Approximate Mechanism Design Without Money

Ariel D. Procaccia
Microsoft Israel R&D Center
arielpro@gmail.com

Moshe Tennenholtz
Microsoft Israel R&D Center and Technion, IIT
moshet@microsoft.com

ABSTRACT

The literature on algorithmic mechanism design is mostly concerned with game-theoretic versions of optimization problems to which standard economic money-based mechanisms cannot be applied efficiently. Recent years have seen the design of various truthful approximation mechanisms that rely on enforcing payments. In this paper, we advocate the reconsideration of highly structured optimization problems in the context of mechanism design. We argue that, in such domains, *approximation can be leveraged to obtain truthfulness without resorting to payments*. This stands in contrast to previous work where payments are ubiquitous, and (more often than not) approximation is a necessary evil that is required to circumvent computational complexity.

We present a case study in approximate mechanism design without money. In our basic setting agents are located on the real line and the mechanism must select the location of a public facility; the cost of an agent is its distance to the facility. We establish tight upper and lower bounds for the approximation ratio given by strategyproof mechanisms without payments, with respect to both deterministic and randomized mechanisms, under two objective functions: the social cost, and the maximum cost. We then extend our results in two natural directions: a domain where two facilities must be located, and a domain where each agent controls multiple locations.

Categories and Subject Descriptors

F.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms

Algorithms, Theory, Economics

Keywords

Game theory, Algorithmic mechanism design, Social choice

1. INTRODUCTION

The vibrant field of algorithmic mechanism design, which originated in the work of Nisan and Ronen [25], deals with game-theoretic versions of (often Internet-related) optimization problems such as task scheduling and resource allocation. In these settings the problem input is distributed among selfish agents; the agents might lie about their private information if this serves their own ends, resulting in a deterioration in the quality of the outcome. A *mechanism* is a function that selects an outcome, and possibly also a payment scheme, given the reported types of the agents. The goal is then to design mechanisms that encourage truthfulness while optimizing an objective function.

It has been observed [8] that there are two major classes of problems in algorithmic mechanism design. The first class contains problems for which there exist optimal truthful mechanisms, but the problem is computationally intractable. Typical examples include the line of work on combinatorial auctions (see, e.g., [17, 15, 11]), where the objective function is usually the maximization of the social welfare, that is, the sum of agents' utilities. For this objective function a truthful optimal mechanism is given by the (now) well known Vickrey-Clarke-Groves (VCG) mechanism (see, e.g., [24]). VCG uses payments in order to align the interests of individual agents with the interests of society. Unfortunately, it turns out that an approximation of the social welfare is insufficient to guarantee truthfulness using VCG. Therefore, researchers have focused on designing *truthful yet efficient approximation mechanisms*; by approximation we refer to the standard multiplicative sense, that is, an α-approximation mechanism always returns a solution that is within an α-factor of the optimal solution. In other words, researchers circumvent the computational hardness by resorting to approximation, and at the same time enforce tailor-made payments to guarantee truthfulness. Papers about scheduling on related machines (see, e.g., [2, 1, 10]) also fall into the first class, although in the scheduling domain the objective is usually to minimize the makespan.

The second (significantly smaller) class of problems involves optimization problems which are not necessarily intractable, but for which there is no optimal truthful mechanism. The prominent problem in this class is scheduling on unrelated machines (see, e.g., [25, 16, 8]). In such domains one might investigate the optimal approximation ratio achievable by any truthful mechanism, regardless of computational feasibility.

The assumption underlying essentially all[1] previous work

---

[1] There are three exceptions, discussed in the sequel.

on truthful approximation mechanisms is the existence of money, or, in other words, the ability to make payments. This assumption is explicit in Nisan and Ronen's very definition of mechanism [25], but is easily challenged when it comes to computational settings. In particular, in Internet domains payments are notoriously difficult to implement, mainly due to security and banking issues. Moreover, Schummer and Vohra [31] note that "there are many important environments where money cannot be used as a medium of compensation", due to ethical considerations (for instance, in political decision making) or legal considerations (e.g., in the context of organ donations). It is therefore natural to ask whether it is possible to design truthful mechanisms without payments; such mechanisms are known as *strategyproof* in the social choice literature.

Our Agenda, or: What is Approximate Mechanism Design Without Money? We consider game-theoretic optimization problems where returning the optimal solution is *not* strategyproof. Our main conceptual contribution is the explicit suggestion that *approximation can be used to obtain strategyproofness without resorting to payments*; In other words, we propose achieving strategyproofness, without using money, by sacrificing the optimality of the solution. In essence, this agenda is reminiscent of the second class of problems discussed above, in the sense that approximation is seen to enable truthfulness rather than hinder it. However, our rejection of money stands in contrast to the existing work in algorithmic mechanism design, where payments are ubiquitous.

The contrast with previous work becomes even more striking when one considers (as we do in this paper) *computationally tractable* optimization problems where *there is an optimal, computationally efficient, truthful, payment-based mechanism*, but *there is no optimal truthful mechanism without money*. Crucially, this type of problems does not fall into either of the two classes mentioned above. We therefore have a new class of problems that has previously been disregarded, and, we suggest, should be reconsidered.

Importantly, our agenda only applies to optimization problems where there exist reasonable strategyproof mechanisms without payments. In particular, we must escape social choice impossibility results such as the well known Gibbard-Satterthwaite Theorem [14, 29] and its variations, e.g., the important paper of Barberà and Peleg regarding continuous preferences [6]. Hence, we consider highly structured domains where these results do not hold.

Our Results. This paper presents a case study in approximate mechanism design without money. In the basic domain that we study, each agent i has a *location* $x_i \in \mathbb{R}$. Given the locations of all the agents, a mechanism selects the location $y \in \mathbb{R}$ of a facility. The *cost* of agent i is simply the distance $|y - x_i|$. For example, x_i might be the location of the house of agent i on a street, and y might be the location of a grocery store or a public library. This type of preferences (more accurately, a slight generalization thereof) is known as *single peaked*. Single peaked preferences and their extensions have been extensively studied in the social choice literature, starting with the work of Moulin [23]; see the surveys by Barberà [5] and Sprumont [32], and the references therein. We use the terminology of facility location problems, but the facility is simply an abstraction of a public good, and the same domain can also be (and has been) used to represent

political policies, economic decisions, locating mirrors in a network, etc. Under many of these interpretations payments may be infeasible, for the reasons discussed above.

We study the foregoing, basic setting in Section 2. We observe that choosing the median location is a *group strategyproof* (i.e., even coalitions of agents cannot gain by lying) mechanism that minimizes the social cost, that is, the sum of the agents' costs. However, if the goal is to minimize the maximum cost, selecting the optimal facility location—the average of the leftmost and rightmost locations—is no longer strategyproof. With respect to this objective function, we give a deterministic group strategyproof mechanism (without money) that yields an approximation ratio of 2, and provide a matching lower bound that holds even against (individually) strategyproof deterministic mechanisms (without money). Further, we give a group strategyproof randomized mechanism with an approximation ratio of $3/2$, and provide a matching strategyproof lower bound. These results are summarized in Table 1.

We subsequently study two natural extensions of the basic setting. In both settings, the optimal solution is not strategyproof even with respect to the social cost, and we resort to strategyproof approximation mechanisms, some straightforward and some nontrivial. Section 3 deals with a setting where two facilities must be located; the cost of an agent is its distance to the nearest one. Our main result of Section 3 is a randomized strategyproof $5/3$-approximation mechanism for the maximum cost objective function. This result is notable since the mechanism (Mechanism 2) incorporates several new ideas in order to achieve strategyproofness, and, unlike other mechanisms, the difficult part of its analysis (Theorem 3.5) is the proof of strategyproofness.

Section 4 is concerned with a setting where only one facility must be located, but each agent is associated with multiple locations (e.g., a real estate agent), and is interested in optimizing the objective function with respect to its own multiset of locations (whereas the designer is interested in optimizing over the entire multiset of locations). In this section, our main results are a randomized strategyproof mechanism that yields a 2-approximation to the social cost when there are two agents that control the same number of locations, and a randomized group strategyproof mechanism that has a tight approximation ratio of $3/2$ for the maximum cost. Due to space constraints and the sheer number of results we do not list them all here, but rather refer the impatient reader to Tables 2 and 3.

