Facility Location with Variable and Dynamic Populations

Yuho Wada, Tomohiro Ono, Taiki Todo, and Makoto Yokoo
Dept. Informatics, Kyushu University
Fukuoka, Japan
{wada@agent., ono@agent., todo@, yokoo@}inf.kyushu-u.ac.jp

Abstract

Facility location is a well-studied problem in social choice literature, where agents' preferences are restricted to be single-peaked. When the number of agents is treated as a variable (e.g., not observable a priori), a social choice function must be defined so that it can accept any possible number of preferences as input. Furthermore, there exist cases where multiple choices must be made continuously while agents dynamically arrive/leave. Under such variable/dynamic populations, a social choice function needs to give each agent an incentive to sincerely report her existence (e.g., participation/no-hiding). In this paper we investigate facility location models with variable/dynamic populations. For a static (one-shot), variable population model, we provide a necessary and sufficient condition for a social choice function to satisfy participation, as well as truthfulness, anonymity, and Pareto efficiency. The condition is given as a further restriction on the well-known median voter schemes. For a dynamic model, we first propose an online social choice function, which is optimal for the total sum of the distances between the choices in the previous and current periods, among any Pareto efficient functions. We then define a generalized class of online social choice functions and compare their performances both theoretically and experimentally.

1 Introduction

Facility location is a well-studied problem in the literature of social choice and known as a special case of voting (Moulin 1980). In the problem, each agent (voter) is located at a point on an interval, which represents the set of social alternatives. Under the realization of a social alternative as the outcome of a social choice function, an agent's cost is defined as the distance between the outcome and its location. In other words, agents' preferences are restricted to single-peaked ones. This restriction on their preferences guarantees the existence of a Condorcet winner. Actually the alternative most preferred by the *median voter*, i.e., the agent whose locations is the $\lfloor (n+1)/2 \rfloor$-th smallest among n agents, is a Condorcet winner.

In the literature of mechanism design, truthfulness is one of the most important properties that must be preserved by a social choice function. It requires that no agent can benefit by misreporting her location to the social choice function, regardless of other agents' reports. Clarifying the necessary and sufficient condition for a social choice function to satisfy truthfulness has greatly attracted considerable attention of researchers. The *median voter schemes*, which contains the above median rule as a special case, are the only social choice functions that satisfy truthfulness, Pareto efficiency, and anonymity (Moulin 1980).

The traditional model of facility location problems assumes that a social choice function knows the number n of agents a priori, or at least, the parameter n is given as a fixed constant. In practice, however, the number of agents is sometimes not previously observable. For instance, in a poll on a social networking service, it is somewhat unrealistic to assume that its organizer knows the exact number of votes beforehand. It may be even possible that agents in the market can partially manipulate the population. In such a case, the behavior of a social choice function for different sets of agents must be carefully analyzed.

One such situation is the case where agents' participation is voluntary, like the above online poll. An agent would choose not to participate in a social choice function if she benefits by doing so. Indeed, the development of the Internet makes it quite easy for people to get much information about an online poll before it opens. To appropriately reflect agents' preference in the outcome, a social choice function should guarantee that participation by reporting true location is weakly better than not participating. This property is known as *participation* and has been studied in the literature of social choice (Fishburn and Brams 1983). However, to the best of our knowledge, there has been no work that addressed a necessary and sufficient condition for a social choice function to satisfy participation as well as truthfulness in facility location problems.

Another possibility is to face the dynamic participation of agents, where the outcome space is divided into periods and agents arrive and depart over time. Since it is natural to assume that their arrivals and departures are also private information, a social choice function needs to be defined for any number of agents participating in the market during a period. At the same time, the truthfulness property also becomes stronger, which requires that participating at their true arrival periods by reporting their true locations is a dominant strategy.

In this paper we study such variable population in both static and dynamic facility location models. The static model is an extension of Moulin's model (1980). The dynamic model is an application of the concept of online mechanism design, originally proposed by Hajiaghayi et al. (2004) to facility location. For both of our models, our analysis is based on a perspective of mechanism design, where these four properties, i.e., truthfulness, participation, Pareto efficiency and anonymity, are treated as desiderata.

