

Facility Location with Double-peaked Preferences

Aris Filos-Ratsikas

Aarhus University, Denmark
filosra@cs.au.dk

Minming Li

City University of Hong Kong, HK
minming.li@cityu.edu.hk

Jie Zhang

University of Oxford, UK
jie.zhang@cs.ox.ac.uk

Qiang Zhang

University of Warsaw, Poland
qzhang@mimuw.edu.pl

Abstract

We study the problem of locating a single *facility* on a real line based on the reports of self-interested agents, when agents have *double-peaked preferences*, with the peaks being on opposite sides of their locations. We observe that double-peaked preferences capture real-life scenarios and thus complement the well-studied notion of single-peaked preferences. We mainly focus on the case where peaks are equidistant from the agents' locations and discuss how our results extend to more general settings. We show that most of the results for single-peaked preferences do not directly apply to this setting; this makes the problem essentially more challenging. As our main contribution, we present a simple truthful-in-expectation mechanism that achieves an approximation ratio of $1+b/c$ for both the social and the maximum cost, where b is the distance of the agent from the peak and c is the minimum cost of an agent. For the latter case, we provide a $3/2$ lower bound on the approximation ratio of any truthful-in-expectation mechanism. We also study deterministic mechanisms under some natural conditions, proving lower bounds and approximation guarantees. We prove that among a large class of reasonable mechanisms, there is no deterministic mechanism that outperforms our truthful-in-expectation mechanism.

1 Introduction

We study the problem of locating a single *facility* on a real line, based on the input provided by selfish agents who wish to minimize their costs. Each agent has a *location* $x_i \in \mathbb{R}$ which is her private information and is asked to report it to some central authority, which then decides where to locate the facility, aiming to optimize some function of the agents' reported locations. This model corresponds to problems such as finding the ideal location for building a primary school or a bus stop along a street, so that the total distance of all agents' houses from the location is minimized, or so that no agent's house will lie too far away from that location.

In our setting, we assume that agents have *double-peaked preferences*, i.e. we assume that each agent i has two unique most preferred points or *peaks*, located at some distances

from x_i on opposite sides, where her cost is minimum. Traditionally, preferences in facility location problems are assumed to be *single-peaked*, i.e. each agent's location is her most preferred point on the line and her cost increases linearly (at the same rate) to the left and the right of that peak. Sometimes however, single-peaked preferences do not model real-life scenarios accurately. Take for instance the example mentioned above, where the government plans to build a primary school on a street. An agent with single-peaked preferences would definitely want the school built next to her house, so that she wouldn't have to drive her children there everyday. However, it is quite possible that she is also not very keen on the inevitable drawbacks of having a primary school next to her house either, like unpleasant noise or trouble with parking. On the other hand, a five-minute walking distance is sufficiently far for those problems to no longer be a factor but also sufficiently close for her children to be able to walk to school. There are two such positions, (symmetrically) in each direction, and those would be her two peaks.

Our primary objective is to explore double-peaked preferences in facility location settings similar to the ones studied extensively for single-peaked preferences throughout the years (Procaccia and Tennenholtz 2009; Schummer and Vohra 2002; Lu, Wang, and Zhou 2009; Lu et al. 2010; Alon et al. 2010; Fotakis and Tzamos 2010; Escoffier et al. 2011; Fotakis and Tzamos 2012; Dokow et al. 2012; Feldman and Wilf 2013). For that reason, following the literature we assume that the cost functions are the same for all agents and that the cost increases linearly, at the same rate, as the output moves away from the peaks. The straightforward extension to the double-peaked case is piecewise-linear cost functions, with the same slope in all intervals, which gives rise to the natural model of *symmetric* agents, i.e. the peaks are equidistant from the agent's location. Note that this symmetry is completely analogous to the single-peaked case (for facility location problems, e.g. see (Procaccia and Tennenholtz 2009)), where agents have exactly the same cost on two points equidistant from their peaks. Our lower bounds and impossibility results naturally extend to non-symmetric settings, but some of our mechanisms do not. We discuss those extensions in Section 5.

Our model also applies to more general spaces, beyond the real line. One can imagine for instance that the goal is

to build a facility on the plane where for the same reasons, agents would like the facility to be built at some distance from their location, in *every direction*. This translates to an agent having infinitely many peaks, located on a circle centered around her location. In that case of course, we would no longer refer to agents' preferences as double-peaked but the underlying idea is similar to the one presented in this paper. We do not explore such extensions here; we leave that for future work.

Agents are self-interested entities that wish to minimize their costs. We are interested in mechanisms that ensure that agents are not incentivized to report anything but their actual locations, namely *strategyproof* mechanisms. We are also interested in *group strategyproof* mechanisms, i.e., mechanisms that are resistant to manipulation by coalitions of agents. Moreover, we want those mechanisms to achieve some good performance guarantees, with respect to our goals. If our objective is to minimize the sum of the agent's costs, known as the *social cost*, then we are looking for strategyproof mechanisms that achieve a social cost as close as possible to that of the optimal mechanism, which need not be strategyproof. The prominent measure of performance for mechanisms in computer science literature is the approximation ratio (Dughmi and Gosh 2010; Guo and Conitzer 2010; Ashlagi et al. 2010; Caragiannis, Filos-Ratsikas, and Procaccia 2011), i.e., the worst possible ratio of the social cost achieved by the mechanism over the minimum social cost over all instances of the problem. The same holds if our objective is to minimize the *maximum cost* of any agent. In the case of *randomized mechanisms*, i.e., mechanisms that output a probability distribution over points in \mathbb{R} , instead of a single point, as a weaker strategyproofness constraint, we require *truthfulness-in-expectation*, i.e., a guarantee that no agent can reduce her expected cost from misreporting.