Related Work. The origins of the agenda of approximate mechanism design without money can be traced to the paper of Dekel et al. [9] on incentive compatible learning, a line of work that was followed up in recent papers [20, 21]. It turns out that the study of incentives in general learning-theoretic domains reduces to simpler settings where strategyproof approximation mechanisms without money can be designed. There are some mathematical connections between our work and that of Dekel et al. [9], on which we elaborate in Section 4. One of the main contributions of this paper is that we properly crystallize and explicitly advocate approximate mechanism design without money.

Our agenda is reminiscent of the line of work on the frugality of mechanisms (see, e.g., [3, 12] in the context of buying an *s-t* path). This body of research deals with designing truthful mechanisms that have to pay as little as possible. One way to see our work is as taking the concept of frugal-

ity to the limit by requiring that the mechanism make no payments.

The domain that we study in Section 3, in which two facilities must be located on the real line, was previously studied by Miyagawa [22]. He gave an interesting characterization of strategyproof, Pareto-optimal, and continuous mechanisms in this setting, where the last property informally means that the facility moves continuously with the locations of the agents. Unfortunately, continuity is incompatible with approximation, hence we cannot technically utilize this result.

Incentives aside and taking an algorithmic point of view, the problems that we deal with are the one-dimensional Euclidean k-median and k-center problems, when the objective functions are the social cost and the maximum cost, respectively, and $k = 1$ (Section 2 and 4) or $k = 2$ (Section 3). This may sound discouraging, but recall that we deliberately focus on relatively simple, structured problems, and the domains that we deal with are extremely well-studied in the social choice literature. The k-median and k-center problems were extensively investigated, especially in the context of clustering, and can be approximated using sophisticated algorithms (see, e.g., [7, 4]).

Further afield, there is a body of work that deals with mechanism design without money, but so far this was pursued by relatively few computer scientists; for a survey see the book chapter by Schummer and Vohra [31]. A prominent example is the work on strategyproof mechanisms for stable matchings [13, 28]. There are a few papers that deal with the game-theoretic properties of existing Internet mechanisms that do not require payments, e.g., the recent work on interdomain routing by Levin et al. [19]. Finally, our work is remotely related to work on strategyproof mechanisms for cost sharing in facility location problems (see, e.g., [18, 26]).

2. THE BASIC SETTING

Let $N = \{1, \ldots, n\}$ be a set of *agents*. Each agent $i \in N$ has a location $x_i \in \mathbb{R}$. We refer to the collection $\mathbf{x} = \langle x_1, \ldots, x_n \rangle$ as the *location profile*.[2]

A (deterministic) *mechanism* in this simple setting is a function $f : \mathbb{R}^n \to \mathbb{R}$, that is, a function that maps a given location profile to a location of a *facility*. If the facility is located at y, the *cost* of agent $i \in N$ is $\text{cost}(y, x_i) = |y - x_i|$.

A randomized mechanism is a function f from $\mathbb{R}^n$ to probability distributions over $\mathbb{R}$. In other words, a randomized mechanism allows us to randomly specify the location of the facility for every given location profile. If $f(\mathbf{x}) = P$, where P is a probability distribution, the cost of agent $i \in N$ is defined as the expected distance from the location of i, i.e., $\text{cost}(P, x_i) = \mathbb{E}_{y \sim P} |y - x_i|$.

A mechanism f is *strategyproof* if an agent can never benefit from reporting a false location, regardless of the strategies of the other agents. In the current setting, this means that for all $\mathbf{x} \in \mathbb{R}^n$, for all $i \in N$, and for all $x_i' \in \mathbb{R}$, $\text{cost}(f(\mathbf{x}), x_i) \leq \text{cost}(f(x_i', \mathbf{x}_{-i}), x_i)$, where $\mathbf{x}_{-i} = \langle x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n \rangle$ is the vector of the locations of all agents in $N \setminus \{i\}$.

A mechanism is *group strategyproof* if for any location profile $\mathbf{x}$ and any coalition $S \subseteq N$, there is no joint deviation x_S' of the agents in S such that all the agents in S gain, that

is, for all $\mathbf{x} \in \mathbb{R}^n$, for all $S \subseteq N$, and for all $\mathbf{x}_S \in \mathbb{R}^{|S|}$, there exists $i \in S$ such that $\text{cost}(f(\mathbf{x}), x_i) \leq \text{cost}(f(\mathbf{x}_S', \mathbf{x}_{-S}), x_i)$. Notice that it is possible to define (strong) group strategyproofness by asking that it cannot be the case that all the deviating agents do not lose and at least one gains. Some of our group strategyproofness results do not hold under this stronger definition. However, our (weaker) notion of strategyproofness is very common in social choice, since in settings where payments (and in particular, side payments) cannot be made, an agent that does not strictly gain has no incentive to become a member of the deviating coalition.

In this paper, we shall be interested in strategyproof mechanisms that also do well with respect to optimizing one of two objective functions: minimizing the social cost, or minimizing the maximum cost.

The *social cost* of a facility location $y \in \mathbb{R}^n$ with respect to the profile $\mathbf{x} \in \mathbb{R}^n$ is $\text{sc}(y, \mathbf{x}) = \sum_{i \in N} \text{cost}(y, x_i)$; the social cost of a distribution P with respect to $\mathbf{x}$ is $\text{sc}(P, \mathbf{x}) = \mathbb{E}_{y \sim P}[\text{sc}(y, \mathbf{x})]$. The *maximum cost* of a y with respect to $\mathbf{x}$ is $\text{mc}(y, \mathbf{x}) = \max_{i \in N} \text{cost}(y, x_i)$, whereas the maximum cost of P with respect to $\mathbf{x}$ is $\text{mc}(P, \mathbf{x}) = \mathbb{E}_{y \sim P}[\text{mc}(y, \mathbf{x})]$.

2.1 Social Cost

We warm up by tackling an easy question: is there a strategyproof mechanism that minimizes the social cost? The solution is very simple: choose the median location in $\mathbf{x}$, which we shall denote by $\text{med}(\mathbf{x})$. Indeed, assume that n is odd, $n = 2k + 1$. Any point that is to the left of the median has higher social cost than that of the median since it is further away from at least $k + 1$ locations and closer to at most k locations, and the same holds for any point to the right of the median. If n is even, $n = 2k$, and without loss of generality $x_1 \leq x_2 \leq \cdots \leq x_n$, then any point in the interval $[x_k, x_{k+1}]$ is an optimal facility location. In this case, when we refer to the median $\text{med}(\mathbf{x})$ we mean the leftmost point of the optimal interval, i.e., the kth order statistic.

As noted in Section 1, the structure of the preferences of our agents is known in the social choice literature as *single peaked*: the peak, or bliss point, of agent i is at x_i, and the closer a location is to x_i, the more preferred it is. It has long been known that, when agents have single peaked preferences, the selection of the kth order statistic for some $k \in \{1, \ldots, n\}$ is group strategyproof [23]; this is also very easy to verify. In particular, selecting the median peak is group strategyproof. Hence, in our basic setting, the social cost can in fact be minimized using a group strategyproof mechanism.

2.2 Maximum Cost

The second objective function that we consider is minimizing the maximum cost. Here the situation becomes nontrivial, even in the basic setting presented above. We will first investigate deterministic mechanisms, and then turn our attention to randomized mechanisms.

Deterministic Mechanisms. For a location profile $\mathbf{x} \in \mathbb{R}^n$, denote the leftmost location in $\mathbf{x}$ by $\text{lt}(\mathbf{x}) = \min_{i \in N} x_i$, and the rightmost location by $\text{rt}(\mathbf{x}) = \max_{i \in N} x_i$. Furthermore, denote the center of the interval $[\text{lt}(\mathbf{x}), \text{rt}(\mathbf{x})]$ by $\text{cen}(\mathbf{x}) = (\text{lt}(\mathbf{x}) + \text{rt}(\mathbf{x}))/2$. Given $\mathbf{x}$, the solution that minimizes the maximum cost is $\text{cen}(\mathbf{x})$. Unfortunately, this solution is not (even individually) strategyproof. Indeed, if $N = \{1, 2\}$, $x_1 = 0$ and $x_2 = 1$, agent 2 can move the optimal solution to its own location by reporting $x_2' = 2$.