For the static model, we give a necessary and sufficient condition for a social choice function to satisfy participation, as well as truthfulness, anonymity, and Pareto efficiency. The condition is given as a further refinement of the parameters of a median voter scheme. A quite similar approach by Todo et al. (2011), addresses a property called false-name-proofness, instead of participation. The condition obtained in our paper is strictly weaker than that obtained by them. Note that, however, we also clarified that participation and false-name-proofness themselves have no inclusion relation by presenting a social choice function that is false-name-proof and violates participation.

For the dynamic model, we propose an online social choice function, the so-called dynamic target rule, which satisfies all four desiderata. It is inspired from well-known target rules for a static model (Ching and Thomson forthcoming; Klaus 2001), which is actually the social choice function obtained by the above condition by Todo et al. (2011). For the static model, the target rule works as follows: it has a pre-defined parameter p and returns, for a given profile of reported locations, the closest alternative to p among the Pareto efficient ones if p is not Pareto efficient, and returns p otherwise. The idea of the dynamic target rule for the dynamic model is that, a target rule (in the static sense) is used for each period, where the outcome at the current period t is set as the parameter for the next period $t + 1$.

Although Pareto efficiency is a very popular qualitative measure for evaluating social choice functions, quantitative analysis has also been used in mechanism design. Guaranteeing the worst-case performance is one such example. We analyze the dynamic target rule from this perspective and show that it is optimal among online Pareto efficient social choice functions in terms of a new measure called *replacement cost*, which is defined as the total sum of the distances between the choices in previous and current periods, where the choice at the very first period 0 is commonly fixed as α for all those rules. On the other hand, it performs poorly in terms of another well-known measure called *social cost*. Therefore, we propose a class of online social choice functions and show that on average they perform reasonably well in terms of social cost, although there is no dominance between any two of them.

2 Related Works

Moulin (1980) proposes a necessary and sufficient condition for a social choice function to satisfy truthfulness, anonymity, and Pareto efficiency, which can also be considered as a characterization of median voter schemes. Procaccia and Tennenholtz (2013) initiated the worst-case analysis of a social choice function, which is also known as approximate mechanism design. Besides them, there are quite a few works on voting in dynamic situation. Tennenholtz (2004) studied a voting model where agents arrive dynamically, but it focused on analyzing social rankings instead of social choice. Parkes and Procaccia (2013) studied the dynamic preferences represented by Markov decision processes, without considering the dynamic arrival of agents.

Median voter schemes have also been widely studied in the literature of social choice. Ching and Thomson (forthcoming) proposed a class of social choice functions called *target rules* as a subclass of the median voter schemes, and give a characterization of them by Pareto efficiency and another consistency property called population monotonicity. However, they did not consider any strategic manipulation by agents. Arribillaga and Massó (2016) compared a superclass of median voter schemes, called *generalized median voter schemes*, in terms of manipulability by preference misreports. However, they focused on a fixed population model.

Fishburn and Brams (1983) and Moulin (1988) initiated discussion on participation in social choice. However, their analysis is for general voting situations with unrestricted preferences, instead of single-peaked preferences. Participation has also recently been attracting much attention. Bochet and Gordon (2012) dealt with the property in locating multiple facilities. Brandl et al. (2017) studied the participation property in the problem of assigning indivisible objects, which requires that reporting that all the items are indifferent is not beneficial for each agent. Some properties that are relative to participation, including partially hiding information, have also been studied in various social choice and mechanism design environments, such as voting (Brandl et al. 2015; Brandl, Brandt, and Hofbauer 2015; Brandt, Geist, and Peters 2016) and assignment (Atlamaz and Klaus 2007; Guo and Deligkas 2013; Todo, Sun, and Yokoo 2014; Fujita et al. 2015).

In social choice literature, some works focused on manipulations in situations with variable populations, especially known as false-name manipulations (Todo, Iwasaki, and Yokoo 2011; Bu 2013; Lesca, Todo, and Yokoo 2014; Sonoda, Todo, and Yokoo 2016; Ono, Todo, and Yokoo 2017). Conitzer (2008) also tackled a property called *anonymity-proofness* as a combination of false-name-proofness and participation and proposed a randomized anonymity-proof voting rule for general preferences; but it did not focus on participation. In general, to the best of our knowledge, there has been no work focusing on the participation property in facility location with variable or dynamic populations.

