incentive compatible if and only if there are $\alpha_0 \geq \alpha_1 \geq \ldots \geq \alpha_n$ in $[0, 1]$ such that, for all profiles $\mathbf{P}$,*

$$\mathcal{A}(\mathbf{P})_1 = \text{med}(p_{1,1}, p_{2,1}, \ldots, p_{n,1}, \alpha_0, \alpha_1, \ldots, \alpha_n),$$
$$\mathcal{A}(\mathbf{P})_2 = \text{med}(p_{1,2}, p_{2,2}, \ldots, p_{n,2}, 1 - \alpha_0, 1 - \alpha_1, \ldots, 1 - \alpha_n).$$

The numbers α_k are known as *phantoms*. Each mechanism described by Theorem 3.1 can be understood as taking the coordinate-wise median of the reported distributions, after inserting $n + 1$ phantom voters (whose report is fixed and independent of the input profile).

One can check that $\alpha_0, \ldots, \alpha_n$ define a neutral mechanism if and only if the phantom placements are symmetric, that is if and only if $\{\alpha_0, \ldots, \alpha_n\} = \{1 - \alpha_0, \ldots, 1 - \alpha_n\}$. Note that there are $n + 1$ phantoms but only n voters, so that the phantoms can outweigh the voters. For example, when $\alpha_k = 1/2$ for all $k \in \{0, \ldots, n\}$ then the mechanism is just the constant mechanism returning $(1/2, 1/2)$. However, if we take $\alpha_0 = 1$ and $\alpha_n = 0$, then these two phantoms "cancel out" and there are only $n - 1$ phantoms left. In fact, one can check that the mechanism is Pareto-optimal if and only if $\alpha_0 = 1$ and $\alpha_n = 0$ [26].

A particularly interesting example is the *uniform phantom mechanism*, obtained when placing the phantoms uniformly over the interval $[0, 1]$, so that $\alpha_k = 1 - k/n$ for each $k \in \{0, \ldots, n\}$. This placement of phantom voters appears in a paper by Caragiannis et al. [11]. They were aiming for mechanisms whose output is close to the mean, and they prove that the uniform phantom mechanism yields an aggregate that is closer to the mean than that obtained from any other phantom placements, in the worst case over inputs. The uniform phantom mechanism has other attractive properties, including being proportional in the sense of Definition 2.3.

Proposition 3.2. *For $m = 2$, the uniform phantom mechanism is the unique (anonymous and continuous) budget aggregation mechanism $\mathcal{A}$ that is both incentive compatible and proportional.*

Proof. Theorem 3.1 gives us that $\mathcal{A}$ is incentive compatible if and only if it can be written in terms of phantom medians. We therefore need only to consider the additional requirement of proportionality. The uniform phantom mechanism is proportional, because if $\mathbf{P}$ consists of $n - k$ voters reporting $(1, 0)$ and k voters reporting $(0, 1)$, then $\mathbf{A}(\mathbf{P})_1 = \alpha_k = (n - k)/n$, as required.

[1]Our preference model using ℓ_1 imposes slightly more structure than just single-peakedness, namely that voters are indifferent between points that are equidistant to their peak. However, this restriction does not enlarge the class of incentive compatible mechanisms, at least if we impose continuity [25].

For uniqueness, suppose $\alpha_0, \dots, \alpha_n$ are phantom positions that induce a proportional mechanism. Let $k \in \{0, \dots, n\}$. We show that $\alpha_k = 1 - k/n$. Let $\mathbf{P}$ be a profile consisting of only single-minded voters with $n_1 = n - k$ voters reporting $\hat{\mathbf{p}}_\mathbf{i} = (1, 0)$. Then α_k is the median, and proportionality requires that $\alpha_k = n_1/n = (n-k)/n = 1 - k/n$. □

Another natural way to place the phantoms is one that takes the coordinate-wise median. When $n+1$ is even, this is achieved by placing half the phantoms at 0 and the other half at 1, outputting precisely the median of the reported values on each coordinate. When $n+1$ is odd, we place $n/2$ phantoms at 0, $n/2$ phantoms at 1, and we place a single phantom at $1/2$ to preserve neutrality. This mechanism outputs the point between the left and right medians that is closest to $1/2$. The resulting mechanism returns an aggregate $\mathbf{p}$ that minimizes the sum of distances between the reports $(\mathbf{p_1}, \mathbf{p_2}, \dots, \mathbf{p_n})$ and $\mathbf{p}$. We will generalize this mechanism for larger m in Section 6.

4 A CLASS OF INCENTIVE COMPATIBLE MECHANISMS FOR HIGHER DIMENSIONS

For $m = 2$, we have a complete picture of incentive-compatible mechanisms, thanks to Moulin's characterization. For $m \geq 3$, it is less clear how to construct examples of incentive-compatible mechanisms. One could try to take a generalized median in each alternative independently, but the result of such a mechanism would not respect the normalization constraint.

However, there is a way of extending the idea of generalized medians to the higher-dimensional setting. The basic idea is that if a coordinate-wise generalized median violates the normalization constraint, then we can adjust the placement of the phantoms, increasing or decreasing the sum of the generalized medians as needed. Such a procedure might, in principle, give voters incentives to manipulate in order to affect the phantom placements. However, our class of *moving phantom mechanisms* manages to avoid this problem.

Definition 4.1. Let $\mathcal{F} = \{f_k : k \in \{0, \dots, n\}\}$ be a family of functions, or *phantom system*, where $f_k : [0,1] \to [0,1]$ is a continuous, weakly increasing function with $f_k(0) = 0$ and $f_k(1) = 1$ for each k, and we have $f_0(t) \geq f_1(t) \geq \cdots \geq f_n(t)$ for all $t \in [0,1]$. Then, the *moving phantom mechanism* $\mathcal{A}^\mathcal{F}$ is defined so that for all profiles $\mathbf{P}$ and all $j \in [m]$,

$$\mathcal{A}^\mathcal{F}(\mathbf{P})_j = \text{med}(f_0(t^*), \dots, f_n(t^*), \hat{p}_{1,j}, \dots, \hat{p}_{n,j}), \tag{1}$$

where t^* is chosen so that $t^* \in \{t : \sum_{j \in [m]} \text{med}(f_0(t), \dots, f_n(t), \hat{p}_{1,j}, \dots, \hat{p}_{n,j}) = 1\}$.

For brevity, we write $\mathcal{F}(t) = (f_0(t), \dots, f_n(t))$ and abbreviate the median in (1) to $\text{med}(\mathcal{F}(t), \mathbf{P}_{i \in [n], j})$.

Let us examine the definition. Each f_k represents a phantom, and the phantom system $\mathcal{F}$ represents a "movie" in which all phantoms continuously increase from 0 to 1, with the function argument t defining an instantaneous snapshot of the phantom positions. The moving phantom mechanism $\mathcal{A}^\mathcal{F}$ defined by $\mathcal{F}$ identifies a particular snapshot in time, t^*, for which the sum of generalized medians over all coordinates is exactly 1. One can check that at least one such t^* exists, and that the output of the mechanism is independent of which of these t^* is chosen.

Proposition 4.2. *The moving phantom mechanism $\mathcal{A}^\mathcal{F}$ is well-defined for every phantom system $\mathcal{F}$ satisfying the conditions of Definition 4.1.*

Proof. First note that the function $t \mapsto \sum_{j \in [m]} \text{med}(\mathcal{F}(t), \mathbf{P}_{i \in [n], j})$ is continuous and increasing in t, because f_k is continuous and increasing, and these properties are preserved under taking the median and sum. This implies that, provided the set $\{t : \sum_{j \in [m]} \text{med}(\mathcal{F}(t), \mathbf{P}_{i \in [n], j}) = 1\}$ is non-empty, the aggregate $\mathcal{A}^\mathcal{F}(\mathbf{P})$ does not depend on the choice of t^*.

When $t = 0$, $\sum_{j \in [m]} \text{med}(\mathcal{F}(t), \mathbf{P}_{i \in [n], j}) = 0$, since all $n+1$ phantom entries are 0. When $t = 1$, $\sum_{j \in [m]} \text{med}(\mathcal{F}(t), \mathbf{P}_{i \in [n], j}) = m > 1$, since all $n+1$ phantom entries are 1. By the Intermediate Value Theorem, using continuity, there exists $t \in [0,1]$ with $\sum_{j \in [m]} \text{med}(\mathcal{F}(t), \mathbf{P}_{i \in [n], j}) = 1$. □

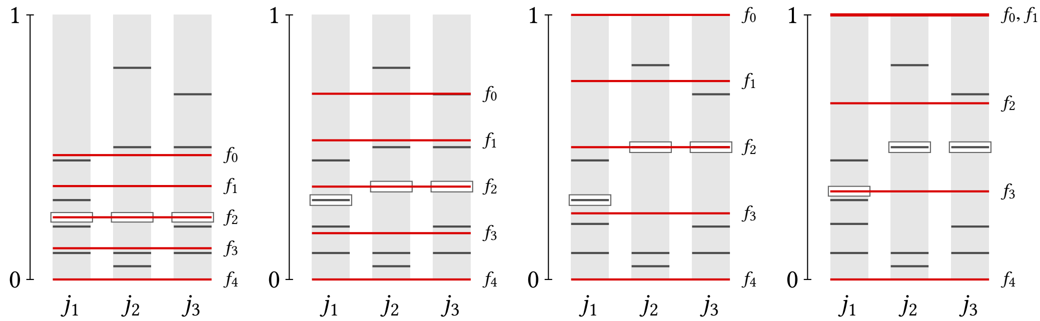

Fig. 1. A moving phantom mechanism operating on an instance with $n = 4$ and $m = 3$.

To build intuition, we consider an example moving phantom mechanism in Figure 1. There are three alternatives, each occupying a column on the horizontal axis, and four voters. Voter reports are indicated by gray horizontal line segments, with their magnitude $\hat{p}_{i,j}$ indicated by their vertical position. The phantom placements are indicated by the red lines and labeled $f_0, \ldots, f_4$. For each alternative, the median of the four agent reports and the five phantoms is indicated by a rectangle.

The four snapshots shown in Figure 1 display increasing values of t. Observe that the position of each phantom (weakly) increases from left to right, as does the median on each alternative. Although the vertical axis is not labeled, for simplicity of presentation, normalization here occurs in the second image from the left. In the leftmost image, the sum of the highlighted entries is less than 1, while in the two rightmost images it is more than 1.

For simplicity, the definition of moving phantom mechanisms treats the number of voters n as fixed. To allow n to vary, it is necessary to define a family of phantom systems, one for each n. In the next two sections, we give two examples of such families, but for this section we keep the presentation simple by considering only a fixed n.

