

Approximately Strategy-proof Mechanisms for (Constrained) Facility Location

Xin Sui

Department of Computer Science
University of Toronto, Toronto, Canada
xsui@cs.toronto.edu

Craig Boutilier

Department of Computer Science
University of Toronto, Toronto, Canada
cebly@cs.toronto.edu

ABSTRACT

Mechanism design for facility location (or selection of alternatives in a metric space) has been studied for decades. While strategy-proof, efficient mechanisms exist for unconstrained, one-dimensional, single-facility problems, guarantees of strategy-proofness and efficiency often break when allowing: (a) multiple dimensions; (b) multiple facilities; or (c) constraints on the feasible placement of facilities. We address these more general problems, providing several possibility/impossibility results with respect to individual and group strategy-proofness in both constrained and unconstrained problems. We also bound the incentive for manipulation in median-like mechanisms in settings where strategy-proofness is not possible. We complement our results with empirical analysis of both electoral and geographic facility data, showing that the odds of successful manipulation, and more importantly, the gains and impact on social welfare, are small in practice (much less than worst-case theoretical bounds).

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence, Multi-agent Systems

General Terms

Algorithms, Economics, Theory

Keywords

Social choice; Facility location; Single-peaked preferences; Mechanism design

1. INTRODUCTION

Mechanism design deals with the design of protocols to elicit individual preferences while achieving some social objective (e.g., maximizing social welfare) [17]. An important property of mechanisms is *strategy-proofness*, which requires that no agent can gain (or induce a more preferred outcome) by misreporting her preferences to the mechanism. While classic results preclude the possibility of strategy-proof mechanisms for arbitrary social choice functions [12,

23], a number of important classes of problems admit powerful strategy-proof mechanisms.

A well-studied class of problems are those in which agent preferences are *single-peaked* [4]. In this setting, an ordering over outcomes is given. Each agent has a single most-preferred outcome—her *peak* or *ideal point*—and her preferences for other outcomes are dictated by the “distance” from this peak. In this setting, the *median mechanism* and its generalizations [19, 2] are strategy-proof and have a number of interesting properties. This approach is sometimes referred to as *mechanism design without money* [21, 25]. Special cases include the restriction to preferences defined on metric spaces, and generalizations include extensions to multiple facilities, multiple dimensions, and constraints on feasible outcomes (i.e., where the preference space is richer than the allowable outcome space). These are often called *facility location problems (FLPs)*, referring to the choice of one or more facilities to serve multiple agents, whose preferences are dictated by the distances from their “ideal points.” However, the problem is much more general, encompassing voting (in which candidates fall on some political spectrum), product design and customer segmentation (where products lie in some feature space), and other domains.

In this work, we address mechanism design in the multi-dimensional case when multiple facilities can be chosen, addressing both *unconstrained FLPs*—in which facilities can be placed at any point in some (metric) space—and *constrained FLPs*—in which some outcomes in the preference space are not feasible (i.e., the outcome space is constrained). In particular, we consider cases in which strategy-proofness cannot be achieved, and analyze *approximately strategy-proof mechanisms*. If one can bound the potential gain an agent (or group) can obtain by misreporting their preferences, the cost of determining an optimal misreport may outweigh the benefits of misreporting, rendering such mechanisms “practically strategy-proof” [14, 16].

In unconstrained problems, individual strategy-proofness can be achieved using *generalized median mechanisms*, or *GMMs* [19, 2] for single-FLPs, and *quantile mechanisms (QMs)* for multi-FLPs [28]. Unfortunately, group strategy-proofness is unachievable in general. Our first contribution is to provide an impossibility result showing that the incentive for any group of agents to misreport is unbounded (w.r.t. arbitrary preference profiles). Then we give profile-specific bounds on the incentive to misreport. Second, we analyze constrained FLPs, defining a new family of *closest candidate mechanisms (CCMs)*. CCMs use QMs to determine *tentative* locations, then project these to the nearest feasi-

ble locations using some distance function. While CCMs are not strategy-proof in general, we are able to bound the incentive for individuals and groups to misreport. Finally, we empirically evaluate the performance of our mechanisms using real-world preference data drawn from electoral and geographic facility domains. We evaluate the probability of agents (or groups) successfully manipulating choices, and more importantly show that their expected gain, and impact on social welfare, is quite small in practice. This suggests that the mechanisms analyzed here, namely, GMMs, QMs and CCMs, may be “sufficiently strategy-proof” for practical purposes.

2. BACKGROUND AND NOTATION

We begin by describing FLPs, single-peaked preferences, mechanism design for FLPs, and related work.

2.1 Facility Location Problems

We focus on *multi-dimensional, multi-facility location problems*, where the goal is to select q homogeneous facilities in some m -dimensional space \mathbb{R}^m (or some bounded subspace thereof). Such an outcome is represented by a *location vector* $\mathbf{x} = (x_1, \dots, x_q)$, where $x_j \in \mathbb{R}^m$. We also have a set of *agents* $N = \{1, \dots, n\}$, each with a type $t_i \in T_i$ determining her cost $c_i(\mathbf{x}, t_i)$ associated with any location vector $\mathbf{x} \in (\mathbb{R}^m)^q$. Given vector \mathbf{x} , agent i uses the facility with least cost, that is, $c_i(\mathbf{x}, t_i) = \min_{j \leq q} c_i(x_j, t_i)$, where the latter function $c_i : \mathbb{R}^m \times T_i \rightarrow \mathbb{R}$ is i 's cost function for the use of a specific (single) facility.

FLPs readily capture the placement of q facilities in some geographical space, in which each agent uses the closest facility. However, it models many other problems as well. Voting is one example, where candidates lie in an m -dimensional space, representing their position on various political issues, and the aim is election of q candidates to a legislative body [6, 20]. In product design, a vendor may wish to offer a set of q related products, each described by an m -dimensional feature vector, with consumer preferences over these options leading them to select their most preferred product.

In some cases, facilities can be placed anywhere in the option space—we call these *unconstrained FLPs*. In *constrained FLPs*, outcomes can only be placed at a restricted finite set of feasible locations, $\mathbb{D} = \{d_1, \dots, d_l\}$. Such restrictions often apply in voting (a finite set of candidates under consideration), in facility placement (geographic constraints), product design (selecting from an existing assortment), and other forms of FLPs.

In FLPs, it is natural to assume agent preferences are *single-peaked*. This means each agent has a single “ideal” location, and its cost for any chosen location increases as it “moves away” from this *peak* or *ideal point*.

DEFINITION 1. [2] Agent i 's preference over \mathbb{R}^m is single-peaked if there is a most preferred option $\tau(t_i)$ s.t., $\forall \alpha, \beta \in \mathbb{R}^m$ satisfying $\|\tau(t_i) - \beta\|_1 = \|\tau(t_i) - \alpha\|_1 + \|\alpha - \beta\|_1$, we have $c_i(\alpha, t_i) \leq c_i(\beta, t_i)$, where $\|\cdot\|_1$ is the L_1 -norm.