Double-peaked preferences in practice

Single-peaked preferences were introduced in (Black 1957 reprint at 1986) as a way to avoid *Condorcet cycles* in majority elections. Moulin (1980) characterized the class of strategyproof mechanisms in this setting, proving that *median voter schemes* are essentially the only strategyproof mechanisms for agents with single-peaked preferences. Double-peaked preferences have been mentioned in social choice literature (e.g. see (Cooter 2002)), to describe settings where preferences are not single-peaked, voting cycles do exist and majority elections are not possible. In broader social choice settings, they can be used to model situations where e.g. a left-wing party might prefer a more conservative but quite effective policy to a more liberal but ineffective one on a left-to-right political axis. In fact, Egan (2013) provides a detailed discussion on double-peaked preferences in political decisions. He uses a 1964-1970 survey about which course of action the United States should take with regard to the Vietnam war as an example where the status quo (keep U.S. troops in Vietnam but try to terminate the war) was ranked last by a considerable fraction of the population when compared to a left-wing policy (pull out entirely) or a right-wing policy (take a stronger stand). This demonstrates that in a scenario where the standard approach would be to as-

sume that preferences are single-peaked, preferences can instead be double-peaked. Egan provides additional evidence for the occurrence of double-peaked preferences supported by experimental results based on surveys on the U.S. population, for many different problems (education, health care, illegal immigration treatment, foreign oil treatment e.t.c.). More examples of double-peaked preferences in real-life scenarios are presented in (Rosen 2004). The related work demonstrates that although they might not be as popular as their single-peaked counterpart, double-peaked preferences do have applications in settings more general than the street example described earlier. On the other hand, the primary focus of this paper is to study double-peaked preferences on facility location settings and therefore the modelling assumptions follow the ones of the facility location literature.

Our results

Our main contribution is a truthful-in-expectation mechanism (M1) that achieves an approximation ratio of $1 + b/c$ for the social cost and $\max\{1 + b/c, 2\}$ for the maximum cost, where b is the distance between an agent's location and her peak and c is her minimum cost. We also prove that no truthful-in-expectation mechanism can do better than a $3/2$ approximation for the maximum cost proving that at least for the natural special case where $b = c$, Mechanism M1 is not far from the best possible. For deterministic mechanisms, we prove that no mechanism in a wide natural class of strategyproof mechanisms can achieve an approximation ratio better than $1 + b/c$ for the social cost and $1 + 2b/c$ for the maximum cost and hence cannot outperform Mechanism M1. To prove this, we first characterize the class of strategyproof, anonymous and position invariant mechanisms for two agents by a single mechanism (M2). Intuitively, anonymity requires that all agents are handled equally by the mechanism while position invariance essentially requires that if we shift an instance by a constant, the location of the facility should be shifted by the same constant as well. This is a quite natural condition and can be interpreted as a guarantee that the facility will be located *relatively* to the reports of the agents and independently of the underlying network (e.g. the street).

We prove that the approximation ratio of Mechanism M2 for the social cost is $\Theta(n)$, where n is the number of agents and conjecture that no deterministic strategyproof mechanism can achieve a constant approximation ratio in this case. For the maximum cost, the ratio of Mechanism M2 is $\max\{1 + 2b/c, 3\}$ which means that the mechanism is actually the best in the natural class of anonymous and position invariant mechanisms. For any deterministic strategyproof mechanism, we prove a lower bound of 2 on the approximation ratio, proving that at least for the natural case of $b = c$, Mechanism M2 is also not far from optimal. Finally, we prove an impossibility result; there is no group strategyproof, anonymous and position invariant mechanism for the problem. This is in contrast with the single-peaked preference setting, where there is a large class of group strategyproof mechanisms that satisfy those properties. Our results are summarized in Table 1.

Most of our results appear without proofs, due to lack of

space; the proofs appear in the full version.

2 Preliminaries

Let $N = \{1, 2, \dots, n\}$ be a set of *agents*. We consider the case where agents are located on a line, i.e., each agent $i \in N$ has a location $x_i \in \mathbb{R}$. We will occasionally use x_i to refer to both the position of agent i and the agent herself. We will call the collection $\mathbf{x} = \langle x_1, \dots, x_n \rangle$ a *location profile* or an *instance*.