Moving phantom mechanisms satisfy some important basic properties. They are all anonymous and neutral. Here neutrality is a design choice—one could imagine defining moving phantom mechanisms for which the movement of the phantoms depends on the alternative. All moving phantom mechanisms are also, again by design, continuous.

Given a profile, we can efficiently approximate the output of a moving phantom mechanism, assuming oracle access to its defining functions $\mathcal{F}$, by performing a binary search on t. In principle, the precise time $t^* \in [0, 1]$ at which the output of the mechanism is normalized may have many decimal digits, and for badly-behaved $\mathcal{F}$ it may even be irrational. For the same reason, the mechanism may return an irrational division, so the precise computation of the output may not be possible. However, for the mechanisms studied in the following sections, we can show that t^* has few digits and the output is always rational, so polynomial-time computation is possible.

We now show our main result in this section, that every moving phantom mechanism is incentive compatible. Before proving the result formally, we provide some intuition. If i changes her report from $\mathbf{p_i}$ to $\hat{\mathbf{p}}_\mathbf{i}$, the effect on the aggregate can be decomposed into two parts. First, we can think of holding the phantoms fixed at the snapshot dictated by the truthful instance, while changing i's report to $\hat{\mathbf{p}}_\mathbf{i}$. Second, we can think of repositioning the phantoms to the snapshot required to guarantee normalization of the aggregate vector after i reports $\hat{\mathbf{p}}_\mathbf{i}$. To prove incentive compatibility, we show that any change that the aggregate division undergoes in the first stage can only be bad for voter i, pushing the aggregate away from $\mathbf{p_i}$. Change in the second stage can push the aggregate towards $\mathbf{p_i}$, helping voter i, but the magnitude of this change is upper bounded by the magnitude of the harmful change in the first stage.

Theorem 4.3. *Every moving phantom mechanism is incentive compatible.*

Proof. Let $\mathcal{F}$ define a moving phantom mechanism $\mathcal{A}^{\mathcal{F}}$. Consider some report $\hat{\mathbf{p}}_\mathbf{i} \neq \mathbf{p_i}$, and fix the reports of all other voters $\mathbf{P_{-i}}$. Let t^* determine the phantom placement for reports $(\mathbf{p_i}, \mathbf{P_{-i}})$ and $\hat{t}^*$ for reports $(\hat{\mathbf{p}}_\mathbf{i}, \mathbf{P_{-i}})$.

Consider the effect of i's misreport from $\mathbf{p_i}$ to $\hat{\mathbf{p}}_\mathbf{i}$ while *holding the phantom placement fixed at* $\mathcal{F}(t^*)$. Then, because phantom placements are fixed on each alternative, any change that voter i can cause on alternative j by misreporting must be away from her preference $p_{i,j}$. For each $j \in [m]$,

- if $p_{i,j} \leq \mathcal{A}^f(\mathbf{p_i}, \mathbf{P_{-i}})_j < \hat{p}_{i,j}$, then we must have $\text{med}(\mathcal{F}(t^*), \hat{p}_{i,j}, \mathbf{P}_{-i,j}) \geq \text{med}(\mathcal{F}(t^*), p_{i,j}, \mathbf{P}_{-i,j})$;
- if $\hat{p}_{i,j} \leq \mathcal{A}^f(\mathbf{p_i}, \mathbf{P_{-i}})_j < p_{i,j}$, then we must have $\text{med}(\mathcal{F}(t^*), \hat{p}_{i,j}, \mathbf{P}_{-i,j}) \leq \text{med}(\mathcal{F}(t^*), p_{i,j}, \mathbf{P}_{-i,j})$;
- if $\hat{p}_{i,j}$ and $p_{i,j}$ lie on the same side of $\mathcal{A}^f(\mathbf{p_i}, \mathbf{P_{-i}})_j$, then $\text{med}(\mathcal{F}(t^*), \hat{p}_{i,j}, \mathbf{P}_{-i,j}) = \text{med}(\mathcal{F}(t^*), p_{i,j}, \mathbf{P}_{-i,j})$.

Let $y_j = \text{med}(\mathcal{F}(t^*), \hat{p}_{i,j}, \mathbf{P}_{-i,j}) - \text{med}(\mathcal{F}(t^*), p_{i,j}, \mathbf{P}_{-i,j})$ denote the change caused on alternative j by voter i's misreport, subject to holding the phantom placement fixed at $\mathcal{F}(t^*)$. By the above, the ℓ_1 distance from i's preferred division has increased by $\sum_{j\in[m]} |y_j|$ as a result of i's misreport.

Next, we consider the change that results from moving the phantoms from $\mathcal{F}(t^*)$ to $\mathcal{F}(\hat{t}^*)$. Assume that $\sum_{j\in[m]} y_j \geq 0$ (otherwise, a very similar argument applies). Then we have that $\sum_{j\in[m]} \text{med}(\mathcal{F}(t^*), \hat{p}_{i,j}, \mathbf{P}_{-i,j}) \geq 1$, which implies that $\hat{t}^* \leq t^*$ since the sum is monotonic in t (see the proof of Proposition 4.2). This produces aggregate division $\mathcal{A}^{\mathcal{F}}(\hat{\mathbf{p}}_\mathbf{i}, \mathbf{P_{-i}})$ with $\mathcal{A}^{\mathcal{F}}(\hat{\mathbf{p}}_\mathbf{i}, \mathbf{P_{-i}})_j = \text{med}(\mathcal{F}(\hat{t}^*), \hat{p}_{i,j}, \mathbf{P}_{-i,j}) \leq \text{med}(\mathcal{F}(t^*), \hat{p}_{i,j}, \mathbf{P}_{-i,j})$ for all j, and $\sum_{j\in[m]}(\text{med}(\mathcal{F}(t^*), \hat{p}_{i,j}, \mathbf{P}_{-i,j}) - \mathcal{A}^{\mathcal{F}}(\mathbf{P})_j) = \sum_{j\in[m]} y_j$. That is, the ℓ_1 distance between taking generalized medians with phantoms defined by t^* and doing so with phantoms defined by $\hat{t}^*$, conditioned on voter i reporting $\hat{\mathbf{p}}_\mathbf{i}$, is at most $\sum_{j\in[m]} y_j$.

Therefore, $d(\mathcal{A}^{\mathcal{F}}(\hat{\mathbf{p}}_\mathbf{i}, \mathbf{P_{-i}}), \mathbf{p_i}) \geq d(\mathcal{A}^{\mathcal{F}}(\mathbf{p_i}, \mathbf{P_{-i}}), \mathbf{p_i}) + \sum_j |y_j| - \sum_j y_j \geq d(\mathcal{A}^{\mathcal{F}}(\mathbf{p_i}, \mathbf{P_{-i}})$. □

In addition to incentive compatibility, moving phantom mechanisms satisfy a natural monotonicity property that says that if some voter increases her report on alternative j, and decreases her report on all other alternatives, then the aggregate weight on alternative j should not decrease.

Definition 4.4. A budget aggregation mechanism $\mathcal{A}$ satisfies *monotonicity* if, for all $\mathbf{p_i}$, $\mathbf{p'_i}$ with $p_{i,j} > p'_{i,j}$ for some j and $p_{i,k} \leq p'_{i,k}$ for all $k \neq j$,

$$\mathcal{A}(\mathbf{p_i}, \mathbf{P_{-i}})_j \geq \mathcal{A}(\mathbf{p'_i}, \mathbf{P_{-i}})_j.$$

Theorem 4.5. *Every moving phantom mechanism satisfies monotonicity.*

Proof. Let $\mathbf{p_i}$, $\mathbf{p'_i}$ be such that $p_{i,j} > p'_{i,j}$ for some j and $p_{i,k} \leq p'_{i,k}$ for all $k \neq j$. Let t^* determine the phantom placement for reports $(\mathbf{p_i}, \mathbf{P_{-i}})$ and t'^* for reports $(\mathbf{p'_i}, \mathbf{P_{-i}})$.

Suppose that $t'^* < t^*$. We have

$$\mathcal{A}(\mathbf{p'_i}, \mathbf{P_{-i}})_j = \text{med}(\mathcal{F}(t'^*), p'_{i,j}, P_{-i,j}) \leq \text{med}(\mathcal{F}(t^*), p_{i,j}, P_{-i,j}) = \mathcal{A}(\mathbf{p_i}, \mathbf{P_{-i}})_j$$

where the inequality holds because $p_{i,j} > p'_{i,j}$ and $f_k(t^*) \geq f_k(t'^*)$ for all $k \in \{0, \dots, n\}$.

Next, suppose that $t'^* > t^*$. Then

$$\begin{aligned}\mathcal{A}(\mathbf{p'_i}, \mathbf{P_{-i}})_j = 1 - \sum_{k\neq j} \mathcal{A}(\mathbf{p'_i}, \mathbf{P_{-i}})_k &= 1 - \sum_{j'\neq j} \text{med}(\mathcal{F}(t'^*), p'_{i,j'}, P_{-i,j'}) \\ &\leq 1 - \sum_{j'\neq j} \text{med}(\mathcal{F}(t^*), p_{i,j'}, P_{-i,j'}) = \mathcal{A}(\mathbf{p_i}, \mathbf{P_{-i}})_j\end{aligned}$$

where the inequality holds because $p_{i,j'} < p'_{i,j'}$ for all $j' \neq j$ and $f_k(t^*) \leq f_k(t'^*)$ for all $k \in \{0, \dots, n\}$. □

Before we move on to particular moving phantom mechanisms, let us end this section with a tantalizing open question: Does there exist an (anonymous, neutral, continuous) incentive compatible budget aggregation mechanism that is *not* a moving phantom mechanism? We have not been able to construct any example, and have found that some mechanisms that on first sight seem to have nothing to do with medians end up having an equivalent description as a moving phantom mechanism. For the simpler two-alternative case, we already have a characterization of all incentive compatible mechanisms (Theorem 3.1). This class can equivalently be described in terms of moving phantoms, and so the answer to our question for $m = 2$ is *no*.

