The Surprising Power of Hiding Information in Facility Location

Safwan Hossain
Vector Institute, University of Toronto
safwan.hossain@mail.utoronto.ca

Evi Micha
University of Toronto
emicha@cs.toronto.edu

Nisarg Shah
University of Toronto
nisarg@cs.toronto.edu

Abstract

Facility location is the problem of locating a public facility based on the preferences of multiple agents. In the classic framework, where each agent holds a single location on a line and can misreport it, strategyproof mechanisms for choosing the location of the facility are well-understood.

We revisit this problem in a more general framework. We assume that each agent may hold several locations on the line with different degrees of importance to the agent. We study mechanisms which elicit the locations of the agents and different levels of information about their importance. Further, in addition to the classic manipulation of misreporting locations, we introduce and study a new manipulation, whereby agents may *hide* some of their locations. We argue for its novelty in facility location and applicability in practice. Our results provide a complete picture of the power of strategyproof mechanisms eliciting different levels of information and with respect to each type of manipulation. Surprisingly, we show that in some cases hiding locations can be a strictly more powerful manipulation than misreporting locations.

1 Introduction

Approximate mechanism design without money is a paradigm introduced by Procaccia and Tennenholtz (2009), which sits at the intersection of computer science and economics, and reasons about ways to prevent strategic manipulations by agents without monetary transfers. They illustrate this through the canonical facility location problem, where n agents are located on the real line, and a mechanism elicits their locations to decide where to build a public facility. However, the agents are strategic, and may manipulate their reports to bring the facility closer to their location. To prevent such manipulations, one may seek a *strategyproof* mechanism under which no agent can benefit by manipulating, regardless of what the others do. However, imposing this constraint comes at a price. Given an objective the designer wants to minimize (such as the maximum distance of the facility to any agent), she may only be able to *approximately* minimize it subject to strategyproofness. In the last decade, research on facility location has exploded, and many variants have been studied such as: locating multiple facilities (Escoffier et al. 2011), locating a facility in multiple dimensions (Sui, Boutilier, and Sandholm 2013), exploring different types of agent preferences (Filos-Ratsikas et al. 2017) and objectives (Feldman and Wilf 2013), and strategic opening of facilities (Chen et al. 2019).

However, this literature has mainly focused on a single type of manipulation: agents *misreporting* their location. In certain contexts however, agents may not be able to lie about their location, but can still manipulate by *hiding* their location. For example, to decide where the facility should be built, a survey may request residents to provide their home address, or school boards to provide the school address. Such reports can often be easily verified, either through external methods or by requiring participants to upload proof. In such cases, agents cannot lie about their location, as it may be tantamount to fraud; however, they can choose to not participate, thus hiding their location. Furthermore, even if the principal can detect agents hiding data, such detection might be very time-consuming or violate privacy concerns.

When each agent holds a single location, the hiding manipulation is very restrictive: an agent can either participate (and reveal the correct location) or not participate (and hide the location). The desideratum of incentivizing agents to participate is known as *individual rationality*, and is already widely studied (Nisan et al. 2007). However, when each agent holds *multiple* locations, she may choose to reveal any subset of these locations, making the hiding manipulation much more complex.

In facility location, agents holding multiple locations arise naturally. In the aforementioned example, residents may report both their home and work address, or a school district may report the addresses of multiple schools under its purview. This has been somewhat explored in facility location (Dekel, Fischer, and Procaccia 2010; Filos-Ratsikas et al. 2017). We also note additional motivation from a different line of literature. Recent explorations of strategic interactions in machine learning have revealed that research on facility location provides great insight into designing strategyproof algorithms for tasks such as linear regression (Chen et al. 2018; Hossain and Shah 2019), where the training data may come from strategic sources and assuming that each

data source provides a single data point is highly unrealistic.

In our model, we assume that each strategic agent holds multiple points with potentially different weights, and is interested in minimizing the weighted sum of their distances to the facility (a.k.a. her cost). This immediately raises a number of questions.

- How powerful is the hiding manipulation compared to the more commonly studied misreporting manipulation?
- How do we characterize strategyproof mechanisms with respect to such manipulations?
- What is the price of imposing strategyproofness in terms of natural objectives that we may care about?

Our Results

This work focuses on answering such questions. We consider two natural objectives: *social cost*, which is the sum of costs to the agents, and *unweighted social cost*, which is the sum of distances of all points to the facility (disregarding the weights).

In addition to eliciting the points, our mechanisms also elicit information about their weights. For *full information* mechanisms, which ask agents to report the exact weights, we show that the PROJECT-AND-FIT mechanism introduced by Dekel, Fischer, and Procaccia (2010), with appropriate generalization to our setting, is strategyproof with respect to both types of manipulations, providing a 3 approximation to social cost and $2m - 1$ approximation to unweighted social cost. Both approximations are essentially optimal. While this may suggest a deeper connection between families of strategyproof mechanisms with respect to the two manipulations, we show that the families are incomparable as there exist mechanisms that are strategyproof with respect to one manipulation but not the other.

For *ordinal* mechanisms, which ask agents to report only a ranking of points by weight rather than the exact weights, we show that only constant mechanisms are strategyproof with respect to hiding; for misreporting however, the family of strategyproof mechanisms is strictly larger. This indicates that hiding is strictly more powerful than misreporting in this case. We show that imposing strategyproofness with respect to either manipulation results in infinite approximation to both objectives, but without it, ordinal mechanisms can achieve $\Theta(m)$-approximation to social cost and 1-approximation to unweighted social cost.