Intuitively, if a point α lies within the “bounding box” of the ideal point $\tau(t_i)$ and β , then α is at least as preferred as β . This does not restrict i 's relative preference for α and β if neither lies within the other's bounding box w.r.t. $\tau(t_i)$.

2.2 Mechanism Design

In FLPs, a *mechanism* is a function $f : \prod_i T_i \rightarrow (\mathbb{R}^m)^q$ that maps a type profile $\mathbf{t} = (t_1, \dots, t_n)$ to a location vector \mathbf{x} . Once agents reveal their types, the mechanism selects an outcome $f(\mathbf{t})$ based on these reports. A critical property of a mechanism is (*additive approximate*) *strategy-proofness*:

DEFINITION 2. A mechanism f is ε -strategy-proof if:

$$c_i(f(t_i, \mathbf{t}_{-i}), t_i) \leq c_i(f(t'_i, \mathbf{t}_{-i}), t_i) + \varepsilon, \quad \forall i, t_i, t'_i, \mathbf{t}_{-i}.$$

where \mathbf{t}_{-i} is the type profile of all agents but i .

We say the mechanism is *strategy-proof* if this holds for $\varepsilon = 0$. Here the LHS is i 's cost if she reports her type t_i truthfully (given the fixed reports of others), and the first term of the RHS is her cost if she misreports type t'_i . This concept extends to groups by requiring that if any coalition of agents misreports their preferences (in any coordinated fashion), at least one member of the group is not strictly better off.

DEFINITION 3. A mechanism f is ε -group strategy-proof if, for a subset of agents $S \subseteq N$, there is some agent $i \in S$ such that:

$$c_i(f(t_S, \mathbf{t}_{-S}), t_i) \leq c_i(f(t'_S, \mathbf{t}_{-S}), t_i) + \varepsilon, \quad \forall S, t_S, t'_S, \mathbf{t}_{-S}$$

where \mathbf{t}_{-S} is the type profile of all agents in $N \setminus S$.

Similarly, we use the term *group strategy-proof* to mean 0-group strategy-proof. This definition requires that *each* agent in a manipulating coalition S has some gain by participating, which is sensible in settings with non-transferable utility (as is the case in many social choice problems).

Let f be any mechanism, S be a coalition with (fixed) true type profile \mathbf{t}_S , and \mathbf{t}_{-S} be the (fixed) reports of the other agents. We define the *gain* of $i \in S$ for a (coalitional) misreport \mathbf{t}'_S to be $G(i, S, \mathbf{t}'_S) = c_i(f(\mathbf{t}_S, \mathbf{t}_{-S}), t_i) - c_i(f(\mathbf{t}'_S, \mathbf{t}_{-S}), t_i)$; the *maximum gain* of i to be $G(i, S) = \max_{\mathbf{t}'_S} G(i, S, \mathbf{t}'_S)$; and *incentive for S to misreport* to be $G(S) = \max_{i \in S} G(i, S)$. We say a misreport \mathbf{t}'_S is *viable* iff $G(i, S, \mathbf{t}'_S) \geq 0$ for each $i \in S$ and $G(i, S, \mathbf{t}'_S) > 0$ for some $i \in S$.

2.3 Related Work

The study on unconstrained FLPs dates to Black [4], who proposed the *median mechanism* for single FLPs in one dimension (1D). Moulin [19] provides an important characterization: *generalized median mechanisms (GMMs)*, which allowing phantom peaks,¹ comprise the class of all strategy-proof mechanisms. Sui et al. [28] extend GMMs to the multi-facility case with *quantile mechanisms (QMs)*, in which strategy-proofness is guaranteed if each facility is selected independently using a (form of) GMM.

DEFINITION 4. [19] Let b_1, \dots, b_{n+1} be $n + 1$ phantom peaks in $\mathbb{R} \cup \{-\infty, +\infty\}$. The generalized median mechanism locates the facility at the median position of the union of the reported agent peaks and the phantom peaks.

DEFINITION 5. [28] Let $\mathbf{p} = (p_1, \dots, p_q)$ be a vector such that $0 \leq p_1 \leq \dots \leq p_q \leq 1$. The \mathbf{p} -quantile mechanism for multi-FLPs locates the j th facility at the p_j th quantile of the reported peaks.

¹Phantom peaks are the peaks of a set of (pre-existing) hypothetical voters, which are used in Moulin's [19] definition of GMMs and his characterization of strategy-proof mechanisms.

GMMs and QMs offer (group) strategy-proofness in 1D:

REMARK 1. *In one-dimensional, unconstrained facility location problems, GMMs (single facility) and QMs (multiple facilities) are individual and group strategy-proof.*

The extension to multiple dimensions has also been studied. Barberà et al. [2] generalize Moulin’s result to m -dimensions, showing that a mechanism is strategy-proof iff it is a multi-dimensional GMM, which selects a location by choosing its coordinate in each dimension independently. As m -dimensional GMMs are not group strategy-proof, this characterization also serves as an impossibility result. QMs can be generalized to multiple dimensions similarly:

REMARK 2. *In m -dimensional, unconstrained facility location problems, GMMs (single facility) and QMs (multiple facilities) are strategy-proof. However, no (anonymous) group strategy-proof mechanisms exist in this setting.*²

FLPs have also been studied in specific preference domains, including those over metric spaces. Border and Jordan [5] offer characterizations when preferences are *separable star-shaped* (incl. quadratic costs). Massó and Moreno de Barreda [18] consider symmetric, single-peaked preferences (e.g., L_1 or L_2 costs), showing *disturbed GMMs* encompass all strategy-proof mechanisms. FLPs have been also studied on tree- and circle-based extensions of single-peakedness [24, 8], while recent work develops approximation ratios when agents have L_2 preferences [21, 15, 11, 9]. Despite the rich literature on unconstrained FLPs, there is little work on constrained FLPs. Dokow et al. [8] characterize the class of strategy-proof mechanisms on a discrete line, which are similar to the results of Border and Jordan [5], and Massó and Moreno de Barreda [18]. Another exception is the work of Barberà et al. [3], who characterize the class of strategy-proof mechanisms for constrained FLPs. They show that a mechanism is strategy-proof in a constrained setting iff: a) it is a GMM; and b) it satisfies the *intersection property*, which requires that the mechanism must be coordinated in each dimension to guarantee a feasible location. However, their result will, as they say, “anticipate impossibility theorems in most applications.”

3. UNCONSTRAINED FACILITY LOCATION

We begin with an analysis of *approximate group strategy-proofness* for unconstrained FLPs. In 1D problems, GMMs and QMs are group strategy-proof for single- and multi-facility problems, respectively. When we move to multiple dimensions, these mechanisms remain (individual) strategy-proof, but unfortunately, not *group* strategy-proof (indeed, as discussed above, no general group strategy-proof mechanisms exist).