Theorem 4.6. *For $m = 2$, moving phantom mechanisms are the only budget aggregation mechanisms that satisfy anonymity, neutrality, continuity, and incentive compatibility.*

Proof. Certainly all moving phantom mechanisms satisfy these properties. For the other direction, we know from Theorem 3.1 that any mechanism $\mathcal{A}$ satisfying these properties can be described as a generalized median with phantoms $\alpha_0, \dots, \alpha_n$ satisfying, due to neutrality, $\{\alpha_0, \dots, \alpha_n\} = \{1-\alpha_0, \dots, 1-\alpha_n\}$. We show that $\mathcal{A}$ is equivalent to a moving phantom mechanism. Define $\mathcal{A}^{\mathcal{F}}$ using a phantom system $\mathcal{F}$ for which there exists a $t^* \in [0, 1]$ with $f_k(t^*) = \alpha_k$ for every $k \in \{0, \dots, n\}$. Then, for every preference profile $\mathbf{P}$, we have that $\mathcal{A}^{\mathcal{F}}(\mathbf{P})_1 = \text{med}(\mathcal{F}(t^*), \mathbf{P}_{i\in[n],j}) = (\alpha_0, \dots, \alpha_n, \mathbf{P}_{i\in[n],j})$, and $\mathcal{A}^{\mathcal{F}}(\mathbf{P})_2 = 1 - \mathcal{A}^{\mathcal{F}}(\mathbf{P})_1$, matching the output of $\mathcal{A}$. □

5 THE INDEPENDENT MARKETS MECHANISM

We have seen that uniform phantoms is uniquely proportional for $m = 2$. By a similar argument to the proof of Proposition 3.2, it is easy to see that any family of functions $\mathcal{F}$ that generates uniform phantoms at some snapshot will be proportional, and will reduce to the uniform phantom mechanism for $m = 2$. However, this leaves a large class of moving phantom mechanisms to choose from. In this section, we identify a particular moving phantom mechanism that generalizes the uniform phantom mechanism for arbitrary m. Its output can be interpreted as a market equilibrium.

Definition 5.1. The independent markets mechanism ($\mathcal{A}^{IM}$) is the moving phantom mechanism defined by the phantom system $f_k(t) = \min\{t(n-k), 1\}$ for each $k \in \{0, \dots, n\}$.

To visualize the phantom placement, observe that for any $t \leq 1/n$, phantoms are being placed at $0, t, 2t, \dots, nt$. Once t reaches $1/n$, phantoms continue to grow in the same manner, but the higher phantoms get capped at 1.[2] This is actually the mechanism that we displayed in Figure 1. Note that, when $t = 1/n$, the phantom placement is uniform on $[0, 1]$ (as is the case in the third panel of Figure 1); thus, $\mathcal{A}^{IM}$ reduces to the uniform phantom mechanism for $m = 2$.

Example 5.2. Let us consider a simple numerical example. Let $n = m = 3$, and suppose voter reports are $\mathbf{p_1} = (0, 0.5, 0.5)$, $\mathbf{p_2} = (0.5, 0.5, 0)$, and $\mathbf{p_3} = (0.9, 0, 0.1)$. Consider the placement of the $n + 1 = 4$ phantoms when $t = 0.6$. They are placed at $f_0(t) = 0.6$, $f_1(t) = 0.4$, $f_2(t) = 0.2$, $f_3(t) = 0$. On the first alternative,

$$\text{med}\{f_0(t), f_1(t), f_2(t), f_3(t), p_{1,1}, p_{2,1}, p_{3,1}\} = \text{med}\{0.6, 0.4, 0.2, 0, 0, 0.5, 0.9\} = 0.4.$$

Similarly, it is easy to check that the generalized median on the second alternative is 0.4 and on the third alternative is 0.2. Because these are normalized, $t^* = 0.6$ is a valid choice of t^*, and the outcome $\mathcal{A}^{IM}(\mathbf{P}) = (0.4, 0.4, 0.2)$.

[2] As written, $f_n(1) = 0$, but Definition 4.1 requires $f_k(1) = 1$ for all k. This detail does not matter here, since normalization is always achieved without moving phantom n, but one could write f_n in a different form to satisfy Definition 4.1 without it changing the behavior of the mechanism.

5.1 Market Interpretation

Why do we call this mechanism the independent markets mechanism? To explain this, we first establish a connection between the market clearing price in a simple single-good market and the median of some familiar-looking numbers.

Suppose we are selling a single divisible good, of which a total amount of $x \in [0, \infty)$ is available. Each of n voters has a budget of 1, and a value $v_i \in [0, \infty)$ per unit of the good. At a price $\pi \geq 0$ per unit of the good, the *demand* of voter i, $D_i(\pi)$ is given by the following function:

$$D_i(\pi) = \begin{cases} \infty & \pi = 0, \\ \frac{1}{\pi} & 0 < \pi < v_i, \\ 0 & \pi \geq v_i \text{ and } \pi > 0. \end{cases}$$

Thus, each voter demands as much of the good as their budget of 1 allows at price π, as long as the price per unit is lower than their value per unit. The *market clearing price* c is the price at which the supply of the good (x) equals the total demand. Formally,

$$c = \sup\{\pi : \textstyle\sum_{i \in [n]} D_i(\pi) > x\}, \tag{2}$$

where the supremum is necessary because, due to discontinuities in the demand function, supply and demand may never be exactly equal.

It turns out that the market clearing price c is equal to the median of the n voter values v_i and the $n + 1$ "phantom values" which are uniformly distributed on the interval $[0, n/x]$. To the best of our knowledge, this connection has not previously been appreciated in the literature.

Lemma 5.3. *In the market defined above, the market clearing price c equals*

$$\text{med}(0, 1/x, \ldots, (n-1)/x, n/x, v_1, \ldots, v_n).$$

Proof. We distinguish the cases that the median is a phantom entry or a voter entry. Suppose that the median is a/x for some a. Then we can partition the (real and phantom) entries, with the exception of the phantom at a/x, into sets A and B with $|A| = |B| = n$, where A consists only of entries less than or equal to a/x, and B consists only of entries greater than or equal to a/x.

The set B contains $n - a$ phantoms, so $n - (n - a) = a$ voter reports. At any price $\pi < a/x$, each voter $i \in B$ has demand $D_i(\pi) = 1/\pi > x/a$. The total demand of all voters in B is therefore greater than x. At price $\pi = a/x$, each voter $i \in B$ has demand $D_i(\pi) = 1/\pi = x/a$ (if $v_i > a/x$) or $D_i(\pi) = 0$ (if $v_i = a/x$), and each voter $i \notin B$ has demand 0. Therefore the total demand of all voters is at most x, so the market clearing price is a/x.

Next, suppose that the generalized median is $a/x < y < (a+1)/x$ for some $a \leq n - 1$ (note that the generalized median cannot be greater than n/x, because it cannot be higher than the largest phantom value). Then we can partition the (real and phantom) entries, not including a single voter with $v_i = y$ (one such voter must exist because the median coincides with some entry, and no phantom entry lies at y), into sets A and B each of size n, where A consists only of entries less than or equal to y, and B consists only of entries greater than or equal to y.

Again, B contains $n - a$ phantom reports, so a voter reports. At all prices $\pi < y$, each of these a voters, as well as voter i with $\pi_i = y$, has demand $1/\pi > 1/y$. The total demand is thus greater than $(a+1)/y > x$. At price $\pi = y$, the total demand is at most $a/y < x$ (since the number of voters with $v_i > y$ is at most the number voter reports in set B). The market clearing price is therefore y. □

The "market" connection to independent markets is now clear: For each alternative j, we set up a market in which we sell an amount x of a good; this amount is the same across markets. Voter $i \in [n]$ has value $\hat{p}_{i,j}$ for the good sold in market j, and has a budget of 1 in each market. The markets are "independent" because, while each voter is engaged in every market, the budget of

1 for each market can only be used to buy the good sold in that market. Using Lemma 5.3, we can derive the market clearing prices in each of these markets. If we write $t = n/x$, then these prices correspond exactly to the output of $\mathcal{A}^{IM}$ with the phantoms as placed at time t. Changing the phantom placement by varying t to normalize the output is equivalent to varying the amount x of the good sold in each market until the clearing prices across markets sum to 1. While we prevent phantoms from moving above 1 in the definition of independent markets—complying with Definition 4.1—the exact positions of these phantoms do not affect the clearing price since all reports are at most 1.

Returning to Example 5.2, we can verify the outcome using the market interpretation, by setting the quantity of goods to be sold in each market to $x^* = n/t^* = 5$. In the market corresponding to alternative 1, the market clears at price $\pi_1 = 0.4$, at which price voters 2 and 3 demand $1/\pi_1 = 2.5$ goods each, matching supply, and voter 1 demands nothing as $p_{1,1} = 0 < \pi$. It can be checked that the market prices also match the independent markets outcome for alternatives 2 and 3.

The market system we have described yields an incentive-compatible aggregator, since it corresponds to a moving phantom mechanism. There are other market-based aggregation mechanisms described in the literature, most famously the parimutuel consensus mechanism of Eisenberg and Gale [15]. That mechanism differs from ours in that voters have only a single budget of 1 which they can use in all of the markets. (The supply of goods can be fixed at $x = n$, which guarantees that prices are normalized, because total spending is fixed.) For the case $m = 2$, it does not matter whether markets are independent or not, and our mechanism is equivalent to the one of Eisenberg and Gale [15]. It follows that the parimutuel consensus mechanism is incentive compatible for $m = 2$ (in our ℓ_1 sense). However, for $m \geq 3$, the mechanism is manipulable,[3] and hence cannot be represented as a moving phantom mechanism. We point the reader to the work of Garg et al. [18] for a detailed overview of other settings in which market mechanisms have been used in the context of public decision making.

5.2 Voting Game Interpretation

We have seen descriptions of $\mathcal{A}^{IM}$ as a moving phantom mechanism and as clearing prices of a market system. We next give a game-theoretic description: the independent markets mechanism can be seen as the unique Nash equilibrium outcome of a voting game inspired by range voting [31]. The game works as follows. Each voter receives one "credit" per alternative, and chooses how much of that credit to place on the alternative. That is, each voter chooses a vector $\mathbf{s}_i \in [0, 1]^m$. The outcome of the game is the division $\mathbf{p}$ where p_j is proportional to $\sum_{i \in [n]} s_{i,j}$, the amount of credits spent on alternative j. Voters choose their spending so as to obtain an outcome that is as close to their ideal point $\hat{\mathbf{p}}_i$ as possible, according to ℓ_1 distance.

Theorem 5.4. *The voting game defined above has a unique outcome* $\mathbf{p}$ *that can be obtained in Nash equilibrium, and it is equal to the output of the Independent Markets mechanism.*

The idea of the proof is to set $\mathbf{s}_{i,j}$ equal to the amount that agent i spends in the market for alternative j, under the market setup that we described earlier. We show that when every agent casts "vote" $\mathbf{s_i}$, the system is at (its unique) equilibrium, and that the spending is proportional to the independent markets outcome.

