

Facility Location with Double-peaked Preferences

Aris Filos-Ratsikas
University of Oxford, UK
Aris.Filos-Ratsikas@cs.ox.ac.uk

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

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

Qiang Zhang
University of Warsaw
csqzhang@gmail.com

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 truthful mechanisms, there is no deterministic mechanism that outperforms our truthful-in-expectation mechanism. In order to obtain this result, we first characterize mechanisms for two agents that satisfy two simple properties; we use the same characterization to prove that no mechanism in this class can be *group-strategyproof*.

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 [18, 20, 16, 15, 1, 12, 9, 13, 6, 11]. 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 [18]), 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 [7, 14, 2, 4], 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.

1.1 Double-peaked preferences in practice

Single-peaked preferences were introduced in [3] as a way to avoid *Condorcet cycles* in majority elections. Moulin [17] 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 [5]), 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 [8] 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 assume 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 [19]. 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.

1.2 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.

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 $SC_{f(\mathbf{x})}(\mathbf{x})$ 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 [10] 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

3.1 Social cost

Next, we will calculate the approximation ratio of the mechanism for the social cost. 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 $x_m = \text{median}(x_1, \dots, x_n)$, 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 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 $f_{\text{M1}}(\mathbf{x}) = x_k - b$ or $f_{\text{M1}}(\mathbf{x}) = x_k + b$. The cost of the optimal location is $SC_{\text{opt}}(\mathbf{x}) = nc + b$ whereas the cost of Mechanism M1 is $SC_{\text{M1}}(\mathbf{x}) = nc + (1/2)(n-1)b + (1/2)(n-1)b = nc + (n-1)b$. 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

3.2 Maximum cost

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.

Theorem 3. *Mechanism M1 has an approximation ratio of $\max\{1+b/c, 2\}$ for the maximum cost.*

Proof. Let $\mathbf{x} = \langle x_1, \dots, x_n \rangle$ be an arbitrary instance and let x_m be the median agent. We will consider two cases, based on the location of $f_{\text{opt}}(\mathbf{x})$ with respect to $x_m - b$ (or symmetrically $x_m + b$).

Case 1: $f_{\text{opt}}(\mathbf{x}) < x_m - b$ (or $f_{\text{opt}}(\mathbf{x}) > x_m + b$).

Let $\delta = (x_m - b) - f_{\text{opt}}(\mathbf{x})$. For the same reason as in the proof of Theorem 2, for every agent i , it holds that $\text{cost}(x_i, x_m - b) \leq \text{cost}(x_i, f_{\text{opt}}(\mathbf{x})) + \delta$ and also that $\text{cost}(x_i, x_m + b) \leq \text{cost}(x_i, f_{\text{opt}}(\mathbf{x})) + (2b + \delta)$.

The maximum cost of Mechanism M1 is

$$\begin{aligned} MC_{\text{M1}}(\mathbf{x}) &= \frac{1}{2} \max_{i \in N} \text{cost}(x_i, x_m - b) + \frac{1}{2} \max_{i \in N} \text{cost}(x_i, x_m + b) \\ &\leq \frac{1}{2} \left(\max_{i \in N} \text{cost}(x_i, f_{\text{opt}}(\mathbf{x})) + \delta \right) + \frac{1}{2} \left(\max_{i \in N} \text{cost}(x_i, f_{\text{opt}}(\mathbf{x})) + (2b + \delta) \right) \\ &= MC_{\text{opt}}(\mathbf{x}) + b + \delta \leq 2MC_{\text{opt}}(\mathbf{x}) + b - c \end{aligned}$$

since $MC_{\text{opt}}(\mathbf{x}) \geq c + \delta$. The approximation is at most $2 + \frac{b-c}{MC_{\text{opt}}(\mathbf{x})}$, which is at most $1 + \frac{b}{c}$ if $b \geq c$ (since $MC_{\text{opt}}(\mathbf{x}) \geq c$) and at most 2 if $b < c$ (since $MC_{\text{opt}}(\mathbf{x}) > 0$).

Case 2: $x_m - c \leq f_{\text{opt}}(\mathbf{x}) \leq x_m + b$.

Now, let $\delta = f_{\text{opt}}(\mathbf{x}) - (x_m - b)$. Again, it holds that $\text{cost}(x_i, x_m - b) \leq \text{cost}(x_i, f_{\text{opt}}(\mathbf{x})) + \delta$ and also that $\text{cost}(x_i, x_m + b) \leq \text{cost}(x_i, f_{\text{opt}}(\mathbf{x})) + (2b - \delta)$.

The maximum cost of Mechanism M1 is

$$\begin{aligned} MC_{\text{M1}}(\mathbf{x}) &= \frac{1}{2} \max_{i \in N} \text{cost}(x_i, x_m - b) + \frac{1}{2} \max_{i \in N} \text{cost}(x_i, x_m + b) \\ &\leq \frac{1}{2} \left(\max_{i \in N} \text{cost}(x_i, f_{\text{opt}}(\mathbf{x})) + \delta \right) + \frac{1}{2} \left(\max_{i \in N} \text{cost}(x_i, f_{\text{opt}}(\mathbf{x})) + (2b - \delta) \right) \\ &= MC_{\text{opt}}(\mathbf{x}) + b \end{aligned}$$

and the approximation ratio is at most $1 + \frac{b}{c}$ (since $MC_{\text{opt}}(\mathbf{x}) \geq c$).

For the matching lower bound, consider an instance $\mathbf{x} = \langle x_1, \dots, x_n \rangle$ on which $x_1 + b < x_2 - b$ and $x_i = x_2$ for all $i \notin \{1, 2\}$. It is $f_{\text{opt}}(\mathbf{x}) = \frac{x_1 + x_2}{2}$, i.e. the middle of the interval between x_1 and x_2 , whereas Mechanism M1 selects equiprobably among $x_2 - b$ and $x_2 + b$. Let $d = f_{\text{opt}}(\mathbf{x}) - (x_1 + b)$. Then $MC_{\text{opt}}(\mathbf{x}) = c + d$, whereas $MC_{\text{M1}}(\mathbf{x}) = c + \frac{1}{2}2d + \frac{1}{2}(2d + 2b) = b + c + 2d$. The approximation ratio is $1 + \frac{b+d}{c+d}$ which is $1 + \frac{b}{c}$ as d goes to 0 and 2 as d goes to infinity. \square \square

Next, we provide a lower bound on the approximation ratio of any truthful-in-expectation mechanism.

Theorem 4. *Any truthful-in-expectation mechanism has an approximation ratio of at least $\frac{3}{2}$ for the maximum cost.*

First, we state a couple of lemmas which are in essence very similar to those used in the proof of the single-peaked preferences case in [18]. Let $\mathbf{x} = \langle x_1, x_2 \rangle$ be an instance such that $x_1 + b < x_2 - b$ and let $\lambda = (x_2 - b) - (x_1 + b)$. Let f be a truthful-in-expectation mechanism and let \mathcal{D} be the distribution that $y = f(\mathbf{x})$ follows on instance \mathbf{x} .