Alternatively, one can try to bound the incentive for any group of agents to misreport, showing GMMs and QMs to be *approximately* group strategy-proof. To quantify this claim, one must make assumptions about agent cost functions. Here we assume that cost is proportional to the L_2 distance from the ideal point, i.e., $c_i(x_j, t_i) = \|x_j - t_i\|_2$.³ Note that an agent’s cost is fully determined by her peak

²Anonymity is critical, as dictatorial mechanisms belong to the class of GMMs and offer group strategy-proofness.

³Barberà et al.’s [2] characterizations do not preclude the existence of group strategy-proof mechanisms when specific




Figure 1: A two-dimensional counter example showing the incentive for a group of agents to misreport can be unbounded.

$\tau(t_i)$ under the L_2 -norm, and we will equate her type t_i with this peak for convenience.

We first give an impossibility result, showing that the incentive for group manipulation can be arbitrarily large.

THEOREM 1. *GMMs and QMs are not ε -group strategy-proof for any fixed $\varepsilon > 0$.*

PROOF. We give a counter-example for two-dimensional, two-facility location under QMs. The result applies directly to GMMs since QMs are a specific instance of GMMs.

Let $\mathbf{p} = (p^1, p^2)$ be a two-dimensional, quantile matrix used for a QM, in which the facility is located at the coordinate of the p^1 th peak in the first dimension, and at the coordinate of the p^2 th peak in the second dimension. Consider the following two cases:

- I. $p^1 + p^2 \leq 1$. Consider the following peak profile $\mathbf{t} = ((a, 0), \dots, (a, 0), (0, a), \dots, (0, a), (a, a), \dots, (a, a))$,
 $\underbrace{\hspace{10em}}_{np^2 \text{ copies}} \quad \underbrace{\hspace{10em}}_{np^1 \text{ copies}} \quad \underbrace{\hspace{10em}}_{n(1-p^1-p^2) \text{ copies}}$

where $a > 0$ is a positive real number (as shown in Fig. 1). The QM will locate the facility at position $(0, 0)$, and the costs of the agents are: a , for those at $(a, 0)$ and $(0, a)$; and $\sqrt{2}a$, for those at (a, a) . However, if all n agents are manipulators, then there exists a viable misreport in which all agents report $(a/2, a/2)$, which will then be selected. The cost under this misreport is $\sqrt{2}/2a$ for each agent, and the gains are: $a - \sqrt{2}/2a \approx 0.293a$, for those at $(a, 0)$ and $(0, a)$; and $\sqrt{2}/2a \approx 0.707a$, for those at (a, a) . As a can be arbitrarily large, so are the gains due to manipulation.

- II. $p^1 + p^2 > 1$. Consider the following peak profile $\mathbf{t} = ((a, 0), \dots, (a, 0), (0, a), \dots, (0, a), (0, 0), \dots, (0, 0))$.
 $\underbrace{\hspace{10em}}_{n(1-p^1) \text{ copies}} \quad \underbrace{\hspace{10em}}_{n(1-p^2) \text{ copies}} \quad \underbrace{\hspace{10em}}_{n(p^1+p^2-1) \text{ copies}}$

The QM will locate the facility at (a, a) , and the agents costs are: a , for those at $(a, 0)$ and $(0, a)$; and $\sqrt{2}a$, for those at $(0, 0)$. As above, a viable manipulation exists in which each manipulator misreports $(a/2, a/2)$, and again the gain of the manipulators is arbitrarily large as $a \rightarrow \infty$.

This demonstrates that the (additive) incentive for manipulation is unbounded for GMMs and QMs. \square

cost functions are used (e.g., L_2 -norm). However, it is still meaningful to study the group manipulation of GMMs and QMs due to their simplicity and intuitive nature, their (individual) strategyproofness, and flexibility (e.g., the fact that they can be optimized or tuned for specific prior distributions over preferences). Similar remarks apply to the negative results of Barberà et al. [3] for constrained FLPs.

We note the unboundedness of the gain is in an additive sense; the relative gain in this example is of course bounded. While this observation serves as a negative result, it is an *a priori* worst-case analysis, allowing arbitrary preference profiles. In practice, the incentive for a group of agents to misreport depends on the actual ideal points of the sincere agents and the manipulators. We provide a profile-specific, *a posteriori* bound that relies on knowledge of voter preference, \mathbf{t}_S and \mathbf{t}_{-S} , and hence can be used when some rough idea of the configuration of preferences (peaks) is known.

We begin with single-facility case, providing an upper bound on the incentive for a group of agents to misreport.

DEFINITION 6. Let $S \subseteq N$ be a set of manipulators. A misreport \mathbf{t}'_S is Pareto optimal if there is no other misreport \mathbf{t}''_S such that $c_i(f(\mathbf{t}''_S, \mathbf{t}_{-S})) \leq c_i(f(\mathbf{t}'_S, \mathbf{t}_{-S}))$ for all $i \in S$ and for some $i^* \in S$, we have $c_{i^*}(f(\mathbf{t}''_S, \mathbf{t}_{-S})) < c_{i^*}(f(\mathbf{t}'_S, \mathbf{t}_{-S}))$.

Intuitively, a misreport is Pareto optimal if there is no other misreport such that no manipulator is worse off and at least one is strictly better off.

When bounding the incentive for a group of agents to misreport, we can focus on Pareto optimal misreports without loss of generality (since a Pareto improvement to a non-Pareto optimal misreport will improve the lot of the manipulators and can only increase the upper bound on this incentive). The following lemma provides a necessary condition for a misreport to be Pareto optimal under GMMs or QMs. Let $S \subseteq N$ be a set of manipulators, and $x = f(\mathbf{t}_S, \mathbf{t}_{-S})$ be the chosen location given truthful reports. We use superscript k to index dimensions, and define $I^k = [\min_{i \in S} t_i^k, \max_{i \in S} t_i^k]$ as the bounding interval of the manipulators' peaks in the k th dimension. We have:

LEMMA 1. Let \mathbf{t}'_S be a Pareto optimal misreport and $x' = f(\mathbf{t}'_S, \mathbf{t}_{-S})$ be the location chosen given misreport \mathbf{t}'_S under a GMM or QM. Then we have $x'^k \in I^k$ if $x^k \in I^k$ and $x'^k = x^k$ otherwise.