Lastly, our negative results hold even when agents are not allowed to manipulate their weight information, and our positive results hold even if they are allowed to.

Related Work

Much of the facility location literature works under the assumption that each agent has single-peaked preferences over possible locations of the facility (Moulin 1980; Schummer and Vohra 2002; Alon et al. 2009; Procaccia and Tennenholtz 2009). In our model, preferences are generated by a weighted sum of ℓ_1 distances to multiple points, and thus are still single-peaked. However, our setting differs from prior work in two key aspects. First, prior work assumes that when agents manipulate they are allowed to report *any* single-peaked preferences, whereas in our model they can manipulate in limited ways. That said, the PROJECT-AND-FIT mechanism we study is inspired by results on strategyproofness in the single-peaked (or more accurately, single-plateau) domain (Moulin 1980; 1984).

The most closely related work to ours is that of Dekel, Fischer, and Procaccia (2010). Among other results, they show that PROJECT-AND-FIT is strategyproof with respect to misreporting when agents care about all their points equally. Theorem 1 extends this to the case where agents have weights for points, can manipulate the weights, and can also misreport or hide points. Their work also establishes that PROJECT-AND-FIT gives a 3-approximation to social cost, which is tight for strategyproof mechanisms with respect to misreporting. We extend their result to our weighted domain and different strategy spaces, while also giving asymptotically tight bounds for unweighted social cost approximation. Finally, they establish their results in a linear regression framework. Our negative results carry over to this more general domain. Our positive results (Theorems 1 and 2) also hold in the more general setting of Dekel, Fischer, and Procaccia (2010), but we omit the details for ease of exposition.

The hiding manipulation has been very well studied in the kidney exchange problem (Roth, Sönmez, and Ünver 2004; Ashlagi et al. 2015). This setting has patient-donor pairs, where a patient needs a kidney, a donor is willing to donate one, but they are not a match. Centralized exchanges ask hospitals to report their patient-donor pairs, so that perhaps the donors and patients of two distinct pairs are a match for each other. But hospitals can hide and internally match some of their pairs to increase the total number of its matched patients. Our work brings the idea of hiding manipulation from this literature to facility location, where it can be combined with complex preference structures and compared to the misreporting manipulation. We note that hiding parts of preferences is also well-studied in unweighted division and assignment problems (Fadaei and Bichler 2017).

2 Model

For a natural number $k \in \mathbb{N}$, define $[k] = \{1, \dots, k\}$. Also, define the extended real line $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, \infty\}$. Let $N = [n]$ be a set of agents. Each agent i holds m_i points denoted by $x_{i,j} \in \overline{\mathbb{R}}$, for $j \in [m_i]$; let D_i denote the (multi)set of points held by agent i. In addition, the agent has a weight function $w_i : D_i \to \mathbb{R}_{\geq 0}$ such that $\sum_{x_{i,j} \in D_i} w_i(x_{i,j}) = 1$; here, $w_i(x_{i,j})$ indicates the relative importance of point $x_{i,j}$ to agent i.[1] In our model, D_i and w_i form the private information held by agent i. Let us define $m = \max_{i \in N} m_i$.

Agent preferences. The outcome of the facility location problem is a single location $x \in \mathbb{R}$ where a public facility will be placed. For this outcome, the cost to agent i is $c_i(x) = \sum_{x_{i,j} \in D_i} w_i(x_{i,j}) \cdot |x - x_{i,j}|$. Note that these preferences are single-peaked (Moulin 1980).

Mechanisms. Often, it may not be feasible or practical to

[1] We treat points $x_{i,j}$ as "labeled" points. Hence, it is possible to have different weights for two points at the same location, i.e., for $j \neq j'$, we can have $w_i(x_{i,j}) \neq w_i(x_{i,j'})$ even when $x_{i,j} = x_{i,j'}$.

ask agents to submit full preference information, and mechanisms may ask instead for partial information. Formally, a mechanism M specifies how each agent i should submit a *response* ρ_i given her private information (D_i, w_i). An instance I consists of the private information of the agents and the responses submitted by them. Let $M(I) \in \mathbb{R}$ denote the location chosen by M on instance I. We consider mechanisms that elicit four different levels of information about agent preferences.

- *Full information mechanisms:* These ask each agent i to report all her points and their weights, i.e., $\rho_i = (D_i, w_i)$.
- *Ordinal mechanisms:* These mechanisms still ask each agent i to report all her points, but instead of reporting their weights, they ask her to report a ranking of the points by their weight. Formally, $\rho_i = (D_i, \sigma_i)$, where σ_i is a linear order over D_i with the property that for all $a, b \in D_i$, $w_i(a) > w_i(b)$ implies $a \succ_{\sigma_i} b$.[2] When this property holds, we say that the ranking is consistent with the weights.
- *Weightless mechanisms:* These mechanisms ask each agent i to report only her points and do not elicit any weight information, i.e., $\rho_i = D_i$.
- *Anonymous mechanisms:* These mechanisms, like weightless mechanisms, also ask each agent i to report only her points (i.e. $\rho_i = D_i$). However, the mechanism only observes $\cup_{i \in N} \rho_i$. That is, the mechanism receives anonymized points, and cannot determine which agent submitted any given point.

Some of our results concern *constant mechanisms* which simply choose a constant location regardless of input. That is, for a constant mechanism M, $M(I) = M(I')$ for all pairs of instances I, I'.

Objective functions. In this work, we consider two objective functions that we may wish to minimize.