A *deterministic mechanism* is a function $f : \mathbb{R}^n \mapsto \mathbb{R}$ that maps a given location profile to a point in \mathbb{R} , the location of the *facility*. We assume that agents have *double-peaked preferences*, symmetric with respect to the origin. We discuss how our results extend to non-symmetric agents in Section 5. Given any instance \mathbf{x} and a location $y \in \mathbb{R}$, the cost of agent i is

$$\text{cost}(y, x_i) = \begin{cases} c + |x_i - b - y| & \text{if } y \leq x_i \\ c + |x_i + b - y| & \text{if } y > x_i \end{cases}$$

where c and b are positive constants. We will say that y *admits* a cost of $\text{cost}(y, x_i)$ for agent i on instance \mathbf{x} . For a mechanism that outputs $f(\mathbf{x})$ on instance \mathbf{x} , the cost of agent i is $\text{cost}(f(\mathbf{x}), x_i)$. Intuitively, each agent has two most favorable locations, i.e., $x_i - b$ and $x_i + b$, which we refer to as the *peaks* of agent i . Note that these peaks are actually the troughs of the curve of the cost function, but much like most related work, we refer to them as peaks. The parameter $c > 0$ is the minimum cost incurred to an agent when the facility is built on one of her peaks.¹ Note that the special case, where $b = c$ corresponds to the natural setting where the incurred minimum cost of an agent is interpreted as the distance she needs to cover to actually reach the facility. This case is particularly appealing, since the bounds we obtain are clean numbers, independent of b and c . The bounds for the natural case can be obtained directly by letting $b = c$ in all of our results.

A *randomized mechanism* is a function $f : \mathbb{R}^n \mapsto \Delta(\mathbb{R})$, where $\Delta(\mathbb{R})$ is the set of probability distributions over \mathbb{R} . It maps a given location profile to probabilistically selected locations of the facility. The expected cost of agent i is $\mathbb{E}_{y \sim \mathcal{D}} [\text{cost}(y, x_i)]$, where \mathcal{D} is the probability distribution of the mechanism outputs.

We will call a deterministic mechanism f *strategyproof* if no agent would benefit by misreporting her location, regardless of the locations of the other agents. This means that for every $\mathbf{x} \in \mathbb{R}^n$, every $i \in N$ and every $x'_i \in \mathbb{R}$, $\text{cost}(f(\mathbf{x}), x_i) \leq \text{cost}(f(x'_i, \mathbf{x}_{-i}), x_i)$, where $\mathbf{x}_{-i} = \langle x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n \rangle$. A mechanism is *truthful-in-expectation* if it guarantees that every agent always minimizes her expected cost by reporting her location truthfully. Throughout the paper we will use the term *strategyproofness* when referring to deterministic mechanisms and the term *truthfulness* when referring to randomized mechanisms.

¹It is not hard to see by our results that if we let an agent's cost be zero on her peaks, then in very general settings, no deterministic strategyproof mechanism can guarantee a finite approximation ratio.

A mechanism is *group strategyproof* if there is no coalition of agents, who by jointly misreporting their locations, affect the outcome in a way such that the cost of none of them increases and the cost of at least one of them strictly decreases. In other words, there is no $S \subseteq N$ such that for some misreports x'_S of agents in S and some reports \mathbf{x}_{-S} of agents in $N \setminus S$, $\text{cost}(f(x'_S, \mathbf{x}_{-S}), x_i) \leq \text{cost}(f(\mathbf{x}), x_i)$ for all $i \in S$, and $\text{cost}(f(x'_S, \mathbf{x}_{-S}), x_j) < \text{cost}(f(\mathbf{x}), x_j)$ for at least one $j \in S$.

Given an instance \mathbf{x} and a location $y \in \mathbb{R}$, the *social cost* and the *maximum cost* of y are defined respectively as:

$$SC_y(\mathbf{x}) = \sum_{i=1}^n \text{cost}(y, x_i) \quad , \quad MC_y(\mathbf{x}) = \max_{i \in N} \text{cost}(y, x_i).$$

We will say that y *admits* a social cost of $SC_y(\mathbf{x})$ or a maximum cost of $MC_y(\mathbf{x})$. We will call $y \in \mathbb{R}$ an *optimal location* (for the social cost), if $y \in \arg \min_y SC_y(\mathbf{x})$. The definition for the maximum cost is analogous. Let $SC_{\text{opt}}(\mathbf{x})$ and $MC_{\text{opt}}(\mathbf{x})$ denote the social cost and the maximum cost of an optimal location respectively, on instance \mathbf{x} . For a mechanism f that outputs $f(\mathbf{x})$ on instance \mathbf{x} , we will call $SC_f(\mathbf{x})$ the social cost of the mechanism and we will denote it by $SC_f(\mathbf{x})$; and analogously for the maximum cost.

We are interested in strategyproof mechanisms that perform well with respect to the goal of minimizing either the social cost or the maximum cost. We measure the performance of the mechanism by comparing the social/maximum cost it achieves with the optimal social/maximum cost, on any instance \mathbf{x} .

The approximation ratio of mechanism f , with respect to the social cost, is given by

$$r = \sup_{\mathbf{x}} \frac{SC_f(\mathbf{x})}{SC_{\text{opt}}(\mathbf{x})}.$$

The approximation ratio of mechanism f , with respect to maximum cost, is defined similarly.

For randomized mechanisms, the definitions are similar and the approximation ratio is calculated with respect to the expected social or maximum cost, i.e., the expected sum of costs of all agents and expected maximum cost of any agent, respectively.

Finally we consider some properties which are quite natural and are satisfied by many mechanisms (including the optimal mechanism). A mechanism f is *anonymous*, if for every location profile \mathbf{x} and every permutation π of the agents, $f(x_1, \dots, x_n) = f(x_{\pi(1)}, \dots, x_{\pi(n)})$. We say that a mechanism f is *onto*, if for every point $y \in \mathbb{R}$ on the line, there exists a location profile \mathbf{x} such that $f(\mathbf{x}) = y$. Without loss of generality, for anonymous mechanisms, we can assume $x_1 \leq \dots \leq x_n$.