---

[2]Some works on single peaked preferences restrict the locations to an interval; our results hold in that model as well.

A trivial, group strategyproof solution would be to choose any kth order statistic for some $k \in \{1, \ldots, n\}$. For reasons that will become apparent in the sequel, we choose the first order statistic, i.e., $\mathrm{lt}(\mathbf{x})$. Notice that any point in the interval $[\mathrm{lt}(\mathbf{x}), \mathrm{rt}(\mathbf{x})]$ would give a 2-approximation to the maximum cost. We have therefore obtained the following straightforward result.

THEOREM 2.1. $f(\mathbf{x}) = \mathrm{lt}(\mathbf{x})$ *is a group strategyproof 2-approximation mechanism for the maximum cost.*

Given the simplicity of our group strategyproof mechanism, it may be somewhat surprising that no (even individually) strategyproof mechanism can do better, as the following theorem asserts.

THEOREM 2.2. *Let* $N = \{1, \ldots, n\}$, $n \geq 2$. *Any deterministic strategyproof mechanism* $f : \mathbb{R}^n \to \mathbb{R}$ *has an approximation ratio of at least 2 for the maximum cost.*

PROOF. We first deal with the case where $N = \{1, 2\}$, and subsequently touch on extending the proof to an arbitrary n.

Assume for contradiction that $f : \mathbb{R}^n \to \mathbb{R}$ is a strategyproof mechanism and has an approximation ratio smaller than 2 for the maximum cost. Consider the location profile $\mathbf{x}$ where $x_1 = 0$ and $x_2 = 1$. Assume without loss of generality that $f(\mathbf{x}) = 1/2 + \epsilon$, $\epsilon \geq 0$. Now, consider the profile where $x_1 = 0$ and $x_2' = 1/2 + \epsilon$. The optimum is the average of the two locations, namely $1/4 + \epsilon/2$, which has a maximum cost of $1/4 + \epsilon/2$. If the mechanism is to achieve an approximation ratio better than 2, the facility must be placed in $(0, 1/2 + \epsilon)$. In that case, given the profile $\mathbf{x}$, agent 2 can benefit by reporting $x_2' = 1$, thus moving the solution to $1/2 + \epsilon$, in contradiction to strategyproofness.

In order to extend this result to an arbitrary n, simply locate all the agents $N \setminus \{1, 2\}$ at $1/2$ in each one of the profiles described above. All the arguments given above go through smoothly. $\square$

Randomized Mechanisms. We presently turn to randomized mechanisms; we shall demonstrate that randomization allows us to break the deterministic lower bound of 2, given by Theorem 2.2. Indeed, we focus on the following mechanism.

MECHANISM 1. *Given* $\mathbf{x}$, *return* $lt(\mathbf{x})$ *with probability* $1/4$, $rt(\mathbf{x})$ *with probability* $1/4$, *and* $cen(\mathbf{x})$ *with probability* $1/2$.

It is possible to demonstrate that Mechanism 1 is group strategyproof. Moreover, the mechanism gives an approximation ratio of $3/2$, well below the deterministic lower bound.

THEOREM 2.3. *Mechanism 1 is a group strategyproof $3/2$-approximation mechanism for the maximum cost.*

The proof of Theorem 2.3 is based on the observation that if the interval over which the mechanism randomized contracts, then the agents at the boundaries must be members of the deviating coalition.

PROOF OF THEOREM 2.3. By scaling the distances, we can assume without loss of generality that $\mathrm{lt}(\mathbf{x}) = 0$ and $\mathrm{rt}(\mathbf{x}) = 1$. We shall first prove the claim about the approximation ratio.

The optimum cost is $1/2$, whereas the expected cost of the algorithm is

$$\frac{1}{4} \cdot 1 + \frac{1}{4} \cdot 1 + \frac{1}{2} \cdot \frac{1}{2} = \frac{3}{4} \quad .$$

The approximation ratio is therefore $3/2$.

We now turn to proving group strategyproofness. Let $S \subseteq N$ be a coalition. We must demonstrate that the agents in S cannot all gain by deviating.

A crucial observation is that, given $\mathbf{x} \in \mathbb{R}^n$, the only deviations that affect the outcome of the mechanism are the ones that modify the locations of the extreme agents $\mathrm{lt}(\mathbf{x})$ and $\mathrm{rt}(\mathbf{x})$. The location of $\mathrm{lt}(\mathbf{x})$ can always be pushed to the left and the location of $\mathrm{rt}(\mathbf{x})$ can always be pushed to the right. However, $\mathrm{lt}(\mathbf{x})$ can be pushed to the right only if the leftmost agent is a member of the deviating coalition S, that is, $\mathrm{argmin}_{i \in N} x_i \cap S \neq \emptyset$. Similarly, $\mathrm{rt}(\mathbf{x})$ can be pushed to the right only if the rightmost agent is a member of S.

Let $\mathbf{x} \in \mathbb{R}^n$, and let $\mathbf{x}' \in \mathbb{R}^n$ where, for every $i \neq S$, $x_i' = x_i$. Further, let $\Delta_1 = \mathrm{lt}(\mathbf{x}) - \mathrm{lt}(\mathbf{x}')$, and $\Delta_2 = \mathrm{rt}(\mathbf{x}') - \mathrm{rt}(\mathbf{x})$. We consider four cases.

Case 1: $\Delta_1 \geq 0$ *and* $\Delta_2 \geq 0$. Let $i \in S$; clearly $x_i \in [\mathrm{lt}(\mathbf{x}), \mathrm{rt}(\mathbf{x})]$. Denoting Mechanism 1 by f, we have:

$$\mathrm{cost}(f(\mathbf{x}'), x_i) = \frac{1}{4} \cdot (x_i - \mathrm{lt}(\mathbf{x}) + \Delta_1) + \frac{1}{4} \cdot (\mathrm{rt}(\mathbf{x}) - x_i + \Delta_2)$$
$$+ \frac{1}{2} \cdot \left| x_i - \frac{\mathrm{lt}(\mathbf{x}) - \Delta_1 + \mathrm{rt}(\mathbf{x}) + \Delta_2}{2} \right|$$
$$\geq \frac{1}{4} \cdot (x_i - \mathrm{lt}(\mathbf{x})) + \frac{1}{4} \cdot (\mathrm{rt}(\mathbf{x}) - x_i)$$
$$+ \frac{1}{2} \cdot \left| x_i - \frac{\mathrm{lt}(\mathbf{x}) + \mathrm{rt}(\mathbf{x})}{2} \right|$$
$$= \mathrm{cost}(f(\mathbf{x}), x_i) \quad .$$

Case 2: $\Delta_1 < 0$ *and* $\Delta_2 \geq 0$. In this case, it must be true that the leftmost agent, which is located at 0, is a member of S. It is obvious that this agent cannot benefit from the deviation, and in fact must strictly lose, since the leftmost point, the center, and possibly the rightmost point are all moving further away from the agent's location at 0.

Case 3: $\Delta_1 \geq 0$ *and* $\Delta_2 < 0$. The case is symmetric to Case 2.