PROOF. Suppose the lemma does not hold. Then for each dimension k , one of the following two situations must arise:

- I. $x'^k \notin I^k$ and $x^k \in I^k$. Note $x'^k \notin I^k$ means either $x'^k < \min_i t_i^k$ or $x'^k > \max_i t_i^k$, and w.l.o.g., we assume it is the former case. Recall that we have $f^k(\mathbf{t}'_S, \mathbf{t}_{-S}) = x'^k < x^k = f^k(\mathbf{t}_S, \mathbf{t}_{-S})$, which means there must be some manipulator whose misreport lies to the left of (is less than) $\min_i t_i^k$ in the k th dimension. We can construct another misreport \mathbf{t}''_S such that $f^k(\mathbf{t}''_S, \mathbf{t}_{-S}) = \min_i t_i^k$ and $f^{\tilde{k}}(\mathbf{t}''_S, \mathbf{t}_{-S}) = x^{\tilde{k}}, \forall \tilde{k} \neq k$, and each manipulator i strictly gains in the k th dimension without losing in any other dimension. This means $c_i(f(\mathbf{t}''_S, \mathbf{t}_{-S})) < c_i(f(\mathbf{t}'_S, \mathbf{t}_{-S}))$, which contradicts our assumption that \mathbf{t}'_S is a Pareto optimal misreport.
- II. $x'^k \neq x^k$ and $x^k \notin I^k$. Similarly $x^k \notin I^k$ means either $x^k < \min_i t_i^k$ or $x^k > \max_i t_i^k$, and w.l.o.g., we assume it is the former case. Since QMs locate the facility at a specified quantile, we must have $x'^k < x^k$. We can construct another misreport \mathbf{t}''_S such that $f^k(\mathbf{t}''_S, \mathbf{t}_{-S}) = x^k$ and $f^{\tilde{k}}(\mathbf{t}''_S, \mathbf{t}_{-S}) = x^{\tilde{k}}, \forall \tilde{k} \neq k$, and each manipulator i strictly gains in the k th dimension without losing in any other dimension. This too contradicts the Pareto optimality of \mathbf{t}'_S .




Figure 2: A two-dimensional example showing that a viable misreport must induce a location contained in $C^\perp(S)$ (the shaded area).

□

Lemma 1 shows that, when bounding the incentive to misreport, we can focus our attention on those dimensions in which the coordinate of the facility selected under truthful reporting lies within the corresponding bounding interval—for those dimensions where this is not true, the manipulators can safely leave their reports on those dimensions unchanged (i.e., report sincerely).

Before describing our bound, we first introduce some notation. Let $S \subseteq N$ be a set of manipulators and x be the chosen location under truthful reporting. We define $C(i, x) = \{\bar{x} \in \mathbb{R}^m : \|\bar{x} - t_i\|_2 \leq \|t_i - x\|_2\}$ to be the circle centered at t_i with radius $\|t_i - x\|_2$. Let $C(S) = \bigcap_{i \in S} C(i, x)$ denote the intersection of these circles. Let I^k be the bounding interval as defined in Lemma 1, and $C^\perp(S) = \{\bar{x} \in \mathbb{R}^m : \bar{x}^k \in C^k(S) \text{ if } x^k \in I^k \text{ and } \bar{x}^k = x^k \text{ otherwise}\}$ be the projection of $C(S)$ onto the subspace of \mathbb{R}^m in which we fix the coordinates of x in those dimensions k not contained in the bounding intervals to x^k . We have the following theorem:

THEOREM 2. For single-facility location under GMMs/QMs, the incentive for a set of manipulators S to misreport is:

$$\varepsilon_S = \max_{g \in C^\perp(S)} \left[\max_{i \in S} (\|x - t_i\|_2 - \|g - t_i\|_2) \right].$$

PROOF. (Sketch) Let \mathbf{t}'_S be any group misreport and $g = f(\mathbf{t}'_S, \mathbf{t}_{-S})$ be the induced location of the facility. The first thing to note is that for the misreport \mathbf{t}'_S to be viable, the induced location g must be contained in $C(S)$, otherwise there will be some manipulator who is strictly worse-off (as shown in Fig. 2).

By Lemma 1, we need only consider the projection of $C(S)$ onto the subspace $C^\perp(S)$ (as defined above). For each location g , the gain of manipulator i is $\|x - t_i\|_2 - \|g - t_i\|_2$. If we take the maximum over all manipulator, and all possible locations g , we obtain the stated bound. □

In the multi-facility case, we provide an upper bound on the incentive to misreport by considering each facility independently. Formally, let $S \subseteq N$ be a set of manipulators, and $\mathbf{x} = f(\mathbf{t}_S, \mathbf{t}_{-S})$ be the chosen location vector under truthful reporting. For each facility j with location $x_j \in \mathbf{x}$, we define $S_j = \{i \in S : j = \arg \min_{j' \leq q} \|x_{j'} - t_i\|_2\}$ as the set of manipulators whose closest facility is j under truthful report. We also define $C(i, x_j)$ and $C(S_j)$ similarly as in the single-facility case, and $D_j = \{i \in S_j : \exists j' \text{ s.t. } C(i, x_j) \cap C(S_{j'}) \neq \emptyset\}$ as the set of manipulators in S_j whose circles intersect with $C(S_{j'})$ for some other facility j' . Intuitively, D_j denotes the set of manipulators in

S_j who may deviate from using facility j to use other facilities. Then we have:

THEOREM 3. *For multi-facility location under GMMs/QMs, the incentive for a set of manipulators S to misreport is $\varepsilon_S = \max_j \varepsilon_{S_j}$, where:*

$$\varepsilon_{S_j} = \max_{g \in C^\perp(S_j \setminus D_j)} \left[\max_i (c_i(\mathbf{x}, t_i) - \|g - t_i\|_2) \right]$$

PROOF. (Sketch) Let \mathbf{t}'_S be any group misreport and $g = f_j(\mathbf{t}'_S, \mathbf{t}_{-S})$ be the induced location of facility j . Among the manipulators whose closest facility is j , we have to preclude those who may deviate from using j to other facilities, which we denote by D_j .

By Lemma 1, we can focus only on the projection of $C(S_j \setminus D_j)$ onto subspace in which the coordinates of x_j is not contained in the bounding intervals. If we take the maximum over all manipulators, over all possible locations g for each facility, and over all facilities, we obtain the stated bound. \square

4. CONSTRAINED FACILITY LOCATION

We now turn our attention to constrained FLPs, in which the feasible locations are a strict subset of those over which agent preferences range. For instance, in political settings, agent i 's ideal point may correspond to a "fictitious" candidate who agrees with i on every issue, while selection is limited to those "actual" candidates who have agreed to stand for election. We assume a finite set $\mathbb{D} = \{d_1, \dots, d_l\} \subset \mathbb{R}^m$ of *feasible locations*.