Lemma 2. *On instance \mathbf{x} , at least one of $\mathbb{E}_{y \sim \mathcal{D}} [\text{cost}(x_1, y)] \geq c + \frac{\lambda}{2}$ and $\mathbb{E}_{y \sim \mathcal{D}} [\text{cost}(x_2, y)] \geq c + \frac{\lambda}{2}$ holds.*

Proof. Obviously, $\text{cost}(x_1, y) + \text{cost}(x_2, y) \geq 2c + \lambda$ for any choice of y , hence $\mathbb{E}_{y \sim \mathcal{D}} \left[\sum_{i=1}^2 \text{cost}(x_i, y) \right] = \sum_{i=1}^2 \mathbb{E}_{y \sim \mathcal{D}} [\text{cost}(x_i, y)] \geq 2c + \lambda$. Therefore, it must be that $\mathbb{E}_{y \sim \mathcal{D}} [\text{cost}(x_i, y)] \geq c + \frac{\lambda}{2}$ for at least one of $i = 1$ or $i = 2$. \square

Lemma 3. *Let $f_{\text{opt}}(\mathbf{x})$ be the outcome of the optimal mechanism on instance \mathbf{x} . If $\mathbb{E}_{y \sim \mathcal{D}} [|y - f_{\text{opt}}(\mathbf{x})|] = \Delta$, then the maximum cost of the mechanism on this instance is $\mathbb{E}[MC_f(\mathbf{x})] = c + \frac{\lambda}{2} + \Delta$.*

Proof. Since $\forall y, MC_f(\mathbf{x}) = c + |y - f_{\text{opt}}(\mathbf{x})| + \lambda/2$, it holds

$$\mathbb{E}_{y \sim \mathcal{D}} [MC_f(\mathbf{x})] = \mathbb{E}_{y \sim \mathcal{D}} \left[c + \frac{\lambda}{2} + |y - f_{\text{opt}}(\mathbf{x})| \right] = c + \frac{\lambda}{2} + \mathbb{E}_{y \sim \mathcal{D}} [|y - f_{\text{opt}}(\mathbf{x})|] = c + \frac{\lambda}{2} + \Delta.$$

\square \square

We can now prove the theorem.

Proof. Consider an instance with two agents on which $x_1 + b < x_2 - b$ and $(x_2 - b) - (x_1 + b) = \lambda$. It holds that $f_{\text{opt}}(\mathbf{x}) = \frac{x_1 + x_2}{2}$. Assume there is a truthful-in-expectation mechanism M on which $y = f(x_1, x_2)$ follows a distribution \mathcal{D} on this instance. According to Lemma 2, at least one of $\mathbb{E}_{y \sim \mathcal{D}} [\text{cost}(y, x_1)] \geq c + \lambda/2$ and $\mathbb{E}_{y \sim \mathcal{D}} [\text{cost}(y, x_2)] \geq c + \lambda/2$ holds. W.l.o.g., assume the second inequality is true (if the first inequality is true then we can make a symmetric argument with agent x_1 deviating).

Next, consider the instance $\mathbf{x}' = \langle x'_1, x'_2 \rangle$ with $x'_1 = x_1$ and $x'_2 = x_2 + \lambda$. Let $f_{\text{opt}}(\mathbf{x}') = (x'_1 + x'_2)/2 = x_2 - b$. Let \mathcal{D}' be the distribution that y' follows on instance \mathbf{x}' . By strategyproofness, $\mathbb{E}_{y' \sim \mathcal{D}'} [\text{cost}(y', x_2)] \geq \mathbb{E}_{y \sim \mathcal{D}} [\text{cost}(y, x_2)] \geq c + \lambda/2$, since x'_2 could be a deviation of agent 2 on instance \mathbf{x} . This implies $\mathbb{E}_{y' \sim \mathcal{D}'} [|y' - (x_2 - b)|] \geq \frac{\lambda}{2}$. To see this, assume otherwise for contradiction, we would have

$$\begin{aligned} \mathbb{E}_{y' \sim \mathcal{D}'} [\text{cost}(y', x_2)] &= \mathbb{E}_{y' \sim \mathcal{D}'} [c + \min\{|y' - (x_2 - b)|, |y' - (x_2 + b)|\}] \\ &\leq c + \min\{\mathbb{E}_{y' \sim \mathcal{D}'} [|y' - (x_2 - b)|], \mathbb{E}_{y' \sim \mathcal{D}'} [|y' - (x_2 + b)|]\} \\ &< c + \frac{\lambda}{2}. \end{aligned}$$

Equivalently, we have $\mathbb{E}_{y' \sim \mathcal{D}'} [|y' - f_{\text{opt}}(\mathbf{x}')|] \geq \frac{\lambda}{2}$. Hence, by applying Lemma 3, we know that the maximum cost of the mechanism on the second instance is $\mathbb{E}_{y' \sim \mathcal{D}'} [MC_M(\mathbf{x}')] \geq c + \lambda + \frac{\lambda}{2} = c + \frac{3\lambda}{2}$. Since the optimal mechanism locates the facility on $f_{\text{opt}}(\mathbf{x}')$, its maximum cost is $c + \lambda$. Therefore, the approximation ratio is at least $\frac{c + 3\lambda/2}{c + \lambda}$. As λ grows to infinity, the approximation ratio approaches $\frac{3}{2}$. To generalize the proof to more than two agents, place every other agent on $f_{\text{opt}}(\mathbf{x}) + b$ on instance \mathbf{x} . Since λ is large enough, the maximum cost is still calculated with respect to x_2 and all the arguments still hold. \square \square

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	$1 + \frac{b}{c}$	$3/2$	$3/2$	$3/2$

4 Deterministic Mechanisms

We now turn our attention to deterministic mechanisms. We will start by stating and proving the following lemma, which will be very useful throughout the paper. Variations of the instances used here will appear in several of our proofs.

When $n = 2$, we call the following instance² the *primary instance*.

Primary instance: $\mathbf{x} = \langle x_1, x_2 \rangle$ with $x_1 + 2b + \epsilon = x_2 - b$, where ϵ is an arbitrarily small positive quantity.

Lemma 4. *On the primary instance, there is no anonymous, position invariant and strategyproof mechanism such that $f(\mathbf{x}) \in [x_1 + b, x_2 - b]$.*

Proof. For contradiction, suppose there exists an anonymous, position invariant strategyproof mechanism M that outputs $f(\mathbf{x}) \in [x_1 + b, x_2 - b]$. Let's denote $\delta_1 = f(\mathbf{x}) - (x_1 + b)$, $\delta_2 = (x_2 - b) - f(\mathbf{x})$. Throughout the proof we start with the primary instance and construct some other instances to prove the lemma. There are 2 cases to be considered.

Case 1: $0 < \delta_1 \leq \frac{1}{2}(b + \epsilon)$, i.e., $x_1 + b < f(\mathbf{x}) \leq \frac{x_1 + x_2}{2}$ (symmetrically, for the case $0 < \delta_2 \leq \frac{1}{2}(b + \epsilon)$). See Figure 1.

- Instance I : $\mathbf{x}^I = \langle x_1^I, x_2^I \rangle$, where $x_1^I = x_1 + \delta_1, x_2^I = x_2$.
- Instance II : $\mathbf{x}^{II} = \langle x_1^{II}, x_2^{II} \rangle$, where $x_1^{II} = x_1 + \delta_1, x_2^{II} = x_2 + \delta_1$.

²We note that this is actually a family of different instances, for each possible value of x_1 on the real line. Any one of those instances can be used for our purposes and hence we use the term loosely.




Figure 1: Case 1 of Lemma 4. On instance I, agent x_2^I can report x_2^{II} and decrease her cost. Instances II and the Primary instance are position equivalent.