A property that requires special mention is that of *position invariance*, which is a very natural property as discussed in the introduction. This property was independently defined by (Feigenbaum, Sethuraman, and Ye 2013) where it was referred to as *shift invariance*. One can view position invariance as an analogue to *neutrality* in problems like the one studied here, where there is a continuum of outcomes instead of a finite set.

Definition 1. A mechanism f satisfies position invariance, if for all location profiles $\mathbf{x} = \langle x_1, \dots, x_n \rangle$ and $t \in \mathbb{R}$, it holds $f(x_1 + t, x_2 + t, \dots, x_n + t) = f(\mathbf{x}) + t$. In this case, we will call such a mechanism position invariant. We will refer to instances \mathbf{x} and $\langle x_1 + t, x_2 + t, \dots, x_n + t \rangle$ as position equivalent.

Note that position invariance implies the onto condition. Indeed, for any location profile \mathbf{x} , with $f(\mathbf{x}) = y$, we have $f(x_1 + t, x_2 + t, \dots, x_n + t) = y' = y + t$ for any $t \in \mathbb{R}$, so every point $y' \in \mathbb{R}$ is a potential output of the mechanism.

3 A truthful-in-expectation mechanism

We start the exposition of our results with our main contribution, a truthful-in-expectation mechanism that achieves an approximation ratio of $1 + b/c$ for the social cost and $\max\{1 + b/c, 2\}$ for the maximum cost.

Mechanism M1. Given any instance $\mathbf{x} = \langle x_1, \dots, x_n \rangle$, find the median agent $x_m = \text{median}(x_1, \dots, x_n)$, breaking ties in favor of the agent with the smallest index. Output $f(\mathbf{x}) = x_m - b$ with probability $\frac{1}{2}$ and $f(\mathbf{x}) = x_m + b$ with probability $\frac{1}{2}$.

Theorem 1. Mechanism M1 is truthful-in-expectation.

Proof. First, note that the median agent does not have an incentive to deviate, since her expected cost is already minimum, neither does any agent i for which $x_i = x_m$. Hence, for the deviating agent i it must be either $x_i < x_m$ or $x_i > x_m$. We consider three cases when $x_i < x_m$. The proof for the case $x_i > x_m$ is symmetric. Observe that for agent i to be able to move the position of the facility, she has to report $x'_i \geq x_m$ and change the identity of the median agent. Let x'_m be the median agent in the new instance $\langle x'_i, x_{-i} \rangle$, after agent i 's deviation. If $x'_m = x_m$, then obviously agent x_i does not gain from deviating, so we will assume that $x'_m > x_m$.

Case 1: $x_i + b \leq x_m - b$ (symmetrically $x_i - b \geq x_m + b$).

In this case, the cost of agent i is calculated with respect to $x_i + b$ for both possible outcomes of the mechanism. Since $x'_m - b > x_m - b$ and $x'_m + b > x_m + b$, it holds that $|(x_i + b) - (x'_m - b)| > |(x_i + b) - (x_m - b)|$ and $|(x_i + b) - (x'_m + b)| > |(x_i + b) - (x_m + b)|$ and agent i can not gain from misreporting.

Case 2: $x_m - b < x_i + b \leq x_m$ (symmetrically $x_m \leq x_i - b < x_m + b$).

Again, the cost of agent i is calculated with respect to $x_i + b$ for both outcomes of the mechanism. This time, it might be that $|(x_i + b) - (x'_m - b)| < |(x_i + b) - (x_m - b)|$ but since $(x'_m - b) - (x_m - b) = (x'_m + b) - (x_m + b)$, it will also hold that $|(x_i + b) - (x'_m + b)| > |(x_i + b) - (x_m + b)|$ and also $|(x_i + b) - (x_m - b)| - |(x_i + b) - (x'_m - b)| = |(x_i + b) - (x'_m + b)| - |(x_i + b) - (x_m + b)|$. Hence, the expected cost of agent i after misreporting is at least as much as it was before.

Case 3: $x_m < x_i + b \leq x_m + b$ (symmetrically $x_m - b \leq x_i - b < x_m$).

The cost of agent i before misreporting is calculated with respect to $x_i - b$ when the outcome is $x_m - b$ and with respect to $x_i + b$ when the outcome is $x_m + b$. For any misreport

$x'_i < x_i + b$, this is still the case (for $x'_m - b$ and $x'_m + b$ respectively) and since $(x'_m - b) - (x_m - b) = (x'_m + b) - (x_m + b)$, her expected cost is not smaller than before. For any misreport $x'_i > x_i + b$, her cost is calculated with respect to $x_i + b$ for both possible outcomes of the mechanism and for the same reason as in Case 2, her expected cost is at least as much as it was before misreporting. \square

Next, we will calculate the approximation ratio of the mechanism. In order to do that, we will need the following lemma.

Lemma 1. Let $\mathbf{x} = \langle x_1, \dots, x_m, \dots, x_n \rangle$, where x_m = median (x_1, \dots, x_n) , breaking ties in favor of the smallest index. There exists an optimal location for the social cost in $[x_m - b, x_m + b]$.