- *Social cost:* This is simply the sum of costs to the agents, i.e., for all $x \in \mathbb{R}$,

$$\mathrm{sc}(x) = \sum_{i \in N} \sum_{x_{i,j} \in D_i} w_i(x_{i,j}) \cdot |x - x_{i,j}|$$

- *Unweighted social cost:* The unweighted social cost is the sum of costs to the individual points, disregarding the weights placed by the agents on the points. Formally, for all $x \in \mathbb{R}$,

$$\mathrm{usc}(x) = \sum_{i \in N} \sum_{x_{i,j} \in D_i} |x - x_{i,j}|$$

Unweighted social cost can be seen as social cost of the individual points. In the example from the introduction, where each school district reports the locations of its schools, unweighted social cost will give equal importance to all schools, ignoring any weights placed by the districts

(Worst-case) Approximation ratio. In this work, we are interested in the (worst-case) approximation that a mechanism provides to the two objectives, assuming agents submit honest reports. Formally, the approximation ratio of mechanism M for objective obj (where $\mathrm{obj} = \mathrm{sc}$ for social cost, and $\mathrm{obj} = \mathrm{usc}$ for unweighted social cost) is defined as $\sup_I \frac{\mathrm{obj}(M(I))}{\min_{x^*} \mathrm{obj}(x^*)}$, where supremum is taken over all instances I. Achieving 1-approximation may not be possible when we either do not have access to full information or want to satisfy other desiderata such as strategyproofness.

Strategic behavior. We assume that each agent i is strategic and seeks to minimize her own cost c_i. To that end, she may submit a strategic response ρ_i' instead of the honest response ρ_i requested by the mechanism. A strong desideratum to prevent manipulations is strategyproofness.

Definition 1. A mechanism M is called *strategyproof* if for every $(D_i, w_i)_{i \in N}$, every possible set of agent reports $\vec{\rho'} = (\rho_1', \ldots, \rho_n')$, and every agent $i \in N$, it holds that $c_i(M(\rho_i, \vec{\rho'}_{-i})) \leq c_i(M(\rho_i', \vec{\rho'}_{-i}))$, where ρ_i is the honest response of agent i given (D_i, w_i). In words, an agent should not be able to gain by manipulating regardless of the reports submitted by the other agents.

The definition of strategyproofness is clearly sensitive to the space of manipulations ρ_i' that agent i is allowed to submit. In this work, we consider two types of manipulations.

- *Misreporting:* This is the standard manipulation studied in facility location, where the agent may misreport her points. Specifically, agent i may submit $D_i' = (x_{i,j}')_{j \in [m_i]}$ as part of her strategic response ρ_i'. Note that $|D_i'| = |D_i| = m_i$, and the agent still submits weight $w_i(x_{i,j})$ for each manipulated point $x_{i,j}'$ to a full information mechanism (or the corresponding ranking to an ordinal mechanism).
- *Hiding:* This is a new type of manipulation that we study, where the agent may hide some of her points. Specifically, agent i may submit D_i' as part of her strategic response ρ_i', where $D_i' \subseteq D_i$. Note that the agent is only allowed to hide a subset of points, and not allowed to change points. Also, the agent now submits re-normalized weight $w_i(x_{i,j}) / \sum_{a \in D_i'} w_i(a)$ for each point $x_{i,j} \in D_i'$ that she reveals to a full information mechanism (or the corresponding ranking to an ordinal mechanism).[3]

Note that in both cases, we assume that the agent does not manipulate the part of her response that conveys weight information. This makes our strategyproofness definition weaker, and thus all our negative results stronger. In our positive result (Theorem 1), the mechanism constructed is strategyproof *even when* agents are allowed to manipulate the part of their response that conveys weight information. Thus, all our results hold regardless of whether the agents can manipulate their weight information.

3 Full Information Mechanisms

We begin by considering the full information case, where the mechanism asks the agents to submit both their points

[2] The agent can break ties among points with equal weight.

[3] When an agent only reveals k zero-weight points, we assume she reports weight $1/k$ for each point. When an agent hides all her points, she submits nothing and the mechanism pretends the agent was not present.

and their weights. This case was studied by Dekel, Fischer, and Procaccia (2010) for the misreporting manipulation, in the special case where agents have uniform weights over their points, i.e., $w_i(x_{i,j}) = 1/m_i$ for each $i \in N$ and $j \in [m_i]$. For this case, they introduce a mechanism called PROJECT-AND-FIT and argue that it is strategyproof. We generalize their mechanism to our setting (presented as Algorithm 1), where agents may have non-uniform weights over their points, and show that the generalized PROJECT-AND-FIT is strategyproof not only with respect to misreporting but also with respect to hiding. This mechanism first computes a location x_i^* most preferred by agent i (breaking ties appropriately), and then returns the median of all x_i^*, denoted by $\text{median}(\{x_i^* : i \in N\})$.[4] Note that x_i^* is simply a weighted median of agent i's points, satisfying $\sum_{j:x_{i,j} \geq x_i^*} w_i(x_{i,j}) \geq 1/2$ and $\sum_{j:x_{i,j} \leq x_i^*} w_i(x_{i,j}) \geq 1/2$.

Algorithm 1: Mechanism PROJECT-AND-FIT

Input: $\rho_i = (D_i, w_i)$ for each $i \in N$
Output: $x^* \in \mathbb{R}$
1 Project: $S_i^* \leftarrow \text{argmin}_x \, c_i(x), \forall i \in N$
2 Tie-break: $x_i^* \leftarrow \text{argmin}_{x \in S_i^*} |x|$
3 Fit: $x^* \leftarrow \text{median}(\{x_i^* : i \in N\})$

Theorem 1. PROJECT-AND-FIT *is strategyproof with respect to both misreporting and hiding, even when agents can manipulate their weight information.*

The proof effectively makes the same argument that Dekel, Fischer, and Procaccia (2010) make, and is given in the full version [5]. Informally, the reason PROJECT-AND-FIT is strategyproof is that by misreporting her points, misreporting their weights and/or by hiding points, an agent effectively only changes the x_i^* computed by the mechanism for her. Because median (the fit step) is strategyproof, the agent would want the mechanism to compute her correct x_i^*, and thus cannot gain by any manipulation.