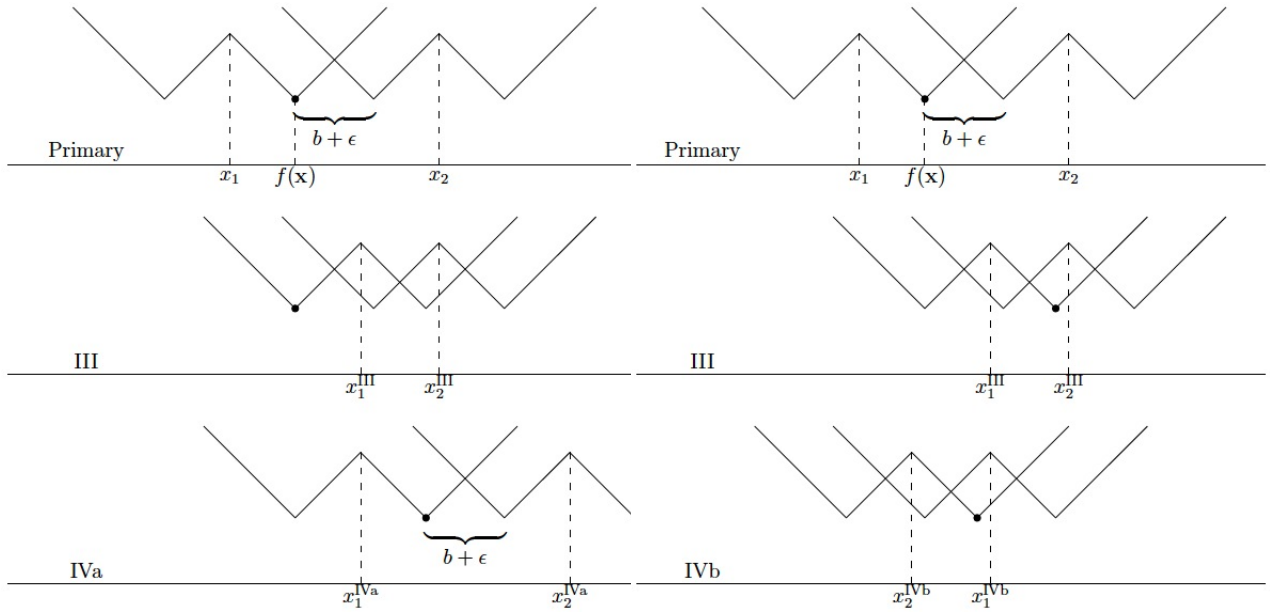


Figure 2: Case 2(a) and 2(b) of Lemma 4. If $f(x_1^{III}, x_2^{III}) = x_1^{III} - b$ on instance III, agent x_2^{III} can misreport x_2^{IVa} and decrease her cost, since instances IVa and the Primary instance are position equivalent. If $f(x_1^{III}, x_2^{III}) = x_1^{III} + b$ on instance III, then on the Primary instance, agent x_2 can misreport x_2^{IVb} and decrease her cost. By anonymity, instances III and IVb are position equivalent.

First, we know that on instance I, strategyproofness requires that $f(\mathbf{x}^I) = x_1^I - b$ or $x_1^I + b$, otherwise agent x_1^I could misreport x_1 and move the facility to position $x_1^I + b$, minimizing her cost. Second, since instance II and the primary instance are position equivalent (II is a shifted version of the primary instance to the right by δ_1), by position invariance we know $f(\mathbf{x}^{II}) = x_1^{II} + b + \delta_1 = x_1^I + b + \delta_1$. Now let's consider agent x_2^I in instance I. If agent x_2^I misreports x_2^I , she can move the facility from $x_1^I - b$ or $x_1^I + b$ to $x_1^I + b + \delta_1 \leq x_2^I - b$; in either case the facility is closer to her left peak. Therefore x_2^I can manipulate the mechanism, which violates strategyproofness. Hence $f(\mathbf{x}) \notin (x_1 + b, \frac{x_1 + x_2}{2}]$.

Case 2: $\delta_1 = 0$, i.e., $f(\mathbf{x}) = x_1 + b$ (symmetrically, for the case $\delta_2 = 0$). See Figure 2.

- Instance III : $\mathbf{x}^{III} = \langle x_1^{III}, x_2^{III} \rangle$, where $x_1^{III} = x_1 + 2b, x_2^{III} = x_2$.

First, we know that on instance III, strategyproofness requires that $f(\mathbf{x}^{III}) = x_1^{III} - b$ or $x_1^{III} + b$, otherwise agent x_1^{III} could misreport x_1 and move the facility to position $x_1^{III} - b$, minimizing her cost. Now, there are two subcases.

- (a) $f(\mathbf{x}^{III}) = x_1^{III} - b$. See left hand side of Figure 2, where we have instance IVa: $x_1^{IVa} = x_1 + 2b, x_2^{IVa} = x_2 + 2b$.

Obviously, instance IVa is position equivalent to the primary instance, so $f(x_1^{IVa}, x_2^{IVa}) = x_1^{IVa} + b$. Note that the cost of agent x_2^{IVa} is $b + c + \epsilon$. If agent x_2^{IVa} misreports x_2^{IVa} , then her cost becomes $b + c - \epsilon$, which is smaller than when reporting truthfully. So $f(\mathbf{x}^{III}) \neq x_1^{III} - b$ if Mechanism M is strategyproof.

- (b) $f(\mathbf{x}^{III}) = x_1^{III} + b$. See right hand side of Figure 2, where we have instance IVb: $x_1^{IVb} = x_1^{III}, x_2^{IVb} = x_1^{III} - (b + \epsilon)$.

Obviously, instance IVb is position equivalent to instance III, so $f(x_1^{IVb}, x_2^{IVb}) = x_2^{IVb} + b$ as implied by position invariance and anonymity. Note that on instance III, the cost of agent x_2^{III} is $b + c - \epsilon$. If agent x_2^{III} misreports x_2^{IVb} , then her cost becomes $c + 2\epsilon$, which is smaller than when reporting truthfully (since ϵ is arbitrarily small). So $f(\mathbf{x}^{III}) \neq x_1^{III} + b$ if Mechanism M is strategyproof.

Hence, in order for mechanism M to be strategyproof, it must be $f(\mathbf{x}) \neq x_1 + b$.

In all, there is no anonymous, position invariant and strategyproof mechanism such that $f(\mathbf{x}) \in [x_1 + b, x_2 - b]$. □ □

4.1 Group strategyproofness

As we mentioned in the introduction, under the reasonable conditions of position invariance and anonymity, there is no group strategyproof mechanism for the problem. We will prove this claim by using Lemma 4 and the following lemma.

When $n = 2k + 1, k \in \mathbb{Z}^+$, let P be the instance obtained by locating $k + 1$ agents on x_1 and k agents on x_2 on the primary instance. Similarly, let S be the instance obtained by locating k agents on x_1 and $k + 1$ agents on x_2 on the primary instance. Formally, let $\mathbf{x}^P = \langle x_1^P, \dots, x_n^P \rangle$, where $x_1^P = \dots = x_{k+1}^P, x_{k+2}^P = \dots = x_n^P$, and $(x_n^P - b) - (x_1^P + b) = b + \epsilon$ and $\mathbf{x}^S = \langle x_1^S, \dots, x_n^S \rangle$, where $x_1^S = \dots = x_k^S, x_{k+1}^S = \dots = x_n^S$, and $(x_n^S - b) - (x_1^S + b) = b + \epsilon$, where ϵ is the same quantity as in the primary instance.

Lemma 5. *When $n = 2k + 1$, any position invariant and group strategyproof mechanism that outputs $f(\mathbf{x}^P) = x_1^P + b$ on instance P, must output $f(\mathbf{x}) = x_1 + b$ on any instance $\mathbf{x} = \langle x_1, \dots, x_n \rangle$, where $x_1 = \dots = x_{k+1}$, $x_{k+2} = \dots = x_n$ and $(x_n - b) - (x_1 + b) = 2b$. Similarly, any position invariant and group strategyproof mechanism that outputs $f(\mathbf{x}^S) = x_n^S - b$ on instance S, must output $f(\mathbf{x}) = x_n - b$ on any instance $\mathbf{x} = \langle x_1, \dots, x_n \rangle$, where $x_1 = \dots = x_k$, $x_{k+1} = \dots = x_n$ and $(x_n - b) - (x_1 + b) = 2b$.*

Proof. We prove the first part of the lemma. The proof of the second part is symmetric. Note that the difference between instances \mathbf{x}^P and \mathbf{x} is that the distance between the two groups of agents is $b + \epsilon$ in \mathbf{x}^P while it is $2b$ in \mathbf{x} . First, we argue that $f(\mathbf{x}) \in [x_1 + b, x_n - b]$. Indeed, if that was not the case, if $f(\mathbf{x}) < x_1 + b$, by the onto condition implied by position invariance, all agents could jointly misreport some different positions and move the facility to $x_1 + b$. This point admits a smaller cost for all agents; specifically the cost of agents x_1, \dots, x_{k+1} is minimized while the cost of agents x_{k+2}, \dots, x_n is reduced and group strategyproofness is violated. Using a symmetric argument, we conclude that it can't be $f(\mathbf{x}) > x_n - b$ either.