Proof. Let $\mathbf{s_{-i}}$ denote the spending profile of all voters other than i, and let $\mathbf{q}$ denote the aggregate division in which the weight on an alternative is proportional to the number of credits spent on that alternative. Suppose that $\mathbf{s_i}$ is a best response for voter i. We first show that for every alternative j,

[3] Let $\mathbf{p_1} = (0, 0.5, 0.5)$, $\mathbf{p_2} = (0.5, 0.5, 0)$. Parimutuel consensus yields prices $(1/3, 1/3, 1/3)$, at distance $2/3$ from $\mathbf{p_1}$. If voter 1 instead reports $\hat{\mathbf{p}}_1 = (0, 0, 1)$, the price vector is $(0.25, 0.25, 0.5)$, at distance 0.5 from $\mathbf{p_1}$.

$q_j < p_{i,j}$ implies $s_{i,j} = 1$ and $q_j > p_{i,j}$ implies $s_{i,j} = 0$. That is, voters will always prefer to increase their spending on alternatives that they consider "undervalued", and decrease their spending on "overvalued" alternatives.

To prove this, suppose that $q_j < p_{i,j}$ and $s_{i,j} < 1$. Since $\mathbf{q}$ and $\mathbf{p_i}$ are both normalized, there must exist some alternative j' for which $q_{j'} > p_{i,j'}$. By marginally increasing $s_{i,j}$ while holding $s_{i,\tilde{j}}$ constant for all $\tilde{j} \neq j$, voter i can move the aggregate division from $\mathbf{q}$ to $\mathbf{q}'$, where $q'_j = q_j + \epsilon < p_{i,j}$, and $q'_{\tilde{j}} = q_{\tilde{j}} - \epsilon_{\tilde{j}}$ for all $\tilde{j} \neq j$, with $\sum_{\tilde{j} \neq j} \epsilon_{\tilde{j}} = \epsilon$ and $q'_{j'} > p_{i,j'}$.

We can now show that $d(\mathbf{q}, \mathbf{p_i}) > d(\mathbf{q}', \mathbf{p_i})$, which implies that $\mathbf{s_i}$ is not a best response for i.

$$\begin{aligned} d(\mathbf{q}, \mathbf{p_i}) - d(\mathbf{q}', \mathbf{p_i}) &= |q_j - p_{i,j}| - |q'_j - p_{i,j}| + |q_{j'} - p_{i,j'}| - |q'_{j'} - p_{i,j'}| + \sum_{\tilde{j} \neq j, j'} (|q_{\tilde{j}} - p_{i,\tilde{j}}| - |q'_{\tilde{j}} - p_{i,\tilde{j}}|) \\ &\geq \epsilon + \epsilon_{j'} - \sum_{\tilde{j} \neq j, j'} \epsilon_{\tilde{j}} = 2\epsilon_{j'} > 0 \end{aligned}$$

Next, we show that there exists an equilibrium of the voting game that produces the same outcome as independent markets. To this end, consider the market interpretation. For every alternative, some amount $x^* = n/t^*$ of a divisible good is sold to voters with a budget of 1 credit each. The number of credits that each voter spends in each market defines some spending profile $\mathbf{s}$.[4] Importantly, whenever $p_{i,j} > \mathcal{A}^{IM}(\mathbf{P})_j$, voter i spends their full budget on alternative j (i.e. $s_{i,j} = 1$), and whenever $p_{i,j} < \mathcal{A}^{IM}(\mathbf{P})_j$, voter i spends nothing on alternative j (i.e. $s_{i,j} = 0$). Further, the amount spent on each alternative is $x^* \cdot \mathcal{A}^{IM}(\mathbf{P})_j$ – the amount of goods sold multiplied by the price per unit – which is proportional to the aggregate division, $\mathcal{A}^{IM}(\mathbf{P})$. Therefore, the induced spending profile does produce the same aggregate division as independent markets when aggregated under the rules of the voting game. It remains to show that this spending profile is an equilibrium.

To do so, consider the spending vector $\mathbf{s_i}$ of some voter i, and suppose they have a better response $\mathbf{s}'_\mathbf{i}$. Denote the aggregate division when i spends $\mathbf{s_i}$ by $\mathbf{q}$ (this division is $\mathcal{A}^{IM}(\mathbf{P})$, but we use $\mathbf{q}$ for short), and the aggregate division when i spends $\mathbf{s}'_\mathbf{i}$ by $\mathbf{q}'$. There must exist either some alternative j with $q_j < p_{i,j}$ for which $q_j < q'_j$, and/or with $q_j > p_{i,j}$ for which $q_j > q'_j$. Suppose without loss of generality that the former case holds (if instead only the latter case holds, a very similar argument applies). Because $q_j < p_{i,j}$, we know that $s_{i,j} = 1$. Therefore, the only way for

$$q_j = \frac{s_{i,j} + \sum_{i' \neq i} s_{i',j}}{\sum_{j'} s_{i,j'} + \sum_{i' \neq i} \sum_{j'} s_{i',j'}} < \frac{s'_{i,j} + \sum_{i' \neq i} s_{i',j}}{\sum_{j'} s'_{i,j'} + \sum_{i' \neq i} \sum_{j'} s_{i',j'}} = q'_j$$

is for $\sum_{j'} s_{i,j'} > \sum_{j'} s'_{i,j'}$. But, because $s_{i,j'} = 0$ for all j' with $q_{j'} > p_{i,j'}$, the reduction in i's spending must come from alternatives with $q_{j'} \leq p_{i,j'}$. That is, $q_{j'} \leq p_{i,j'}$ for all alternatives with $q'_{j'} < q_{j'}$. Therefore

$$\begin{aligned} d(\mathbf{q}, \mathbf{p_i}) - d(\mathbf{q}', \mathbf{p_i}) &= \sum_{q'_{j'} < q_{j'}} (|q_{j'} - p_{i,j'}| - |q'_{j'} - p_{i,j'}|) + \sum_{q'_{j'} \geq q_{j'}} (|q_{j'} - p_{i,j'}| - |q'_{j'} - p_{i,j'}|) \\ &= \sum_{q'_{j'} < q_{j'}} (q'_{j'} - q_{j'}) + \sum_{q'_{j'} \geq q_{j'}} (|q_{j'} - p_{i,j'}| - |q'_{j'} - p_{i,j'}|) \\ &\leq \sum_{q'_{j'} < q_{j'}} (q'_{j'} - q_{j'}) + \sum_{q'_{j'} \geq q_{j'}} (q'_{j'} - q_{j'}) = 0 \end{aligned}$$

which contradicts that $\mathbf{s}'_\mathbf{i}$ is a better response than $\mathbf{s_i}$ for voter i.

[4]This spending profile is not unique, since for voters with $p_{i,j}$ equal to the clearing price for alternative j, there is some flexibility as to which voters pay for and get assigned goods, and which do not.

Next we show that the voting game has a *unique* equilibrium aggregate division.[5] We know that the independent markets aggregate $\mathbf{q}$ is an equilibrium with spending profile $\mathbf{s}$. For contradiction, suppose that there is some other equilibrium aggregate $\mathbf{q}'$ with spending profile $\mathbf{s}'$. Then there is an alternative j for which $q'_j > q_j$ and an alternative j' for which $q'_{j'} < q_{j'}$. Thus, there are weakly fewer voters with $p_{i,j} \geq q'_j$ than with $p_{i,j} > q_j$. Because only voters with $p_{i,j} \geq q'_j$ can have $s'_{i,j} > 0$, and all voters with $p_{i,j} > q_j$ have $s_{i,j} = 1$, we know that $\sum_i s_{i,j} > \sum_i s'_{i,j}$. But, because $q'_j > q_j$, global spending across all alternatives must be lower under $\mathbf{s}'$ than under $\mathbf{s}$; $\sum_i \sum_{\tilde{j}} s_{i,\tilde{j}} > \sum_i \sum_{\tilde{j}} s'_{i,\tilde{j}}$. By an identical argument, $\sum_i s_{i,j'} < \sum_i s'_{i,j'}$, implying that $\sum_i \sum_{\tilde{j}} s_{i,\tilde{j}} < \sum_i \sum_{\tilde{j}} s'_{i,\tilde{j}}$, contradicting the previous sentence. Hence, the equilibrium aggregate division $\mathbf{q} = \mathcal{A}^{IM}(\mathbf{P})$ is unique. □

To illustrate the voting game interpretation, consider again Example 5.2. Define spending vectors $\mathbf{s_1} = (0, 1, 1), \mathbf{s_2} = (1, 1, 0), \mathbf{s_3} = (1, 0, 0)$. These vectors sum to $(2, 2, 1)$, which is proportional to $\mathcal{A}^{IM}(\mathbf{P}) = (0.4, 0.4, 0.2)$. It can be checked that no voter wishes to unilaterally change their spending vector. For example, if voter 1 increases the amount she spends on alternative 1, she will increase the first coordinate of the outcome while decreasing the second and third coordinates, all of which increase the distance between the outcome and her preferred budget.

5.3 Other Properties of Independent Markets

Under the voting game interpretation, it is easy to see that independent markets satisfies proportionality. When voters are single-minded, it is a dominant strategy for them to spend their entire budget on alternatives they favor and spend nothing on all other alternatives. The amount spent on each alternative is therefore n_j, where n_j is the number of voters that favor alternative j.

Proposition 5.5. *$\mathcal{A}^{IM}$ satisfies proportionality.*

Independent markets satisfies several other desirable properties. One of these is participation, which says that voters should always prefer (truthfully) participating to not contributing a report.

Definition 5.6. A budget aggregation mechanism $\mathcal{A}$ satisfies *participation* if, for all $\mathbf{p_i}$ and all preference profiles $\mathbf{P}_{-i}$, $d(\mathcal{A}(\mathbf{p_i}, \mathbf{P_{-i}}), \mathbf{p_i}) \leq d(\mathcal{A}(\mathbf{P_{-i}}), \mathbf{p_i})$.