Barberà et al.'s [3] characterization of strategy-proofness in constrained FLPs (see Sec. 2.3) suggests that strategy-proofness is not attainable in most practical settings. So as above, we turn our attention to *approximately* strategy-proof mechanisms. Focusing on QMs (since they apply to single- and multi-FLPs), we deal with constraints by defining *closest candidate mechanisms (CCMs)* (CCMs), and assume L_2 -distance for costs and "projection:"

MECHANISM 1. *Let $\mathbb{D} = \{d_1, \dots, d_l\}$ be a set of feasible locations, and f' a (multi-dimensional) QM. A closest candidate mechanism (CCM) f , based on QM f' , selects a location vector, given reports \mathbf{t} , as follows: (i) let $f'(\mathbf{t}) = \tilde{\mathbf{x}} = \{\tilde{x}_1, \dots, \tilde{x}_q\}$; (ii) return location vector $\mathbf{x} = \{x_1, \dots, x_q\}$, where $x_j = \arg \min_{d \in \mathbb{D}} \|d - \tilde{x}_j\|_2$.*

In other words, the mechanism runs a QM on the reported peaks and replaces any infeasible location $x'_j \notin \mathbb{D}$ by the nearest feasible location in \mathbb{D} . While not strategy-proof in general, CCMs are in fact (group) strategy-proof in 1D:

THEOREM 4. *CCMs are group strategy-proof for 1D FLPs under the L_2 norm.*

PROOF. (Sketch) We provide a sketch assuming $q = 2$. The analysis is easily generalized to $q = 1$ or $q > 2$.

Let $S \subseteq N$, and $\tilde{\mathbf{x}} = \{\tilde{x}_1, \tilde{x}_2\}$ be the location vector chosen by the QM if all agents report truthfully, and $\mathbf{x} = \{x_1, x_2\}$ be the projected location vector into \mathbb{D} . Let $\tilde{\mathbf{x}}' = \{\tilde{x}'_1, \tilde{x}'_2\}$ be the vector chosen by the QM if agents in S jointly misreport, and $\mathbf{x}' = \{x'_1, x'_2\}$ be its projection. W.l.o.g., assume $x_1 < x_2$ and $x'_1 < x'_2$. Consider four cases:

- I. $x_1 \geq x'_1$ and $x_2 > x'_2$: Both x_2 and x'_2 are feasible, so $\tilde{x}'_2 \leq (x'_2 + x_2)/2 \leq \tilde{x}_2$. Since QM chooses each




Figure 3: An example where a manipulator can benefit by changing the outcome from d_1 to d_2 .

location using quantiles, suppose some i , with peak $t_i > \tilde{x}_2$, misreports to the left of \tilde{x}_2 . Then $i \in S$, and i 's cost now is $c_i(\mathbf{x}', t_i) = t_i - x'_2 > t_i - x_2 = c_i(\mathbf{x}, t_i)$.

- II. $x_1 < x'_1$ and $x_2 > x'_2$: As above, there must be some $i \in S$, with peak $t_i > \tilde{x}_2$, who misreports to the left of \tilde{x}_2 . So i 's cost now is $c_i(\mathbf{x}', t_i) = t_i - x'_2 > t_i - x_2 = c_i(\mathbf{x}, t_i)$.
- III. $x_1 < x'_1$ and $x_2 \leq x'_2$: Symmetric to cases I and II.
- IV. $x_1 \geq x'_1$ and $x_2 \leq x'_2$: There must some $i \in S$, with type $\tilde{x}_1 < t_i < \tilde{x}_2$, who misreports to the left of \tilde{x}_1 . Then i 's cost is $c_i(\mathbf{x}', t_i) = \min\{t_i - x'_1, x'_2 - t_i\} \geq \{t_i - x_1, x_2 - t_i\} = c_i(\mathbf{x}, t_i)$.

This establishes group strategy-proofness. \square

One can show that CCMs in the multi-facility case are a straightforward extension of the family of *disturbed GMMs* [18] in the 1D setting, which characterize all strategy-proof mechanisms when agents have symmetric single-peaked preferences (of which L_1 - and L_2 -preferences are a special case). CCMs also satisfy the *intersection property* in 1D, a sufficient condition for a mechanism to be strategy-proof with constraints, hence it is consistent with Barberà et al.'s characterization result.

Evaluating incentives to misreport in multi-dimensional spaces is more involved. Our main results, Thms. 5 and 6 below, require two preliminary lemmas. The first addresses single-agent misreports. We begin with some notation. For each feasible $d \in \mathbb{D}$, define its *electoral zone* to be

$$\mathbb{Z}_d = \{x \in \mathbb{R}^m, d = \arg \min_{d' \in \mathbb{D}} \|d' - x\|_2\}.$$

Let \mathbb{C}_d be the *potential deviation area* for d if a single manipulator changes her report in all but one dimension, i.e., $\mathbb{C}_d = \{x \in \mathbb{R}^m, x^k = x'^k \text{ for some } x' \in \mathbb{Z}_d \text{ and all dimension } k\}$. We also denote $\mathbb{C}_d^k = \{x \in \mathbb{R}^m, x^k = x'^k \text{ for some } x' \in \mathbb{Z}_d\}$. Then we have the following result:

LEMMA 2. *For any two feasible locations $d_1, d_2 \in \mathbb{D}$, an agent i can gain from a misreport that changes the location of a facility from d_1 to d_2 only if $\mathbb{C}_1 \cap \mathbb{Z}_2 \neq \emptyset$.*

PROOF. (Sketch) Our proof is for 2D, but the analysis can be generalized to higher dimensions. Consider two feasible locations d_1 and d_2 (see Fig. 3). Let g be one of the chosen locations under a QM f' , and d_1 be its projected feasible location under CCM f (note we must have $g \in \mathbb{Z}_1$, otherwise f will not project g to d_1). Suppose there exists a location profile \mathbf{t} in which an agent i (with true peak x) will use facility d_1 , but has a positive incentive to change it




Figure 4: The incentive is bounded if some manipulator can benefit from changing the outcome from d_1 to d_2 .

to d_2 . Then we can construct another profile \mathbf{t}' such that if i misreports x' , the selected location for g under the QM f' will be $g' \in \mathbb{C}_1$. Since f projects g' to the closest feasible location, which is d_2 instead of d_1 , agent i gains by misreporting. However, f will project g' to d_2 only if there is no other feasible location closer to g' , i.e., only if g' is in the electoral zone of d_2 . This implies $\mathbb{C}_1 \cap \mathbb{Z}_2 \neq \emptyset$. \square

This lemma ensures an agent can profitably change a facility only if she can move the corresponding quantile-location into the electoral zone of another feasible outcome. The next lemma bounds the gain an agent can realize by changing one of the CCM's outcomes from one feasible location to another. For each pair of feasible locations $d_1, d_2 \in \mathbb{D}$, we define $\mathbf{K}_{1,2} = \{k : \mathbb{Z}_2 \cap \mathbb{C}_1^k \neq \emptyset\}$. For any two points $x, y \in \mathbb{R}^m$, let $\mathbf{B}(x, y)$ be the minimum bounding box containing x and y . Then we have:

LEMMA 3. *Let $d_1, d_2 \in \mathbb{D}$. The maximum gain any agent can realize by replacing d_1 with d_2 in a CCM is:*

$$G(d_1, d_2) = \begin{cases} \|d_2 - d_1\|_2 & \text{if } \exists x \in \mathbb{C}_1 \cap \mathbb{Z}_2 \text{ s.t.} \\ & \mathbf{B}(d_2, x) \cap \mathbb{Z}_1 \neq \emptyset \\ \max_{k' \in \mathbf{K}_{1,2}} \sqrt{\sum_{k \neq k'} |d_1^k - d_2^k|^2} & \text{otherwise.} \end{cases}$$

PROOF. (Sketch) We prove the lemma for 2D case, but the analysis can be generalized to higher dimensions.

