Games and Economic Behavior 124 (2020) 478–490

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

The capacity constrained facility location problem ☆

Haris Aziz [a], Hau Chan [b], Barton E. Lee [a,*], David C. Parkes [c]

[a] UNSW Sydney and Data61 CSIRO, Australia
[b] University of Nebraska-Lincoln, USA
[c] Harvard University, USA

ARTICLE INFO

Article history:
Received 6 July 2019
Available online 2 October 2020

Keywords:
Facility location
Median mechanisms
Strategy-proofness
Characterization
Worst-case approximation

ABSTRACT

We initiate the study of the capacity constrained facility location problem from a mechanism design perspective. In the capacity constrained setting, the facility can serve only a subset of the population, assumed to be the k-closest with respect to agents' true locations (this can be justified as the essentially unique equilibrium outcome of a first-come-first game induced by the facility location). The main result is a complete characterization of dominant-strategy incentive compatible (DIC) mechanisms via the family of generalized median mechanisms (GMMs). Thus, the framework we introduce surprisingly provides a new characterization of GMMs, and is responsive to gaps in the current social choice literature highlighted by Border and Jordan (1983) and Barberà et al. (1998). We also provide algorithmic results and study the performance of DIC mechanisms in optimizing welfare. Adopting a worst-case approximation measure, we attain tight lower bounds on the approximation ratio of any DIC mechanism.

1. Introduction

A common economic problem is deciding where a public facility should be located to serve a population of agents with heterogeneous preferences. For example, a government needs to decide the location of a public hospital or library. More abstractly, the 'location' may represent a type or quality of a service. For example, a government may have a fixed hospital location but must decide on the type of service the hospital will provide, perhaps whether the facility's services are targeted to those suffering from acute, moderate, or mild severity of a certain illness. In such problems, agents may benefit by misreporting their preferences, and this can be problematic for a decision maker trying to find a socially optimal solution. This leads to the mechanism design problem of providing optimal, or approximately optimal, solutions while also being dominant-strategy incentive compatible (DIC) or strategyproof, i.e., so that no agent can profit from misreporting their preferences regardless of what others report.[1] In this paper, we refer to this problem as the *facility location problem*.

The facility location problem has been studied extensively, and when it is not capacity constrained then all agents can benefit from the facility it is modeled as a public good (non-rivalrous and non-excludable). This problem has been explored in several classic papers (Black, 1948; Border and Jordan, 1983; Gibbard, 1973, 1977; Moulin, 1980; Satterthwaite, 1975),

☆ This paper was previously circulated under the title "Mechanism Design without Money for Common Goods." The abstract of this paper appeared in the *Proceedings of the 15th Conference on Web and Internet Economics (WINE 2019)*. We owe thanks to Richard Holden, Herve Moulin, Toby Walsh, UNSW seminar participants, and the anonymous referees.

* Corresponding author.
E-mail addresses: haris.aziz@unsw.edu.au (H. Aziz), hchan3@unl.edu (H. Chan), barton.e.lee@gmail.com (B.E. Lee), parkes@eecs.harvard.edu (D.C. Parkes).

[1] We focus on the setting where the use of money is not permitted, for example because it would be unlawful, unethical, or otherwise unfair.

https://doi.org/10.1016/j.geb.2020.09.001
0899-8256/

and more recently in the field of algorithmic mechanism design (Feldman et al., 2016; Nisan and Ronen, 2001; Procaccia and Tennenholtz, 2013).

To the best of our knowledge, an unexplored setting for the mechanism design problem is where the public facility is also capacity constrained.[2] These kinds of capacity constraints, which limit the number of agents who can benefit from a facility's services, are ubiquitous in practice. Consider a hospital that is capacity constrained by the number of beds and doctors, for example, or a library with limited seating. Capacity constraints introduce a form of *rivalry* to the facility, since once the facility reaches its capacity limit then additional agents are prevented from using the facility.

A number of new strategic challenges arise for the mechanism designer when the public facility is capacity constrained but still non-excludable. For example, when the mechanism designer chooses a location for the facility we cannot stipulate which agents will be served. Instead, we assume that only the subset of agents with their true location closest to the facility will be served (up to the capacity constraint). This can be justified as the essentially unique equilibrium outcome of a first-come-first game induced by the facility location. In this way, an agent's utility and whether they are served depends not only on the facility location but also on the capacity constraint and the true locations of the other agents. This introduces a new technical challenge, because it results in agents having interdependent utilities and requires the designer to consider mechanisms that are strategyproof in this broader game-theoretic context.

In this paper, we initiate the study of the capacity constrained facility location problem from the viewpoint of mechanism design. In our model, n agents are located in the $[0, 1]$ interval, and there is a single facility to be located, this facility is able to serve at most k agents, where k is some positive integer. When $k \geq n$ the problem is equivalent to the classic problem. Agent locations are privately known, and the mechanism induces a reporting game. Once the facility location is decided, we assume the k-closest agents with respect to their true locations are served. We adopt as a design goal that of maximizing social welfare, while also seeking strategyproofness. As with much of the literature since Gibbard-Satterthwaite (Gibbard, 1973; Satterthwaite, 1975), we focus on the case where, conditional on being served, an agent has single-peaked and symmetric preferences over the facility location. This symmetry assumption is also adopted in Border and Jordan (1983). Strategyproofness requires that an agent never benefits from misreporting their location to the mechanism regardless of what other agents report, and regardless of other agents' true locations. Unlike the classic, unconstrained problem, the social welfare optimal mechanism is not DIC except when the capacity constraint is trivial, i.e., $k = 1$ or n.[3] As a result, we also follow the approach of Procaccia and Tennenholtz (2013) and consider the approximate mechanism design problem. We adopt the worst-case approximation measure, and ask what is the best approximation achievable with DIC mechanisms, and how does this vary as a function of the capacity constraint?

The main theoretical contribution is a complete characterization of DIC mechanisms via the family of generalized median mechanisms (GMMs), which therefore provides a new characterization of GMMs, and also closes a gap in the current literature on the facility location problem without capacity constraints. Border and Jordan (1983) has provided a partial characterization of strategyproof mechanisms for this classical setting via the family of GMMs. Border and Jordan show that a mechanism is strategyproof and unanimity respecting[4] if and only if it is a GMM, and that the family of GMMs is strictly smaller than the complete family of strategyproof mechanisms.[5] This has left a gap in the literature to characterize the complete family of strategyproof mechanisms and to understand the difference between strategyproof mechanisms that are GMMs and those that are not. Fig. 2 schematically illustrates this gap.

Our Contributions: We introduce a new mechanism problem, the capacity constrained facility location problem. This problem is a natural variant of the classic facility problem where the facility also faces a capacity constraint.

Our main theoretical contribution is a complete characterization of DIC mechanisms for the capacity constrained facility location problem. We show that a mechanism is DIC if and only if it belongs to the established family of mechanisms of GMMs, which appear in Moulin (1980) and Border and Jordan (1983). Thus, the framework we introduce also provides an interesting new characterization of GMMs, and contributes a novel perspective to a "major open question" (Barberà et al., 1998) posed by Border and Jordan (1983) regarding the relationship between GMMs and the *larger* space of DIC mechanisms in the non-capacity constrained facility location problem. In particular, our characterization result shows that any mechanism that is DIC in the non-capacity constrained problem and is not a GMM must not be DIC in the capacity constrained setting. That is, the strategyproofness of these (non-GMM) mechanisms is not robust to the addition of capacity constraints (further discussion is provided in Section 1.1).

We also provide algorithmic results and study the performance of DIC mechanisms in optimizing social welfare. We adopt a worst-case approximation measure, and provide a lower bound on the approximation ratio of any DIC mechanism. A lower bound that is greater than 1 can be viewed as an impossibility result. We show that at best the approximation ratio

[2] The purely algorithmic problem of locating multiple capacity constrained facilities when agents are not strategic has been studied (Brandeau and Chiu, 1989; Pál et al., 2001; Vygen, 2004).