Secondly, we argue that $f(\mathbf{x}) \notin (x_1 + b, x_n - b]$. Indeed, assume that was not the case. Then agents x_{k+2}^P, \dots, x_n^P on instance \mathbf{x}^P could jointly misreport x_{k+2}, \dots, x_n and move the facility from $f(\mathbf{x}^P) = x_1^P + b$ to $f(\mathbf{x})$. Since by assumption $f(\mathbf{x}) \in (x_1 + b, x_n - b]$, the cost of each deviating agent is smaller than her cost before deviating. Group strategyproofness is then violated and hence it must be that $f(\mathbf{x}) = x_1 + b$. By position invariance, it must be that $f(\bar{\mathbf{x}}) = \bar{x}_1 + b$ on any instance $\bar{\mathbf{x}}$ which is position equivalent to instance \mathbf{x} . \square \square

Theorem 5. *There is no group strategyproof mechanism that is anonymous and position invariant.*

Proof. When $n = 2$, on the primary instance, according to Lemma 4, there is no anonymous, position invariant and strategyproof mechanism such that $f(\mathbf{x}) \in [x_1 + b, x_2 - b]$. In addition, if the facility was placed on some point $f(\mathbf{x}) < x_1 + b$ (the argument for $f(\mathbf{x}) > x_2 - b$ is symmetric), for any mechanism that satisfies position invariance which implies onto, agents 1 and 2 could jointly misreport some positions \hat{x}_1 and \hat{x}_2 such that $f(\hat{x}_1, \hat{x}_2) = x_1 + b$. Obviously, this deviation admits the minimum possible cost for agent 1 and a reduced cost for agent 2, violating group strategyproofness.

The proof can easily be extended to the case when n is even. On the primary instance, simply place $\frac{n}{2}$ agents on x_1 and $\frac{n}{2}$ agents on x_2 . By considering deviations of coalitions of agents coinciding on x_1 or x_2 instead of deviations of agents x_1 and x_2 respectively, all the arguments still hold. However, additional care must be taken when n is odd.

When $n = 2k + 1$, $k \in \mathbb{Z}^+$, we denote by P_J the instance after placing $k + 1$ agents on x_1^J and k agents on x_2^J on instance J, where J is either instance I, II, III, IVa or IVb of the proof of Lemma 4. Similarly, let S_J be the instance after placing k agents on x_1^J and $k + 1$ agents on x_2^J on instance J. Finally, let X_1^K be the group of agents i for which $x_i^K = x_1^K$ on instance K, where K is either $P, P_I, P_{II}, P_{III}, P_{IVa}, P_{IVb}, S, S_I, S_{II}, S_{III}, S_{IVa}$ or S_{IVb} , and let X_2^K be the group of agents i for which $x_i^K = x_n^K$ on instance K.

For contradiction, assume that there exists an anonymous, position invariant, and group strategyproof mechanism. On instance P, by group strategyproofness, it must be $f(\mathbf{x}^P) \in [x_1^P + b, x_n^P - b]$ and by the same arguments used in case 1 of Lemma 4 (with instances P_I, P_{II} instead of I, II and

with X_1^P and X_2^P instead of x_1 and x_2)³, it must be $f(\mathbf{x}^P) \notin (x_1^P + b, x_n^P - b)$. Hence, it must be that either $f(\mathbf{x}^P) = x_1^P + b$ or $f(\mathbf{x}^P) = x_n^P - b$. Assume w.l.o.g. that $f(\mathbf{x}^P) = x_1^P + b$; the other case can be handled symmetrically.

Following the arguments of case 2 of Lemma 4 (using instance P_{III} instead of III), group strategyproofness and position invariance imply that $f(\mathbf{x}^{P_{\text{III}}}) = x_1^{P_{\text{III}}} - b$ or $f(\mathbf{x}^{P_{\text{III}}}) = x_1^{P_{\text{III}}} + b$ on instance P_{III} . By the arguments of subcase (a) (using instance P_{IVa} instead of IVa), it can't be that $f(\mathbf{x}^{P_{\text{III}}}) = x_1^{P_{\text{III}}} - b$, so it must be that $f(\mathbf{x}^{P_{\text{III}}}) = x_1^{P_{\text{III}}} + b$. However, we can not simply apply the argument used in subcase (b) to get a contradiction, because since there is a different number of agents on x_1^P and x_2^P , instances P and P_{IVa} are no longer position equivalent.

Next, consider instance S and observe that $f(\mathbf{x}^S) \notin (x_1^S + b, x_n^S - b)$ by the same arguments as above and $f(\mathbf{x}^S) \in [x_1^S + b, x_n^S - b]$ by group strategyproofness and position invariance. Hence it is either $f(\mathbf{x}^S) = x_1^S + b$ or $f(\mathbf{x}^S) = x_n^S - b$. Assume first that $f(\mathbf{x}^S) = x_1^S + b$. By the same arguments as above (using instances S_{III} and S_{IVa}), on instance S_{III} , it must be that $f(\mathbf{x}^{S_{\text{III}}}) = x_1^{S_{\text{III}}} + b$. Now, observe that if $X_2^{P_{\text{III}}}$ misreport $\bar{x}_i^{P_{\text{III}}} = x_i^{P_{\text{III}}} - 2b - 2\epsilon$, then we get instance P_{IVb} which is position equivalent to instance S_{III} and hence it must be that $f(\mathbf{x}^{P_{\text{IVb}}}) = x_n^{P_{\text{IVb}}} + b$. The cost of $X_2^{P_{\text{III}}}$ before misreporting was $b + c - \epsilon$ while it becomes $c + 2\epsilon$ after misreporting. This violates strategyproofness, which means that on instance S , it must be $f(\mathbf{x}^S) = x_n^S - b$.

Let's denote instance T by $\mathbf{x}^T = \langle x_1^T, \dots, x_n^T \rangle$, where $x_1^T = \dots = x_k^T = x_{k+1}^T - 2b = x_{k+2}^T - 4b = \dots = x_n^T - 4b$. Let X_1^T be the set of agents i for which $x_i^T = x_1^T$, and let X_2^T be the set of agents j for which $x_j^T = x_n^T$, and let x_t be agent x_{k+1}^T . On instance T , it must be either $f(\mathbf{x}^T) = x_1^T + b$ or $f(\mathbf{x}^T) = x_n^T - b$, otherwise agent x_t could misreport $x'_t = x_1^T$ and then by Lemma 5 and the fact that on instance P it is $f(\mathbf{x}^P) = x_1^P + b$, it should be $f(x_1^T, \dots, x'_t, \dots, x_n^T) = x_t - b$, which admits a cost of c for agent x_t . Similarly, by Lemma 5 and the fact that on instance S , it is $f(\mathbf{x}^S) = x_n^S - b$, if agent x_t misreports $x''_t = x_n^T$, then it should be $f(x_1^T, \dots, x''_t, \dots, x_n^T) = x_t + b$. If $f(\mathbf{x}^T) = x_n^T - b$, agent x_t could form a coalition with agents X_1^T and by misreporting x'_t , move the facility to $x_t - b$, a choice that would admit the same cost for her, but a strictly smaller cost for every other member of the coalition. If $f(\mathbf{x}^T) = x_1^T + b$ then agent x_t could form a coalition with agents X_2^T and by misreporting x''_t , move the facility to $x_t + b$, a choice that would admit the same cost for her, but a strictly smaller cost for every other member of the coalition. In each case, there is a coalition of agents that can benefit from misreporting and group strategyproofness is violated. This completes the proof. \square \square

4.2 Strategyproofness

We now turn our attention to (simple) strategyproofness. It is well known [17] that when agents have single-peaked preferences, outputting the location of the k th agent (k th order statistic), results in a strategyproof mechanism. This is not the case however, for double-peaked preferences and any choice of k .