Social Cost Objective

Next, we analyze the worst-case approximation ratio of this mechanism for social cost and unweighted social cost objectives. For this, we need the following technical result; its proof is given in the full versions.

Lemma 1. *Let $x_M \in \mathbb{R}$ and $\alpha \in (0, 1]$. If*

$$\left[\sum_{\substack{i,j:\\ x_{i,j} \leq x_M}} w_i(x_{i,j}) \geq \alpha n\right] \bigwedge \left[\sum_{\substack{i,j:\\ x_{i,j} \geq x_M}} w_i(x_{i,j}) \geq \alpha n\right],$$

then $\text{sc}(x_M)/\text{sc}(x^*) \leq (1-\alpha)/\alpha$. Similarly, if

$$\Big[|\{x_{i,j} : x_{i,j} \leq x_M\}| \geq \alpha n\Big] \wedge \Big[|\{x_{i,j} : x_{i,j} \geq x_M\}| \geq \alpha n\Big],$$

then $\text{usc}(x_M)/\text{usc}(x^*) \leq (m-\alpha)/\alpha$.

Dekel, Fischer, and Procaccia (2010) show that for the uniform weight case, PROJECT-AND-FIT gives a 3-approximation to social cost. We show that this remains true in our more general setting. Our proof, given in the full version, draws ideas from their proof and uses Lemma 1.

Theorem 2. PROJECT-AND-FIT *gives a 3-approximation to social cost in the worst case.*

Dekel, Fischer, and Procaccia (2010) also show that 3 is the best possible approximation ratio to social cost by any strategyproof mechanism with respect to misreporting. Hence, their result continues to hold in our case without the uniform weight assumption.

What about mechanisms that are strategyproof with respect to hiding? At first glance, it may seem that hiding is a significantly weaker manipulation than misreporting. For instance, the strategy space is infinite for misreporting, but finite for hiding. Thus, one might expect it to be easier to achieve strategyproofness with respect to hiding than it is with respect to misreporting. Nonetheless, we show that 3 is also the best approximation ratio to social cost by any strategyproof mechanism with respect to hiding. Note that this negative result holds even if agents cannot manipulate their weight information, and continues to hold if they can.

Theorem 3. *For any $\epsilon > 0$, there is no full information mechanism that is strategyproof with respect to hiding and provides a $3 - \epsilon$ approximation to social cost in the worst case, even when there are only two agents.*

Proof. This proof leverages some of the ideas from the proof of Theorem 5.3 by Dekel, Fischer, and Procaccia (2010), but also introduces new ideas to make the proof work for hiding rather than misreporting. Fix $\epsilon > 0$. Suppose for contradiction that there exists a full information mechanism M that is strategyproof with respect to hiding and achieves $3 - \epsilon$ approximation to social cost.

First, we construct another full information mechanism $\widehat{M}$ which is also strategyproof with respect to hiding and has no greater approximation ratio to social cost than M does. Later, we show that $\widehat{M}$ cannot provide $3 - \epsilon$ approximation.

Construction of $\widehat{M}$: Mechanism $\widehat{M}$, on a given instance I, first constructs an instance $\widehat{I}$ by removing all zero-weight points from I, and then returns $M(\widehat{I})$. Let us argue that this is strategyproof. Suppose for contradiction that there exists a pair of instances I and I' which only differ because in I', agent i hides some of her points from I, and $c_i(\widehat{M}(I')) < c_i(\widehat{M}(I))$. However, since $\widehat{M}(I) = M(\widehat{I})$ and $\widehat{M}(I') = M(\widehat{I'})$, we also have $c_i(M(\widehat{I'})) < c_i(M(\widehat{I}))$. Given that $\widehat{I'}$ can be obtained from $\widehat{I}$ with agent i hiding points, this contradicts strategyproofness of M. To see that the worst-case approximation ratio of $\widehat{M}$ is no worse than that of M, note that the approximation ratio of $\widehat{M}$ on instance I is precisely the approximation ratio of M on instance $\widehat{I}$ since zero-weight points do not change social cost.

The benefit of constructing $\widehat{M}$ is that we know its output does not change when zero-weight points are added or

[4] For n points, where n is even, median should either always return the $(n/2)^{\text{th}}$ smallest point or always return the $(n/2+1)^{\text{th}}$ smallest point.

[5] The full version is available at: http://www.cs.toronto.edu/~nisarg/papers/hiding.pdf

removed from an instance, and this helps us derive a lower bound on its worst-case approximation ratio.

Claim 1. *Let $q \in \mathbb{N} \cup \{0\}$. Then, there exists an instance with two agents, $I_q = (D_i, w_i)_{i\in[2]}$, satisfying $D_i = \{x_i\}$ and $w_i(x_i) = 1$ for each $i \in [2]$, and $x_1 - x_2 = 2^q$, such that either $\widehat{M}(I_q) \geq x_1 - 1/2$ or $\widehat{M}(I_q) \leq x_2 + 1/2$.*

The proof of the claim is given in the full version. Now, we derive a contradiction to the assumption that $\widehat{M}$ provides $3 - \epsilon$ approximation to social cost.