[3] For the classical setting, if the objective of the mechanism designer is to maximize social welfare, then the standard median mechanism is both strategyproof and social welfare optimal (Black, 1948).

[4] Unanimity respecting means that if there is a unanimously most preferred facility location, then the mechanism must locate the facility at this location.

[5] We note that in a slightly different setting, where the single-peaked preferences are possibly asymmetric, GMMs provide a complete characterization of strategyproof and "peak only" mechanisms that only receive reports of peaks (Proposition 3 of Moulin (1980)). Example 2.2 in the present paper provides an example of a mechanism that is strategyproof in the Border and Jordan (1983) setting but not the Moulin (1980) setting. Massó and de Barreda (2011) provides a complete characterization of strategyproof mechanisms when agents have symmetric single-peaked preferences.

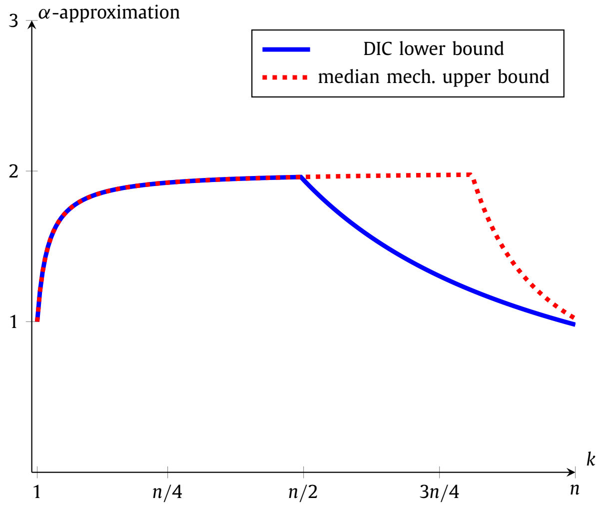

Fig. 1. Worst-case approximation ratio as a function of the capacity constraint, k.

of a DIC mechanism is $2\frac{k}{k+1}$ when $k \leq \lceil (n-1)/2 \rceil$, and $\max\{\frac{n-1}{k+1}, 1\}$ otherwise. Interestingly, this lower bound is achieved by the standard median mechanism (which is also DIC) when $k \leq \lceil (n-1)/2 \rceil$ or $k = n$, and hence the median mechanism is optimal among all DIC mechanisms in those ranges. Fig. 1 illustrates these approximation results.

1.1. Related literature

Facility location problems. Facility location problems have been extensively studied. The standard setting considers the problem of locating a single facility along the real line when agents have single-peaked preferences over the facility's location.[6] Moulin (1980) focuses on characterizing mechanisms that are strategyproof and mechanisms that are strategyproof, anonymous, and efficient; however, the literature has extended this problem in a variety of directions. For example, Nehring and Puppe (2006, 2007) explore a setting where agents' preferences satisfy a more general notion of single-peakedness; Miyagawa (1998, 2001) and Ehlers (2002, 2003) consider an extension whereby multiple facilities must be located; Barberà et al. (1998) study a setting where the space of feasible facility locations is restricted; and Procaccia and Tennenholtz (2013) introduce computational and approximation properties into the mechanism design problem. Our paper contributes to this literature by considering a natural extension of the facility location problem where the facility is capacity constrained.

Most closely related to our paper is Procaccia and Tennenholtz (2013), where agents with single-peaked preferences are located along the real line and the problem of locating a (non-capacity constrained) public facility is studied with the goal of minimizing two distinct objective functions: the total social cost and the maximum cost across agents. In this paper, we focus on minimizing the total social cost in the presence of a capacity constrained facility. In contrast to the setting studied by Procaccia and Tennenholtz (2013), agents have interdependent utilities in our model because of the capacity constraints of the facility, since their utility for a choice by the mechanism depends on the true locations of others and whether they will gain access.

Characterization of strategyproof (DIC) mechanisms. Another large body of literature is concerned with characterizing DIC mechanisms for the unconstrained facility location problem. In one-dimensional space and for symmetric and single-peaked preferences, Border and Jordan (1983) characterize a general class of DIC mechanisms which have become to be known as *generalized median mechanisms* (GMM), and in addition, show that when the property of unanimity is enforced every DIC mechanism is a GMM. Border and Jordan (1983) also consider the problem in higher dimensions. These results differ slightly from the characterization results of Moulin (1980) since the setting studied in Moulin (1980) does not restrict the single-peaked preferences to be symmetric. Characterizing DIC mechanisms that need not satisfy unanimity was posed as an open problem; as stated by Border and Jordan (1983) "*[the characterization] leaves several open problems. The most obvious question is: what happens if the unanimity assumption is dropped?*" This has become known as a "*major open question*" (Barberà et al., 1998), with only partial progress towards a resolution (Barberà and Jackson, 1994; Barberà et al., 1998; Ching, 1997; Peremans et al., 1997; Weymark, 2011). In particular, the results of Border and Jordan (1983) for facility location in a one-dimensional space leave two gaps in regard to GMMs:

[6] We do not review the computer science and operations research literature on facility location problems that assumes complete information and hence does not require a mechanism design approach (see Brandeau and Chiu (1989)). Furthermore, this literature, when incorporating capacity constraints, typically focuses on the problem of locating multiple capacity constrained facilities that collectively have sufficient capacity to serve all agents (Charikar et al., 2002; Cygan et al., 2012; Pál et al., 2001; Vygen, 2004).

(1) there exist non-unanimity respecting DIC mechanisms that are not GMM, and
(2) there exist DIC mechanisms that are GMMs but do not respect unanimity.

The results of Massó and de Barreda (2011) close the first gap by providing a complete characterization of DIC mechanisms. Their results imply that any non-unanimity respecting DIC mechanism that is not a GMM must belong to a family of "disturbed minmax" mechanisms that is obtained by modifying each GMM to admit a particular kind of discontinuity. Our characterization of DIC mechanisms via the family of GMM, although considered in a different setting where the facility is capacity constrained, applies more generally to mechanisms that are not unanimity respecting. Hence, we contribute a novel perspective to these gaps in characterization, showing that a mechanism is DIC for all possible capacity constraints $k \le n$ if and only if it is a GMM. This means that any mechanism in gap (1) is not DIC when the facility is capacity constrained with $k < n$. Furthermore, the unanimity property is sufficient to ensure that a mechanism that is DIC in the non-capacity constrained setting remains DIC when capacity constraints are present. Our results complement those of Massó and de Barreda (2011) by showing that the family of disturbed minmax mechanisms are not DIC in the capacity constrained setting, while the family of GMMs are DIC in the capacity constrained setting.

Facility location problems with capacity constraints. To the best of our knowledge, the facility location problem with capacity constraints has not been studied within the mechanism design literature. However, the settings studied in Jackson and Nicolò (2004) and Cantala (2004) share some similarities with our setting and, hence, deserve greater discussion.

Jackson and Nicolò (2004) study a non-capacity constrained facility location problem where agents have preferences over the facility's location and also the number of other agents that are served. In their setting, the mechanism designer chooses both the location of the facility and assigns a subset of agents to be served the facility.[7] This differs from the setting that we study where the mechanism designer does not assign a subset of agents to be served; instead, the facility's location and capacity constraint determines the subset of agents that will be served.[8] Thus, Jackson and Nicolò's results are applicable to settings where the designer has coercive power to both force and exclude agents from accessing the facility, while our results are applicable to settings where the designer has no such power. Our results are complementary and show that the scope for designing strategyproof mechanisms is significantly enlarged when the designer does not have the power to assign agents to the facility.

Cantala (2004) studies the related problem of locating a non-capacity constrained facility when agents have an outside option. These outside options play an important role. In particular, if a facility is located sufficiently far from an agent's ideal location, then the agent will prefer not to be served by the facility.[9] Cantala's results focus on a notion of efficiency that is specific to his setting. In particular, due to agents having an outside option, an efficient mechanism must locate the facility at some agent's ideal location. This implication does not hold in the standard setting of Moulin (1980) and does not hold in our setting.[10] For this reason, Cantala's characterization of all strategyproof and efficient mechanisms via a subfamily of GMMs, which are called *extreme minimax rules*, does not apply.[11] Towards the end of Section 2, we further discuss how our approximation results connect to the notion of Pareto optimality.