Case 4: $\Delta_1 < 0$ *and* $\Delta_2 < 0$. In this case, the leftmost agent, located at 0, and the rightmost agent, located at 1, must both be members of S. We shall demonstrate that they cannot both gain from the deviation.

$$\mathrm{cost}(f(\mathbf{x}'), 0) = \frac{1}{4} \cdot \Delta_1 + \frac{1}{4} \cdot (1 - \Delta_2) + \frac{1}{2} \cdot \frac{\Delta_1 + 1 - \Delta_2}{2}$$
$$= \mathrm{cost}(f(\mathbf{x}), 0) + \frac{\Delta_1 - \Delta_2}{2} \quad .$$

Similarly,

$$\mathrm{cost}(f(\mathbf{x}'), 1) = \mathrm{cost}(f(\mathbf{x}), 1) + \frac{\Delta_2 - \Delta_1}{2} \quad .$$

We conclude that

$$\mathrm{cost}(f(\mathbf{x}'), 0) + \mathrm{cost}(f(\mathbf{x}'), 1) = \mathrm{cost}(f(\mathbf{x}), 0) + \mathrm{cost}(f(\mathbf{x}), 1) \quad ,$$

and hence it holds that either $\mathrm{cost}(f(\mathbf{x}'), 0) \geq \mathrm{cost}(f(\mathbf{x}), 0)$ or $\mathrm{cost}(f(\mathbf{x}'), 1) \geq \mathrm{cost}(f(\mathbf{x}), 1)$. $\square$

While the theorem implies that randomization allows us to drop the feasible strategyproof approximation ratio from

Objective Function	Deterministic	Randomized
Social Cost	UB: 1 GSP LB: 1 SP	
Maximum Cost	UB: 2 GSP (Thm 2.1) LB: 2 SP (Thm 2.2)	UB: 3/2 GSP (Thm 2.3) LB: 3/2 SP (Thm 2.4)

Table 1: A summary of the results of Section 2. UB and LB stand for upper bound and lower bound, respectively. SP and GSP stand for strategyproof and group strategyproof, respectively.

2 to 3/2, we can also show that this is as far as randomization can take us.

THEOREM 2.4. *Let $N = \{1, \dots, n\}$, $n \geq 2$. Any randomized strategyproof mechanism has an approximation ratio of at least 3/2 for the maximum cost.*

In order to prove the theorem, we require two straightforward lemmata.

LEMMA 2.5. *Let $N = \{1, 2\}$, and let $\mathbf{x} \in \mathbb{R}^2$. Let P be a probability distribution over $\mathbb{R}$ such that*

$$\mathbb{E}_{y \sim P}\left[\left|y - \frac{x_1 + x_2}{2}\right|\right] = \Delta \quad.$$

Then the expected maximum cost is

$$\Delta + \frac{|x_1 - x_2|}{2} \quad.$$

PROOF. For every $y \in \mathbb{R}$, we have that the maximum cost is $\left|y - \frac{x_1 + x_2}{2}\right| + \frac{|x_1 - x_2|}{2}$. Therefore, the expected maximum cost is

$$\mathbb{E}_{y \sim P}\left[\left|y - \frac{x_1 + x_2}{2}\right| + \frac{|x_1 - x_2|}{2}\right] = \Delta + \frac{|x_1 - x_2|}{2} \quad.$$
□

LEMMA 2.6. *Let $N = \{1, 2\}$, and let $x_1, x_2 \in \mathbb{R}$. Let P be a probability distribution over $\mathbb{R}$. Then there exists $i \in N$ such that*

$$\mathbb{E}_{y \sim P}[|y - x_i|] \geq \frac{|x_1 - x_2|}{2} \quad.$$

PROOF. Let Y be a random variable distributed according to P, and let X_1 and X_2 be random variables defined by $X_1 = |Y - x_1|$, $X_2 = |Y - x_2|$. Then

$$\mathbb{E}[X_1] + \mathbb{E}[X_2] = \mathbb{E}[X_1 + X_2] \geq |x_1 - x_2| \quad.$$

The lemma directly follows. □

We are now ready to prove the theorem.

PROOF OF THEOREM 2.4. As in the proof of Theorem 2.2, we first deal with the case $N = \{1, 2\}$, and then extend the proof to more agents.

Let f be a randomized mechanism. Consider the profile $\mathbf{x} \in \mathbb{R}^2$ where $x_1 = 0$ and $x_2 = 1$. We have that $f(\mathbf{x}) = P$, where P is a distribution over $\mathbb{R}$. By Lemma 2.6, there exists $i \in N$, without loss of generality x_2, such that $\text{cost}(P, x_2) \geq 1/2$.

Now, consider the profile where $x_1 = 0$, $x_2' = 2$. By strategyproofness, the expected distance from 1 must still be at least 1/2, otherwise agent 2 gains from deviating from x_2 to x_2'. By Lemma 2.5 (with $\Delta = 1/2$), the expected maximum cost is therefore at least 3/2, whereas the optimum has a

cost of 1; it follows that the approximation ratio of f is at least 3/2.

In order to extend the proof to an arbitrary number of agents n, we simply set the locations of the additional agents to be 1/2; the proof works as before. □

2.3 Discussion

Table 1 summarizes the results of Section 2. Our results in this section are completely tight. As we move on to significantly more involved settings, obtaining tight bounds inevitably becomes much more difficult.

Interestingly, if payments are allowed, it is possible to obtain a truthful optimal solution even for the maximum cost, by using VCG-like payments: each agent $i \in N$ pays the distance between the optimal facility location when $\mathbf{x}$ is reported and the optimal facility location when $\mathbf{x}_{-i}$ is reported.

3. EXTENSION I: TWO FACILITIES

In this section we investigate a first natural extension to the setting examined in Section 2: locating two facilities instead of just one. A deterministic mechanism is now a function $f : \mathbb{R}^n \to \mathbb{R}^2$, that is, the mechanism returns the locations $\mathbf{y} \in \mathbb{R}^2$ of both facilities given a location profile. If $\mathbf{y} = \langle y_1, y_2 \rangle$, the cost of an agent is its distance to the nearest facility: $\text{cost}(\mathbf{y}, x_i) = \min\{|y_1 - x_i|, |y_2 - x_i|\}$. We usually assume that $y_1 \leq y_2$.

Similarly, a randomized mechanism returns a probability distribution P over $\mathbb{R}^2$, and the cost of an agent is its expected distance to the nearest facility. We redefine $\text{sc}(\mathbf{y}, \mathbf{x})$ and $\text{mc}(\mathbf{y}, \mathbf{x})$ in the obvious way according to the new definition of cost given above.

3.1 Social Cost

As before, we shall first look into minimizing the social cost in a strategyproof way. Let us first consider the algorithmic problem of locating two facilities in a way that minimizes the social cost, disregarding incentives. This problem is quite simple, although this may not be immediately apparent. Indeed, given a location profile $\mathbf{x} \in \mathbb{R}^n$, let the optimal facility locations be $y_1, y_2 \in \mathbb{R}$, $y_1 \leq y_2$. Informally, we can associate with y_1 a multiset of locations $L(\mathbf{x}) \subsetneq \{x_1, \dots, x_n\}$ (for "left") whose cost is computed with respect to y_1, and similarly associate with y_2 a multiset of locations $R(\mathbf{x}) \subsetneq \{x_1, \dots, x_n\}$ (for "right") whose cost is computed with respect to y_2, such that for all $x_i \in L$, $x_j \in R$, $x_i \leq x_j$. Now, y_1 is the median of $L(\mathbf{x})$ and y_2 is the median of $R(\mathbf{x})$. Hence, it is sufficient to optimize over the $n - 1$ possible choices of $L(\mathbf{x})$ and $R(\mathbf{x})$.

Despite the algorithmic simplicity of the problem, and in contrast to the single facility setting, minimizing the social cost in the two facility setting is not strategyproof. Intu-

itively, the reason is that it is impossible to elicit the structure of L and R in a strategyproof way. The next theorem in fact establishes a lower bound of $3/2 - \mathcal{O}(1/n)$.

THEOREM 3.1. *Let $N = \{1, \ldots, n\}$, $n \geq 3$. In the two facility setting, any deterministic strategyproof mechanism $f : \mathbb{R}^n \to \mathbb{R}^2$ has an approximation ratio of at least $3/2 - \mathcal{O}(1/n)$ for the maximum cost.*

PROOF. Let $n \geq 3$. We construct a location profile $\mathbf{x} \in \mathbb{R}^n$ as follows: $x_1 = -1$, $x_2 = 1$, and $x_i = 0$ for all $i \in N \setminus \{1, 2\}$. The optimal solution has a social cost of 1. Let f be a mechanism, and let $f(\mathbf{x}) = \langle y_1, y_2 \rangle \in \mathbb{R}^2$. If $|y_1| \geq \frac{2}{n-2}$, and $|y_2| \geq \frac{2}{n-2}$, then $\mathrm{sc}(f(\mathbf{x}), \mathbf{x}) \geq 2$, hence the mechanism's approximation ratio is at least 2.