Theorem 5.7. *$\mathcal{A}^{IM}$ satisfies participation.*

Proof. Under $\mathcal{A}^{IM}$, adding new a voter who agrees with the aggregate division does not change the aggregate: $\mathcal{A}^{IM}(\mathcal{A}^{IM}(\mathbf{P_{-i}}), \mathbf{P_{-i}}) = \mathcal{A}^{IM}(\mathbf{P_{-i}})$. This can be seen by noting that for the new voters, it is an equilibrium spending strategy to spend nothing, thus not changing the spending profile. Then by incentive compatibility, $d(\mathcal{A}^{IM}(\mathbf{p_i}, \mathbf{P_{-i}}), \mathbf{p_i}) \leq d(\mathcal{A}^{IM}(\mathcal{A}^{IM}(\mathbf{P_{-i}}), \mathbf{P_{-i}}), \mathbf{p_i}) = d(\mathcal{A}^{IM}(\mathbf{P_{-i}}), \mathbf{p_i})$. □

Another desirable property is reinforcement, which states that the aggregate division should not change when two preference profiles that each agree with that aggregate are combined. For $\mathbf{P} = (\hat{\mathbf{p}}_1, \hat{\mathbf{p}}_2, \ldots, \hat{\mathbf{p}}_{\mathbf{n_P}})$ and $\mathbf{Q} = (\hat{\mathbf{q}}_1, \hat{\mathbf{q}}_2, \ldots, \hat{\mathbf{q}}_{\mathbf{n_Q}})$, let $\mathbf{P} \cup \mathbf{Q} = (\hat{\mathbf{p}}_1, \ldots, \hat{\mathbf{p}}_{\mathbf{n_P}}, \hat{\mathbf{q}}_1, \ldots, \hat{\mathbf{q}}_{\mathbf{n_Q}})$.

Definition 5.8. A budget aggregation mechanism $\mathcal{A}$ satisfies reinforcement if, for all preference profiles $\mathbf{P}$ and $\mathbf{R}$ with $\mathcal{A}(\mathbf{P}) = \mathcal{A}(\mathbf{R})$, it holds that $\mathcal{A}(\mathbf{P} \cup \mathbf{R}) = \mathcal{A}(\mathbf{P}) = \mathcal{A}(\mathbf{R})$.

Theorem 5.9. *Independent markets satisfies reinforcement.*

Proof. Let $\mathbf{P}$ and $\mathbf{R}$ be profiles with $\mathcal{A}^{IM}(\mathbf{P}) = \mathcal{A}^{IM}(\mathbf{R}) = \mathbf{q}$. We utilize the market interpretation of independent markets. Suppose that for profile $\mathbf{P}$, market prices are normalized when x_P^* goods

[5] In contrast, the exact spending profile is clearly not unique. For instance, in an instance with only a single voter, that voter can enforce their belief exactly as long as they spend on each outcome in proportion to their belief, with no restriction on the magnitude of their spending.

are sold and for profile **R**, market prices are normalized when x_R^* goods are sold. Now consider the combined profile $\mathbf{P} \cup \mathbf{R}$ when $x_P^* + x_R^*$ goods are sold. For every alternative j, and every price $\pi \in [0, 1]$, the total demand is equal to the total demand in profile P at price π plus the total demand in profile R at price π, since a voter's demand depends only on their valuation and the price π. Likewise, the total supply is, by definition, the sum of the supplies in each individual instance. Therefore, the market clearing prices in the combined instance when $x_P^* + x_R^*$ goods are sold are equal to the (normalized) aggregate vector **q**. □

We next check that the independent markets outcome is always rational and can be described in polynomially many bits, thus ensuring that it can be computed efficiently as suggested in Section 4. Our argument proceeds by showing that the outcome is a solution of a linear program, similar to a proof of rationality for the parimutuel consensus mechanism [32, Thm. 5.1]. Consider the outcome **p** of the independent markets, and write $N_j = \{i \in N : p_j < \hat{p}_{i,j}\}$ for the set of voters that purchase good $j \in [m]$ since their value is lower than its price. Now, if x is the supply of each good, then the amount xp_j of money spent on j equals the budget of the demanders, which is $|N_j|$. Introducing a variable $z \equiv 1/x$, we can write this as $p_j = z \cdot |N_j|$. Thus, **p** is the solution of maximizing ϵ subject to

$$\begin{aligned}
\hat{p}_{i,j} &\leq p_j - \epsilon && \text{for } j \in [m] \text{ and } i \in N_j, \\
\hat{p}_{i,j} &\geq p_j && \text{for } j \in [m] \text{ and } i \in N \setminus N_j, \\
p_j &= z \cdot |N_j| && \text{for } j \in [m], \\
\textstyle\sum_{j \in [m]} p_j &= 1, \quad p_j \geq 0, \quad z \geq 0 && \text{for } j \in [m].
\end{aligned}$$

Using standard encoding techniques, one can also calculate the independent markets mechanism using an ILP with binary variables encoding "$i \in N_j$."

6 PARETO OPTIMALITY AND SOCIAL WELFARE

The independent markets mechanism has several natural interpretations, and it satisfies proportionality. However, it is not Pareto-optimal. If voter 1 reports $(0.8, 0.2, 0)$ and voter 2 reports $(0.8, 0, 0.2)$, then independent markets returns $(0.6, 0.2, 0.2)$, which is dominated by $(0.8, 0.1, 0.1)$. On this example, independent markets even fails to be *range-respecting*, which requires that $\min_{i \in [n]} \hat{p}_{i,j} \leq \mathcal{A}(\mathbf{P})_j \leq \max_{i \in [n]} \hat{p}_{i,j}$ for all $j \in [m]$.

The failure to be range-respecting can be fixed, if desired, by changing the positions of phantoms 0 and n. One can show that a moving phantom mechanism $\mathcal{A}^{\mathcal{F}}$ is range-respecting if and only if $f_0(t) = 1$ and $f_n(t) = 0$ for all $t \in [0, 1]$ except for an initial period where phantom 0 moves from 0 to 1 while all other phantoms remain at 0, and a period at the end where phantom n moves from 0 to 1 while all other phantoms are at 1. This mirrors a result in Section 3; if the outer two phantoms are at 0 and 1, the $n - 1$ remaining phantoms cannot outweigh the n voter reports.

While many moving phantom mechanisms are range-respecting, it is much more difficult to find a mechanism in this class which is Pareto-optimal. Usually, it is possible to construct a profile in which the mechanism returns a vector **p** all of whose entries are phantom reports, and then a Pareto improvement can be obtained by perturbing this vector in the directions where the majority of voter reports lie. Such a perturbation is not possible if the phantoms lie at 0 or 1, which turns out to be the only escape. As we prove below, no mechanism $\mathcal{A}^{\mathcal{F}}$ can be Pareto-optimal if there is any time point t when two phantoms are both strictly between 0 and 1.

This condition is extremely restrictive, and a moment's thought reveals that there is only one legal phantom system which avoids having two interior phantoms: All phantoms start at 0, and then, one by one, one of the phantoms is moved to 1. At each t, at most one phantom lies strictly between 0 and 1 while travelling. We call this phantom system $\mathcal{F}^*$. It can be formalized as

$$f_k(t) = \begin{cases} 0 & 0 \le t \le \frac{k}{n+1}, \\ t(n+1) - k & \frac{k}{n+1} < t < \frac{k+1}{n+1}, \\ 1 & \frac{k+1}{n+1} \le t \le 1. \end{cases}$$

Below we will show that $\mathcal{A}^{\mathcal{F}^*}$ precisely corresponds to the budget aggregation mechanism that maximizes voter welfare, breaking ties in favor of the maximum entropy division. It will immediately follow that $\mathcal{A}^{\mathcal{F}^*}$ is indeed Pareto-optimal. Combined with Theorem 6.1 below, which shows that all other moving phantom mechanisms are Pareto-inefficient, this implies that the welfare-maximizing mechanism is the unique Pareto-optimal moving phantom mechanism.

6.1 Characterizing Pareto Optimality

The proof of Theorem 6.1 shows, by induction, that each phantom needs to move all the way to 1 before the next phantom can leave its position at 0. In case this does not happen, based on the approximate phantom positions, we construct a profile where the mechanism is Pareto-inefficient. These constructions are of two kinds: an easier case when the interior phantoms are low (lying below $\frac{1}{n(n-1)}$), and a more involved case when one of the phantoms has moved higher. In both cases, our constructions utilize two types of alternatives. More voters report "high" probabilities on alternatives of the first type than on alternatives of the second type. The constructions work so that if two phantoms simultaneously take values between 0 and 1, then the mechanism outputs middling values on all alternatives. Social welfare can be improved by increasing the output on alternatives of the first type, and decreasing the output on alternatives of the second type. By incorporating enough symmetry between voters, we guarantee that social welfare gains are shared equally, so obtain a Pareto improvement.

Theorem 6.1. *A moving phantom mechanism $\mathcal{A}^{\mathcal{F}}$ cannot be Pareto-optimal for any $m \ge n^2$ unless $\mathcal{A}^{\mathcal{F}} = \mathcal{A}^{\mathcal{F}^*}$.*

Proof. We first show that any Pareto optimal moving phantom mechanism $\mathcal{A}^{\mathcal{F}}$ for which there exists a t with $f_0(t) < 1$ and $f_1(t) > 0$ can be equivalently expressed as a moving phantom mechanism that does not have such a t. Suppose that $\mathcal{A}^{\mathcal{F}} = \text{med}(f_0(t^*), \dots, f_n(t^*), \hat{p}_{1,j}, \dots, \hat{p}_{n,j}) = f_0(t^*) < \bar{p}_{1,j}$ for some j. Then $\mathcal{A}^{\mathcal{F}}$ is not Pareto optimal, because increasing $\mathcal{A}^{\mathcal{F}}(\mathbf{P})_j$ and decreasing $\mathcal{A}^{\mathcal{F}}(\mathbf{P})_{j'}$ for any coordinate with $\mathcal{A}^{\mathcal{F}}(\mathbf{P})_{j'} > \bar{p}_{1,j'}$ is a Pareto improvement (such a coordinate must exist, because $\sum_j \bar{p}_{1,j} \le 1$). Therefore, for all preference profiles $\mathbf{P}$, $\mathcal{A}^{\mathcal{F}} = \text{med}(f_0(t^*), \dots, f_n(t^*), \hat{p}_{1,j}, \dots, \hat{p}_{n,j}) \ge \bar{p}_{1,j}$. This implies that $f_0(t^*) \ge \bar{p}_{1,j}$ and so the exact position of $f_0(t^*)$ has no effect on the mechanism. It would be equivalent to move phantom f_0 to position 1 before moving phantom f_1.

A very similar argument can be used to show that there cannot exist a t for which $f_{n-1}(t) < 1$ and $f_n(t) > 0$. For the rest of the proof, we focus on the intermediate phantoms. Suppose that there exists some index $1 \le k \le n-2$ for which $f_k(t) < 1$ and $f_{k+1}(t) > 0$ for some t. If no such k exists, then phantom system $\mathcal{F} = \mathcal{F}^*$.