Lemma 6. *Given any instance $\mathbf{x} = \langle x_1, \dots, x_n \rangle$, any mechanism that outputs $f(\mathbf{x}) = x_i - b$, for $i = 2, \dots, n$ or any mechanism that outputs $f(\mathbf{x}) = x_i + b$, for $i = 1, \dots, n - 1$, is not strategyproof.*

Proof. We prove the case when $f(\mathbf{x}) = x_1 + b$. The arguments for the other cases are similar. Consider any instance where $x_2 = x_1 + b$ and $x_i - b > x_2 + b$ for $i = 3, \dots, n$. Since $f(\mathbf{x}) =$

³Here we assume for convenience that any mechanism outputs the same location in $[x_1^P + c, x_n^P - c]$ on the primary instance and instance P . This is without loss of generality because the argument for any output in $[x_1^P + b, x_n^P - b]$ is exactly the same.

$f(x_1, x_2, x_3, \dots, x_n) = x_1 + b$, the cost of agent 2 is $\text{cost}(f(\mathbf{x}), x_2) = b + c$. If agent 2 misreports $x'_2 = x_1 - b$, the outcome will be $f(\mathbf{x}') = f(x_1, x'_2, x_3, \dots, x_n) = x'_2 + c = x_1$, and $\text{cost}(f(\mathbf{x}'), x_2) = c < b + c = \text{cost}(f(\mathbf{x}), x_2)$. Agent 2 has an incentive to misreport, therefore the mechanism is not strategyproof. \square \square

This only leaves two potential choices among k th order statistics, either $f(\mathbf{x}) = x_1 - b$ or $f(\mathbf{x}) = x_n + b$. Consider the following 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. $f(\mathbf{x}) = x_1 - b$ or always on the right peak of agent n , i.e. $f(\mathbf{x}) = x_n + b$.*

From now on, we will assume that M2 locates the facility on $x_1 - b$ on any instance \mathbf{x} . The analysis for the other case is similar.

Theorem 6. *Mechanism M2 is strategyproof.*

Proof. Obviously, agent 1 has no incentive to misreport, since her cost is already minimum. For any other agent $i, i = 2, \dots, n$, the cost is $\text{cost}(f(\mathbf{x}), x_i) = c + x_i - x_1$. For any misreport $x'_i \geq x_1$, the facility is still located on $x_1 - b$ and every agent's cost is the same as before. For some misreport $x'_i < x_1$, the facility moves to $f(x_1, \dots, x'_i, \dots, x_n) = x'_i - b < x_1 - b$, i.e. further away from any of agent i 's peaks and hence this choice admits a larger cost for her. \square \square

In the following, we prove that for the case of two agents, Mechanism M2 is actually the only strategyproof mechanism that satisfies anonymity and position invariance. We start with the following lemma.

Lemma 7. *For any instance $\mathbf{x} = \langle x_1, x_2 \rangle$, where $x_1 + b < x_2 - b$, if an anonymous, position invariant and strategyproof mechanism outputs $f(x_1, x_2) = x_1 - b$, then it must output $f(x'_1, x'_2) = x'_1 - b$ for any instance $\mathbf{x}' = \langle x'_1, x'_2 \rangle$ with $x'_1 \leq x'_2$.*

Proof. Let $(x_2 - b) - (x_1 + b) = \gamma$. For any other instance $\mathbf{x}' = \langle x'_1, x'_2 \rangle$ with $x'_1 \leq x'_2$, we assume $x'_1 = x_1$ without loss of generality due to position invariance. We will consider potential deviations of agent 2 of instance \mathbf{x} . Note that position invariance and anonymity suggest that we only need to consider deviations $x'_2 \geq x_1$ because any instance on which $x'_2 < x_1$ is position equivalent to some instance $\langle \hat{x}_1, \hat{x}_2 \rangle$ with $\hat{x}_1 = x_1$ and $\hat{x}_2 > x_1$.

On instance \mathbf{x} , the cost of agent 2 is $2b + c + \gamma$, so for any deviation of x_2 , her cost must be at least $2b + c + \gamma$, as required by strategyproofness. This implies that if agent 2 misreports x'_2 , then on the resulting instance it must be either $f(x_1, x'_2) \in (-\infty, x_1 - b]$ or $f(x_1, x'_2) \in [x_2 + 3b + \gamma, +\infty)$.

First, assume $f(x_1, x'_2) \in [x_2 + 3b + \gamma, +\infty)$. Let instance I be $\mathbf{x}^I = \langle x_1^I, x_2^I \rangle$, where $x_1^I = x_1, x_2^I = f(x_1, x'_2) + b$. On instance I it must be either $f(x_1^I, x_2^I) = f(x_1, x'_2)$ or $f(x_1^I, x_2^I) = f(x_1, x'_2) + 2b$, otherwise agent 2 can deviate from x_2^I to x'_2 and move the facility to $x'_2 - b$, minimizing her cost and violating strategyproofness. However then, on instance I, agent 1 can misreport $\bar{x}_1^I = x_2^I - 2b - \gamma$ and move the facility to $x_2^I - 3b - \gamma$, reducing her cost and violating strategyproofness. This follows from the fact that the resulting instance $\langle \bar{x}_1^I, x_2^I \rangle$ and instance \mathbf{x} are position equivalent and that $f(x_1, x_2) = x_1 - b$ on instance \mathbf{x} . Hence, it can not be that $f(x_1, x'_2) \in [x_2 + 3b + \gamma, +\infty)$ on instance \mathbf{x}' .

Second, assume $f(x_1, x'_2) \in (-\infty, x_1 - b)$. Then, since $x'_2 > x_1$, agent 2 can deviate from x'_2 to x_2 and move the facility to $x_1 - b$, i.e. closer to her actual position. Hence, it can't be that $f(x_1, x'_2) \in (-\infty, x_1 - b)$ on instance \mathbf{x}' either.

In conclusion, it must be $f(x_1, x'_2) = f(x'_1, x'_2) = x'_1 - b$ for any instance $\mathbf{x}' = (x'_1, x'_2)$ with $x'_1 \leq x'_2$. \square

Theorem 7. *When $n = 2$, the only strategyproof mechanism that satisfies position invariance and anonymity is Mechanism M2.*

Proof. For contradiction, suppose there exists an anonymous, position invariant strategyproof mechanism M which is different from Mechanism M2 and consider the primary instance used in Lemma 4.

We first argue that it must be that $f(x_1, x_2) \in [x_1 + b, x_2 - b]$. Assume on the contrary that $f(x_1, x_2) < x_1 + b$ (the argument for $f(x_1, x_2) > x_2 - b$ is symmetric). Then, consider the instance $\langle x'_1, x'_2 \rangle$, where $x'_2 = x_2$ and $x'_1 = f(x_1, x_2) - b$. On this instance, it must be that $f(x'_1, x'_2) = x'_1 - b$ or $f(x'_1, x'_2) = x'_1 + b$, otherwise agent 1 can deviate from x'_1 to x_1 and move the facility to $x'_1 + b$, minimizing her cost. In addition, if $f(x'_1, x'_2) = x'_1 - b$, then according to Lemma 7, the unique strategyproof mechanism that satisfies position invariance and anonymity is Mechanism M2, and thus our assumption is violated. So let's assume $f(x'_1, x'_2) = x'_1 + b$. Then, on the primary instance, agent 2 could report $\hat{x}_2 = x_2 + (x_1 + b - f(x_1, x_2))$ and by position invariance (since instances $\langle x'_1, x'_2 \rangle$ and $\langle x_1, \hat{x}_2 \rangle$ are position equivalent), it should be $f(x_1, \hat{x}_2) = x_1 + b$. This would give agent 2 an incentive to misreport, violating strategyproofness. Hence, on the primary instance it must be that $f(x_1, x_2) \in [x_1 + b, x_2 - b]$.