Proof. Assume for contradiction that this is not the case. Then, for any optimal location y , it must be that either $y < x_m - b$ or $y > x_m + b$.

Assume first that $y < x_m - b$. Since x_m is the median agent, it holds that for at least $\lceil n/2 \rceil$ agents, $x_i - b \geq x_m - b$, that is $x_m - b$ admits a smaller cost for at least $\lceil n/2 \rceil$ agents when compared to y . Let X_1 be the set of those agents. On the other hand, for each agent $x_i < x_m$, $x_m - b$ may admit a smaller or larger cost than y , depending on her position with respect to y . In the worst case, the cost is larger for every one of those agents, which happens when $x_i + b \leq y$ for every agent with $x_i < x_m$. Let X_2 be the set of those agents. Now observe that for any two agents $x_a \in X_1$ and $x_b \in X_2$, it holds that $\text{cost}(x_a, y) - \text{cost}(x_a, x_m - b) = \text{cost}(x_b, x_m - b) - \text{cost}(x_b, y)$. Since $|X_1| \geq |X_2|$, it holds that $SC_{x_m - b}(\mathbf{x}) \leq SC_y(\mathbf{x})$. Since it can not be that $SC_{x_m - b}(\mathbf{x}) < SC_y(\mathbf{x})$, $x_m - b$ is an optimal location and we get a contradiction.

Now assume $y > x_m + b$. Let $y = x_m + b$. If the number of agents is odd, then we can use an exactly symmetric argument to prove that $SC_{x_m + b} \leq SC_y$. If the number of agents is even, the argument can still be used, since our tie-breaking rule selects agent $x_{n/2}$ as the median. Specifically, $x_m + b$ admits a smaller cost for exactly $n/2$ of the agents (including agent $x_{n/2}$) and in the worst case, y admits a smaller cost for $n/2$ agents as well. If X_1 and X_2 are the sets of those agents respectively, then again it holds that $\text{cost}(x_a, y) - \text{cost}(x_a, x_m + b) = \text{cost}(x_b, x_m + b) - \text{cost}(x_b, y)$ for $x_a \in X_1$ and $x_b \in X_2$ and we get a contradiction as before. \square

We now proceed to proving the approximation ratio of Mechanism M1.

Theorem 2. Mechanism M1 has an approximation ratio of $1 + \frac{b}{c}$ for the social cost.

Proof. Consider an arbitrary instance $\mathbf{x} = \langle x_1, \dots, x_n \rangle$ and let x_m be the median agent. By Lemma 1, there exists an optimal location $y \in [x_m - b, x_m + b]$. Let $\delta = y - (x_m - b)$. For every agent i , it holds that $\text{cost}(x_i, x_m - b) \leq \text{cost}(x_i, y) + \delta$. To see this, first observe that $|(x_i - b) - (x_m - b)| \leq |(x_i - b) - y| + \delta$ and that $|(x_i + b) - (x_m - b)| \leq |(x_i + b) - y| + \delta$. If the cost of an agent admitted by y

and $x_m - b$ is calculated with respect to the same peak, then $\min(|(x_i - b) - (x_m - b)|, |(x_i + b) - (x_m - b)|) \leq \min(|(x_i - b) - y|, |(x_i + b) - y|) + \delta$ and the inequality holds. If the cost is calculated with respect to different peaks for y and $x_m - b$, it must be that $\text{cost}(x_i, x_m - b) = c + |(x_i - b) - (x_m - b)|$ and $\text{cost}(x_i, y) = c + |x_i + b - y|$, because $x_m - b < y$. Since $|(x_i - b) - (x_m - b)| \leq |(x_i + b) - (x_m - b)| \leq |(x_i + b) - y| + \delta$, the inequality holds. Similarly, we can prove that $\text{cost}(x_i, x_m + b) \leq \text{cost}(x_i, y) + (2b - \delta)$ for every agent i . Hence, we can upper bound the cost of Mechanism M1 by $\frac{1}{2} \sum_{i=1}^n \text{cost}(x_i, x_m - b) + \frac{1}{2} \sum_{i=1}^n \text{cost}(x_i, x_m + b) \leq \frac{1}{2} \sum_{i=1}^n (\text{cost}(x_i, y) + \delta) + \frac{1}{2} \sum_{i=1}^n (\text{cost}(x_i, y) + 2b - \delta) = SC_y(\mathbf{x}) + nb = SC_{\text{opt}}(\mathbf{x}) + nb$. The approximation ratio then becomes $1 + \frac{nb}{SC_{\text{opt}}(\mathbf{x})}$, which is at most $1 + \frac{b}{c}$, since $SC_{\text{opt}}(\mathbf{x})$ is at least nc .