We next show that if $f_k(t) < \frac{1}{n(n-1)}$, it must be the case that $x_{k+1}(t) = 0$. Define an instance with $m = n^2$ alternatives. Voter $i \in [n]$ reports $\hat{p}_{i,j} = \frac{1}{n^2-kn-1}$ for alternatives $j \in \{((i-1)(n-1)+1, \dots, (i-1)(n-1) + n^2 - n - kn + k) \mod n(n-1)\}$ and for alternatives $j \in \{n(n-1) + (i, i+1, i+n-k-2) \mod n\}$, and $\hat{p}_{i,j} = 0$ for all other alternatives. Note that $\sum_{j=1}^{m} \hat{p}_{i,j} = 1$. Further, note that among alternatives 1, ..., $n(n-1)$, each voter makes $n^2 - n - kn + k = (n-1)(n-k)$ non-zero reports and each alternative has $n-k$ non-zero reports, while among alternatives $n(n-1)+1, \dots, n^2$, each voter makes $n-k-1$ non-zero reports and each alternative has $n-k-1$ non-zero reports. Therefore, if

$f_k(t) < \frac{1}{n^2-kn-1}$, the generalized median on alternative j is $f_k(t)$ for $j \in \{1, \dots, n(n-1)\}$ and the median on alternative j is $f_{k+1}(t)$ for $j \in \{n+1, \dots, 2n\}$.

Suppose that there exists t for which $f_k(t) < \frac{1}{n(n-1)} < \frac{1}{n^2-kn-1}$ and $f_{k+1}(t) > 0$. Then, since f is increasing and continuous, and the aggregate division $\mathcal{A}^{\mathcal{F}}(\mathbf{P})$ is normalized, it will necessarily be the case that for all $j \in \{1, ..., n(n-1)\}$, $\mathcal{A}^f(\mathbf{P})_j = f_k(t^*) < \frac{1}{n(n-1)}$, and for all $j \in \{1, ..., n(n-1)\}$, $\mathcal{A}^f(\mathbf{P})_j = f_{k+1}(t^*) > 0$, with $n(n-1)f_k(t^*) + nf_{k+1}(t^*) = 1$. But this is not Pareto optimal. Consider, for some small enough ϵ, increasing $\mathcal{A}^{\mathcal{F}}(\mathbf{P})_j$ by ϵ on alternatives $j \in \{1, \dots, n(n-1)\}$, and decreasing $\mathcal{A}^{\mathcal{F}}(\mathbf{P})_j$ by $\epsilon(n-1)$ on alternatives $j \in \{n(n-1)+1, \dots, n^2\}$. For every voter i, there are $n^2-n-kn+k$ alternatives on which the aggregate moves ϵ closer to i's report, $kn-k$ alternatives for which the aggregate moves ϵ farther from i's report, $n-k-1$ alternatives on which the aggregate moves $\epsilon(n-1)$ further from i's report, and $k+1$ alternatives for which the aggregate moves $\epsilon(n-1)$ closer to i's report. Summing these up, the change moves the aggregate $(2n-2)\epsilon$ closer to $\hat{\mathbf{p}}_\mathbf{i}$ in ℓ_1 distance.

We now show that if $f_k(t) < 1$ then it must be the case that $f_{k+1}(t) = 0$. For contradiction suppose otherwise. Let $\bar{t} = \sup\{t : f_{k+1}(t) = 0\}$ be the final snapshot at which $f_{k+1}(t) = 0$. By assumption, $f_k(\bar{t}) < 1$. We define an instance similar to that above. Let $\delta > 0$ (we will determine the exact value of δ later). For every voter i, let $\hat{p}_{i,j} = \frac{1-f_k(\bar{t})-\delta}{n(n-1)-1}$ for every $j \in \{(i-1)(n-1)+1, \dots, (i-1)(n-1)+n-1\}$ and $\hat{p}_{i,j} = \frac{1-f_k(\bar{t})}{n(n-1)-1}$ for every $j \in \{((i-1)(n-1)+n, \dots, (i-1)(n-1)+n^2-n-kn+k) \mod n(n-1)\}$. However, for $j = 1$, for every voter i with $p_{i,1} = \frac{1-f_k(\bar{t})-\delta}{n(n-1)-1}$ we instead set $p_{i,1} = f_k(\bar{t}) + \delta$, overriding the earlier setting. Because we know that $f_k(\bar{t}) \geq \frac{1}{n(n-1)}$, we have that $\frac{1-f_k(\bar{t})-\delta}{n(n-1)-1} < \frac{1-\frac{1}{n(n-1)}}{n(n-1)-1} \leq \frac{1}{n(n-1)} \leq x_k$, therefore the new value of $\hat{p}_{i,1}$ is higher than the one it replaces. To set δ, choose some value that guarantees $\sum_{j=1}^{n(n-1)} \hat{p}_{i,j} < 1$ for all i. In particular, by the previous observation, it is sufficient to set δ so that

$$f_k(\bar{t}) + \delta + (n-2)\frac{1-f_k(\bar{t})-\delta}{n(n-1)-1} + (n-k-1)(n-1)\frac{1-f_k(\bar{t})}{n(n-1)-1} < 1. \tag{3}$$

To see that such a value of δ exists, note that Equation 3 is continuous in δ and takes value strictly less than 1 when $\delta = 0$:

$$\begin{aligned} f_k(\bar{t}) + (n-2)\frac{1-f_k(\bar{t})}{n(n-1)-1} + (n-k-1)(n-1)\frac{1-f_k(\bar{t})}{n(n-1)-1} &< f_k(\bar{t}) + n(n-2)\frac{1-f_k(\bar{t})}{n(n-1)-1} \\ &< f_k(\bar{t}) + 1 - f_k(\bar{t}) = 1, \end{aligned}$$

where we may assume $n \geq 3$ because the case of $n = 2$ has only a single phantom that is not f_0 or f_n. For every $j \in \{1, \dots, n(n-1)\}$ for which $\hat{p}_{i,j}$ is not explicitly set greater than 0, we set it to 0.

For all i, we evenly distribute the remaining (positive) mass $1 - \sum_{j=1}^{n(n-1)} p_{i,j}$ evenly among $j \in \{n(n-1) + (i, i+1, i+n-k-2) \mod n\}$.

When $t = \bar{t}$, the generalized median on each alternative is $f_k(t) = f_k(\bar{t})$ for alternative 1, $\frac{1-f_k(\bar{t})-\delta}{n(n-1)-1}$ for all $j \in \{2, \dots, n(n-1)$, and $x_{k+1}(\bar{t}) = 0$ for all $j \in \{n(n-1)+1, \dots, n^2\}$. The sum of these generalized medians is $1-\delta$. Therefore t needs to increase to achieve normalization. By the definition of $\bar{t}$, for any $t > \bar{t}$, we have that $f_{k+1}(t) > 0$, and therefore the generalized median on all alternatives $j \in \{n(n-1)+1, \dots, n^2\}$ is greater than 0. It is therefore impossible for $f_k(t)$ to reach $f_k(\bar{t}) + \delta$, because then the sum of generalized medians would exceed 1. It is also impossible for $f_{k+1}(t)$ to reach $\frac{1-f_k(\bar{t})}{n(n-1)-1}$. If it does, then the generalized median on alternative 1 is at least $f_k(\bar{t})$, on $j \in \{2, \dots, n(n-1)\}$ is $\frac{1-f_k(\bar{t})}{n(n-1)-1}$ and on $j \in \{n(n-1)+1, \dots, n^2\}$ is strictly greater than 0. Therefore the aggregate is not normalized. To summarize, we are guaranteed that $\frac{1}{n(n-1)} \leq \mathcal{A}^{\mathcal{F}}(\mathbf{P})_1 < f_k(\bar{t}) + \delta$, $\frac{1-f_k(\bar{t})-\delta}{n(n-1)-1} \leq \mathcal{A}^{\mathcal{F}}(\mathbf{P})_j < \frac{1-f_k(\bar{t})}{n(n-1)-1}$ for all $j \in \{2, \dots, n(n-1)\}$, and $\mathcal{A}^{\mathcal{F}}(\mathbf{P})_j > 0$ for all $j \in \{n(n+1)+1, \dots, n^2\}$.

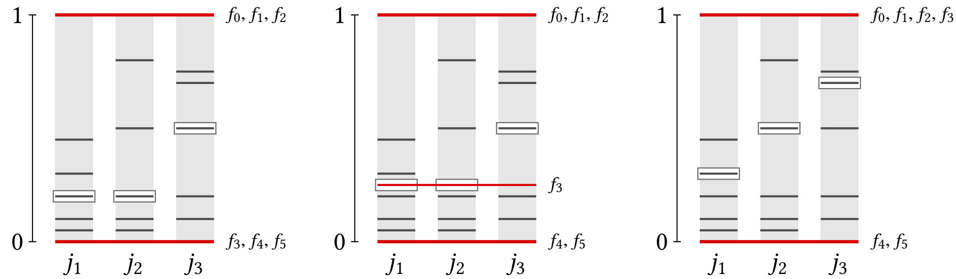

Fig. 2. Snapshots of the phantom system $\mathcal{F}^*$ with $t < t^*$ (left), $t = t^*$ (center), and $t > t^*$ (right) on an instance with $n = 5$, $m = 3$.

Now we can define a Pareto improvement to $\mathcal{A}^{\mathcal{F}}(\mathbf{P})$ of the same form as previously. For some small enough ϵ, increase the aggregate by ϵ on alternatives $j \in \{1, \ldots, n(n-1)\}$, and decrease the aggregate by $\epsilon(n-1)$ on alternatives $j \in \{n(n-1)+1, \ldots, n^2\}$. For voter 1, with $p_{1,1} = f_k(\bar{t}) + \delta$, the new aggregate is ϵ better than $\mathcal{A}^{\mathcal{F}}(\mathbf{P})$ on alternative 1, ϵ worse on alternatives $j \in \{2, \ldots, n-1\}$, with $p_{1,j} = \frac{1-f_k(\bar{t})-\delta}{n(n-1)-1}$, ϵ better on alternatives $j \in \{n, \ldots, n^2 - n - kn + k\}$, with $p_{i,j} = \frac{1-f_k(\bar{t})}{n(n-1)-1}$, and ϵ worse on alternatives $j \in \{n^2 - n - kn + k + 1, \ldots, n(n-1)\}$, with $p_{i,j} = 0$. On the final n alternatives, there are at least $k+1$ alternatives on which the aggregate moves $\epsilon(n-1)$ closer to voter 1's report (on alternatives for which voter 1 reports $p_{i,j} = 0$), and at most $n-k-1$ alternatives for which the aggregate moves $\epsilon(n-1)$ farther from voter 1's report. Summing these up, we get the aggregate has moved $1-(n-2)+(n-k-1)(n-1)-k(n-1)-(n-1)(n-k-1)+(n-1)(k+1) = 2$ towards $\mathbf{p_1}$.