Consider an instance I_q with $D_1 = \{x_1\}$ and $D_2 = \{x_2\}$ constructed in Lemma 1. Let us denote $x_{\widehat{M}} = \widehat{M}(I_q)$. Without loss of generality, assume that $x_{\widehat{M}} \geq x_1 - 1/2$ (the argument for the other case is symmetric). First, we argue that $x_{\widehat{M}} \leq x_1$. Suppose for contradiction that $x_{\widehat{M}} > x_1$. Consider another instance I' which is obtained from I by adding a point at x_1 with zero weight to D_2 (i.e. $D'_1 = D_1 = \{x_1\}$ and $D'_2 = \{x_2, x_1\}$ where $w'_2(x_1) = 0$). Because $\widehat{M}$ is unaffected by zero-weight points, it still returns $x_{\widehat{M}}$. Next, construct an instance I'' where $D''_1 = D_1 = \{x_1\}$ and $D''_2 = \{x_1\}$. For $\widehat{M}$ to have any finite approximation of social cost, it must return x_1 on I'', which violates strategyproofness for agent 2 because I'' can be obtained from I' when agent 2 hides point x_2.

We have thus established $x_{\widehat{M}} \in [x_1 - 1/2, x_1]$. Now, consider a new instance $\tilde{I}$ in which $\tilde{D}_2 = D_2$, $\tilde{D}_1 = \{x_1, x_2\}$, $\tilde{w}_1(x_1) = 1/2 + \epsilon/8$, and $\tilde{w}_1(x_2) = 1/2 - \epsilon/8$. Let $\tilde{x}_{\widehat{M}} = \widehat{M}(\tilde{I})$ denote the output of the mechanism on this instance. We consider three cases, and in each case, we either derive a contradiction to strategyproofness of $\widehat{M}$ or to its $3 - \epsilon$ worst-case approximation ratio.

1. Suppose $\tilde{x}_{\widehat{M}} < x_2$. Then, the cost to agent 1 without hiding x_2 is more than $(1/2+\epsilon/8)\cdot(x_1-x_2)$. In contrast, when the agent hides x_2, the outcome of the mechanism is $x_{\widehat{M}} \in [x_1 - 1/2, x_1]$, and her cost is at most

$$\left(\frac{1}{2} + \frac{\epsilon}{8}\right) \cdot (1/2) + \left(\frac{1}{2} - \frac{\epsilon}{8}\right) \cdot (x_1 - 1/2 - x_2)$$
$$= \left(\frac{1}{2} - \frac{\epsilon}{8}\right)(x_1 - x_2) + \frac{\epsilon}{8} < \left(\frac{1}{2} + \frac{\epsilon}{8}\right)(x_1 - x_2),$$

   where the last inequality holds because $x_1 - x_2 = 2^q \geq 1$. Hence, the agent benefits by hiding x_2, which is a contradition to strategyproofness of $\widehat{M}$.

2. Suppose $\tilde{x}_{\widehat{M}} > x_1$. Then, under $\tilde{I}$, we have $\mathrm{sc}(\tilde{x}_{\widehat{M}}) > (1/2 - \epsilon/8 + 1)(x_1 - x_2)$, whereas $\mathrm{sc}(x_2) \leq (1/2 + \epsilon/8)(x_1 - x_2)$. Hence, the approximation ratio of $\widehat{M}$ is at least $(3/2 - \epsilon/8)/(1/2 + \epsilon/8) > 3 - \epsilon/8$, which is a contradiction.

3. Finally, suppose $\tilde{x}_{\widehat{M}} \in [x_2, x_1]$. Then, noting that agent 1 should not be able to gain by hiding x_2 in $\tilde{I}$, we get

$$\left(\frac{1}{2} + \frac{\epsilon}{8}\right)(x_1 - \tilde{x}_{\widehat{M}}) + \left(\frac{1}{2} - \frac{\epsilon}{8}\right)(\tilde{x}_{\widehat{M}} - x_2)$$
$$\leq \left(\frac{1}{2} + \frac{\epsilon}{8}\right)(x_1 - x_{\widehat{M}}) + \left(\frac{1}{2} - \frac{\epsilon}{8}\right)(x_{\widehat{M}} - x_2),$$

   which implies $\tilde{x}_{\widehat{M}} \geq x_{\widehat{M}} \geq x_1 - 1/2$. Hence, $\mathrm{sc}(\tilde{x}_{\widehat{M}}) \geq (1/2 - \epsilon/8 + 1)(2^q - 1/2)$, whereas $\mathrm{sc}(x_2) = (1/2 + \epsilon/8) \cdot 2^q$. It is easy to check that for a sufficiently large q, $\frac{3/2-\epsilon/8}{1/2+\epsilon/8} \cdot \frac{2^q-1/2}{2^q} > 3 - \epsilon$, which is a contradiction.

This completes the entire proof. □

Our results so far establish that PROJECT-AND-FIT gives the lowest approximation ratio to social cost (which is 3) among all mechanisms that are strategyproof with respect to misreporting or hiding.

Unweighted Social Cost Objective

Next, we show that PROJECT-AND-FIT also gives asymptotically lowest approximation to unweighted social cost among all mechanisms that are strategyproof with respect to misreporting or hiding; however, this approximation ratio is now $\Theta(m)$. Recall that m is the maximum number of points held by any agent. We begin by establishing an upper bound on the approximation ratio of PROJECT-AND-FIT for unweighted social cost.