For the lower bound, consider the location profile $\mathbf{x} = \langle x_1, \dots, x_n \rangle$ with $x_1 = \dots = x_{k-1} = x_k - b = x_{k+1} - 2b = \dots = x_n - 2b$. Note that the argument works both when $n = 2k$ and when $n = 2k + 1$ because Mechanism M1 selects agent x_k as the median agent in each case. The optimal location is $x_1 + c$ whereas Mechanism M1 equiprobably outputs $f_{M1}(\mathbf{x}) = x_k - b$ or $f_{M1}(\mathbf{x}) = x_k + b$. The cost of the optimal location is $SC_{\text{opt}}(\mathbf{x}) = nc + b$ whereas the cost of Mechanism M1 is $SC_{M1}(\mathbf{x}) = nc + (1/2)(n-1)b + (1/2)(n-1)b = nc + (n-1)b$. The approximation ratio then becomes $\frac{nc+(n-1)b}{nc+b} = 1 + \frac{b}{c} \cdot \frac{n-2}{n+(b/c)}$. As the number of agents grows to infinity, the approximation ratio of the mechanism on this instance approaches $1 + b/c$. This completes the proof. \square

We also consider the maximum cost and prove the approximation ratio of Mechanism M1 as well as a lower bound on the approximation ratio of any truthful-in-expectation mechanism. The results are summarized in Table 1.

4 Deterministic Mechanisms

We now turn our attention to deterministic mechanisms. Here, we provide an overview of our results; see the full version for a more detailed exposition. First, consider the following very simple mechanism.

Mechanism M2. *Given any instance $\mathbf{x} = \langle x_1, \dots, x_n \rangle$, locate the facility always on the left peak of agent 1, i.e. $f(\mathbf{x}) = x_1 - b$.*

We note here that the mechanism $f(\mathbf{x}) = x_n + b$ has exactly the same properties as Mechanism M2 and they are identical up to symmetric profiles. The analysis for the latter mechanism would be exactly the same. For that reason, we will consider them to be the same mechanism.

It is not difficult to see that Mechanism M2 is strategyproof. Mechanisms like Mechanism M2 that output the peaks of agents are known as k -th order statistics. In fact, Mechanism M2 is the only strategyproof k -th order statistic mechanism for double-peaked preferences, unlike the single-peaked preference case, where every k -th order statistic mechanism is (group) strategyproof. Note that Mechanism M2 is anonymous and position invariant. In fact, we obtain the following characterization for two agents.

Theorem 3. *When $n = 2$, the only strategyproof mechanism that satisfies position invariance and anonymity is Mechanism M2.*

The approximation guarantees of Mechanism M2 for the social cost and the maximum cost are summarized in Table 1. What stands out is the non-constant approximation ratio of the mechanism for the social cost. We believe that the inability to guarantee constant ratios is a broader characteristic of strategyproof deterministic mechanisms and not just Mechanism M2. Unfortunately, we were not able to prove that; what we were able to prove however is that no anonymous and position invariant mechanism can do better than Mechanism M1. Given the fact that those properties are usually possessed by mechanisms with good approximation guarantees, we can conjecture that no deterministic mechanism outperforms Mechanism M1. To prove the aforementioned bound, we make use of Theorem 3 in a creative way. First, we prove the following lemma.

Lemma 2. *Let M^n be a strategyproof, anonymous and position invariant mechanism for n agents. Then, for any location profile $\mathbf{x} = \langle x_1 = \dots = x_{n/2}, x_{n/2+1} = \dots = x_n \rangle$, it holds that $M^n(\mathbf{x}) = x_1 - b$.*

Proof. Let M^2 be the following mechanism for two agents: On input location profile $\langle \hat{x}_1, \hat{x}_2 \rangle$, output $M^2(\mathbf{x})$, where $\mathbf{x} = \langle x_1 = \dots = x_{n/2}, x_{n/2+1} = \dots = x_n \rangle$, and $x_1 = \hat{x}_1$ and $x_{n/2+1} = \hat{x}_2$. First, we claim that M^2 is strategyproof, anonymous and position invariant. If that is true, then by Theorem 3, M^2 is Mechanism M2 and the lemma follows.

First let $\hat{\mathbf{x}} = \langle \hat{x}_1, \hat{x}_2 \rangle$, $\hat{\mathbf{x}}' = \langle \hat{x}'_1, \hat{x}'_2 \rangle$ be any two position equivalent location profiles. Observe that the corresponding n -agent profiles \mathbf{x} and \mathbf{x}' obtained by placing $n/2$ agents on \hat{x}_1 and \hat{x}'_1 and $n/2$ agents on \hat{x}_2 and \hat{x}'_2 respectively are also position equivalent. Since M^n is position invariant, it must hold that $M^n(\mathbf{x}) = M^n(\mathbf{x}')$ and hence by construction of M^2 , $M^2(\hat{\mathbf{x}}) = M^2(\hat{\mathbf{x}}')$. Since $\hat{\mathbf{x}}$ and $\hat{\mathbf{x}}'$ were arbitrary, Mechanism M^2 is position invariant.

Similarly, let $\hat{\mathbf{x}} = \langle \hat{x}_1, \hat{x}_2 \rangle$, $\hat{\mathbf{x}}' = \langle \hat{x}'_1, \hat{x}'_2 \rangle$ be any two location profiles, such that $\hat{\mathbf{x}}'$ is obtained by $\hat{\mathbf{x}}$ by a permutation of the agents. The outcome of M^n on the corresponding n -agent location profiles (since the number of agents placed on \hat{x}_1 and \hat{x}_2 is the same) is the same and by construction of M^2 , $M^2(\hat{\mathbf{x}}) = M^2(\hat{\mathbf{x}}')$ and since the profiles were arbitrary, the mechanism is anonymous.