Capacity constraints in mechanism design problems. Finally, we note that capacity constraints, or quotas, have been considered in a number of related domains where the use of money is not permitted. Such domains include committee voting (Aziz and Lee, 2018; Bredereck et al., 2018), apportionment (Balinski and Young, 1982), and matching markets (Abdulkadiroğlu and Sönmez, 2003; Kamada and Kojima, 2015).

Outline: Section 2 presents our model and formalizes the objective of the mechanism designer, Section 2.1 then presents our key characterization result of DIC mechanisms. Section 3 explores the performance, i.e., approximation results, of DIC mechanisms. Lastly, we conclude with a discussion in Section 4.

2. Model, basic properties, and definitions

Model: Let $N = \{1, \dots, n\}$ be a finite set of n agents and let $X = [0, 1]$ be the domain of agent locations. Each agent $i \in N$ has a location $x_i \in X$, which is privately known, and the profile of agent locations is denoted by $\boldsymbol{x} = (x_1, x_2, \dots, x_n)$. The profile of all agents except some agent $i \in N$ is denoted by $\boldsymbol{x}_{-i} = (x_1, x_2, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$. There is a single facility to be located in X. A *mechanism* is a function $M : \prod_{i \in N} X \to X$, mapping a profile of locations to a single location. We restrict our attention to deterministic mechanisms. We denote the mechanism's output, or facility location, by $s \in X$.

[7] Bogomolnaia and Nicolò (2005) also consider a similar problem where the designer must locate and assign agents to multiple facilities.

[8] For expositional convenience, we assume that the k-closest agents with respect to their true locations are served. It is also straightforward to show that this corresponds to the unique equilibrium of a first-come-first game induced by the facility's location.

[9] This leads agents to have preferences that are "weakly" singled-peaked: agents have single-peaked preferences on an interval around their ideal location, but for facility locations outside of this interval they are indifferent. In our model, agents' preferences are also weakly single-peaked; however, this arises due to the facility being capacity constrained rather than agents having an outside option.

[10] In Moulin (1980) a Pareto efficient mechanism is only required to locate the facility in the interval that is bounded by the maximum and minimum of agents' ideal locations. Although we do not focus on Pareto efficiency, in our setting, when facilities are capacity constrained with $k < n$, the appropriate definition of efficiency differs from those of Moulin (1980) and Cantala (2004). In particular, an efficient mechanism in our setting is required to locate the facility such that the facility location is bounded by the maximum and minimum of the k-closest agents' ideal locations.

[11] Cantala also provides a characterization for group strategyproof and efficient mechanisms.

The facility faces a *capacity constraint* $k : k \leq n$, which provides a limit on the number of agents that can be served. Given a facility location s, we assume that the *k-closest* agents (with respect to the Euclidean metric and the agents' true locations) are served, breaking ties via some deterministic priority rule when necessary.[12] We denote this priority ordering by the binary relation $\triangleright$, so that $i \triangleright j$ means that agent i has a higher priority than agent j. A *served* agent attains utility $u_i = 1 - d(s, x_i) \geq 0$, where $d(\cdot, \cdot)$ denotes the Euclidean metric. An *unserved* agent attains zero utility, $u_i = 0$. Notice that an agent's utility depends not only on the facility location and their own location but also on the location of other agents and the capacity constraint of the facility. Thus, we will denote an agent's utility by a function $u_i(s, \boldsymbol{x}, k)$, where $s \in X$ is the facility location, k is the facility's capacity constraint, and $\boldsymbol{x}$ is the true location profile of all agents.

Our approximation results rely on the choice of utility function. On the other hand, our characterization results do not rely on this particular form, and nor do they require that locations be restricted to the closed interval $[0, 1]$. We only require that, for any true location, every agent weakly prefers to be served than not, and conditional on being served an agent's utility is symmetric and (strictly) single-peaked (as per Border and Jordan (1983)).

A useful observation is that, for any profile of other agents' locations, an agent has weakly single-peaked preferences over the facility location. This result is stated in Proposition 2.1 and is intuitive. If an agent i is among the k-closest agents for a given facility location $s > x_i$ then moving the facility location closer, to say $s' : s > s' \geq x_i$, never results in agent i being excluded from the k-closest agents. Thus, agent i's utility is guaranteed to weakly increase as the facility location moves closer to the agent's true location.

Proposition 2.1. *For any agent $i \in N$, if $s < s' \leq x_i$ or $x_i \leq s' < s$, then $u_i(s, \boldsymbol{x}) \leq u_i(s', \boldsymbol{x})$.*

Proof. Fix $k \leq n$, fix a profile of agent locations $\boldsymbol{x}$, and let s, s' be two distinct facility locations. Suppose there exists some agent i such that $s < s' \leq x_i$, the case where $x_i \leq s' < s$ is dealt with similarly. For the sake of a contradiction, suppose that $u_i(\boldsymbol{x}, s, k) > u_i(\boldsymbol{x}, s', k)$. Since $u_i(\boldsymbol{x}, s', k)$ equals either $1 - d(s', x_i)$ or 0 and $1 - d(s', x_i) > 1 - d(s, x_i)$, it is immediate that $u_i(\boldsymbol{x}, s, k) = 1 - d(s, x_i)$ and $u_i(\boldsymbol{x}, s', k) = 0$. That is, agent i is served (and hence among the k-closest agents) when the facility is located at s and is not served when the facility is located at the closer location of s'. But for any agent j such that agent i is closer to location s, i.e., $d(x_i, s) \leq d(x_j, s)$, it must also be true that agent i is closer to s', i.e., $d(x_i, s') \leq d(x_j, s')$. This is immediate for any $x_j \geq x_i$. Now suppose that $x_j < x_i$, then it must be that $x_j < s$, since $d(x_i, s) \leq d(x_j, s)$, but then it follows that

$$d(s', x_j) > d(s, x_j) \geq d(x_i, s) > d(x_i, s').$$

This is a contradiction. We conclude that $u_i(\boldsymbol{x}, s, k) \leq u_i(\boldsymbol{x}, s', k)$. □

In this paper, we are interested in strategyproof or dominant-strategy Incentive Compatible (DIC) mechanisms, so that agents do not have an incentive to misreport their location. A mechanism M is DIC if for every agent $i \in N$, we have

$$u_i\Big(M(x_i, \hat{\boldsymbol{x}}_{-i}), \boldsymbol{x}, k\Big) \geq u_i\Big(M(x_i', \hat{\boldsymbol{x}}_{-i}), \boldsymbol{x}, k\Big),$$

for every x_i', for every $\hat{\boldsymbol{x}}_{-i}$, and for every $\boldsymbol{x}_{-i}$. Formally speaking, the DIC definition depends on the capacity constraint k, however we omit this k dependence as this will be clear from the context.

Objective of the mechanism designer: The design goal is to find DIC mechanisms that perform well with respect to *social welfare*, i.e., the sum of agents' utilities, and rather than make distributional assumptions, we measure the performance of a DIC mechanism by its worst-case performance over the domain of preference profiles.