Theorem 4. *In the worst case,* PROJECT-AND-FIT *gives $(2m - 1)$-approximation to unweighted social cost.*

Proof. Fix an instance $I = (D_i, w_i)_{i\in N}$. Let x^* denote an optimal solution for unweighted social cost, and x_M denote the output of PROJECT-AND-FIT. Let x_i^* denote the location computed for agent i in step 2 of PROJECT-AND-FIT.

As the mechanism returns median of all x_i^*, it holds that $|\{x_i^* : x_i^* \leq x_M\}| \geq n/2$. Further, as we noted earlier, x_i^* is a weighted median of points held by agent i. In particular, our tie-breaking in step 2 of the algorithm ensures that it must be one of the points held by agent i.

Hence, we have that $|\{x_{i,j} : x_{i,j} \leq x_M\}| \geq n/2$, and by a symmetric argument, $|\{x_{i,j} : x_{i,j} \geq x_M\}| \geq n/2$. The result now follows by applying Lemma 1. □

Next, we show that no strategyproof mechanism with respect to misreporting can achieve an asymptotically better approximation ratio to unweighted social cost. The proof is in the full version.

Theorem 5. *The worst-case approximation ratio to unweighted social cost of a full information mechanism that is strategyproof with respect to misreporting is at least $m - 1$.*

Finally, we show that no strategyproof mechanism with respect to hiding can achieve an asymptotically better approximation. The proof is provided in the full version.

Theorem 6. *The worst-case approximation ratio to unweighted social cost of a full information mechanism that is strategyproof with respect to hiding is at least $m - 1$.*

Hiding versus Misreporting

So far, our results point out striking similarities between misreporting and hiding manipulations: (a) PROJECT-AND-FIT is strategyproof with respect to both; (b) we prove a lower bound of 3 (resp. $m-1$) for the worst-case approximation to social cost (resp. unweighted social cost) for strategyproof mechanisms with respect to each type of manipulation, and this bound is tight (resp. asymptotically tight).

Could there be a deeper connection between the family of strategyproof mechanisms with respect to hiding and the family of strategyproof mechanisms with respect to misreporting? For the full information case, we show that the two families are at least incomparable, i.e., each family contains a mechanism that is not in the other family.

One reason is that the two manipulations place very different restrictions on what the agents absolutely cannot do. Under misreporting, agents cannot change the number of points they hold, and under hiding, agents cannot expand the support of the set of points they hold. We utilize this to create mechanisms that are strategyproof with respect to one type of manipulation but not with respect to the other.

For example, consider mechanism $M_{\text{misreport}}$ which returns 0 when the total number of points reported by the agents is exactly n, and 1 otherwise. This is clearly strategyproof with respect to misreporting because agents cannot change the number of reported points, and thus cannot influence which of the two constant mechanisms ("return 0" and "return 1") is used. Since each constant mechanism is strategyproof, so is the overall mechanism. However, it is also easy to see that this is not strategyproof for hiding. Consider an instance where each agent except agent 1 has a single point. Agent 1 has two points: one at 0 with weight 1, and one at 1 with weight 0. If the agents report honestly, the mechanism returns 1. But agent 1 can benefit by hiding the point at 1, resulting in the mechanism returning 0.

Interestingly, we could not find an equally trivial mechanism that is strategyproof with respect to hiding but not with respect to misreporting. In the full version, we present a mechanism that leverages PROJECT-AND-FIT as a subroutine to achieve this. This yields the following result.

Theorem 7. *There exists a full information mechanism that is strategyproof with respect to hiding but not with respect to misreporting, and also one that is strategyproof with respect to misreporting but not with respect to hiding.*

4 Ordinal Mechanisms

We now consider mechanisms which do not elicit full information about agents' weights. In particular, we start by studying ordinal mechanisms, which ask agents to report only a ranking of the points by their weight (rather than the exact weights), in addition to reporting the points. That is, the response of each agent i is $\rho_i = (D_i, \sigma_i)$, where σ_i is a ranking of points in D_i by their weight.

We remark that eliciting less information can only constrain the family of strategyproof mechanisms, when we view mechanisms as functions mapping instances to their corresponding outputs. For example, if there is an ordinal strategyproof mechanism M_{ord}, we can construct an equivalent[6] full information strategyproof mechanism M_{full} which elicits full information, converts the reported weights into a ranking of points by their weight, and feeds it as input to M_{ord}. It is easy to see that agents would have no incentive to manipulate under M_{full} as well.

Strategyproof Ordinal Mechanisms

We are interested in studying how limited the family of ordinal strategyproof mechanisms is compared to the family of full information strategyproof mechanisms. First, we consider strategyproofness with respect to hiding. Our next theorem shows that this family is significantly limited, and contains only constant mechanisms. Note that the result holds even when agents cannot manipulate the ranking of points, and continues to hold if they can.

Theorem 8. *The only ordinal mechanisms that are strategyproof with respect to hiding are constant mechanisms.*

Proof. For contradiction, assume that there is a non-constant ordinal mechanism M that is strategyproof with respect to hiding. Then, there are two instances I and I' such that M has different outputs on these instances, i.e., $M(I) \neq M(I')$. Let $x = M(I)$ and $x' = M(I')$. Without loss of generality, assume that both instances have n agents.[7] Let $I = (D_i, w_i)_{i \in N}$ and $I' = (D'_i, w'_i)_{i \in N}$. Let σ_i (resp. σ'_i) denote the ranking of points in D_i (resp. D'_i) induced by the weight function w_i (resp. w'_i).