Finally, for strategyproofness, start with a location profile $\hat{\mathbf{x}}' = \langle \hat{x}'_1, \hat{x}'_2 \rangle$ and let $\mathbf{x}' = \langle x'_1 = \dots = x'_{n/2}, x'_{n/2+1} = \dots = x'_n \rangle$ be the corresponding n -agent location profile. Let $y = M^n(\mathbf{x}')$ and let $\text{cost}(x'_i, y)$ be the cost of agents $x'_1, \dots, x'_{n/2}$ on \mathbf{x}' . For any x_1 , let $\langle x_1, x'_2 = \dots = x'_{n/2}, x'_{n/2+1} = \dots = x'_n \rangle$ be the resulting location profile. By strategyproofness of M^n , agent x'_1 can not decrease her cost by misreporting x_1 on profile \mathbf{x}' and hence her cost on the new profile is at least $\text{cost}(x'_1, y)$. Next, consider the location profile $\langle x_1 = x_2, x'_3 = \dots = x'_{n/2}, x'_{n/2+1} = \dots = x'_n \rangle$ and observe that by the same argument, the cost of agent x'_2 is not smaller on the new profile when compared to $\langle x_1, x'_2 = \dots = x'_{n/2}, x'_{n/2+1} = \dots = x'_n \rangle$ and hence

Table 1: Summary of our results. The lower bounds for deterministic mechanisms hold for anonymous and position invariant strategyproof mechanisms. For (*), an additional lower bound of 2 holds under no conditions. Fields indicated by (-) are not proven yet. For the maximum cost, the approximation ratios are actually $\max\{1 + 2b/c, 3\}$ and $\max\{1 + b/c, 2\}$ respectively. The results for single-peaked preferences are also noted for comparison.

| | Double-peaked | | Single-peaked | |
|---------------------|----------------------|--------------------|----------------------|-------|
| | Ratio | Lower | Ratio | Lower |
| Social cost | | | | |
| Deterministic | $\Theta(n)$ | $1 + \frac{b}{c}$ | 1 | 1 |
| Randomized | $1 + \frac{b}{c}$ | - | 1 | 1 |
| Maximum cost | | | | |
| Deterministic* | $1 + \frac{2b}{c}$ | $1 + \frac{2b}{c}$ | 2 | 2 |
| Randomized | $1 + \frac{b}{c}$ | $3/2$ | $3/2$ | $3/2$ |

her cost is at least $\text{cost}(x', y)$. Continuing like this, we obtain the profile $\langle x_1 = \dots = x_{n/2}, x'_{n/2+1} = \dots = x'_n \rangle$ and by the same argument, the cost of agent $x'_{n/2}$ on this profile is at least $\text{cost}(x', y)$. The location profile $\langle x_1 = \dots = x_{n/2}, x'_{n/2+1} = \dots = x'_n \rangle$ corresponds to the 2-agent location profile $\hat{\mathbf{x}} = \langle \hat{x}_1, \hat{x}_2 \rangle$ and by construction of M^2 , $\text{cost}(\hat{x}'_1, M^2(\hat{\mathbf{x}}')) \leq \text{cost}(\hat{x}'_1, M^2(\hat{\mathbf{x}}))$ and since the choice of x_1 (and hence the choice of \hat{x}_1) was arbitrary, Mechanism M^2 is strategyproof. \square

Theorem 4. *Any strategyproof mechanism that satisfies position invariance and anonymity achieves an approximation ratio of at least $1 + \frac{b}{c}$ for the social cost.*

Proof. Let M^n be a strategyproof, anonymous and position invariant mechanism and consider any location profile $\mathbf{x} = \langle x_1 = \dots = x_{n/2}, x_{n/2} + 1 = \dots = x_n \text{ with } x_{n/2+1} = x_1 + 2b \rangle$. By Lemma 2, $M^n(\mathbf{x}) = x_1 - b$ and the social cost of M^n is $nc + (n/2)2b$ while the social cost of the optimal allocation is only nc . The lower bound follows. \square

Note that the theorem requires the number of agents to be even; we would like to prove a similar bound when n is odd. We conclude the section with our impossibility result about group strategyproof mechanisms. Recall that for the single-peaked preference case, k -th order statistic mechanisms are group strategyproof. This is not the case for double-peaked preferences.

Theorem 5. *There exists no deterministic, anonymous, position invariant and group strategyproof mechanism.*

5 Generalizations and conclusion

As argued in the introduction, double-peaked preferences are often a very realistic model for facility location and our results initiate the discussion on such settings and shed some light on the capabilities and limitations of strategyproof

Table 2: The results for the case when peaks are not required to be symmetric.

| | Non-symmetric | |
|---------------------|----------------------|-------------------------|
| | Ratio | Lower |
| Social cost | | |
| Deterministic | $\Theta(n)$ | $1 + \frac{b_1+b_2}{c}$ |
| Randomized | $\Theta(n)$ | - |
| Maximum cost | | |
| Deterministic | $\Theta(n)$ | $1 + \frac{b_1+b_2}{c}$ |
| Randomized | $\Theta(n)$ | $3/2$ |

mechanisms. We conclude with a discussion about an extension to the main model and some potential future directions.