By the above, we can assume without loss of generality that $|y_1| \leq \frac{2}{n-2}$. Furthermore, assume without loss of generality that $y_2 \leq 0$. We consider a deviation of agent 2 to $x_2' = 3/2$. Let $f(x_2', \mathbf{x}_{-2}) = \langle y_1', y_2' \rangle$. The optimal solution $\langle 0, 3/2 \rangle$ has a social cost of 1, therefore we can assume once again that $|y_1'| \leq \frac{2}{n-2}$. In addition, $\mathrm{cost}(\langle y_1, y_2 \rangle, x_2) \geq 1 - \frac{2}{n-2}$, hence by strategyproofness we have that $|y_2' - x_2| \geq 1 - \frac{2}{n-2}$. It follows that either $y_2' \geq 2 - \frac{2}{n-2}$, or $y_2' \leq \frac{2}{n-2}$. In both cases, we get that

$$\mathrm{sc}(\langle y_1', y_2' \rangle, \langle x_2', \mathbf{x}_{-2} \rangle) \geq \frac{3}{2} - \frac{2}{n-2} = \frac{3}{2} - \mathcal{O}\left(\frac{1}{n}\right) \quad ,$$

hence the approximation ratio is at least $3/2 - \mathcal{O}(1/n)$. $\square$

It can be verified that a group strategy proof $(n-1)$-approximation mechanism is given by choosing $\mathrm{lt}(\mathbf{x})$ and $\mathrm{rt}(\mathbf{x})$ given the location profile $\mathbf{x} \in \mathbb{R}^n$. In brief, the reason is that $\mathrm{lt}(\mathbf{x}) \in L(\mathbf{x})$ and $\mathrm{rt}(\mathbf{x}) \in R(\mathbf{x})$. The gap between this result and the lower bound given by Theorem 3.1 is still huge, and remains our most enigmatic open problem.

3.2 Maximum Cost

Let us now turn to strategyproof mechanisms that approximate the maximum cost. Similarly to the social cost objective, the problem of locating two facilities in a way that minimizes the maximum cost is computationally straightforward. Moreover, we can give a very accurate characterization of the structure of the optimal solution. We shall first require some notations.

Given $\mathbf{x} \in \mathbb{R}^n$, let the *left boundary location* be $\mathrm{lb}(\mathbf{x}) = \max\{x_i : i \in N, x_i \leq \mathrm{cen}(\mathbf{x})\}$, and the *right boundary location* be $\mathrm{rb}(\mathbf{x}) = \min\{x_i : i \in N, x_i \geq \mathrm{cen}(\mathbf{x})\}$. Now, denote $\mathrm{dist}(\mathbf{x}) = \max\{\mathrm{lb}(\mathbf{x}) - \mathrm{lt}(\mathbf{x}), \mathrm{rt}(\mathbf{x}) - \mathrm{rb}(\mathbf{x})\}$. The following lemma is the foundation of the positive results in this subsection.

LEMMA 3.2. *Given $\mathbf{x} \in \mathbb{R}^n$, the optimal placement of two facilities has a maximum cost of $\mathrm{dist}(\mathbf{x})/2$.*

PROOF. As usual, we can assume without loss of generality (by scaling the distances) that $\mathrm{lt}(\mathbf{x}) = 0$, $\mathrm{lt}(\mathbf{x}) = 1$; further, assume without loss of generality that $\mathrm{lb}(\mathbf{x}) \geq 1 - \mathrm{rb}(\mathbf{x})$, that is, $\mathrm{dist}(\mathbf{x})$ is defined by $\mathrm{lb}(\mathbf{x})$. We shall first show that there is a solution with the announced cost. Indeed, let $\mathbf{y}^*$ with $y_1^* = \mathrm{lb}(\mathbf{x})/2$, $y_2^* = (\mathrm{rb}(\mathbf{x}) + 1)/2$. It holds that $\mathrm{mc}(\mathbf{y}^*, \mathbf{x}) = \mathrm{lb}(\mathbf{x})/2 \leq 1/4$.

We argue that any solution must have a cost of at least $\mathrm{lb}(\mathbf{x})/2$. Indeed, consider first a solution $\mathbf{y}$ where $y_1 \leq 1/2$ and $y_2 \leq 1/2$, or $y_1 \geq 1/2$ and $y_2 \geq 1/2$; then $\mathrm{mc}(\mathbf{y}, \mathbf{x}) \geq$

$1/2$, making this solution inferior to $\mathbf{y}^*$. Now, Given that the solution only locates one facility y_1 to the left of $1/2$, we can assume that $y_2 \geq 3/4$, otherwise the cost is at least $1/4$. Any location such that $|y_1 - \mathrm{lb}(\mathbf{x})/2| = \epsilon > 0$ has a cost of at least $\mathrm{lb}(\mathbf{x})/2 + \epsilon$, incurred by its distance to either 0 or $\mathrm{lb}(\mathbf{x})$. We conclude that the maximum cost is at least $\mathrm{lb}(\mathbf{x})/2$. $\square$

Deterministic Mechanisms. Given our experience with the single facility case and Lemma 3.2, it is quite straightforward to obtain a 2-approximation, group strategyproof, deterministic mechanism: given $\mathbf{x} \in \mathbb{R}^n$, simply select the leftmost location $\mathrm{lt}(\mathbf{x})$ and the rightmost location $\mathrm{rt}(\mathbf{x})$. Indeed, the maximum cost of our solution is $\mathrm{dist}(\mathbf{x})$, whereas the maximum cost of the optimal solution, by Lemma 3.2, is $\mathrm{dist}(\mathbf{x})/2$. We have obtained the following theorem.

THEOREM 3.3. *$f(\mathbf{x}) = \langle \mathit{lt}(\mathbf{x}), \mathit{rt}(\mathbf{x}) \rangle$ is a group strategyproof 2-approximation mechanism for the maximum cost in the two facility setting.*

As for lower bounds, notice that, when there are two agents it is possible to obtain an optimal strategyproof solution by selecting the locations of the two agents. When $n \geq 3$, however, we can use a variation on the construction in the proof of Theorem 2.2.

COROLLARY 3.4. *Let $N = \{1, \ldots, n\}$, $n \geq 3$. Any deterministic strategyproof mechanism $f : \mathbb{R}^n \to \mathbb{R}^2$ has an approximation ratio of at least 2 for the maximum cost in the two facility setting.*

PROOF SKETCH. Use the same construction as in the proof of Theorem 2.2 for $n - 1$ agents, and locate an additional agent at, say, 10 in all the location profiles used in the proof. In order to get a 2-approximation, one of the two facilities must always be close to 10, whereas the same arguments as before apply to the second facility and the rest of the agents. $\square$

Randomized Mechanisms. Above we saw that, with respect to deterministic mechanisms, the results from Section 2 carry over quite smoothly to the two facility setting. This is no longer true with respect to randomized mechanisms, for a variety of reasons.

We consider the following mechanism. It is inspired by Mechanism 1, but requires several additional new ideas: randomizing over two equal intervals, unbalanced weights at the edges, and correlation between the two facilities. These "tricks" play a crucial role in satisfying the delicate strategyproofness constraints associated with the two facility setting.

MECHANISM 2. *Given $\mathbf{x} \in \mathbb{R}^n$, compute $\mathit{dist}(\mathbf{x})$. Then return $\mathbf{y}$ according to the following probability distribution: $\langle \mathit{lt}(\mathbf{x}), \mathit{rt}(\mathbf{x}) \rangle$ with probability 1/2, $\langle \mathit{lt}(\mathbf{x}) + \mathit{dist}(\mathbf{x}), \mathit{rt}(\mathbf{x}) - \mathit{dist}(\mathbf{x}) \rangle$ with probability 1/6, and $\langle \mathit{lt}(\mathbf{x}) + \mathit{dist}(\mathbf{x})/2, \mathit{rt}(\mathbf{x}) - \mathit{dist}(\mathbf{x})/2 \rangle$ with probability 1/3.*

The unbalanced weights inevitably harm the mechanism's approximation performance. Nevertheless, we shall demonstrate that Mechanism 2 succeeds in breaking the deterministic lower bound of 2 by a significant margin.