Now, consider an instance $\overline{I}^1$ that is similar to I' except that for agent 1, $\overline{D}^1_1 = D_1 \cup D'_1 \cup \{x, x, x'\}$ and $\overline{w}^1_1(x') = 1$. Let $\overline{\sigma}^1_1 = x' \succ x \succ x \succ \sigma_1 \succ \sigma'_1$ be the ranking that agent 1 chooses to submit; note that this is consistent with the weight function $\overline{w}^1_1$. Then, the output of M on $\overline{I}^1$ must be x', otherwise agent 1 would have an incentive to hide some of her points and return to instance I'. Next, for $k \in \{2, \ldots, n\}$ we similarly create $\overline{I}^k$ from $\overline{I}^{k-1}$ by changing the set of points held by agent k to $\overline{D}^k_k = D_k \cup D'_k \cup \{x, x, x'\}$, setting $\overline{w}^k_k(x') = 1$, and having the agent submit the ranking $\overline{\sigma}^k_k = x' \succ x \succ x \succ \sigma_k \succ \sigma'_k$. By the same argument, the output of M must be $M(\overline{I}^k) = x'$.

Specifically, note that $M(\overline{I}^n) = x'$. In this instance, each agent i holds the set $\overline{D}^n_i = D_i \cup D'_i \cup \{x, x, x'\}$ and submits the ranking $x' \succ x \succ x \succ \sigma_i \succ \sigma'_i$.

Next, construct an instance I^* that is similar to $\overline{I}^n$ except the weight function of each agent i is changed so that $w^*_i(x') = 1/3 + \epsilon$ and $w^*_i(x) = 1/3 - \epsilon/2$ for each point x, where $\epsilon \in (0, 1/6)$; the remaining points still have zero weight. Suppose each agent i still submits the same ranking of points $\sigma^*_i = x' \succ x \succ x \succ \sigma_i \succ \sigma'_i$, which is still consistent with the weights. Since I^* is indistinguishable from $\overline{I}^n$ to the ordinal mechanism M, it must return x' on I^*.

[6] Two mechanisms are called equivalent if they return the same output on each instance.

[7] If they have different number of agents, we can add dummy agents with no points to the instance with fewer agents.

However, note that in I^*, agent 1 strictly prefers outcome x to any other outcome. Hence, she should not be able to obtain outcome x by hiding some of her points. Specifically, construct instance $\widehat{I}^1$ which is similar to I^* except the set of points held by agent 1 is $\widehat{D}_1^1 = D_1$, she has equal weight for all these points, and she submits ranking σ_i over these points. Then, the output of M on $\widehat{I}^1$ should not be x. Now, for $k \in \{2, \ldots, n\}$, we similarly construct instance $\widehat{I}^k$ from $\widehat{I}^{k-1}$ by changing the data of agent k, and obtain that the output of the mechanism cannot be x. However, the final instance $\widehat{I}^n$ is precisely instance I, on which the output of the mechanism is x, which is the desired contradiction. □

As we argued before, eliciting less information only restricts the family of strategyproof mechanisms. Hence, Theorem 8 immediately yields the following negative result for weightless and anonymous mechanisms.

Corollary 1. *The only weightless or anonymous mechanisms that are strategyproof with respect to hiding are constant mechanisms.*

Thus, for hiding, there exist good full information strategyproof mechanisms (such as PROJECT-AND-FIT), but as soon as we drop to the ordinal case, surprisingly, we have nothing but constant mechanisms.

For misreporting, clearly constant mechanisms are also strategyproof. But it is also easy to construct non-constant mechanisms that are strategyproof. In fact, mechanism $M_{\text{misreport}}$ that we constructed in Section 3 is *anonymous* (and therefore, also weightless and ordinal) and yet strategyproof with respect to misreporting. This implies that in the ordinal, weightless, and anonymous cases, the family of mechanisms that are strategyproof with respect to hiding is a strict subset of the family of mechanisms that are strategyproof with respect to misreporting. In this sense, hiding is a stronger manipulation than misreporting in these three cases. In our opinion, this is not only in stark contrast to the full information case (where the two families are incomparable, as shown in Theorem 7), but also counter-intuitive.

Approximation Ratio with SP

In approximating our two objective functions (social cost and unweighted social cost), Theorem 8 implies that no strategyproof (with regard to hiding) ordinal mechanism can provide a finite approximation to either. To see this, consider a constant mechanism which always outputs x. When all agents have a single point x' with $x' \neq x$, the mechanism has infinite approximation ratio for both objective.

Corollary 2. *Any ordinal mechanism that is strategyproof with respect to hiding has infinite worst-case approximation ratio to social cost and unweighted social cost.*

For ordinal mechanisms that are strategyproof with respect to *misreporting*, we do not have a characterization. Nonetheless, we can show that they also cannot provide a finite approximation. The proof follows roughly the same outline as in the proof of Theorem 8, and is given in the full version. The result holds even if agents cannot manipulate their ranking of points, and continues to hold if they can.

Theorem 9. *Any ordinal mechanism that is strategyproof with respect to misreporting has infinite worst-case approximation ratio to social cost and unweighted social cost.*

Approximation Ratio without SP

We saw that ordinal mechanisms that are strategyproof with respect to hiding or misreporting cannot give finite approximation to social cost or unweighted social cost. There are two potential reasons why: it could be due to the enforcement of strategyproofness, or due to ordinal mechanisms not having access to full information about the weights.

For unweighted social cost, the reason is clearly the former. This is because optimizing unweighted social cost does not require knowledge of weights. Hence, without strategyproofness, one can simply achieve an optimal 1-approximation to this objective via an ordinal mechanism.