However, according to Lemma 4, it is impossible for a strategyproof mechanism to output $f(x_1, x_2) \in [x_1 + b, x_2 - b]$ on the primary instance. In conclusion, no mechanism other than Mechanism M2 is strategyproof, anonymous and position invariant. \square \square

4.3 Social Cost

We have seen that Mechanism M2 is strategyproof, but how well does it perform with respect to our goals, namely minimizing the social cost or the maximum cost? In other words, what is the approximation ratio that Mechanism M2 achieves against the optimal choice? First, we observe that the optimal mechanism, which minimizes the social cost, is not strategyproof.

Theorem 8. *The optimal mechanism with respect to the social cost, $f_{\text{opt}}(\mathbf{x}) = \arg \min_y \sum_{i=1}^n \text{cost}(y, x_i)$, is not strategyproof.*

Proof. Consider an instance $\mathbf{x} = \langle x_1, x_2, x_3 \rangle$, such that $x_2 + b < x_3 - b$, $x_2 - b < x_1 + b < x_2$ and $(x_1 + b) - (x_2 - b) = \epsilon$, where ϵ is an arbitrarily small positive quantity. On this instance, the optimal facility location is $x_2 + b$ and the cost of agent x_1 is $c + 2b - \epsilon$. Suppose now that agent x_1 reports $x'_1 < x_2 - 2b$. Moreover, suppose that when there are two locations y_1 and y_2 , with $y_1 < y_2$ that admit the same social cost, the optimal mechanism outputs y_1 . If the mechanism outputs y_2 instead, we can use a symmetric argument on the instance $\mathbf{x}' = \langle x'_1, x'_2, x'_3 \rangle = \langle x'_1, x_2, x_2 + 2b - \epsilon \rangle$ with agent 3 misreporting x_3 .⁴ By this tie-breaking rule, on instance $\mathbf{x} = \langle x'_1, x_2, x_3 \rangle$, the location of the facility is $x_2 - b$ and the cost of agent x_1 is $c + \epsilon$, i.e. smaller than before. Hence, the optimal mechanism is not strategyproof. To extend this to an arbitrary number of agents, let $x_j = x_2$ for every other agent x_j . \square \square

⁴In fact, even if the optimal mechanism outputs a distribution over points that all admit the minimum social cost, the argument still works.




Figure 3: Cases of Theorem 9. The dotted lines indicate the position of agent x_i before moving her from the left side to the right side of y . The cost of every agent with respect to y is the same before and after moving her to the right. The figure depicts only three agents (including x_1 and x_n).

Unfortunately, when considering the social cost, in the extremal case, the approximation ratio of Mechanism M2 is dependent on the number of agents. The approximation ratio is given by the following theorem.

Theorem 9. *For $n > 3$ agents, Mechanism M2 achieves an approximation ratio of $\Theta(n)$ for the social cost.*

Proof. Consider any instance $\mathbf{x}' = \langle x'_1, \dots, x'_n \rangle$ and let $y = f_{\text{opt}}(\mathbf{x}')$. It also holds that $f_{\text{M2}}(\mathbf{x}') = x'_1 - b$. Denote the social costs of the optimal mechanism and Mechanism M2 on instance \mathbf{x}' by $SC_{\text{opt}}(\mathbf{x}')$ and $SC_{\text{M2}}(\mathbf{x}')$, respectively. Let \mathbf{x} be the instance obtained by instance \mathbf{x}' as follows. For every agent $i, i \neq 1$, if $x'_i + b \leq y$, let $x_i = 2y - x'_i$; if $x'_i < y < x'_i + b$, let $x_i = x'_i + 2b$; if $x'_i - b < y < x'_i$, let $x_i = 2y - x'_i + 2b$; otherwise let $x_i = x'_i$. Observe that $|x_i - b - y| = \min(|x'_i - b - y|, |x'_i + b - y|)$ and hence $\text{cost}(x_i, y) = \text{cost}(x'_i, y)$ for every agent i . Similarly, it holds that $(x_i - b) - (x_1 - b) \geq (x'_i - b) - (x'_1 - b)$ and hence (since $x_1 = x'_1$), $\text{cost}(x_i - b, x_1 - b) \geq \text{cost}(x'_i - b, x'_1 - b)$ for all agents i .

We will calculate an upper bound on the approximation ratio on instance x' . To do that, we will calculate an upper bound on the value of the ratio $SC_{\text{M2}}(\mathbf{x})/SC_y(\mathbf{x})$ on instance \mathbf{x} , where $SC_{\text{M2}}(\mathbf{x})$ is the social cost of Mechanism M2 on instance \mathbf{x} and $SC_y(\mathbf{x})$ is the social cost admitted by y on instance \mathbf{x} . By the way the instance was constructed, it holds that $SC_{\text{M2}}(\mathbf{x}) \geq SC_{\text{M2}}(\mathbf{x}')$ and $SC_y(\mathbf{x}) = SC_y(\mathbf{x}') = SC_{\text{opt}}(\mathbf{x}')$ and hence $SC_{\text{M2}}(\mathbf{x})/SC_y(\mathbf{x})$ is an upper bound on $SC_{\text{M2}}(\mathbf{x}')/SC_{\text{opt}}(\mathbf{x}')$.

Let $d_i = (x_i - b) - y$, for $i = 2, \dots, n$ and let $k = \sum_{i \neq 1} d_i$. Finally let $d = |y - (x_1 + b)|$. We consider three cases. (See Figure 3).

Case 1: $x_1 + b < y$.

In this case, the social cost admitted by y is $SC_y(\mathbf{x}) = nc + k + d$, and the cost of Mechanism M2 is $SC_{M2}(\mathbf{x}) = nc + k + (n - 1)(d + 2b)$. The ratio is

$$\frac{SC_{M2}(\mathbf{x})}{SC_y(\mathbf{x})} = g(k, d, n) = \frac{k + (n - 1)(d + 2b) + nc}{k + d + nc}.$$

By calculating the partial derivative of the ratio with respect to k , we have

$$\frac{\partial g(k, d, n)}{\partial k} = \frac{k + d + nc - k - (n - 1)(d + 2b) - nc}{(k + d + nc)^2} = \frac{d - (n - 1)(d + 2b)}{(k + d + nc)^2} < 0.$$

When $k = 0$, i.e. $x_i = y + b$ for all agents $i = 2, \dots, n$, the ratio achieves the maximum value with respect to k . We plug in $k = 0$ to the above formula and get

$$\frac{SC_{M2}(\mathbf{x})}{SC_y(\mathbf{x})} = g(d, n) = \frac{(n - 1)(d + 2b) + nc}{d + nc} = \frac{(n - 1)d + 2(n - 1)b + nc}{d + nc}.$$

Again, we calculate the partial derivative of $g(d, n)$ with respect to d ,

$$\frac{\partial g(d, n)}{\partial d} = \frac{(n - 1)(d + nc) - (n - 1)d - 2(n - 1)b - nc}{(d + nc)^2} = \frac{c(n^2 - 4n + 2)}{(d + nc)^2}.$$

For sufficiently large n ($n \geq 2b/c$), the above value is positive, so the ratio is maximized when d approaches to infinity, which means that agent x_1 is positioned very far away from the rest of the agents coinciding on position x_i . Therefore the ratio goes to $n - 1$.

Case 2: $x_1 < y \leq x_1 + b$, which means $0 \leq d < b$.