For all other voters, the new aggregate is ϵ worse than $\mathcal{A}^{\mathcal{F}}(\mathbf{P})$ on alternatives $j \in \{(i-1)(n-1)+1, \ldots, n-1\}$, with $p_{i,j} = \frac{1-f_k(\bar{t})-\delta}{n(n-1)-1}$, and ϵ better on alternatives $j \in \{((i-1)(n-1)+n, \ldots, (i-1)(n-1)+n^2-n-kn+k) \mod n(n-1)\}$, with $p_{i,j} = \frac{1-f_k(\bar{t})}{n(n-1)-1}$, and ϵ worse on alternatives $j \in \{((i-1)(n-1)+n^2-n-kn+k+1, \ldots, (i-1)(n-1)+n(n-1)) \mod n(n-1)\}$, with $p_{i,j} = 0$. On the final n alternatives, there are at least $k+1$ alternatives on which the aggregate moves $\epsilon(n-1)$ closer to voter i's report (on alternatives for which voter i reports $p_{i,j} = 0$), and at most $n-k-1$ alternatives for which the aggregate moves $\epsilon(n-1)$ farther from voter i's report. Summing these up, we get that the new and old aggregates are equal ℓ_1 distances from $\mathbf{p_i}$.

Finally, our construction uses $m = n^2$ alternatives, but we can extend it to larger m by adding dummy alternatives that no voter puts any weight on. □

On profiles consisting of single-minded voters, $\mathcal{A}^{\mathcal{F}^*}$ selects a division that is also single-minded, following the plurality. Hence, it is not proportional, which gives the following corollary.

Corollary 6.2. *For $m \geq n^2$, no moving phantom mechanism is proportional and Pareto-optimal.*

6.2 Maximizing Social Welfare

Having narrowed down the space of Pareto-optimal moving phantom mechanisms to at most one mechanism, let us examine the behavior of $\mathcal{A}^{\mathcal{F}^*}$ with the assistance of Figure 2, which takes the same form as Figure 1. On every alternative, order the entries $\{p_{i,j}\}$ from largest to smallest. We denote the relabeled entries $\bar{p}_{1,j} \geq \ldots \geq \bar{p}_{n,j}$. At the snapshot of $\mathcal{F}^*$ for which $f_0(t) = \ldots = f_k(t) = 1$ and $f_{k+1}(t) = \ldots = f_n(t) = 0$, the generalized median selects the order statistic $\bar{p}_{n-k,j}$ for all j. We see this in Figure 2 where, in the left image, $k = 2$ and the generalized median is the $n - k = 3$rd highest report on each alternative, and in the right image $k = 3$ and the $n - k = 2$nd highest reports are chosen.

We can think of $\mathcal{F}^*$ as partitioning the phantom "movie" into periods defined by which phantom is moving. Initially, all phantoms are at 0, and the generalized medians are 0 for each $j \in [m]$. Then phantom f_0 moves to 1, and the generalized medians are $\bar{p}_{n,j}$. As phantom f_k moves from 0 to 1, the generalized medians progress from $\bar{p}_{n-k+1,j}$ to $\bar{p}_{n-k,j}$, until all phantoms reach 1 and the generalized medians are uniformly 1. By (a discrete analogue of) the intermediate value theorem, there must exist some value I for which $\sum_{j\in[m]} \bar{p}_{I+1,j} \le 1$ and $\sum_{j\in[m]} \bar{p}_{I,j} \ge 1$, and this transition is made during the period in which phantom $n - I$ is moving. In Figure 2, we have $I = 2$ because the sum of the third-highest entries is less than one (see the left image), while the sum of the second-highest entries is more than one (the right image).

Normalization therefore occurs during the movement of phantom f_{n-I}, and the final value $\mathcal{A}^{\mathcal{F}^*}(\mathbf{P})_j$ lies in the interval $[\bar{p}_{I+1,j}, \bar{p}_{I,j}]$. If $f_{n-I}(t^*)$ lies in this interval, then $\mathcal{A}^{\mathcal{F}^*}(\mathbf{P})_j = f_{n-I}(t^*)$, otherwise $\mathcal{A}^{\mathcal{F}^*}(\mathbf{P})_j$ is equal to the endpoint of the interval closest to $f_{n-I}(t^*)$. This is depicted in the center image of Figure 2, where $f_3(t^*)$ lies between the second and third-highest reports on the first two alternatives, but below the third-highest report on the third alternative.

Finding the exact value of $f_{n-I}(t^*)$, and therefore the output $\mathcal{A}^{\mathcal{F}^*}(\mathbf{P})$, can be thought of as finding the "most equal" division, subject to interval constraints on each alternative. This problem has been studied before, and the (unique) value of $f_{n-I}(t^*)$ can be found in $O(m \log m)$ time by the Divvy algorithm of Gulati et al. [22].

Given a profile $\mathbf{P}$, the *social cost* of an outcome $\mathbf{p}$ is $\sum_{i\in[n]} d(\hat{\mathbf{p}}_i, \mathbf{p})$, and the (utilitarian) social welfare of $\mathbf{p}$ is the negation of the social cost. In general, there may be multiple divisions that maximize social welfare. For example, if $m = 2$, one voter reports $(1, 0)$ and another reports $(0, 1)$, then all divisions have the same social cost of 2. As it turns out, any division that satisfies the upper and lower bound constraints of $\bar{p}_{I,j}$ and $\bar{p}_{I+1,j}$ maximizes social welfare.

Lemma 6.3. *A division* $\mathbf{q}$ *maximizes social welfare if and only if* $\bar{p}_{I+1,j} \le q_j \le \bar{p}_{I,j}$ *for all* j.

Proof. Let $\mathbf{q}$ be a division with $\bar{p}_{I+1,j} \le q_j = \bar{p}_{I+1,j} + \epsilon_j \le \bar{p}_{I,j}$ for all j, with normalization of $\mathbf{q}$ implying that $\sum_{j\in[m]} \epsilon_j = 1 - \sum_{j\in[m]} \bar{p}_{I+1,j}$. Then the social cost of $\mathbf{q}$ is

$$\sum_{j\in[m]} \sum_{i\in[n]} |\bar{p}_{i,j} - q_j| = \sum_{j\in[m]} \left(\sum_{i\in[n]} |\bar{p}_{i,j} - \bar{p}_{I+1,j}| + \sum_{i\ge I+1} \epsilon_j - \sum_{i\le I} \epsilon_j \right)$$

$$= \sum_{j\in[m]} \sum_{i\in[n]} |\bar{p}_{i,j} - \bar{p}_{I+1,j}| + (n - 2I) \left(1 - \sum_{j\in[m]} \bar{p}_{I+1,j} \right)$$

Because this expression does not depend on ϵ_j, all such divisions $\mathbf{q}$ have the same social cost.

We now show that this distance is minimal. Let $\mathbf{q}$ be a division that does not satisfy $\bar{p}_{I+1,j} \le q_j \le \bar{p}_{I,j}$ for some j. Suppose $q_j > \bar{p}_{I,j}$ (the case where $q_j < \bar{p}_{I+1,j}$ can be handled similarly). By the definition of I, there must exist some alternative j' for which $q_{j'} < \bar{p}_{I,j'}$. Now, consider the division $\mathbf{q}'$ defined by $q'_j = q_j - \epsilon > \bar{p}_{I,j}$ and $q'_{j'} = q_{j'} + \epsilon < \bar{p}_{I,j'}$, with $\mathbf{q}'$ and $\mathbf{q}$ equal on all other coordinates. Compare $\mathbf{q}$ and $\mathbf{q}'$ in terms of ℓ_1 distance from the reports. They are indistinguishable on all alternatives other than j and j'. On alternative j, $\mathbf{q}'$ is ϵ closer than $\mathbf{q}$ to all entries $\bar{p}_{i,j}$ with $i \ge I$, and at most ϵ farther from all other entries. On alternative j', $\mathbf{q}'$ is ϵ closer than $\mathbf{q}$ to all entries $\bar{p}_{i,j'}$ with $i \le I$, and at most ϵ farther from all other entries. Therefore, of the $2n$ entries on alternatives j and j', $\mathbf{q}'$ is ϵ closer than $\mathbf{q}$ to at least $n + 1$ of them, and no more than ϵ farther than $\mathbf{q}$ from the other $n - 1$. Therefore, $\mathbf{q}$ does not maximize social welfare. □

As a corollary of Lemma 6.3, we immediately obtain that $\mathcal{A}^{\mathcal{F}^*}$ maximizes social welfare, and is therefore Pareto-optimal. Social welfare-maximizing mechanisms have been considered before; all that is needed is a suitable tiebreaking procedure to select a single division from the set of maximizers. Goel et al. [21] suggest breaking ties by selecting the lexicographically largest welfare maximizer, but this is not neutral. We propose a different way to break ties, which is neutral: select the welfare maximizer $\mathbf{p}$ with largest Shannon entropy $-\sum_{j\in[m]} p_j \log p_j$. Because the set of welfare-maximizers is convex, and Shannon entropy is a convex function, existence and uniqueness of $\mathbf{p}$ is guaranteed.

Theorem 6.4. *For every profile* $\mathbf{P}$, $\mathcal{A}^{\mathcal{F}^*}$ *selects the entropy-maximizing division among those that maximize social welfare.*

Proof. From the earlier discussion, we know that $\mathcal{A}^{\mathcal{F}^*}(\mathbf{P})_j = \text{med}\{\bar{p}_{I+1,j}, \bar{p}_{I,j}, f_{n-I}(t^*)\} \in [\bar{p}_{I+1,j}, \bar{p}_{I,j}]$ for all $j \in [m]$. Therefore, it maximizes social welfare.

It remains to show that, subject to these constraints, $\mathcal{A}^{\mathcal{F}^*}(\mathbf{P})$ maximizes Shannon entropy. Consider any other division $\mathbf{q} \neq \mathcal{A}^{\mathcal{F}^*}(\mathbf{P})$ with $\bar{p}_{I+1,j} \leq q_j \leq \bar{p}_{I,j}$. Then there must exist an alternative j for which $\bar{p}_{I+1,j} \leq \mathcal{A}^{\mathcal{F}^*}(\mathbf{P})_j < q_j \leq \bar{p}_{I,j}$. Further, because $\mathcal{A}^{\mathcal{F}^*}(\mathbf{P})_j < \bar{p}_{I,j}$, it must be the case that $f_{n-I}(t^*) \leq \mathcal{A}^{\mathcal{F}^*}(\mathbf{P})_j = \text{med}\{\bar{p}_{I+1,j}, \bar{p}_{I,j}, f_{n-I}(t^*)\}$. There must also be an alternative j' for which $\bar{p}_{I+1,j'} \leq q_{j'} < \mathcal{A}^{\mathcal{F}^*}(\mathbf{P})_{j'} \leq \bar{p}_{I,j'}$, with $\mathcal{A}^{\mathcal{F}^*}(\mathbf{P})_{j'} \leq f_{n-I}(t^*)$.