The situation is not as clear for social cost. Here, we show that without strategyproofness, the best worst-case approximation ratio to social cost that ordinal mechanisms can provide is $\Theta(m)$. This implies that while we face a significant "price of incomplete information", the unbounded approximation arises as the "price of strategyproofness". We first present the lower bound, with its proof in the full version.

Theorem 10. *Any ordinal mechanism gives $\Omega(m)$ approximation to social cost in the worst case.*

For the upper bound, we construct an ordinal mechanism — MEDIAN-OF-TOPS, given in the full version — which returns the median of top-ranked points of the agents, and show that it achieves $O(m)$ approximation to social cost. The proof of the next result is in the full version.

Theorem 11. MEDIAN-OF-TOPS *achieves $O(m)$ approximation to social cost in the worst case.*

5 Discussion

We considered a facility location setting in which agents hold multiple locations with different weights. We introduced and studied a new type of manipulation, whereby agents can hide some of their points, and compared its power to that of the standard misreporting manipulation, whereby agents change their reported locations.

Our work leaves open multiple directions for future work. An interesting extension would be to use more general preference structures. For instance, our formulation of agent costs uses ℓ_1 distances. Although some of our results extend to more general distances, many of them rely on the ℓ_1 distance. What if a different distance measure is used instead?

The hiding manipulation is also quite realistic in machine learning settings such as linear regression or classification (Perote and Perote-Pena 2004; Meir, Procaccia, and Rosenschein 2012; Chen et al. 2018), where strategic agents provide training datasets and may decide to omit parts of it for their own benefit. Importantly, the negative results in this paper apply to the broader regression setting. Designing learning algorithms robust not only to stochastic noise (Littlestone 1991; Goldman and Sloan 1995) or adversarial noise (Kearns and Li 1993; Chen, Caramanis, and Mannor 2013), but also to "strategic noise" requires investigating ways to address such manipulations.

References

Alon, N.; Feldman, M.; Procaccia, A. D.; and Tennenholtz, M. 2009. Strategyproof approximation mechanisms for location on networks. *arXiv preprint arXiv:0907.2049*.

Ashlagi, I.; Fischer, F.; Kash, I. A.; and Procaccia, A. D. 2015. Mix and match: A strategyproof mechanism for multi-hospital kidney exchange. *Games and Economic Behavior* 91:284–296.

Chen, Y.; Podimata, C.; Procaccia, A. D.; and Shah, N. 2018. Strategyproof linear regression in high dimensions. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 9–26. ACM.

Chen, X.; Li, M.; Wang, C.; Wang, C.; and Zhao, Y. 2019. Truthful mechanisms for location games of dual-role facilities. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, 1470–1478.

Chen, Y.; Caramanis, C.; and Mannor, S. 2013. Robust sparse regression under adversarial corruption. In *International Conference on Machine Learning*, 774–782.

Dekel, O.; Fischer, F.; and Procaccia, A. D. 2010. Incentive compatible regression learning. *Journal of Computer and System Sciences* 76(8):759–777.

Escoffier, B.; Gourves, L.; Thang, N. K.; Pascual, F.; and Spanjaard, O. 2011. Strategy-proof mechanisms for facility location games with many facilities. In *International Conference on Algorithmic DecisionTheory*, 67–81. Springer.

Fadaei, S., and Bichler, M. 2017. Generalized assignment problem: Truthful mechanism design without money. *Operations Research Letters* 45(1):72–76.

Feldman, M., and Wilf, Y. 2013. Strategyproof facility location and the least squares objective. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, 873–890.

Filos-Ratsikas, A.; Li, M.; Zhang, J.; and Zhang, Q. 2017. Facility location with double-peaked preferences. *Autonomous Agents and Multi-Agent Systems* 31(6):1209–1235.

Goldman, S. A., and Sloan, R. H. 1995. Can PAC learning algorithms tolerate random attribute noise? *Algorithmica* 14(1):70–84.

Hossain, S., and Shah, N. 2019. Pure nash equilibria in linear regression. Manuscript.

Kearns, M., and Li, M. 1993. Learning in the presence of malicious errors. *SIAM Journal on Computing* 22(4):807–837.

Littlestone, N. 1991. Redundant noisy attributes, attribute errors, and linear-threshold learning using winnow. In *Proceedings of the 4th Conference on Computational Learning Theory (COLT)*, 147–156.

Meir, R.; Procaccia, A. D.; and Rosenschein, J. S. 2012. Algorithms for strategyproof classification. *Artificial Intelligence* 186:123–156.

Moulin, H. 1980. On strategy-proofness and single peakedness. *Public Choice* 35(4):437–455.

Moulin, H. 1984. Generalized Condorcet-winners for single peaked and single-plateau preferences. *Social Choice and Welfare* 1(2):127–147.

Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V. V. 2007. *Algorithmic game theory*. Cambridge university press.

Perote, J., and Perote-Pena, J. 2004. Strategy-proof estimators for simple regression. *Mathematical Social Sciences* 47(2):153–176.

Procaccia, A. D., and Tennenholtz, M. 2009. Approximate mechanism design without money. In *Proceedings of the 10th ACM conference on Electronic commerce*, 177–186. ACM.

Roth, A. E.; Sönmez, T.; and Ünver, M. U. 2004. Kidney exchange. *The Quarterly Journal of Economics* 119(2):457–488.

Schummer, J., and Vohra, R. V. 2002. Strategy-proof location on a network. *Journal of Economic Theory* 104(2):405–428.

Sui, X.; Boutilier, C.; and Sandholm, T. 2013. Analysis and optimization of multi-dimensional percentile mechanisms. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, 367–374.