In this case, the social cost admitted by y is $SC_y(\mathbf{x}) = k + d + nc$, and the cost of Mechanism M2 is $SC_{M2}(\mathbf{x}) = k + (n - 1)(2b - d) + nc$. So the ratio is

$$\frac{SC_{M2}(\mathbf{x})}{SC_y(\mathbf{x})} = g(k, d, n) = \frac{k + (n - 1)(2b - d) + nc}{k + d + nc}.$$

Again, we calculate the partial derivative of $g(k, d, n)$ with respect to k ,

$$\frac{\partial g(k, d, n)}{\partial k} = \frac{k + d + nc - k - (n - 1)(2b - d) - nc}{(k + d + nc)^2} = \frac{d - (n - 1)(2b - d)}{(k + d + nc)^2}.$$

Since $0 \leq d < b$, $d < 2b - d$ and the above value is negative. To get the maximum value of the ratio, we plug in $k = 0$ and the approximation ratio becomes

$$\frac{SC_{M2}(\mathbf{x})}{SC_y(\mathbf{x})} = g(d, n) = \frac{(n - 1)(2b - d) + nc}{d + nc}.$$

By calculating the partial derivative with respect to d , we get

$$\frac{\partial g(d, n)}{\partial d} = \frac{-(n - 1)(d + nc) - (n - 1)(2b - d) - nc}{(d + nc)^2} = \frac{-(n - 2)nc - (n - 1)2b}{(d + nc)^2} < 0.$$

So the ratio is maximized when $d = 0$, and the ratio is $1 + \frac{(2n-2)b+nc}{nc}$, which (for sufficiently large n) is smaller than the ratio in Case 1.

Case 3: $x_1 - b < y \leq x_1$, which means $b \leq d < 2b$.

In this case, the social cost admitted by y is $SC_y(\mathbf{x}) = k + 2b - d + nc$, and cost of Mechanism M2 is $SC_{M2}(\mathbf{x}) = k + (n-1)(2b-d) + nc$. So the ratio is

$$\frac{SC_{M2}(\mathbf{x})}{SC_y(\mathbf{x})} = g(k, d, n) = \frac{k + (n-1)(2b-d) + nc}{k + 2b - d + nc}.$$

We first calculate partial derivative of the function g with respect to k ,

$$\frac{\partial g(k, d, n)}{\partial k} = \frac{k + 2b - d + nc - k - (n-1)(2b-d) - nc}{(k + 2b - d + nc)^2} = \frac{(2-n)(2b-d)}{(k + 2b - d + nc)^2} < 0.$$

The function is decreasing in k so let's set $k = 0$. The function f becomes

$$\frac{SC_{M2}(\mathbf{x})}{SC_y(\mathbf{x})} = g(d, n) = \frac{(n-1)(2b-d) + nc}{2b-d+nc}.$$

Now, we calculate partial derivative of g with respect to d and get

$$\frac{\partial g(d, n)}{\partial d} = \frac{-(n-1)(2b-d) - n(n-1)c + (n-1)(2b-d) + nc}{(2b-d+nc)^2} \leq 0.$$

So function $g(d, n)$ is decreasing in d . Let's set $d = b$, and the ratio goes to $\frac{(n-1)b+nc}{b+nc}$, which goes to $n-1$ as well.

In all, since $SC_{M2}(\mathbf{x})/SC_y(\mathbf{x}) \leq n-1$ the approximation ratio of Mechanism M2 is at most $n-1$. The approximation ratio is exactly $n-1$ on any instance $\mathbf{x} = \langle x_1, x_2, \dots, x_n \rangle$ with $x_2 = \dots = x_n$ and $x_1 \ll x_2$, i.e. when agent 1 lies on the left, really far away from the other $n-1$ agents. \square \square

Next, we will prove a lower bound of $1 + b/c$ on the approximation ratio of any anonymous, position invariant, and strategyproof mechanism, when the number of agents is even. We will start with the following lemma.

Lemma 8. *Let M^n be a strategyproof, anonymous and position invariant mechanism for n agents, where n is even. 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 $\langle x_1, x_2 \rangle$, output $M^2(\mathbf{x}')$, where $\mathbf{x}' = \langle x'_1 = \dots = x'_{n/2}, x'_{n/2+1} = \dots = x'_n \rangle$, and $x'_1 = x_1$ and $x'_{n/2+1} = x_2$. First, we claim that M^2 is strategyproof, anonymous and position invariant. If that is true, then by Theorem 7, M^2 is Mechanism M2 and the lemma follows.

First let $\mathbf{x} = \langle x_1, x_2 \rangle$, $\hat{\mathbf{x}} = \langle \hat{x}_1, \hat{x}_2 \rangle$ be any two position equivalent location profiles. Observe that the corresponding n -agent profiles \mathbf{x}' and $\hat{\mathbf{x}}'$ obtained by placing $n/2$ agents on x_1 and \hat{x}_1 and $n/2$ agents on x_2 and \hat{x}_2 respectively are also position equivalent. Since M^n is position invariant, it must hold that $M^n(\mathbf{x}') = M^n(\hat{\mathbf{x}}')$ and hence by construction of M^2 , $M^2(\mathbf{x}) = M^2(\hat{\mathbf{x}})$. Since \mathbf{x} and $\hat{\mathbf{x}}$ where arbitrary, Mechanism M^2 is position invariant.

Similarly, let $\mathbf{x} = \langle x_1, x_2 \rangle$, $\hat{\mathbf{x}} = \langle \hat{x}_1, \hat{x}_2 \rangle$ be any two location profiles, such that $\hat{\mathbf{x}}$ is obtained by \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 x_1 and x_2 is the same) is the same and by construction of M^2 , $M^2(\mathbf{x}) = M^2(\hat{\mathbf{x}})$ 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 $\text{cost}(x', y)$ 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 $\text{cost}(x', y)$. 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 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 10. *When the number of agents is even, any strategyproof mechanism that satisfies position invariance and anonymity achieves an approximation ratio of at least $1 + (b/c)$ 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 8, $M^n(\mathbf{x}) = x_1 - b$ 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

4.4 Maximum Cost

First, it is easy to see that the mechanism that outputs the location that minimizes the maximum cost is not strategyproof. On any instance $\langle x_1, x_2 \rangle$ with $x_1 + b < x_2 - b$ the optimal location of the facility is $(x_1 + x_2)/2$. If agent x_2 misreports $x'_2 = 2x_2 - 2b - x_1$ then the location moves to $x_2 - b$, minimizing her cost.

While the approximation ratio of Mechanism M2 for the social cost is not constant, for the maximum cost that is indeed the case. In fact, as we will prove, when the number of agents is even, Mechanism M2 actually achieves the best possible approximation ratio amongst strategyproof mechanisms. We start with the theorem about the approximation ratio of Mechanism M2.

Theorem 11. *For $n \geq 3$, Mechanism M2 achieves an approximation ratio of $1 + \frac{2b}{c}$ for the maximum cost.*

Proof. Let $\mathbf{x} = \langle x_1, \dots, x_n \rangle$ be any instance. We consider three cases, depending on the distance between agents x_1 and x_n .

Case 1: $x_1 + b \leq x_n - b \Rightarrow x_n - x_1 \geq 2b$.

In this case, the cost of the optimal mechanism is $MC_{\text{opt}}(\mathbf{x}) = (x_n - b) - \frac{x_1 + x_n}{2} + c = \frac{x_n - x_1}{2} - b + c$, whereas the cost of Mechanism M2 is $MC_{\text{M2}}(\mathbf{x}) = x_n - x_1 + c$. The approximation ratio is

$$\frac{MC_{\text{M2}}(\mathbf{x})}{MC_{\text{opt}}(\mathbf{x})} = \frac{2(x_n - x_1 + c)}{x_n - x_1 - 2b + 2c} = 2 + \frac{4b - 2c}{x_n - x_1 - 2b + 2c} \leq 1 + \frac{2b}{c}.$$

Hence, in this case, the approximation ratio is at most $1 + \frac{2b}{c}$. For $x_n - b = x_1 + b$ (the instance on which the right peak of the first agent and the left peak of the last agent coincide), the approximation ratio is exactly $1 + \frac{2b}{c}$.

Case 2: $x_1 < x_n - b < x_1 + b \Rightarrow c < x_n - x_1 < 2b$.

The cost of Mechanism M2 in this case is $MC_{\text{M2}}(\mathbf{x}) = x_n - b - (x_1 - b) + c = x_n - x_1 + c$, while the cost of the optimal mechanism is at least c . The approximation ratio is

$$\frac{MC_{\text{M2}}(\mathbf{x})}{MC_{\text{opt}}(\mathbf{x})} \leq \frac{(x_n - x_1 + c)}{c} < 1 + \frac{2b}{c}$$

Case 3: $x_1 - b \leq x_n - b \leq x_1 \Rightarrow x_n - x_1 \leq b$.