For the feasible pair of outcomes $d_1, d_2 \in \mathbb{D}$, we consider the following two cases:

- I. $d_2 \in \mathbb{C}_1^k$ for some k , $\exists x \in \mathbb{C}_1 \cap \mathbb{Z}_2$ and $g \in \mathbf{B}(d_2, x) \cap \mathbb{Z}_1$ (as shown in Fig.4 left). Consider the situation in which a manipulator's true peak coincides with d_2 , which provides the maximum gain for a manipulation that induces location d_2 . We can construct a location profile such that g is one of the quantile-location under truthful report before projection. As we have $g \in \mathbb{Z}_1$, the CCM will project it to d_1 , and the manipulator cost is at most $\|d_2 - d_1\|_2$ (equality if d_1 is the closest facility under truthful report). However, the manipulator can misreport and change the quantile-location for g to x (as $g \in \mathbf{B}(d_2, x)$), inducing a projection to d_2 (as $x \in \mathbb{Z}_2$) and a cost of 0, so her gain is at most $\|d_2 - d_1\|_2$.

- II. $d_2 \notin \mathbb{C}_1^k$ for any k . If $\mathbf{K}_{1,2} = \emptyset$, then by Lemma 2 we have $G(d_1, d_2) = 0$, otherwise the upper bound is demonstrated using the properties of a hyperbola. Given two focal points, the difference of the distances to these two foci from any point on a hyperbola is constant. Let a and b be the semi-major and semi-minor axes, and c the half distance between two foci satisfying $c^2 = a^2 + b^2$.




Figure 5: An example where a manipulator can benefit from changing the outcome from d_1 to d_2 .

Let d_1 and d_2 be two focal points of a hyperbola (see Fig. 4 right). Let the angle between line d_1d_2 and the horizontal axis be α , and the angle between the asymptotes and the semi-major axis be θ . Our goal is to bound the maximum value of $2a$, which is the difference of distances to the two foci on a hyperbola, s.t. the constraint that hyperbola l intersects the horizontal or vertical axis (otherwise no agent can benefit this much). Suppose w.l.o.g., we have $\mathbb{Z}_2 \cap \mathbb{C}_1^2 \neq \emptyset$. The maximum gain is achieved when the angle $\theta > 90^\circ - \alpha$ and the hyperbola intersects the shaded area \mathbb{C}_1^1 . Recall that for an asymptote, we have $\tan(\theta) = b/a$, so we can formulate this as a maximization:

$$\begin{aligned} & \max 2a \\ & \text{s.t. } (d_2^1 - d_1^1)^2 + (d_2^2 - d_1^2)^2 = 4(a^2 + b^2) \\ & \frac{b}{a} > \frac{|d_2^1 - d_1^1|}{|d_2^2 - d_1^2|} \end{aligned}$$

Solving the above maximization, we have $2a = |d_1^2 - d_2^2|$. And if we consider every dimension $k \in \mathbf{K}_{1,2}$, we can get the above bound.

\square

We now describe the main results of this section, and provide upper bounds on the incentives for individuals and groups misreport in CCMs under the L_2 -norm. Unlike the unconstrained case, the bound here applies for any group of manipulators with any preference profile, and is a function of the feasible locations only. The first result is for a single manipulator:

THEOREM 5. *CCMs are ε -strategy-proof in m -dimensional FLPs under the L_2 -norm, where*

$$\varepsilon = \max_{(d_r, d_s) \in \mathbb{D}} G(d_r, d_s)$$

PROOF. (Sketch) For each feasible pair of outcomes $d_r, d_s \in \mathbb{D}$, the gain of any agent when changing the outcome from d_r to d_s is at most $G(d_r, d_s)$ by Lemma 3. Maximizing over all feasible pairs completes the proof. \square

For group misreports, we provide a loose bound:

THEOREM 6. *CCMs are ε -group strategy-proof in multi-dimensional FLPs under the L_2 -norm, where*

$$\varepsilon = \max_{d_r, d_s \in \mathbb{D}} \|d_r - d_s\|_2$$

PROOF. Consider any feasible pair of outcomes $d_r, d_s \in \mathbb{D}$. Let g be one of the chosen locations under a QM f' ,

which is projected to d_r under the CCM f . We can construct a location profile and a manipulator set $S = \{x, y, p\}$ such that: (i) all the manipulators are closer to d_s than to d_r ; and (ii) one of the manipulators p coincides with d_s . In addition, we can also ensure that x and y are “far enough away” so that the bounding box containing x , y and p intersects with Z_{d_r} (as shown in Fig. 5).

A viable group manipulation exists if all three manipulators misreport d_s , and move the selected quantile-location from g to d_s , in which the gain of the each manipulators is at most $\|d_r - d_s\|_2$ (the bound is tight if $q = 1$). Maximizing over all feasible pairs completes the proof. \square

Note that this bound can also be viewed as a negative result, as it naturally holds in any mechanism for constrained FLPs. However, the proof of the worst-case bound is achieved makes strong assumptions about the locations of the peaks of both the sincere agents and the manipulators. Such worst-case bounds are unlikely to arise in practice, as we explore empirically in the next section.

5. EMPIRICAL ANALYSIS

The theoretical bounds derived above offer some insight into the performance of GMMs, QMs and CCMs w.r.t. incentive for manipulation. But the tightness of these bounds in practice depends on the distribution of agent preferences (i.e., their peaks in the underlying space). We evaluate these incentives empirically, using two real-world data sets.

The first uses voting data from the Dublin West constituency in the 2002 Irish General Election.⁴ It consists of 29,989 votes over nine candidates, with each vote a ranking of a subset the candidates. We use the 3800 votes that rank all nine candidates. Since voters simply rank candidates rather than specifying ideal points in some space, we fit this data to a *spatial voter model* [20] using L_2 distance to measure (stochastic) preferences: using an EM algorithm, we place both voters (ideal points or peaks) and candidates (feasible locations) in a 2D space to maximize the likelihood of the observed rankings. We defer discussion of further details to a longer version of this paper.⁵ We use the estimated voter peaks in tests of unconstrained QMs (ignoring the candidates) and constrained CCMs (limiting selection to the nine candidates).