Given a profile of agent locations, $\boldsymbol{x}$, and a capacity constraint, k, we define the *optimal social welfare* by $\Pi^*(\boldsymbol{x}, k) := \max_{s \in X} \sum_{i=1}^{n} u_i(s, \boldsymbol{x}, k)$, and given a mechanism M let $\Pi_M(\boldsymbol{x}, k)$ denote the social welfare attained by the mechanism, i.e.,

$$\Pi_M(\boldsymbol{x}, k) := \sum_{i=1}^{n} u_i(s, \boldsymbol{x}, k) \qquad \text{where } s = M(\boldsymbol{x}).$$

The mechanism M is an α-approximation if

$$\max_{\boldsymbol{x} \in \prod_{i=1}^{n} X} \left\{ \frac{\Pi^*(\boldsymbol{x}, k)}{\Pi_M(\boldsymbol{x}, k)} \right\} \leq \alpha, \tag{1}$$

and the LHS of (1) is the *approximation ratio*. A mechanism (or family of mechanisms) has a *lower bound*, $\bar{\alpha}$, on the approximation ratio if

[12] This outcome can also be shown to be the essentially unique equilibrium of a simple game where agents compete to be served by the facility via a first-come-first-served protocol. Here essentially unique means that in the degenerate cases where multiple equilibria exist they are all payoff equivalent and hence the multiplicity does not affect the incentives of agents.

$$\bar{\alpha} \leq \max_{\boldsymbol{x} \in \prod_{i=1}^{n} X} \left\{ \frac{\Pi^*(\boldsymbol{x}, k)}{\Pi_M(\boldsymbol{x}, k)} \right\}. \tag{2}$$

When this lower bound is greater than 1 then this can be viewed as an impossibility result. We are also interested in algorithms that match the lower bound. For a given capacity constraint, k, we refer to a mechanism M that attains the optimal social welfare for all profiles of agent locations, $\boldsymbol{x}$, as an *optimal mechanism*. Such a mechanism is a 1-approximation. Again, the optimal mechanism definition depends on the capacity constraint k but we omit this dependence as this will be clear from the context. Note that the optimal mechanism need not, and in general will not, be DIC for a given k.

The results that we provide on approximations to social welfare also have a connection with works that consider Pareto optimality as the primary efficiency notion (see, e.g., Cantala (2004); Moulin (1980)). A utility profile $u = (u_1, \dots, u_n)$ Pareto dominates utility profile $u' = (u'_1, \dots, u'_n)$ if $u_i \geq u'_i$ for all agents i and $u_i > u'_i$ for some agent i. Given any $\beta \in [0, 1]$, a utility profile u is *β-Pareto optimal* if there exists no other achievable utility profile u' such that $\beta \cdot u'$ Pareto dominates u. Note that just as maximum welfare implies Pareto optimality, any outcome achieving β fraction of the maximum social welfare — which corresponds to what we call an $\alpha = 1/\beta$-approximation — also satisfies β-Pareto optimality (see, e.g., Aziz (2019)).

Remark 1. When $k = n$ our model reduces to the well-known facility location problem (Black, 1948; Moulin, 1980; Procaccia and Tennenholtz, 2013). Accordingly, this case ($k = n$) is fully resolved: the standard median mechanism that always locates the facility at the median reported location is both optimal and DIC.

To illustrate how the case where $k < n$ differs from the standard $k = n$ setting, consider Example 2.2. The example considers a mechanism that is DIC when $k = n$ but for any capacity constraint $k < n$ is not DIC.

Example 2.2. Let M be the mechanism such that $M(\boldsymbol{x}) = \arg\min_{s \in \{1/4,\ 3/4\}} d(s, x_i)$ for some $i \in N$, tie-breaking in favor of $s = 1/4$ if necessary. That is, the mechanism locates the facility at either location 1/4 or 3/4 depending on which is closest to agent i's report.

First, notice that the mechanism M is DIC when $k = n$. If $k = n$ then every agent i is always served by the facility and hence attains utility $1 - d(s, x_i)$ for any facility location s. It is immediate that agent i can never strictly benefit from misreporting their location.

However, when $k < n$ the mechanism is not DIC. To see this, consider an instance where agent i is located at 3/8 and all other agents are located at 1/4. When agent i truthfully reports, the facility is located at 1/4 and is not served – leading to zero utility. On the other hand, misreporting to $x'_i \in (1/2, 1]$ leads to the facility location 3/4 and agent i is the closest agent to the facility. In this case agent i attains strictly higher utility equal to $1 - d(3/4, 3/8) > 0$. Thus, the mechanism is not DIC for any $k < n$. ◇

2.1. A complete characterization of DIC mechanisms

We begin by defining the family of *generalized median mechanisms (GMMs)*. This family was introduced by Moulin (1980) and Border and Jordan (1983) for the $k = n$ setting, and provides a partial characterization of DIC mechanisms. For the $k = n$ setting, the well-known family of phantom mechanisms (Moulin, 1980) is also attained from the family of GMMs by requiring anonymity and efficiency.[13] The main result of the present paper shows that GMMs provide a complete characterization of mechanisms that are (1) DIC for all $k \leq n$, and (2) DIC for some $k < n$.

Definition 2.3 *(Generalized median mechanism (GMM)).* A mechanism M is said to be a *generalized median mechanism* (GMM) if for each $S \subseteq N$ there are constants $a_S\ :\ S \subseteq T \implies a_S \leq a_T$, such that for all location profiles $\boldsymbol{x}$,

$$M(\boldsymbol{x}) = \max_{S \subseteq N} \left\{ \min_{i \in S} \{x_i,\ a_S\} \right\}. \tag{3}$$

To build some intuition, we highlight three well-known mechanisms that belong to the family of GMMs:

- The *median mechanism*, which always outputs the median of the reported location profile, i.e., the $\lfloor (n+1)/2 \rfloor$-th smallest report, is attained from (3) by setting $a_S = 0$ for all subsets $S \subseteq N$ with $|S| < \lfloor (n+1)/2 \rfloor$ and $a_S = 1$ otherwise.
- The *s-constant mechanism*, which always outputs some location $s \in X$, is attained from (3) by setting $a_S = s$ for all subsets $S \subseteq N$ (including the empty set).
- The *agent i dictatorship mechanism*, which always outputs the location of agent i's report, is attained from (3) by setting $a_S = 1$ for all $S \subseteq N\ : i \in S$, and $a_S = 0$ for all other subsets.

[13] In the setting studied by Moulin (1980) where agents may have asymmetric utility functions, the family of GMMs provides a complete characterization of strategyproof mechanisms (Moulin, 1980, Proposition 3).

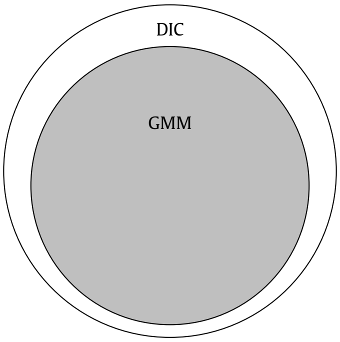

Fig. 2. Setting where $k = n$ (Border and Jordan, 1983).

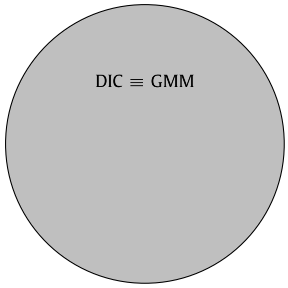

Fig. 3. Setting where $k < n$.

An example of a mechanism that is not a GMM is the dictatorial-style mechanism considered in Example 2.2.

We now state the main characterization result.

Theorem 2.4. *Let M be a mechanism. The following are equivalent:*

(1) *M is a GMM,*
(2) *M is DIC for some $k < n$,*
(3) *M is DIC for every $k \leq n$.*

We present the proof via a series of propositions and utilize a characterization of Border and Jordan (1983). First, we illustrate the contribution of Theorem 2.4, benchmarked against the results of Border and Jordan (1983), where GMMs are shown to be a strict subset of DIC mechanisms when $k = n$. Fig. 2 presents the result of Border and Jordan (1983) and Fig. 3 illustrates our characterization. When considering the capacity constrained problem, with $k < n$, the family of DIC mechanisms coincides precisely with the GMM family.

First, we present a result of Border and Jordan (1983), characterizing the family of GMMs via a property of the mechanism that they call *uncompromising*. Informally speaking, an uncompromising mechanism means that an agent cannot influence the mechanism output in their favor by reporting extreme locations. The most obvious mechanism satisfying this property is the median mechanism.