Objective Function	Deterministic	Randomized
Social Cost	UB: $n-1$ GSP LB: $3/2 - \mathcal{O}(1/n)$ SP (Thm 3.1)	N/A
Maximum Cost	UB: 2 GSP (Thm 3.3) LB: 2 SP (Cor 3.4)	UB: 5/3 SP (Thm 3.5) LB: 3/2 SP (Cor 3.6)

Table 2: A summary of the results of Section 3. UB and LB stand for upper bound and lower bound, respectively. SP and GSP stand for strategyproof and group strategyproof, respectively.

THEOREM 3.5. *Mechanism 2 is a strategyproof 5/3 approximation mechanism for the maximum cost in the two facility setting.*

The proof of the theorem is rather elaborate, and can be found in the full version of the paper [27]. As in the deterministic case, we observe that the lower bound of $3/2$ given in Theorem 2.4 also holds, up to an additive term of ϵ, in our current setting, as long as $n \geq 3$.

COROLLARY 3.6. *Let $N = \{1, \ldots, n\}$, and let $\epsilon > 0$. Any randomized strategyproof mechanism has an approximation ratio of at least $3/2 - \epsilon$ for the maximum cost in the two facility setting.*

PROOF SKETCH. We use the same construction as in the proof of Theorem 2.4 for $n-1$ agents, and add an additional agent located at a large enough value $v(\epsilon)$ that depends on ϵ. Now, in order to obtain a small approximation ratio, the expected distance of the right facility y_2 from $v(\epsilon)$ must be small, hence the probability that y_2 is relevant to the first $n-1$ agents can be made arbitrarily small. We conclude that the arguments of the proof of Theorem 2.4 work here as well, up to an arbitrarily small additive term. $\square$

3.3 Discussion

Table 2 summarizes the results of Section 3. A truly intriguing gap is the one between the trivial $n-1$ strategyproof upper bound for the social cost, and the lower bound of $3/2$. This problem seems straightforward at first, but has proved quite elusive. We conjecture that it is possible to obtain a deterministic lower bound of $\Omega(n)$. In fact, we can easily prove an $\Omega(n)$ lower bound for *group strategyproof* mechanisms.

Another gap is between our randomized upper bound of $5/3$ for the maximum cost, and the lower bound of $3/2$. Moreover, it is unclear whether Mechanism 2 is group strategyproof.

A natural, asked for way to further extend the results of this section is to consider a setting with more than two facilities. The computational problems involved are still tractable when the number of facilities is constant. However, the intuitions behind the positive results given in this section (that is, Theorems 3.3 and 3.5) already collapse even with respect to three facilities.

Note that the preferences of the agents in the setting of Section 3 are *not* single peaked. This fact is meaningful with respect to the generality of our agenda.

4. EXTENSION II: MULTIPLE LOCATIONS PER AGENT

Another natural extension of the setting of Section 2 is the one in which each agent controls multiple locations. Let w_i be the number of locations controlled by agent $i \in N$. We denote the set of locations that agent i controls by $\mathbf{x}_i = \langle x_{i1}, \ldots, x_{iw_i} \rangle$, and the location profile is now $\mathbf{x} = \langle \mathbf{x}_1, \ldots, \mathbf{x}_n \rangle$.

A deterministic mechanism in the multiple locations setting is a function $f : \mathbb{R}^{w_1} \times \cdots \mathbb{R}^{w_n} \to \mathbb{R}$, which locates a single facility given the multiple locations reported by each agent. As in Section 2, a randomized mechanism returns a probability distribution over $\mathbb{R}$.

As before, we will be interested in minimizing the social cost or the maximum cost, but now the cost of an agent depends on the objective function. If the objective function is minimizing the social cost, given a facility location y, the cost of an agent is the sum of distances to its locations: $\text{cost}(y, \mathbf{x}_i) = \text{sc}(y, \mathbf{x}_i) = \sum_{j=1}^{w_i} |y - x_{ij}|$. If the goal is minimizing the maximum cost, then the cost of an agent is the maximum distance to its locations: $\text{cost}(y, \mathbf{x}_i) = \text{mc}(y, \mathbf{x}_i) = \max_{j \in \{1, \ldots, w_i\}} |y - x_{ij}|$.

The same goes for randomized mechanisms, with respect to expected costs. Notice that, when the individual costs are defined as above, optimizing the social cost is in fact equivalent to minimizing the sum of distances to all the locations controlled by all the agents, that is, choosing y that minimizes $\sum_{i \in N} \sum_{j \in \{1, \ldots, w_i\}} |y - x_{ij}|$. Optimizing the maximum cost implies minimizing the maximum distance with respect to all the locations controlled by all the agents, i.e., minimizing $\max_{i \in N} \max_{j \in \{1, \ldots, w_i\}} |y - x_{ij}|$.

4.1 Social Cost

As in Section 3, when moving from the basic setting to this more elaborate setting, optimization of the social cost is no longer strategyproof. To see this, consider a simple example with two agents. Let $\mathbf{x}_1 = \langle 0, 1, 1 \rangle$ and $\mathbf{x}_2 = \langle 0, 0 \rangle$. The optimal solution is the median of all the locations, which is 0; we have that $\text{cost}(0, \mathbf{x}_1) = 2$. However, by reporting $\mathbf{x}'_1 = \langle 1, 1, 1 \rangle$, agent 1 can move the median of all the locations to 1; notice that $\text{cost}(1, \mathbf{x}_1) = 1$, hence agent 1 benefits from misreporting its locations.

Deterministic Mechanisms. Dekel, Fischer and Procaccia [9] have in fact investigated our current setting (that is, optimizing the social cost when each agent controls multiple locations) with respect to deterministic mechanisms, in the context of incentive compatible regression learning. Some of their results (Section 4 of [9]) deal with a discrete setting where one wishes to optimize the social cost under the absolute loss function, when the function class is the class of constant functions; it can be verified, although it is not immediately obvious, that the two settings are equivalent. Note that the results of Dekel et al. are stated under the assumption that the agents all control the same number of points, that is, $w_i = w_j$ for all $i, j \in N$, but they also hold when this is not the case.

Dekel et al. designed a deterministic group strategyproof 3-approximation mechanism for the social cost in the multiple locations setting. Furthermore, they showed a matching lower bound of 3 for deterministic strategyproof mechanisms. Interestingly, their lower bound holds even when there are only two agents that control the same number of locations.

Randomized Mechanisms. Dekel et al. [9] did not discuss randomized mechanisms. We design a simple randomized mechanism that succeeds in breaking the deterministic lower bound given by Dekel et al. [9].

MECHANISM 3. *Given* $\mathbf{x}$*, return* $med(\mathbf{x}_i)$ *with probability* $w_i/(\sum_{j\in N} w_j)$.

This mechanism is strategyproof. Indeed, for each agent $i \in N$, agent i has single peaked preferences with a peak at $med(\mathbf{x}_i)$. Consider a situation where i lies; if it is not selected by the mechanism, the lie does not make a difference; if i is selected, then it can only be worse off.

However, somewhat counterintuitively and in contrast to the group strategyproof mechanism given by Dekel et al., Mechanism 3 is not group strategyproof; this is demonstrated by following example.

EXAMPLE 4.1 (MECHANISM 3 IS NOT GSP). Let $N = \{1, 2\}$, and set $\mathbf{x}_1 = \langle -3, -2, 1 \rangle$ and $\mathbf{x}_2 = \langle -1, 2, 3 \rangle$. The medians are $med(\mathbf{x}_1) = -2$, $med(\mathbf{x}_2) = 2$, and each is selected by Mechanism 3 with probability $1/2$. Hence, denoting Mechanism 3 by f, we have that for both agents $i \in N$,

$$cost(f(\mathbf{x}), \mathbf{x}_i) = \frac{1}{2} \cdot (1 + 3) + \frac{1}{2} \cdot (1 + 4 + 5) = 7 \quad .$$

On the other hand, consider the location profile $\mathbf{x}'$ where both agents report all their locations to be at 0. Then $f(\mathbf{x}')$ selects 0 with probability one. Hence, for all $i \in N$, $cost(\mathbf{x}_i, f(\mathbf{x}')) = 6$. This means that both agents strictly benefit from the deviation from $\mathbf{x}$ to $\mathbf{x}'$.

We now turn to establishing the approximation guarantees provided by Mechanism 3.