Putting these together, we have that $q_{j'} < f_{n-I}(t^*) < q_j$. We also know that $\bar{p}_{I+1,j} < q_j$ and $q_{j'} < \bar{p}_{I,j'}$. Therefore, adjusting q_j to $q_j - \epsilon$ and $q_{j'}$ to $q_{j'} + \epsilon$, for ϵ small enough that none of the above strict inequalities are violated, both (1) decreases $|q_j - q_{j'}|$, which it is easy to check increases Shannon entropy, and (2) respects social-welfare maximization. Therefore, $\mathbf{q}$ is not the unique entropy-maximizing division among social welfare maximizers, so $\mathcal{A}^{\mathcal{F}^*}$ is. □

Similar to independent markets, $\mathcal{A}^{\mathcal{F}^*}$ satisfies participation. The argument proceeds along exactly the same lines. For an individual voter, participating and reporting the existing aggregate has no effect, while reporting truthfully is at least as beneficial, by incentive compatibility.

Theorem 6.5. $\mathcal{A}^{\mathcal{F}^*}$ *satisfies participation.*

The utilitarian mechanism also satisfies reinforcement, which follows from the additivity of social welfare.

Theorem 6.6. $\mathcal{A}^{\mathcal{F}^*}$ *satisfies reinforcement.*

Proof. Let $\mathbf{P}$ and $\mathbf{R}$ be profiles with $\mathcal{A}^{\mathcal{F}^*}(\mathbf{P}) = \mathcal{A}^{\mathcal{F}^*}(\mathbf{R}) = \mathbf{q}$. Let $\mathbf{q}'$ be some division with higher entropy than $\mathbf{q}$. Then it must be the case that $d(\mathbf{q}', \mathbf{P}) > d(\mathbf{q}, \mathbf{P})$ and $d(\mathbf{q}', \mathbf{R}) > d(\mathbf{q}, \mathbf{R})$. Since the ℓ_1 distance is additive, $d(\mathbf{q}', \mathbf{P} \cup \mathbf{R}) > d(\mathbf{q}, \mathbf{P} \cup \mathbf{R})$.

Now let $\mathbf{q}'$ be some division with $d(\mathbf{q}', \mathbf{P} \cup \mathbf{R}) < d(\mathbf{q}, \mathbf{P} \cup \mathbf{R})$. Then $d(\mathbf{q}', \mathbf{P}) + d(\mathbf{q}', \mathbf{R}) < d(\mathbf{q}, \mathbf{P}) + d(\mathbf{q}, \mathbf{R})$. But this is impossible, since $\mathbf{q}$ maximizes social welfare for profiles $\mathbf{P}$ and $\mathbf{R}$.

Therefore, any division $\mathbf{q}' \neq \mathbf{q}$ either has strictly lower social welfare on profile $\mathbf{P} \cup \mathbf{R}$, or equal social welfare but lower entropy. So $\mathcal{A}^{\mathcal{F}^*}((\mathbf{P} \cup \mathbf{R})) = \mathbf{q}$. □

7 CONCLUSION

We considered the problem of aggregating budget proposals for participatory budgeting. Inspired by the generalized median mechanisms of Moulin [26], we introduced the broad class of moving phantom mechanisms and proved that all mechanisms in this class are incentive compatible under ℓ_1 voter preferences. We analyzed two moving phantom mechanisms in detail: one that maximizes social welfare while violating the natural fairness notion of proportionality, and another that satisfies proportionality while violating Pareto optimality.

In addition to the properties discussed in the main body of the paper, some other contrasts can be drawn between these two mechanisms. For example, suppose that n is large, and $n/2$ voters report $(1, 0)$ while the other $n/2$ report $(0, 1)$. Then, under $\mathcal{A}^{\mathcal{F}^*}$, a single voter can completely dictate the outcome. Under independent markets, we can show that the ability of a single voter to affect the outcome vanishes as $n \to \infty$.

We have implemented both mechanisms. Some preliminary simulation results suggest that even when voters are not single-minded, independent markets is "more proportional" than $\mathcal{A}^{\mathcal{F}^*}$, in the sense that it better reflects all voters' opinions and not just that of the majority. However, independent markets also has a tendency to shift the aggregate towards the uniform division, relative to what we might otherwise expect. This deserves more investigation.

There are many other moving phantom mechanisms that we have not considered. It would be interesting to investigate what other properties can be achieved by investigating other phantom systems. And finally, as we mentioned earlier, when there are only two outcomes, we know that all (anonymous, neutral, and continuous) incentive compatible budget aggregation mechanisms can be represented as moving phantom mechanisms. It remains an open question whether this continues to hold when the number of outcomes m is more than 2.

ACKNOWLEDGMENTS

Dominik Peters was supported by EPSRC and by ERC under grant 639945 (ACCORD).

REFERENCES

[1] N. Alon, M. Feldman, A. D. Procaccia, and M. Tennenholtz. 2009. Strategyproof Approximation Mechanisms for Location on Networks. arXiv:0907.2049. (2009).

[2] H. Aziz, A. Bogomolnaia, and H. Moulin. 2017. Fair mixing: the case of dichotomous preferences. arXiv:1712.02542. (2017).

[3] H. Aziz, B.E. Lee, and N. Talmon. 2018. Proportionally representative participatory budgeting: Axioms and algorithms. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 23–31.

[4] S. Barberà, F. Gul, and E. Stacchetti. 1993. Generalized Median Voter Schemes and Committees. *Journal of Economic Theory* 61, 2 (1993), 262–289.

[5] S. Barberà and M. Jackson. 1994. A characterization of strategy-proof social choice functions for economies with pure public goods. *Social Choice and Welfare* 11, 3 (1994), 241–252.

[6] S. Barberà, J. Masso, and A. Neme. 1997. Voting under Constraints. *Journal of Economic Theory* 76, 2 (1997), 298–321.

[7] G. Benade, A. D. Procaccia, S. Nath, and N. Shah. 2017. Preference Elicitation for Participatory Budgeting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*. 376–382.

[8] A. Bogomolnaia, H. Moulin, and R. Stong. 2005. Collective Choice Under Dichotomous Preferences. *Journal of Economic Theory* 122, 2 (2005), 165–184.

[9] K. C. Border and J. S. Jordan. 1983. Straightforward elections, unanimity and phantom voters. *Review of Economic Studies* 50, 1 (1983), 153–170.

[10] Y. Cabannes. 2004. Participatory Budgeting: A Significant Contribution to Participatory Democracy. *Environment and Urbanization* 16, 1 (2004), 27–46.

[11] I. Caragiannis, A. D. Procaccia, and N. Shah. 2016. Truthful Univariate Estimators. In *Proceedings of the 33rd International Conference on Machine Learning (ICML)*. 127–135.

[12] R. T. Clemen. 1989. Combining forecasts: A review and annotated bibliography. *International Journal of Forecasting* 5 (1989), 559–583.

[13] V. Conitzer, M. Brill, and R. Freeman. 2015. Crowdsourcing Societal Tradeoffs. In *Proceedings of the 14th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. 1213–1217.

[14] V. Conitzer, R. Freeman, M. Brill, and Y. Li. 2016. Rules for Choosing Societal Tradeoffs.. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*. 460–467.

[15] E. Eisenberg and D. Gale. 1959. Consensus of Subjective Probabilities: The Pari-Mutuel Method. *The Annals of Mathematical Statistics* 30, 1 (1959), 165–168.

[16] B. Fain, A. Goel, and K. Munagala. 2016. The Core of the Participatory Budgeting Problem. In *Proceedings of the 12th Conference on Web and Internet Economics (WINE)*. 384–399.

[17] S. French. 1985. Group consensus probability distributions: A critical survey. *Bayesian Statistics* 2 (1985), 183–202.
[18] N. Garg, A. Goel, and B. Plaut. 2018. Markets for Public Decision-making. arXiv:1807.10836. (2018).
[19] N. Garg, V. Kamble, A. Goel, D. Marn, and K. Munagala. 2019. Iterative Local Voting for Collective Decision-making in Continuous Spaces. *Journal of Artificial Intelligence Research* 64 (2019), 315–355.
[20] C. Genest and J.V. Zidek. 1986. Combining Probability Distributions: A Critique and an Annotated Bibliography. *Statist. Sci.* 1, 1 (1986), 114–148.
[21] A. Goel, A. K. Krishnaswamy, S. Sakshuwong, and T. Aitamurto. 2016. Knapsack Voting: Voting mechanisms for Participatory Budgeting. Manuscript. (2016).
[22] A. Gulati, G. Shanmuganathan, X. Zhang, and P.J. Varman. 2012. Demand Based Hierarchical QoS Using Storage Resource Pools.. In *USENIX Annual Technical Conference*. 1–13.
[23] M.D. Intriligator. 1973. A probabilistic model of social choice. *The Review of Economic Studies* 40, 4 (1973), 553–560.
[24] T. Lu and C. Boutilier. 2011. Budgeted Social Choice: From Consensus to Personalized Decision Making. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*. 280–286.
[25] J. Massó and I. M. de Barreda. 2011. On strategy-proofness and symmetric single-peakedness. *Games and Economic Behavior* 72, 2 (2011), 467–484.
[26] H. Moulin. 1980. On strategy-proofness and single-peakedness. *Public Choice* 35 (1980), 437–455.
[27] Participatory Budgeting Project. 2019. Participatory Budgeting Project. https://www.participatorybudgeting.org/what-is-pb/. (2019). Accessed: 2019-02-12.
[28] H. Peters, H. van der Stel, and T. Storcken. 1992. Pareto optimality, anonymity, and strategy-proofness in location problems. *International Journal of Game Theory* 21, 3 (1992), 221–235.
[29] A.D. Procaccia and M. Tennenholtz. 2009. Approximate Mechanism Design Without Money. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*. Stanford, CA, USA, 177–186.
[30] A. Schneider and B. Goldfrank. 2002. Budgets and ballots in Brazil: Participatory budgeting from the city to the state. (2002). IDS Working Paper 149.
[31] W. D. Smith. 2000. Range Voting. http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.32.9150. (2000).
[32] V. V. Vazirani. 2007. Combinatorial Algorithms for Market Equilibria. In *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani (Eds.). Cambridge University Press, Chapter 5, 103–134.