Formally, a mechanism M is *uncompromising* if, for every profile of locations $\boldsymbol{x}$, and each agent $i \in N$, if $M(\boldsymbol{x}) = s$ then

$$x_i > s \implies M(x_i', \boldsymbol{x}_{-i}) = s \qquad \text{for all } x_i' \geq s \qquad \text{and,} \tag{4}$$

$$x_i < s \implies M(x_i', \boldsymbol{x}_{-i}) = s \qquad \text{for all } x_i' \leq s. \tag{5}$$

Lemma 2.5 *(Border and Jordan (1983)). A mechanism M is uncompromising if and only if it is a GMM.*

Note that Lemma 2.5, although proved in the setting where $k = n$, does not rely on any strategic properties of the mechanism and so applies more generally to our setting with $k \leq n$.

We now prove our first proposition towards the characterization result.

Proposition 2.6. *Every GMM is DIC for any $k \leq n$.*

Proof. Fix $k \leq n$ and let M be a GMM. For the sake of a contradiction, suppose that M is not DIC. That is, for some agent i with location x_i, there exist a profile of other agent locations $\boldsymbol{x}_{-i}$, and reports $\hat{\boldsymbol{x}}_{-i}$ such that for some $x_i' \neq x_i$

$$u_i(M(x_i', \hat{\boldsymbol{x}}_{-i}), \boldsymbol{x}, k) > u_i(M(x_i, \hat{\boldsymbol{x}}_{-i}), \boldsymbol{x}, k). \tag{6}$$

Define $s' = M(x_i', \hat{\boldsymbol{x}}_{-i})$ and $s = M(x_i, \hat{\boldsymbol{x}}_{-i})$. It is immediate from (6) that $s \neq x_i$ and $s \neq s'$. Without loss of generality we assume that $x_i > s$. By assumption, M is a GMM and hence, by Lemma 2.5, satisfies the uncompromising property. It follows that $x_i' < s$, since otherwise $x_i' \geq s$ and (4) would imply $s' = s$ contradicting (6).

Case 1: Suppose $s < s'$. Then $x_i' < s'$ and the uncompromising property (5) implies that

$$M(x_i'', \hat{\boldsymbol{x}}_{-i}) = s' \qquad \text{for all } x_i'' \leq s'.$$

If $x_i'' \in [s, s']$, the uncompromising property implies that $M(x_i'', \hat{\boldsymbol{x}}_{-i}) = M(x_i, \hat{\boldsymbol{x}}_{-i})$, i.e., $s' = s$, which contradicts (6). Thus, we conclude that $x_i'' < s$.

Now consider a new instance where agent i has true location $y_i = \varepsilon \in (0, s)$, all other agents have true location $y_j = 0$ but collectively report $\hat{\boldsymbol{x}}_{-i}$. If agent i reports $y_i = \varepsilon$, then the facility location is s' and i attains utility $1 - d(s', \varepsilon)$. If instead agent i reports $y'_i = x_i$ then the facility location is $s < s'$ and i attains strictly higher utility $1 - d(s, \varepsilon)$. Thus, the mechanism is not DIC — a contradiction.

Case 2: Suppose $s > s'$. Since $x_i > s > s'$, it follows from the single-peaked property (Proposition 2.1) that $u_i(s, \boldsymbol{x}, k) \geq u_i(s', \boldsymbol{x}, k)$. This contradicts (6). □

We now prove our second proposition towards the characterization result. Proposition 2.7 shows that the DIC requirement is more restrictive for $k < n$ than for $k = n$ — meaning that the capacity constraints induce new strategic concerns for the mechanism designer.

Proposition 2.7. *If a mechanism M is DIC, for some $k < n$, then it is DIC for $k = n$. The converse is not true.*

Proof. We prove the contrapositive. Suppose that M is not DIC for $k = n$. That is, for some agent i with location x_i there exists a report x'_i, a profile of other agent reports $\hat{\boldsymbol{x}}_{-i}$, and a profile of other agent locations $\boldsymbol{x}_{-i}$ such that

$$u_i(M(x'_i, \hat{\boldsymbol{x}}_{-i}), \boldsymbol{x}, n) > u_i(M(x_i, \hat{\boldsymbol{x}}_{-i}), \boldsymbol{x}, n). \tag{7}$$

Let $s' = M(x'_i, \hat{\boldsymbol{x}}_{-i})$ and $s = M(x_i, \hat{\boldsymbol{x}}_{-i})$. When $k = n$ all agents are served and so (7) simplifies to

$$1 - d(s', x_i) > 1 - d(s, x_i). \tag{8}$$

Now we consider the same profile of reports but for an arbitrary $k < n$. Further, suppose all agents have location equal to x_i and agent i has the highest priority ($\triangleright$), i.e., after tie-breaking. The mechanism output is independent of agent true locations and so we still attain $M(x'_i, \hat{\boldsymbol{x}}_{-i}) = s'$ and $M(x_i, \hat{\boldsymbol{x}}_{-i}) = s$. Furthermore, since i has the highest priority (recall that the priority is distance-based but in this instance all agents are equidistant for every facility location) they are always served for every facility location. In particular, the utility from reporting truthfully is $1 - d(s, x_i)$ and misreporting is $1 - d(s', x_i)$ — the latter provides strictly higher utility, as per (8). We conclude that the mechanism is not DIC, and since $k < n$ was chosen arbitrarily it holds for all $k < n$.

The final statement in the proposition was shown in Example 2.2. □

We now prove our third and final proposition, which completes the characterization result.

Proposition 2.8. *If a mechanism M is DIC, for some $k < n$, then it is a GMM.*

Proof. Let M be a mechanism that is DIC for some $k < n$.

First, consider an instance where an arbitrary agent i has location x_i, and the other agents report $\hat{\boldsymbol{x}}_{-i}$. If i reports truthfully, the mechanism outputs some location that we denote as s, i.e.,

$$s := M(x_i, \hat{\boldsymbol{x}}_{-i}). \tag{9}$$

If $s = x_i$, then consider an alternate location and profile of other agents' reports so that the equality does not hold. If no such location and report profile exists then the mechanism always coincides with agent i's report; that is, the mechanism is the agent i dictatorship mechanism, which is a GMM.

Now suppose $s \neq x_i$, and without loss of generality assume $s < x_i$. By assumption M is DIC, for some $k < n$, and so it must be that for all x'_i

$$u_i(s, \boldsymbol{x}, k) \geq u_i(M(x'_i, \hat{\boldsymbol{x}}_{-i}), \boldsymbol{x}, k), \tag{10}$$

where $\boldsymbol{x}$ denotes the true location profile of all agents.

Recall that a mechanism is a GMM if and only if it is uncompromising (Lemma 2.5). We now show that deviation by agent i will satisfy the uncompromising property, i.e., for any $x'_i \geq s$ $M(x'_i, \hat{\boldsymbol{x}}_{-i}) = s$. To do so, we analyze different cases and sequential refine the possible values of $M(x'_i, \hat{\boldsymbol{x}}_{-i})$, we then derive a contradiction and conclude that $M(x'_i, \hat{\boldsymbol{x}}_{-i}) = s$.