THEOREM 4.2. *Mechanism 3 is a strategyproof mechanism in the multiple locations setting. Moreover, if* $n = 2$*, the mechanism yields an approximation ratio of* $2 + \frac{|w_1 - w_2|}{w_1 + w_2}$ *for the social cost.*

In particular, the theorem implies that Mechanism 3 gives a 2-approximation when there are two agents that control the same number of points, which is a setting where the deterministic lower bound of 3 (given in [9]) holds. The proof of the theorem is relegated to the full version of the paper [27]. Extending the theorem beyond two agents remains an open problem.

We now construct an example that serves two purposes. First, the example shows that Mechanism 3 does not provide an approximation ratio better than $3 - 2/n$ when there are n agents, even when the agents control the same number of locations, and thus does not significantly beat the deterministic lower bound of 3 when the number of agents is large. Second, the example demonstrates the tightness of the upper bound given in Theorem 4.2, that is, when there are two agents with w_1 and w_2 points, the mechanism does not obtain an approximation ratio better than $2 + \frac{|w_1 - w_2|}{w_1 + w_2}$.

EXAMPLE 4.3 (LOWER BOUNDS FOR MECHANISM 3). We first establish that, when $N = \{1, \ldots, n\}$, given $\epsilon > 0$ there is $w \in \mathbb{N}$ large enough such that even when each agent controls exactly w locations, the approximation ratio given by Mechanism 3 is at least $3 - \frac{2}{n} - \epsilon$.

Let $w = 2k + 1$, where k is to be chosen later. Construct a location profile $\mathbf{x}$ as follows. For agent 1, we have $x_1, \ldots, x_{k+1} = 0$, and $x_{k+2}, \ldots, x_{2k+1} = 1$. For all $j \neq i$ and all $l = 1, \ldots, w$, $x_{jl} = 1$. Notice that $med(\mathbf{x}_1) = 0$. With probability $1/n$ the algorithm returns 0, and has a social cost of $(n-1)(2k+1) + k$. With probability $(n-1)/n$ the algorithm selects 1 and has a social cost of $k + 1$. The ratio is

$$\frac{\frac{1}{n} \cdot ((n-1)(2k+1) + k) + \frac{n-1}{n} \cdot (k+1)}{k+1} = 3 - \frac{2}{n} - \frac{1}{k+1} \quad .$$

To prove the claim, choose $k > 1/\epsilon - 1$.

Interestingly, the same example also shows a lower bound of $2 + \frac{|w_1 - w_2|}{w_1 + w_2} - \epsilon$ in the setting of Theorem 4.2, by choosing $w_1 = 2k + 1$, $w_2 = (n-1)(2k+1)$. The analysis is as above with respect to agent 1, whereas agent 2 replaces agents $2, \ldots, n$ above. In this case,

$$2 + \frac{|w_1 - w_2|}{w_1 + w_2} = 2 + \frac{n-2}{n} = 3 - \frac{2}{n} \quad .$$

4.2 Maximum Cost

Our last object of interest is mechanisms for minimizing the maximum cost, in the setting where each agent $i \in N$ controls w_i locations. Similarly to Section 3, we shall demonstrate that the results of Section 2 can be leveraged to obtain tight or nearly tight results in the current setting.

A crucial observation is that, given an agent $i \in N$, its location profile $\mathbf{x}_i \in \mathbb{R}^{w_i}$, and a facility location $y \in \mathbb{R}$,

$$mc(y, \mathbf{x}_i) = |y - cen(\mathbf{x}_i)| + \frac{rt(\mathbf{x}_i) - lt(\mathbf{x}_i)}{2} \quad . \quad (1)$$

Hence, when $cost(y, \mathbf{x}_i) = mc(y, \mathbf{x}_i)$, the preferences of the agents are single peaked with the peak at $cen(\mathbf{x}_i)$, and, moreover, their utility depends only on the distance $|y - cen(\mathbf{x}_i)|$.

Deterministic Mechanisms. In previous settings we have seen that it is straightforward to obtain a deterministic strategyproof 2-approximation mechanism for the maximum cost. The reason (implicitly underlying the result of Section 3) was that returning any location between $lt(\mathbf{x})$ and $rt(\mathbf{x})$ yields a 2-approximation. The same logic also delivers in our current setting.

Given $\mathbf{x} \in \mathbb{R}^{w_1} \times \cdots \times \mathbb{R}^{w_n}$, we define the vector

$$multicen(\mathbf{x}) = \langle cen(\mathbf{x}_1), \ldots, cen(\mathbf{x}_n) \rangle \quad .$$

This is the vector of the centers of the agents' location profiles, or, in other words, the vector of the peaks of the agents' preferences. Hence, choosing the leftmost center, $lt(multicen(\mathbf{x}))$, is a group strategyproof solution. Moreover, we have that

$$lt(\mathbf{x}) \leq lt(multicen(\mathbf{x})) \leq rt(\mathbf{x}) \quad ,$$

so

$$mc(lt(multicen(\mathbf{x}), \mathbf{x}) \leq rt(\mathbf{x}) - lt(\mathbf{x}) \quad ,$$

whereas the optimal solution has a maximum cost of at least $(rt(\mathbf{x}) - lt(\mathbf{x}))/2$. We have proved:

Objective Function	Deterministic	Randomized
Social Cost	UB: 3 GSP (Dekel et al. [9]) LB: 3 SP (Dekel et al. [9])	UB: $2 + \frac{|w_1 - w_2|}{w_1 + w_2}$ SP ($n = 2$, Thm 4.2) LB: N/A
Maximum Cost	UB: 2 GSP (Thm 4.4) LB: 2 SP (Thm 2.2)	UB: 3/2 GSP (Thm 4.5) LB: 3/2 SP (Thm 2.4)

Table 3: A summary of the results of Section 4. UB and LB stand for upper bound and lower bound, respectively. SP and GSP stand for strategyproof and group strategyproof, respectively.

THEOREM 4.4. $f(\mathbf{x}) = lt(multicen(\mathbf{x}))$ *is a group strategyproof 2-approximation mechanism for the maximum cost in the multiple location setting.*

Since in the current setting we can have that $w_i = 1$ for all $i \in N$, any lower bound from Section 2 holds here as well. In particular, Theorem 2.2 provides a tight lower bound of 2.

Randomized Mechanisms. In order to obtain randomized mechanisms for the maximum cost in the multiple location setting we once again leverage the techniques of Section 2. Consider the following Mechanism.

MECHANISM 4. *Given* $\mathbf{x} \in \mathbb{R}^{w_1} \times \cdots \mathbb{R}^{w_n}$, *return the point* $lt(multicen(\mathbf{x}))$ *with probability 1/4,* $rt(multicen(\mathbf{x}))$ *with probability 1/4, and*

$$cen(multicen(\mathbf{x})) = (lt(multicen(\mathbf{x})) + rt(multicen(\mathbf{x})))/2$$

with probability 1/2.

The following theorem establishes that the mechanism has some very desirable properties.

THEOREM 4.5. *Mechanism 4 is a group strategyproof 3/2-approximation mechanism for the maximum cost in the multiple location setting.*

PROOF. It can easily be verified that, using (1), the group strategyproofness of the mechanism follows from exactly the same arguments as in the proof of Theorem 2.3. Therefore, we concentrate on establishing the announced approximation ratio.