Non-symmetric peaks

Although the symmetric case is arguably the best analogue of the single-peaked preference setting, it could certainly make sense to consider a more general model, where the cost functions do not have the same slope in every interval and hence the peaks are not equidistant from the location of an agent. Let b_1 and b_2 be the distances from the left and the right peaks respectively. Clearly, all our lower bounds still hold, although one could potentially prove even stronger bounds by taking advantage of the more general setting. The main observation is that Mechanism M1 is no longer truthful-in-expectation, because its truthfulness depends heavily on the peaks being equidistant. On the other hand, mechanism M2 is still strategyproof and the approximation ratio bounds extend naturally. A summary of the results for the non-symmetric setting is depicted in Table 2.

Future work

Starting from randomized mechanisms, we would like to obtain lower bounds that are functions of b and c , to see how well Mechanism M1 fares in the general setting. For deterministic mechanisms, we would like to get a result that would clear up the picture. Characterizing strategyproof, anonymous and position invariant mechanisms would be ideal, but proving a lower bound that depends on n on the ratio of such mechanisms (for the social cost) would also be quite helpful. The techniques used in our characterization for two agents and our lower bounds (available at the full version) seem to convey promising intuition for achieving such a task. Finally, it would be interesting to see if we can come up with a “good” randomized truthful-in-expectation mechanism for the extended model, when peaks are not assumed to be symmetric.

Acknowledgements

Aris Filos-Ratsikas and Jie Zhang were supported by the Sino-Danish Center for the Theory of Interactive Computation, funded by the Danish National Research Foundation and the National Science Foundation of China (under the grant 61061130540), and by the Center for research

in the Foundations of Electronic Markets (CFEM), supported by the Danish Strategic Research Council. Jie Zhang was also supported by ERC Advanced Grant 321171 (AL-GAME). Minming Li was partly supported by a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China [Project No. CityU 117913]. Qiang Zhang was supported by FET IP project MULTIPLEX 317532.

References

- Alon, N.; Feldman, M.; Procaccia, A. D.; and Tennenholtz, M. 2010. Strategyproof approximation of the minimax on networks. *Mathematics of Operations Research* 35(3):513–526.
- Ashlagi, I.; Fischer, F.; Kash, I.; and Procaccia, A. D. 2010. Mix and match. In *Proceedings of the 11th ACM conference on Electronic commerce (ACM-EC)*, 305–314. ACM.
- Black, D. 1957 (reprint at 1986). *The theory of committees and elections*. Kluwer Academic Publishers.
- Caragiannis, I.; Filos-Ratsikas, A.; and Procaccia, A. D. 2011. An improved 2-agent kidney exchange mechanism. In *Proceedings of the 7th Workshop of Internet and Network Economics (WINE)*, 37–48. Springer.
- Cooter, R. D. 2002. *The strategic constitution*. Princeton University Press.
- Dokow, E.; Feldman, M.; Meir, R.; and Nehama, I. 2012. Mechanism design on discrete lines and cycles. In *Proceedings of the 13th ACM Conference on Electronic Commerce (ACM-EC)*, 423–440.
- Dughmi, S., and Gosh, A. 2010. Truthful assignment without money. In *Proceedings of the 11th ACM conference on Electronic commerce (ACM-EC)*, 325–334.
- Egan, P. J. 2013. "Do something" politics and double-peaked policy preferences. *Journal of Politics* 76(2):333–349.
- Escoffier, B.; Gourvès, L.; Kim Thang, N.; Pascual, F.; and Spanjaard, O. 2011. Strategy-proof mechanisms for facility location games with many facilities. *Algorithmic Decision Theory* 67–81.
- Feigenbaum, I.; Sethuraman, J.; and Ye, C. 2013. Approximately optimal mechanisms for strategyproof facility location: Minimizing l_p norm of costs. *arXiv preprint arXiv:1305.2446*.
- Feldman, M., and Wilf, Y. 2013. Strategyproof facility location and the least squares objective. In *Proceedings of the 14th ACM conference on Electronic Commerce (ACM-EC)*, 873–890.
- Fotakis, D., and Tzamos, C. 2010. Winner-imposing strategyproof mechanisms for multiple facility location games. In *Proceeding of the 5th International Workshop of Internet and Network Economics (WINE)*, 234–245.
- Fotakis, D., and Tzamos, C. 2012. On the power of deterministic mechanisms for facility location games. *arXiv preprint arXiv:1207.0935*.
- Guo, M., and Conitzer, V. 2010. Strategy-proof allocation of multiple items between two agents without payments or priors. In *Ninth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS)*, volume 10, 881–888.
- Lu, P.; Sun, X.; Wang, Y.; and Zhu, Z. A. 2010. Asymptotically optimal strategy-proof mechanisms for two-facility games. In *Proceedings of the 11th ACM Conference on Electronic Commerce (ACM-EC)*, 315–324.
- Lu, P.; Wang, Y.; and Zhou, Y. 2009. Tighter bounds for facility games. In *Proceeding of the 5th International Workshop of Internet and Network Economics (WINE)*, 137–148.
- Moulin, H. 1980. On strategy-proofness and single peakedness. *Public Choice* 35(4):437–455.
- Procaccia, A. D., and Tennenholtz, M. 2009. Approximate mechanism design without money. In *Proceedings of the 10th ACM Conference on Electronic Commerce (ACM-EC)*, 177–186.
- Rosen, H. S. 2004. *Public Finance*. Springer. chapter 6 - Political Economy, 115–117.
- Schummer, J., and Vohra, R. V. 2002. Strategy-proof location on a network. *Journal of Economic Theory* 104(2):405–428.