Case 1: Suppose all other agents have true location s. When agent i truthfully reports x_i the facility location is s and agent i attains zero utility. Now consider some report $x'_i \geq s$, leading to facility location

$$s_{x'_i} := M(x'_i, \hat{\boldsymbol{x}}_{-i}).$$

If $s_{x'_i} \in (\frac{s+x_i}{2}, 1]$ for any $x'_i \geq s$ we attain a contradiction, since this means that agent i would be served from this report and attain strictly more utility than being truthful. We conclude that

$$s_{x'_i} \in [0, s) \cup \{s\} \cup (s, \frac{s+x_i}{2}) \qquad \text{for all } x'_i \geq s.$$

Case 2: Suppose all other agents have true location 1, noting that $s < x_i \leq 1$. In the event that $x_i = 1$ (in which case all agents are equidistant from every facility location), assume agent i has the highest priority in the tie-breaking rule ($\triangleright$). When agent i truthfully reports their location, they are served and attain utility $1 - d(s, x_i)$. To avoid a contradiction of (10), it must be that $s_{x'_i} \leq s$. Thus, we conclude

$$s_{x'_i} \in [0, s) \cup \{s\} \qquad \text{for all } x'_i \geq s.$$

For the sake of a contradiction suppose there exists some $x''_i \geq s$ such that

$$s_{x''_i} \in [0, s). \tag{11}$$

Consider a new instance where agent i's location is $y_i = x''_i$ (note that $x''_i \geq s$), all other agents have location 1, and the other agents report $\hat{\boldsymbol{x}}_{-i}$ (the same profile of reports as per (10)). In the event that $y_i = x''_i = 1$ (in which case all agents are equidistant from every facility location), assume agent i has the highest priority in the tie-breaking rule ($\triangleright$). If agent i reports their location y_i the facility location is $s_{y_i} = s_{x''_i} < s$, as per (11), and they attain utility $1 - d(s_{y_i}, y_i)$. But now misreporting to $y'_i = x_i$ then as per (9) the facility location is s where

$$s_{y_i} < s \leq y_i,$$

leading to utility $1 - d(s, y_i)$. This is a contradiction of the mechanism being DIC, since $d(s, y_i) < d(s_{y_i}, y_i)$; that is, agent i by reporting y'_i instead of their true location y_i attains strictly higher utility. We conclude that $s_{x'_i} = s$ for all $x'_i \geq s$. Thus, the mechanism is uncompromising and hence a GMM. □

3. Approximation of DIC mechanisms

Given the characterization result (Theorem 2.4) of the previous section, there is no distinction between the family of mechanisms that are DIC for some $k < n$, and the family of mechanisms that are DIC for all $k \leq n$: both families are equal to the GMM family. Accordingly, we will now simply refer to a mechanism as being DIC.

3.1. The optimal mechanism is not DIC

We first show that in general for $k < n$, the optimal mechanism is not DIC. Note that this result contrasts with the $k = n$ setting where the median mechanism is both optimal and DIC (Remark 1).

Theorem 3.1. *The optimal mechanism is DIC if and only if $k \in \{1, n\}$.*

Proof. The backward direction of the theorem statement is straightforward: If $k = 1$ then for any $i \in N$ the agent i dictator mechanism, where the mechanism output always coincides with agent i's report, is both optimal and DIC. This is trivial and we do not provide further details. If $k = n$ then the median mechanism is both optimal and DIC. This result has long been known and can be found in Black (1948); Moulin (1980); Procaccia and Tennenholtz (2013).

We now prove the forward direction using the contrapositive. Let $k \notin \{1, n\}$ and partition the agents into $\lfloor n/k \rfloor$ groups of size k, denoted by N_t for $t = 1, 2, \ldots, \lfloor n/k \rfloor$, and one group of size $n - \lfloor n/k \rfloor$, denoted by $N_{\lfloor n/k \rfloor + 1}$. We now identify $\lfloor n/k \rfloor + 1$ locations in $[0, 1]$, let

$$y_t = \frac{t}{\lfloor n/k \rfloor + 1} \qquad \text{for } t = 1, 2, \ldots, \lfloor n/k \rfloor + 1.$$

Consider a scenario such that for each $t = 1, 2, \ldots, \lfloor n/k \rfloor + 1$, all but one agent in N_t is located at y_t and a single agent is located at $y_t - t\,\varepsilon$ for some sufficiently small $\varepsilon > 0$. In each instance denote the single agent located at $y_t - t\,\varepsilon$ by $i_t \in N_t$.

In this scenario it is immediate that the optimal welfare is attained by locating the facility at location y_1, leading to a social welfare of $k - \varepsilon$ and agent i_1 attains utility $1 - \varepsilon$.

Now in a new scenario where agent i_1 has true location at $y_1 - 3\varepsilon$ the optimal mechanism must locate the facility at y_2. In this case agent i_1 attains utility zero. However, if agent i_1 misreports their location to $y_1 - \varepsilon$ then (as shown above) the facility location will be y_1 and they will attain strictly higher utility $1 - 3\varepsilon$. That is, the optimal mechanism is not DIC for $k \notin \{1, n\}$. □

Despite Theorem 3.1 stating a stark impossibility result, we note that absent strategic manipulations by the agents the optimal mechanism can be efficiently computed.

Remark 2. The optimal facility location and welfare can be computed in polynomial time for any $k \leq n$.

We sketch an informal argument. Order the agents $i \in N$ such that $x_i \le x_j$ if and only if $i \le j$. An optimal solution has two features: (1) the facility serves a *contiguous* set of k agents, i.e., if agent i and $i+2$ are served then agent $i+1$ is served, and (2) the facility is located at the median of these k served agents. Given this, a polynomial-time procedure exists by simply comparing the welfare produced by, the at most n, sets of k contiguous agents.

3.2. A lower bound on DIC approximation

Utilizing the characterization result of DIC mechanisms via the family of GMMs, we provide a lower bound on the approximation ratio for all DIC mechanisms.

The lower bound of Theorem 3.2 on the approximation ratio is illustrated in Fig. 1.

Theorem 3.2. *Let $n \ge 2$. A DIC mechanism is at best an α-approximation with $\alpha = 2\frac{k}{k+1}$ when $1 \le k \le \lceil (n-1)/2 \rceil$, and $\alpha = \max\{\frac{n-1}{k+1}, 1\}$ otherwise.*

Proof. Let M be a DIC mechanism, and consider a scenario where all n agents have distinct (true) locations contained in the interval $I = (1/2 - 1/2\varepsilon,\ 1/2 + 1/2\varepsilon)$ for some sufficiently small $\varepsilon > 0$. Denote the profile of agent locations by $\boldsymbol{x}$, and the mechanism's corresponding output by $s = M(\boldsymbol{x})$.

We consider two cases.

Case 1: Suppose $s \notin I$ and without loss of generality assume $s < 1/2 - 1/2\varepsilon$. Now suppose that agents $i = 1, 2, \ldots, n$ sequentially have their true (and reported) locations changed to $x_i = 1$, and consider the sequence of facility locations produced by the mechanism $s_1, s_2, \ldots, s_n$. By the uncompromising property (satisfied by M since it is a GMM) the location of the facility never changes from s. That is, $s_n = s$ despite every agent having location at 1. The optimal social welfare in this scenario is clearly k, however, the mechanism provides welfare of

$$k(1 - d(s, 1)) = k\,s < k(1/2 - 1/2\varepsilon) \to k/2 \qquad \text{as } \varepsilon \to 0.$$

Thus, the approximation ratio is at best $k/(k/2) = 2$.

Case 2: Suppose $s \in I$ and without loss of generality assume $s \le 1/2$. Let λ_1, λ_2 be the number of agents with true location strictly less than s, and strictly above s, respectively. Note that $\lambda_1 + \lambda_2 \in \{n-1, n\}$, since all agents have distinct locations. Similar to Case 1, suppose the λ_1 agents instead had their true (and reported) locations shifted to 0 and the λ_2 agents had true (and reported) locations shifted to 1 – by the uncompromising property the facility location is unchanged.

To attain the bound on the approximation ratio we consider two subcases where $k \le \lceil (n-1)/2 \rceil$ and $k > \lceil (n-1)/2 \rceil$.

In the first subcase ($k \le \lceil (n-1)/2 \rceil$): the optimal welfare is k, since either λ_1 or λ_2 exceeds k meaning that k agents can be served at either 0 or 1. The mechanism's welfare is at most

$$1 + (k-1)(1 - d(s, 0)) < 1 + (k-1)(1/2 + 1/2\varepsilon) \to 1/2 + k/2 \qquad \text{as } \varepsilon \to 0.$$

Thus, the approximation ratio is at best $k/(1/2 + k/2) = 2k/(k+1)$.