Let $\mathbf{x} \in \mathbb{R}^n$. Without loss of generality (by scaling the distances) we assume that $lt(\mathbf{x}) = 0$, $rt(\mathbf{x}) = 1$. We first claim that $lt(multicen(\mathbf{x})) \leq 1/2$. Indeed, let $i \in N$ be the agent that controls 0. Then $lt(\mathbf{x}_i) = 0$, $rt(\mathbf{x}_i) \leq 1$, hence $cen(\mathbf{x}_i) \leq 1/2$. The claim directly follows. Similarly, we have that $rt(multicen(\mathbf{x})) \geq 1/2$. In other words, it holds that $lt(multicen(\mathbf{x}))$ is at least as close to 0 as to 1, whereas $rt(multicen(\mathbf{x}))$ is at least as close to 1 as to 0. Therefore, denoting Mechanism 4 by f, we have:

$$mc(f(\mathbf{x}), \mathbf{x}) = \frac{1}{4} \cdot (1 - lt(multicen(\mathbf{x}))) + \frac{1}{4} \cdot rt(multicen(\mathbf{x}))$$

$$+ \frac{1}{2} \cdot \max \left\{ \frac{lt(multicen(\mathbf{x})) + rt(multicen(\mathbf{x}))}{2}, \right.$$

$$\left. 1 - \frac{lt(multicen(\mathbf{x})) + rt(multicen(\mathbf{x}))}{2} \right\}$$

$$= \max \left\{ \frac{1}{4} + \frac{rt(multicen(\mathbf{x}))}{2}, \right.$$

$$\left. \frac{3}{4} - \frac{lt(multicen(\mathbf{x}))}{2} \right\}$$

$$\leq \frac{3}{4} \ ,$$

where the last inequality follows from $lt(multicen(\mathbf{x})) \geq 0$ and $rt(multicen(\mathbf{x})) \leq 1$. The optimal solution has a cost of 1/2. Therefore, we get an approximation ratio of 3/2. $\square$

Finally, we remark that the randomized lower bound of 3/2 given by Theorem 2.4 holds here as well. We find it quite surprising that the upper bound yielded by the seemingly "generous" Mechanism 4 is tight.

4.3 Discussion

Table 3 summarizes the results of Section 4. The most interesting open question is how the analysis of Mechanism 3 extends to $n > 2$. We conjecture that for any number of agents n, if $w_i = w_j$ for all $i, j \in N$, then the mechanism yields an approximation ratio of $3 - 2/n$. Such a result would be tight by Example 4.3. Moreover, we have no general lower bound that works for randomized mechanisms for the social cost.

The setting investigated in this section has many applications, but we note that, in particular, any results about randomized strategyproof mechanisms for the social cost can be directly applied in the incentive compatible regression learning setting of Dekel et al. [9].

5. OPEN PROBLEMS AND FURTHER DISCUSSION

We believe we have just scratched the surface with respect to approximate mechanism design without money. First, there are many interesting open problems and conjectures that are directly related to the domains investigated above; these problems are summarized in Sections 3.3 and 4.3.

Second, there are several additional, very natural, extensions of the setting of Section 2 that we have not considered above. One example is looking at domains where the space of locations is more involved, either multi-dimensional Euclidean space, or settings where agents have single peaked preferences over graphs [30]. Another example is considering *allotment rules*, namely rules that assign a point $a_i \in [0, 1]$ to each agent, such that $\sum_{i \in N} a_i = 1$; this setting models the division of a task or a good among the agents [5, Section 4.1]. Furthermore, it is possible to consider almost any combination of the extensions, e.g., a domain in which agents control multiple locations (as in Section 4) *and* two facilities must be located (as in Section 3).

Third, in the long run our intention is to apply the ideas discussed in this paper to completely different domains. We have several examples in mind, but here we outline just one. Consider a directed graph where the vertices are agents, and an edge from agent i to agent j means that i recommends j, trusts j, or "votes" for j in some sense. The score of an agent is its indegree in the graph. The mechanism receives the graph as input, and outputs a set of k agents. This

model can be interpreted in the context of social networks, recommendation systems, etc. An agent's strategy is its set of outgoing edges, and its utility is one if it is among the selected agents and zero if not; the objective function is the total score of the selected agents. In ongoing work with colleagues, we have designed some strategyproof approximation mechanisms with a constant approximation ratio for this problem.

6. ACKNOWLEDGMENTS

7. REFERENCES

[1] N. Andelman, Y. Azar, and M. Sorani. Truthful approximation mechanisms for scheduling selfish related machines. In *Proc. of 22nd STACS*, pages 69–82, 2005.

[2] A. Archer and É. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of 42nd FOCS*, pages 482–491, 2001.

[3] A. Archer and É. Tardos. Frugal path mechanisms. *ACM Transactions on Algorithms*, 3(1), article 3, 2007.

[4] S. Arora, P. Raghavan, and S. Rao. Approximation schemes for Euclidean k-medians and related problems. In *Proc. of 30th STOC*, pages 106–113, 1998.

[5] S. Barberà. An introduction to strategy-proof social choice functions. *Social Choice and Welfare*, 18:619–653, 2001.

[6] S. Barberà and B. Peleg. Strategy-proof voting schemes with continuous preferences. *Social Choice and Welfare*, 7:31–38, 1990.

[7] M. Bern and D. Eppstein. Approximation algorithms for geometric problems. In D. Hochbaum, editor, *Approximation Algorithms for NP-Hard Problems*. PWS Publishing, 1996.

[8] G. Christodoulou, E. Koutsoupias, and A. Vidali. A lower bound for scheduling mechanisms. In *Proc. of 18th SODA*, pages 1163–1170, 2007.

[9] O. Dekel, F. Fischer, and A. D. Procaccia. Incentive compatible regression learning. In *Proc. of 19th SODA*, pages 277–286, 2008.

[10] P. Dhangwatnotai, S. Dobzinski, S. Dughmi, and T. Roughgarden. Truthful approximation schemes for single-parameter agents. In *Proc. of 49th FOCS*, pages 15–24, 2008.

[11] S. Dobzinski, N. Nisan, and M. Schapira. Truthful randomized mechanisms for combinatorial auctions. In *Proc. of 38th STOC*, pages 644–652, 2006.

[12] E. Elkind, A. Sahai, and K. Steiglitz. Frugality in path auctions. In *Proc. of 15th SODA*, pages 701–709, 2004.

[13] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *Americal Mathematical Monthly*, 69(1):9–15, 1962.

[14] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41:587–602, 1973.

[15] R. Lavi and C. Swami. Truthful and near-optimal mechanism design via linear programming. In *Proc. of 46th FOCS*, pages 595–604, 2005.

[16] R. Lavi and C. Swami. Truthful mechanism design for multi-dimensional scheduling via cycle monotonicity. In *Proc. of 8th ACM-EC*, pages 252–261, 2007.

[17] D. Lehmann, L. I. O'Callaghan, and Y. Shoham. Truth revelation in rapid, approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):577–602, 2002.

[18] S. Leonardi and G. Schäfer. Cross-monotonic cost-sharing methods for connected facility location. In *Proc. of 5th ACM-EC*, pages 242–243, 2004.

[19] H. Levin, M. Schapira, and A. Zohar. Interdomain routing and games. In *Proc. of 40th STOC*, pages 57–66, 2008.

[20] R. Meir, A. D. Procaccia, and J. S. Rosenschein. Strategyproof classification under constant hypotheses: A tale of two functions. In *Proc. of 23rd AAAI*, pages 126–131, 2008.

[21] R. Meir, A. D. Procaccia, and J. S. Rosenschein. Strategyproof classification with shared inputs. In *Proc. of 21st IJCAI*, 2009. To appear.

[22] E. Miyagawa. Locating libraries on a street. *Social Choice and Welfare*, 18(3):527–541, 2001.

[23] H. Moulin. On strategy-proofness and single-peakedness. *Public Choice*, 35:437–455, 1980.

[24] N. Nisan. Introduction to mechanism design (for computer scientists). In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 9. Cambridge University Press, 2007.

[25] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35(1–2):166–196, 2001.

[26] M. Pál and É. Tardos. Group strategyproof mechanisms via primal-dual algorithms. In *Proc. of 44th FOCS*, pages 584–593, 2003.

[27] A. D. Procaccia and M. Tennenholtz. Approximate mechanism design without money. Manuscript, 2009. Available from: http://www.cs.huji.ac.il/~arielpro/papers/facility.pdf.

[28] A. E. Roth and M. Sotomayor. *Two-Sided Mathing: A Study in Game-Theoretic Modelling and Analysis*. Cambridge University Press, 1991.

[29] M. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.

[30] J. Schummer and R. V. Vohra. Strategy-proof location on a network. *Journal of Economic Theory*, 104(2):405–428, 2004.

[31] J. Schummer and R. V. Vohra. Mechanism design without money. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 10. Cambridge University Press, 2007.

[32] Y. Sprumont. Strategyproof collective choice in economic and political environments. *The Canadian Journal of Economics*, 28(1):68–107, 1995.