The second data set comprises geographic data for facility location [7], with latitude and longitudes of 88 cities in the continental United States (the 48 state capitals unioned with the 50 largest cities). Following [26], we treat these locations as both the ideal points of 88 agents and the feasible locations in constrained FLPs. In other words, the agents reveal their locations (which we assume to be private, but in fact linked to a specific site) and then place a small set of facilities among themselves (the setup is similar to a voting for representatives from *within a group* [1]). This data is used to test CCMs.

To generate unconstrained FLPs from the voting data, we assume $s \in \{2, 5, 8\}$ manipulators and $n \in \{0, 2, 5, 10, 20, 50, 80, 100\}$ sincere voters. For each setting, we randomly sample voter peaks from the 3800 estimated (spatial) positions

⁴ Available from www.dublincountyreturningofficer.com.

⁵ Recent analysis has suggested not only that this data is approximately single-peaked in 2D [29], but that a 2D spatial model using L_2 provides a reasonable explanation of voter preferences [13]. Our EM method is similar to [13].

to generate 1000 type profiles. For each profile, we either enumerate all manipulating coalitions of the required size or randomly sample $t \in \{10, 20, 50, 100, 200\}$ sets of s manipulators (depending on problem size). For each of the 1000 profiles, if *any* of the coalitions has a viable manipulation, we say the profile is manipulable and report the average gain of the coalition members in the coalition that has maximal gain.⁶ We report the following in our results: the *probability of manipulation*, i.e., the proportion of the 1000 profiles that admit a beneficial manipulation for *some* coalition; the *normalized gain* (the utility gain of an agent divided by his current cost) for the coalition with maximal gain, averaged over the 1000 profiles; and the average *loss in social welfare* realized, relative to truthful reporting. To test CCM on constrained FLPs, we use a smaller number of manipulators (1, 2, 4), but otherwise use the same settings as in the unconstrained case.

Fig. 6 shows results on unconstrained problems for a single facility (winning candidate) using the median mechanism (quantile 0.5). Interestingly, the probability of manipulation increases with the number of sincere agents and converges to 1.0 (see the middle figure of Fig. 6). This occurs because we simply measure whether *some* coalition among the set of agents can successfully manipulate. This suggests that there is almost always some group whose peaks “contain” the median position. However, the left figure in Fig. 6 shows that the average normalized gain decreases significantly with the number of sincere agents (e.g., with 2 manipulators, manipulation probability increases from 9.7% to 100%, but normalized gain reduces from 6.2% to 0.33%). Manipulative power is limited by the nearby peaks of sincere voters, and diminishes with more sincere voters. Impact on social welfare is also limited and is very small beyond 10 sincere agents, suggesting that QMs (including the median mechanism) are robust to manipulation in practice (note that manipulation may both increase or decrease total social cost).

We next evaluate CCMs in constrained two-facility FLPs, using the QM $\mathbf{q} = \{0.2, 0.3; 0.8, 0.7\}$ to make the initial selections (which are then projected using CCM). Fig. 7 shows the results on both the voting data set (top) and the geographic data set (bottom). The results for the voting data in constrained FLPs is similar to those for the unconstrained FLPs, except that the probability of manipulation initially increases as the number of sincere agents grows, and then decreases. The initial increase occurs for the same reason as in the unconstrained case, and subsequently decreases because the number of feasible locations is fixed and small, which limits the probability of manipulation as the number of sincere agents increases. For the geographic data set, the probability of manipulation remains high, suggesting that there is always some group that can profitably manipulate a QM. Compared with the results on the voting data set, this occurs, in part, because the number of feasible outcomes increases as the number of agents increases, making it more probable for the manipulators to probe new possibilities. Average normalized gain and loss in social welfare

⁶This set up assumes, somewhat unrealistically, that the members of this worst-case coalition can “discover” each other, and that they generate their misreport with *full knowledge* the reports of the sincere agents, as is common in analysis of manipulation in voting. For an analysis of manipulation in voting under more realistic knowledge assumptions, see [16].




Figure 6: Unconstrained, single-FLPs and GMMs: normalized gain (left), prob. of manipulation (middle), and loss in social welfare (right). The error bars show the standard deviation for each point.

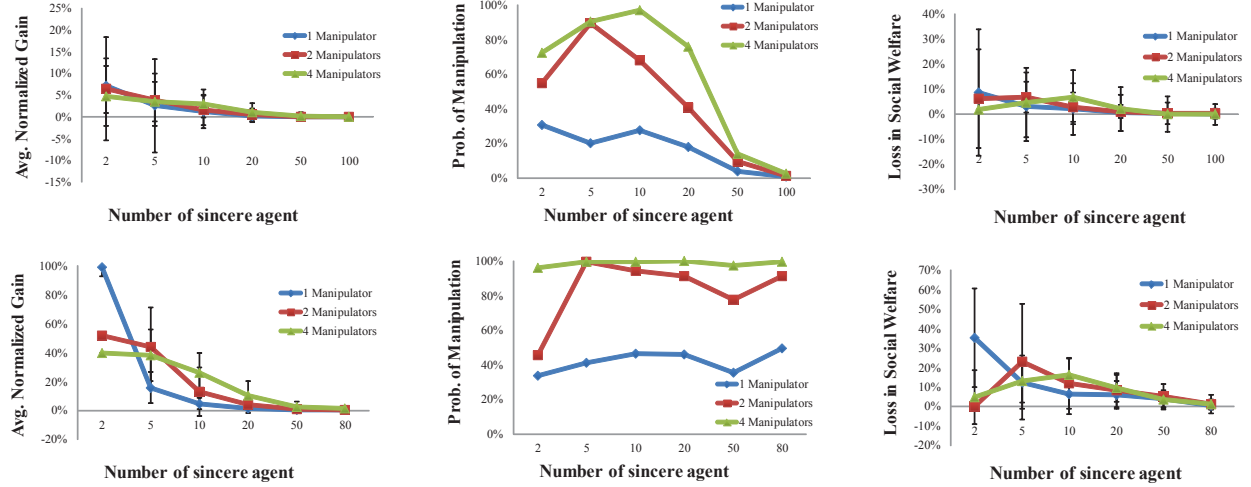


Figure 7: Constrained single-FLPs and CCMS: normalized gain (left), prob. of manipulation (middle), and impact on social cost (right). The results on top are for the voting data set, and the results on bottom are for the geographic data set. The error bars show the standard deviation for each point.

is much higher than in the voting data set (e.g., with 2 of each agent type, average gain in constrained FLPs is 52%, compared to 4.7% in the voting data set). This is largely due to the fact that the agents’ ideal locations and the feasible locations are much more tightly clustered in the geographic data set (since ideal points coincide with feasible locations) than in the voting data set. Despite this, both average gain and impact on social cost drop quickly with the number of sincere agents.