In the second subcase ($k > \lceil (n-1)/2 \rceil$): the optimal welfare is at worst $\lceil (n-1)/2 \rceil$, i.e., when the facility serves either λ_1 or λ_2 agents (whichever is larger) from location 0 or 1. The mechanism's welfare is at most

$$1 + (k-1)(1 - d(0, s)) < k - (k-1)(1/2 + 1/2\varepsilon) \to k/2 + 1/2 \qquad \text{as } \varepsilon \to 0.$$

Thus, the approximation ratio is at best $\lceil (n-1)/2 \rceil/(k/2 + 1/2)$, but

$$\lceil (n-1)/2 \rceil/(k/2 + 1/2) \ge \frac{(n-1)/2}{(k+1)/2} = \frac{n-1}{k+1}.$$

Furthermore since $k > (n-1)/2$ it follows that $\frac{n-1}{k+1} < 2$. Of course, this bound is only meaningful when $n - 1/k + 1 > 1$.

We conclude that when $k \le \lceil (n-1)/2 \rceil$ the approximation ratio is at best $2\frac{k}{k+1}$ and otherwise is at best $\max\{\frac{n-1}{k+1}, 1\}$. $\square$

3.3. An optimized approximation ratio for DIC mechanism

We now analyze the performance of the median mechanism for general $k \le n$. In instances where $k \in \{1, n\}$, the median mechanism is both optimal mechanism and DIC (Theorem 3.1). Furthermore, this mechanism is DIC for all $k \le n$ since the median mechanism is a GMM (Theorem 2.4).

By Theorem 3.3, the median mechanism is optimal among all DIC mechanisms for $k \le \lceil (n-1)/2 \rceil$ since the approximation-ratio matches the lower bound found in Theorem 3.2. These approximation results are also illustrated in Fig. 1.

Theorem 3.3. *The median mechanism is an α-approximation with $\alpha = 2\frac{k}{k+1}$ for $k \le \frac{2}{3}(n+1)$, and $\alpha = \min\{2\frac{k}{k+1}, 1 + 2\frac{n-k+1}{3k-2n-2}\}$ otherwise.*

Proof. Throughout the proof let i_m denote the agent with median location (choose the agent arbitrarily if multiple such agents exist), and let s_m denote the median location. The median mechanism provides welfare

$$\Pi_M(\boldsymbol{x},k) = \max_{N' \in N_k} \sum_{i \in N_k} (1 - d(s_m, x_i)) = 1 + \max_{N' \in N_{k-1,i_m}} \sum_{i \in N_{k-1,i_m}} (1 - d(s_m, x_i)),$$

where N_k is the set of all k-sized subsets of N and N_{k-1,i_m} is the set of all $(k-1)$-sized subsets of $N \backslash \{i_m\}$. This follows since the subset of agents served are always the k-closest to the facility location. Hence, given a facility location, the served subset is welfare maximizing. Furthermore, the median location coincides with at least one agent's location, i.e., agent i_m.

First, we provide an upper bound on the approximation-ratio for all k. The median mechanism locates the facility at the $\lfloor (n+1)/2 \rfloor$-th location and hence there are $\lfloor (n+1)/2 \rfloor - 1$ agents with locations (weakly) below and $\lceil (n+1)/2 \rceil - 1$ with locations (strictly) above. A lower bound on the median mechanism's welfare is attained when the agents below and above the median location at located at 0 and 1, respectively. Thus,

$$\Pi_M(\boldsymbol{x},k) \geq 1 + (k-1) \max\{1 - d(s_m, 0),\ 1 - d(s_m, 1)\},$$

and since either $d(s_m, 0) \leq 1/2$ or $d(s_m, 1) \leq 1/2$ it follows that $\Pi_M(\boldsymbol{x},k) \geq (k+1)/2$. This leads to an upper bound on the approximation-ratio of $k/((k+1)/2) = 2\frac{k}{k+1}$ for all k, since the optimal welfare is always bounded by k.

Now we attain a tighter upper bound for certain values of k. To do so, we bound the median welfare using the optimal welfare. Let s^* be the location of the facility under the optimal mechanism. Let N_m^* denote the set of k agents served under the median mechanism, and let N^* denote the set of k agents served under the optimal mechanism. We have

$$\begin{aligned}
\Pi_M(\boldsymbol{x},k) &= \sum_{i \in N_m^*} \big(1 - d(s_m, x_i)\big) \\
&\geq \sum_{i \in N^*} \big(1 - d(s_m, x_i)\big) \\
&= \sum_{i \in N^*} \Big(1 - d(s_m, x_i) - d(s^*, x_i) + d(s^*, x_i)\Big) \\
&= \Pi^*(\boldsymbol{x},k) - \sum_{i \in N^*} \Big(d(s_m, x_i) - d(s^*, x_i)\Big).
\end{aligned}$$

Clearly, the lower bound is smallest when $s_m \neq s^*$, without loss of generality assume that $s_m < s^*$. Let N_1^*, N_2^* be a partition of N^* such that $|N_1^*|, |N_2^*| \leq \lfloor (n+1)/2 \rfloor$ and all agents in N_1^* have location in $[0, s_m]$ and agent in N_2^* have location in $[s_m, 1]$. Such a partition of N^* exists since the location s_m coincides with the $\lfloor (n+1)/2 \rfloor$ highest location. Using this partition we further bound the median mechanism's welfare:

$$\begin{aligned}
\Pi_M(\boldsymbol{x},k) &\geq \Pi^*(\boldsymbol{x},k) - \sum_{i \in N_1^*} (d(s_m, x_i) - d(s^*, x_i)) - \sum_{i \in N_2^*} (d(s_m, x_i) - d(s^*, x_i)) \\
&\geq \Pi^*(\boldsymbol{x},k) - |N_1^*| \max_{x \in [0, s_m]} \Big(s_m - s^* - 2x\Big) - |N_2^*| \max_{x \in [s_m, 1]} \Big(x_i - s_m - |s^* - x|\Big) \\
&\geq \Pi^*(\boldsymbol{x},k) - |N_1^*|(s_m - s^*) - |N_2^*|(s^* - s_m) \\
&\geq \Pi^*(\boldsymbol{x},k) - (|N_2^*| - |N_1^*|)(s^* - s_m) \\
&\geq \Pi^*(\boldsymbol{x},k) - (|N_2^*| - |N_1^*|).
\end{aligned}$$

We now attain our lower bound by considering the maximum value of $|N_2^*| - |N_1^*|$. For $k \leq \lfloor (n+1)/2 \rfloor$, the value can only be guaranteed to be no larger than k – leading to a trivial zero lower bound for $\Pi_M(\boldsymbol{x},k)$. However, for $k > \lfloor (n+1)/2 \rfloor$ we attain a more useful bound by noting that

$$(|N_2^*| - |N_1^*|) \leq \lfloor (n+1)/2 \rfloor - (k - \lfloor (n+1)/2 \rfloor) = 2\lfloor (n+1)/2 \rfloor - k \leq n + 1 - k.$$

This leads to an approximation-ratio upper bound of

$$\max_{\boldsymbol{x} \in \prod_{i=1}^n X} \left\{ \frac{\Pi^*(\boldsymbol{x},k)}{\Pi^*(\boldsymbol{x},k) - n - 1 + k} \right\}$$

whenever the denominator can be guaranteed to be positive. Noting that $\Pi^*(\boldsymbol{x},k) \geq k/2$, since at least as much welfare is attained by locating the facility at $s = 1/2$, we conclude that the denominator is positive whenever

$$k > \frac{2}{3}(n+1).$$

Thus, for $k > \frac{2}{3}(n+1)$ an upper bound for the approximation-ratio is

$$\begin{aligned}\max_{\boldsymbol{x}\in\prod_{i=1}^{n} X}\left\{\frac{\Pi^*(\boldsymbol{x},k)}{\Pi^*(\boldsymbol{x},k)-n-1+k}\right\} &= \max_{\boldsymbol{x}\in\prod_{i=1}^{n} X}\left\{1+\frac{n+1-k}{\Pi^*(\boldsymbol{x},k)-n-1+k}\right\}\\ &\leq 1+\frac{n+1-k}{k/2-n-1+k}\\ &= 1+2\frac{n+1-k}{3k-2n-2}.\end{aligned}$$

This completes the proof. □

4. Discussion and conclusion

Extensions to multiple facilities. In the present paper we focused on the case of a single facility problem. However, more generally, facility location problems can be extended to consider the problem of locating multiple facilities; see, for example, Bogomolnaia and Nicolò (2005); Ehlers (2002, 2003); Fotakis and Tzamos (2013); Heo (2013); Miyagawa (1998, 2001), and Procaccia and Tennenholtz (2013). Incorporating capacity constraints within the multiple facilities location problem presents a number of challenges.[14] Firstly, the assumption of the k-closest agents being served by the facility can no longer be justified as the unique equilibrium outcome of a first-come-first-served game, since issues of multiple equilibria arise. Furthermore, and even ignoring this issue, the mechanism design problem is drastically more complicated due to the interdependence of agent utilities and the possibility that an agent's report could simultaneously affect the location of multiple facilities. Golowich et al. (2018) have explored the mechanism design problem for multiple facilities without capacity constraints via deep learning. The algorithmic problem of finding optimal facility locations is also more complicated (see Brimberg et al. (2001)).