The cost of Mechanism M2 is $MC_{\text{M2}}(\mathbf{x}) = x_n - x_1 + c$ while the cost of the optimal mechanism is at least c . The approximation ratio is

$$\frac{MC_{\text{M2}}(\mathbf{x})}{MC_{\text{opt}}(\mathbf{x})} \leq \frac{(x_n - x_1 + c)}{c} \leq 1 + \frac{b}{c}.$$

Over the three cases, the worst approximation ratio is at most $1 + \frac{2b}{c}$ and there is actually an instance of the problem with approximation ratio exactly $1 + \frac{2b}{c}$, ensuring that the bound is tight. \square

The lower bound for the case when the number of agents is even follows.

Corollary 1. *When the number of agents is even, any deterministic strategyproof mechanism that satisfies position invariance and anonymity achieves an approximation ratio of at least $1 + \frac{2b}{c}$ for the maximum cost.*

Proof. On instance $\hat{\mathbf{x}}$ of the proof of Theorem 10, there are agents whose cost for any strategyproof, anonymous and position invariant mechanism is $2b + c$, while under the optimal mechanism it is only c . The lower bound on the approximation ratio follows. \square

A worse lower bound that holds for any number of agents (and without using the position invariance property) is proved in the next theorem.

Theorem 12. *Any deterministic strategyproof mechanism achieves an approximation ratio of at least 2 for the maximum cost.*

Proof. Consider an instance $\mathbf{x} = \langle x_1, x_2 \rangle$ with $x_1 + b < x_2 - b$ (the instance can be extended to arbitrarily many agents by placing agents on positions x_1 and x_2 and all the arguments will still hold). The optimal location of the facility is $f_{\text{opt}}(\mathbf{x}) = \frac{x_1 + x_2}{2}$. Assume for contradiction that M is a deterministic strategyproof mechanism with approximation ratio smaller than 2.

First, we argue that it can not be that $f_M(\mathbf{x}) \in [x_2 - b, \infty)$. Let $d = x_2 - b - f_{\text{opt}}(\mathbf{x})$. It holds that $MC_{\text{opt}}(\mathbf{x}) = c + d$. If it was $f_M(\mathbf{x}) \in [x_2 - b, \infty)$, then it would be that $MC_M(\mathbf{x}) \geq c + 2d$

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	$\Theta(n)$	$3/2$

and the approximation ratio would be at least $2 - \frac{c}{d+c}$ which goes to 2 as d grows to infinity (i.e. the agents are placed very far away from each other). Now, for Mechanism M to achieve an approximation ratio smaller than 2, it must be $f_M(\mathbf{x}) \in [f_{\text{opt}}(\mathbf{x}), x_2 - b)$ (or symmetrically $f_M(\mathbf{x}) \in (x_1 + b, f_{\text{opt}}(\mathbf{x}))$).

Now consider the instance $\mathbf{x}' = \langle x'_1, x'_2 \rangle$ with $x'_1 = x_1$ and $x'_2 = f_M(\mathbf{x}) + b$. On this instance, it must be either $f_M(\mathbf{x}') = f_M(\mathbf{x})$ or $f_M(\mathbf{x}') = f_M(\mathbf{x}) + 2b$ (the left or the right peak of agent x'_2), otherwise agent x'_2 could report x_2 and move the facility to $x'_2 - b$, minimizing her cost and violating strategyproofness. For calculating the lower bound, we need the choice that admits the smaller of the two costs, i.e. $f_M(\mathbf{x}') = f_M(\mathbf{x})$. We calculate the approximation ratio on instance \mathbf{x}' .

The optimal choice for the facility is again $f_{\text{opt}}(\mathbf{x}') = (x_1 + x'_2)/2$. Let $\lambda = (x'_2 - b) - f_{\text{opt}}(\mathbf{x}')$. The approximation ratio then is $2 - \frac{c}{\lambda+c}$. We know that $\lambda \geq d/2$, so when d grows to infinity as before, λ also grows to infinity and the approximation ratio goes to 2. This means that there exists an instance for which the approximation ratio of the mechanism is 2, which gives us the lower bound. \square

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

5.1 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.

5.2 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.

References

- [1] Noga Alon, Michal Feldman, Ariel D Procaccia, and Moshe Tennenholtz. Strategyproof approximation of the minimax on networks. *Mathematics of Operations Research*, 35(3):513–526, 2010.
- [2] I. Ashlagi, F. Fischer, I. Kash, and Ariel D. Procaccia. Mix and match. In *Proceedings of the 11th ACM conference on Electronic commerce (ACM-EC)*, pages 305–314. ACM, 2010.
- [3] Duncan Black. *The theory of committees and elections*. Kluwer Academic Publishers, 1957 (reprint at 1986).
- [4] I. Caragiannis, A. Filos-Ratsikas, and Ariel D. Procaccia. An improved 2-agent kidney exchange mechanism. In *Proceedings of the 7th Workshop of Internet and Network Economics (WINE)*, pages 37–48. Springer, 2011.
- [5] Robert D Cooter. *The strategic constitution*. Princeton University Press, 2002.
- [6] Elad Dokow, Michal Feldman, Reshef Meir, and Ilan Nehama. Mechanism design on discrete lines and cycles. In *Proceedings of the 13th ACM Conference on Electronic Commerce (ACM-EC)*, pages 423–440, 2012.
- [7] S. Dughmi and A. Gosh. Truthful assignment without money. In *Proceedings of the 11th ACM conference on Electronic commerce (ACM-EC)*, pages 325–334, 2010.
- [8] Patrick J Egan. ”Do something” politics and double-peaked policy preferences. *Journal of Politics*, 76(2):333–349, 2013.
- [9] B. Escoffier, L. Gourvès, N. Kim Thang, F. Pascual, and O. Spanjaard. Strategy-proof mechanisms for facility location games with many facilities. *Algorithmic Decision Theory*, pages 67–81, 2011.

- [10] Itai Feigenbaum, Jay Sethuraman, and Chun Ye. Approximately optimal mechanisms for strategyproof facility location: Minimizing l_p norm of costs. *arXiv preprint arXiv:1305.2446*, 2013.
- [11] M. Feldman and Y. Wilf. Strategyproof facility location and the least squares objective. In *Proceedings of the 14th ACM conference on Electronic Commerce (ACM-EC)*, pages 873–890, 2013.
- [12] D. Fotakis and C. Tzamos. Winner-imposing strategyproof mechanisms for multiple facility location games. In *Proceeding of the 5th International Workshop of Internet and Network Economics (WINE)*, pages 234–245, 2010.
- [13] D. Fotakis and C. Tzamos. On the power of deterministic mechanisms for facility location games. *arXiv preprint arXiv:1207.0935*, 2012.
- [14] M. Guo and V. Conitzer. 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, pages 881–888, 2010.
- [15] Pinyan Lu, Xiaorui Sun, Yajun Wang, and Zeyuan Allen Zhu. Asymptotically optimal strategy-proof mechanisms for two-facility games. In *Proceedings of the 11th ACM Conference on Electronic Commerce (ACM-EC)*, pages 315–324, 2010.
- [16] Pinyan Lu, Yajun Wang, and Yuan Zhou. Tighter bounds for facility games. In *Proceeding of the 5th International Workshop of Internet and Network Economics (WINE)*, pages 137–148, 2009.
- [17] Hervé Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.
- [18] Ariel D Procaccia and Moshe Tennenholtz. Approximate mechanism design without money. In *Proceedings of the 10th ACM Conference on Electronic Commerce (ACM-EC)*, pages 177–186, 2009.
- [19] Harvey S Rosen. *Public Finance*, chapter 6 - Political Economy, pages 115–117. Springer, 2004.
- [20] James Schummer and Rakesh V Vohra. Strategy-proof location on a network. *Journal of Economic Theory*, 104(2):405–428, 2002.