6. CONCLUSION AND FUTURE WORK

We have studied the mechanism design problem for both unconstrained and constrained FLPs, investigating the degree to which individual and group strategy-proofness can be achieved, and providing bounds on the incentive for individuals and groups to misreport in generalized median, quantile, and our newly proposed closet candidate mechanisms. Empirical analysis of Irish electoral data shows that these mechanisms may perform extremely well in practice, limiting the odds of manipulation and especially the potential gains and impact on social welfare.

Several interesting future directions remain. Exploring the approximate incentive properties of additional mechanisms (beyond GMMs, QMs, CCMS) and cost functions (beyond L_2) is of interest. The exploration of incremental (or multi-stage) mechanisms that trade off social cost, incentives, privacy and communication would be extremely valuable [10, 22, 27]. Finally, preferences are often not fully single-peaked in realistic domains, but are often approxi-

mately so [29]. Extending the theoretical analysis to this setting would be of value. Finally, we are interested in examining the optimization problem facing manipulators when they have only probabilistic knowledge of the potential reports of the sincere agents, as well as the impact of this limited knowledge on the probability of manipulation, average gain/incentive, and loss in social welfare [16]. This would provide a more realistic assessment of the robustness/resistance of GMMs and QMs to group manipulation.

Acknowledgements

We thank the reviewers for helpful comments. This research was supported by NSERC.

REFERENCES

- [1] N. Alon, F. Fischer, A. Procaccia, and M. Tennenholtz. Sum of us: Strategyproof selection from the selectors. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge*, pages 101–110. ACM, 2011.
- [2] S. Barberà, F. Gul, and E. Stacchetti. Generalized median voter schemes and committees. *Journal of Economic Theory*, 61(2):262–289, 1993.
- [3] S. Barberà, J. Massó, and A. Neme. Voting under constraints. *Journal of Economic Theory*, 76, 1997.
- [4] D. Black. On the rationale of group decision-making. *Journal of Political Economy*, 56(1):23–34, 1948.

- [5] K. C. Border and J. S. Jordan. Straightforward elections, unanimity and phantom voters. *The Review of Economic Studies*, 50(1):153–170, 1983.
- [6] J. R. Chamberlin and P. N. Courant. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *The American Political Science Review*, 77(3):718–733, 1983.
- [7] M. S. Daskin. *Network and Discrete Location: Models, Algorithms, and Applications*. John Wiley & Sons, 2011.
- [8] E. Dokow, M. Feldman, R. Meir, and I. Nehama. Mechanism design on discrete lines and cycles. In *Proceedings of the Thirteenth ACM Conference on Electronic Commerce (EC'12)*, pages 423–440, Valencia, Spain, 2012.
- [9] B. Escoffier, L. Gourvès, T. N. Kim, F. Pascual, and O. Spanjaard. Strategy-proof mechanisms for facility location games with many facilities. In *Proceedings of the Second International Conference on Algorithmic Decision Theory (ADT-11)*, pages 67–81, Piscataway, NJ, 2011.
- [10] J. Feigenbaum, A. D. Jaggard, and M. Schapira. Approximate privacy: Foundations and quantification (extended abstract). In *Proceedings of the Eleventh ACM Conference on Electronic Commerce (EC'10)*, pages 167–178, Cambridge, MA, 2010.
- [11] D. Fotakis and C. Tzamos. Winner-imposing strategyproof mechanisms for multiple facility location games. In *Proceedings of the Sixth International Workshop on Internet and Network Economics (WINE-10)*, pages 234–245, Stanford, CA, 2010.
- [12] A. Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973.
- [13] I. C. Gormley and T. B. Murphy. A latent space model for rank data. In *Proceedings of the Twenty-third International Conference on Machine Learning (ICML-06)*, 2006.
- [14] N. Hyafil and C. Boutilier. Regret-based incremental partial revelation mechanisms. In *Proceedings of the Twenty-first National Conference on Artificial Intelligence (AAAI-06)*, pages 672–678, Boston, 2006.
- [15] P. Lu, X. Sun, Y. Wang, and Z. A. Zhu. Asymptotically optimal strategy-proof mechanisms for two-facility games. In *Proceedings of the Eleventh ACM Conference on Electronic Commerce (EC'10)*, pages 315–324, Cambridge, MA, 2010.
- [16] T. Lu, P. Tang, A. D. Procaccia, and C. Boutilier. Bayesian vote manipulation: Optimal strategies and impact on welfare. In *Proceedings of the Twenty-eighth Conference on Uncertainty in Artificial Intelligence (UAI-12)*, Catalina, CA, 2012. 543–553.
- [17] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.
- [18] J. Massó and I. Moreno de Barreda. On strategy-proofness and symmetric single-peakedness. *Games and Economic Behavior*, 72(2):467–484, 2011.
- [19] H. Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.
- [20] K. T. Poole and H. Rosenthal. A spatial model for legislative roll call analysis. *American Journal of Political Science*, 29(2):357–384, 1985.
- [21] A. D. Procaccia and M. Tennenholtz. Approximate mechanism design without money. In *Proceedings of the Tenth ACM Conference on Electronic Commerce (EC'09)*, pages 177–186, Cambridge, MA, 2009.
- [22] T. Sandholm and F. Brandt. On the existence of unconditionally privacy-preserving auction protocols. *ACM Transactions on Information and System Security*, 11(2), 2008. Article 6.
- [23] M. A. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.
- [24] J. Schummer and R. V. Vohra. Strategy-proof location on a network. *Journal of Economic Theory*, 104(2):405–428, 2002.
- [25] J. Schummer and R. V. Vohra. Mechanism design without money. In N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, pages 243–265. Cambridge University Press, 2007.
- [26] L. V. Snyder and M. S. Daskin. Reliability models for facility location: the expected failure cost case. *Transportation Science*, 39(3):400–416, 2005.
- [27] X. Sui and C. Boutilier. Efficiency and privacy tradeoffs in mechanism design. In *Proceedings of the Twenty-fifth AAAI Conference on Artificial Intelligence (AAAI-11)*, pages 738–744, San Francisco, CA, 2011.
- [28] X. Sui, C. Boutilier, and T. Sandholm. Analysis and optimization of multi-dimensional percentile mechanisms. In *Proceedings of the Twenty-third International Joint Conference on Artificial Intelligence (IJCAI-13)*, Beijing, 2013.
- [29] X. Sui, A. Francois-Nienaber, and C. Boutilier. Multi-dimensional single-peaked consistency and its approximations. In *Proceedings of the Twenty-third International Joint Conference on Artificial Intelligence (IJCAI-13)*, Beijing, 2013.