Weakening DIC. Another natural direction to consider is weakening the DIC requirement, which requires that agents must attain maximal utility from reporting their location no matter what other agents report, and no matter other agents' true locations. A weaker, *ex post* notion of incentive compatibility may be interesting to explore, requiring instead that agents attain maximal utility from reporting their location no matter the other agents' true locations, but conditional on the other agents reporting truthfully. It is straightforward to construct *ex post* IC mechanisms that out-perform the median mechanism for certain parameter ranges. However, a complete characterization seems to be a challenging problem, which we leave for future research.

References

Abdulkadiroğlu, A., Sönmez, T., 2003. School choice: a mechanism design approach. Am. Econ. Rev. 93 (3), 729–747.

Aziz, H., 2019. Justifications of welfare guarantees under normalized utilities. ACM SIGecom Exch.

Aziz, H., Chan, H., Lee, B.E., Li, B., Walsh, T., 2020. Facility location problem with capacity constraints: algorithmic and mechanism design perspectives. In: Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI).

Aziz, H., Lee, B.E., 2018. Sub-committee approval voting and generalized justified representation axioms. In: Proceedings of AAAI/ACM Conference on AI, Ethics, and Society, pp. 3–9.

Balinski, M., Young, H., 1982. Fair Representation: Meeting the Ideal of One Man, One Vote. Yale University Press.

Barberà, S., Jackson, M., 1994. A characterization of strategy-proof social choice functions for economies with pure public goods. Soc. Choice Welf. 11 (3), 241–252.

Barberà, S., Massó, J., Serizawa, S., 1998. Strategy-proof voting on compact ranges. Games Econ. Behav. 25, 272–291.

Black, D., 1948. On the rationale of group decision-making. J. Polit. Econ. 56 (1), 23–34.

Bogomolnaia, A., Nicolò, A., 2005. Stable assignment of public facilities under congestion. J. Public Econ. Theory 7 (1), 65–91.

Border, K.C., Jordan, J.S., 1983. Straightforward elections, unanimity and phantom voters. Rev. Econ. Stud. 50 (1), 153–170.

Brandeau, M.L., Chiu, S.S., 1989. An overview of representative problems in location research. Manag. Sci. 35 (6), 645–674.

Bredereck, R., Faliszewski, P., Igarashi, A., Lackner, M., Skowron, P., 2018. Multiwinner elections with diversity constraints. In: Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI), pp. 933–940.

Brimberg, J., Korach, E., Eben-Chaim, M., Mehrez, A., 2001. The capacitated p-facility location problem on the real line. Int. Trans. Oper. Res. 8, 727–738.

Cantala, D., 2004. Choosing the level of a public good when agents have an outside option. Soc. Choice Welf. 22 (3), 491–514.

Charikar, M., Guha, S., Tardos, E., Shmoys, D., 2002. A constant-factor approximation algorithm for the k-median problem. J. Comput. Syst. Sci. 65 (1), 129–149.

Ching, S., 1997. Strategy-proofness and "median voters". Int. J. Game Theory 26, 473–490.

Cygan, M., Hajiaghayi, M.T., Khuller, S., 2012. Lp rounding for k-centers with non-uniform hard capacities. In: Proceedings of the 53rd IEEE Symposium on Foundations of Computer Science (FOCS), pp. 273–282.

Ehlers, L., 2002. Multiple public goods and lexicographic preferences. J. Math. Econ. 37 (1), 1–14.

[14] A subset of the present paper's authors have explored this extension in the paper by Aziz et al. (2020), but without establishing any general characterization results.

Ehlers, L., 2003. Multiple public goods and lexicographic preferences and single-plateaued preference rules. Games Econ. Behav. 43 (1), 1–27.
Feldman, M., Fiat, A., Golomb, I., 2016. On voting and facility location. In: Proceedings of the 17th ACM Conference on Electronic Commerce (ACM-EC), pp. 269–286.
Fotakis, D., Tzamos, C., 2013. On the power of deterministic mechanisms for facility location games. In: Fomin, F.V., Freivalds, R., Kwiatkowska, M., Peleg, D. (Eds.), Automata, Languages, and Programming. In: Lecture Notes in Computer Science, vol. 7965. Springer, Berlin, Heidelberg.
Gibbard, A., 1973. Manipulation of voting schemes: a general result. Econometrica 41 (4), 587–601.
Gibbard, A., 1977. Manipulation of schemes that mix voting with chance. Econometrica 45 (3), 665–681.
Golowich, N., Narasimhan, H., Parkes, D.C., 2018. Deep learning for multi-facility location mechanism design. In: Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI), pp. 261–267.
Heo, E.J., 2013. Strategy-proof rules for two public goods: double median rules. Soc. Choice Welf. 41, 895–922.
Jackson, M., Nicolò, A., 2004. The strategy-proof provision of public goods under congestion and crowding preferences. J. Econ. Theory 115, 278–308.
Kamada, Y., Kojima, F., 2015. Distributional constraints: theory and applications. Am. Econ. Rev. 105 (1), 67–99.
Massó, J., de Barreda, I.M., 2011. On strategy-proofness and symmetric single-peakedness. Games Econ. Behav. 72 (2), 467–484.
Miyagawa, E., 1998. Mechanisms for providing a menu of public goods. PhD dissertation. University of Rochester.
Miyagawa, E., 2001. Locating libraries on a street. Soc. Choice Welf. 18, 527–541.
Moulin, H., 1980. On strategy-proofness and single peakedness. Public Choice 45 (4), 437–455.
Nehring, K., Puppe, C., 2006. The structure of strategy-proof social choice — part I: general characterization and possibility results on median spaces. J. Econ. Theory 135 (1), 269–305.
Nehring, K., Puppe, C., 2007. Efficient and strategy-proof voting rules: a characterization. Games Econ. Behav. 59 (1), 132–153.
Nisan, N., Ronen, A., 2001. Algorithmic mechanism design. Games Econ. Behav. 35 (1–2), 166–196.
Pál, M., Tardos, É., Wexler, T., 2001. Facility location with nonuniform hard capacities. In: Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science (FOCS), pp. 329–338.
Peremans, W., Peters, H., Stel, H.v.d., Storcken, T., 1997. Strategy-proofness on Euclidean spaces. Soc. Choice Welf. 14, 379–401.
Procaccia, A.D., Tennenholtz, M., 2013. Approximate mechanism design without money. In: Proceedings of the 14th ACM Conference on Electronic Commerce (ACM-EC). ACM Press, pp. 1–26.
Satterthwaite, M., 1975. Strategy-proofness and arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions. J. Econ. Theory 10, 187–217.
Vygen, J., 2004. Approximation Algorithms Facility Location Problems. Lecture Notes, Research Institute for Discrete Mathematics. University of Bonn, Germany.
Weymark, J., 2011. A unified approach to strategy-proofness for single-peaked preferences. SERIEs 2 (4), 529–550.