3 Preliminaries

In this section we introduce the general notations and definitions that are commonly used in our two different facility location models to avoid redundantly defining ideas that are essentially identical. Each specified model will be formally defined at the beginning of Sections 4 and 5, respectively.

Let $\mathcal{N}$ be the set of potential agents and let $N \subseteq \mathcal{N}$ be a set of participating agents. A participating agent $i \in N$ has a true type $\theta_i \in \Theta_i$, where Θ_i is the set of potential types of agent i. Let $\theta := (\theta_i)_{i \in N} \in \Theta_N := \times_{i \in N} \Theta_i$ be a profile of the types of a set N of agents, where $\theta_{-i} := (\theta_j)_{j \neq i}$

indicates a profile that consists of all the types in θ except for the agent i's type. Let $\mathcal{O}$ be the set of social alternatives and let $o \in \mathcal{O}$ be a social alternative. A cost function c is shared among all the potential agents, where the cost incurred to an agent i with true type θ_i when a social alternative o is achieved is represented as $c(\theta_i, o) \in \mathbb{R}_{\geq 0}$.

A social choice function $f = (f_N)_{N \subseteq \mathcal{N}}$ is defined as a family of functions, each f_N of which is a mapping from Θ_N to $\mathcal{O}$. This means that, each function f_N takes a profile θ of types jointly reported by the set N of agents as an input, and returns a social alternative o as an outcome. We write f_N as f if it is clear from the context. Here let R be a function that restricts an agent's reportable types, for a given agent's true type. Formally, $R(\theta_i) \subseteq \Theta_i$ indicates the set of types reportable by agent i with true type θ_i. We assume that $\theta_i \in R(\theta_i)$ holds for any i and θ_i, meaning that each agent can always report her true type.

Now we introduce the four desiderata, truthfulness, anonymity, Pareto efficiency, and participation. Truthfulness requires that for each agent, truthfully reporting her type is weakly better than reporting any false type, regardless of the other agents' reports. Anonymity requires that any permutation of agents types does not change the outcome, even though the agents' names are changed. Pareto efficiency requires that there is no other outcome that is weakly preferred to the outcome chosen by a social choice function by all the agents, and strictly preferred to by at least one agent. Finally, participation requires that for each agent, participation by reporting her true type is weakly better than not participating.

Definition 1 (Truthfulness)**.** A social choice function f is *truthful* if for any N, any $\theta \in \Theta_N$, any $i \in N$, and any $\theta'_i \in R(\theta_i)$, it holds that $c(\theta_i, f(\theta)) \leq c(\theta_i, f(\theta'_i, \theta_{-i}))$.

Definition 2 (Anonymity)**.** A social choice function f is *anonymous* if for any N, any N' such that $|N'| = |N|$, any $\theta \in \Theta_N$, and any $\theta' \in \Theta_{N'}$, the existence of permutation $\pi : N \to N'$ of N such that $\theta_i = \theta'_{\pi(i)}$ for every $i \in N$ implies $f(\theta) = f(\theta')$.

Definition 3 (Pareto efficiency)**.** An outcome $o \in \mathcal{O}$ *Pareto dominates* another outcome o' under profile $\theta \in \Theta_N$ if $c(\theta_i, o) \leq c(\theta_i, o')$ for all $i \in N$ and $c(\theta_j, o) < c(\theta_j, o')$ for at least one $j \in N$. A social choice function f is *Pareto efficient* if for any N and any $\theta \in \Theta_N$, there exists no outcome $o \in O$ that Pareto dominates $f(\theta)$.

Definition 4 (Participation)**.** A social choice function f satisfies *participation* if for any N, any $\theta \in \Theta_N$, and any $i \in N$, it holds that $c(\theta_i, f(\theta)) \leq c(\theta_i, f(\theta_{-i}))$.

Let us also define the median function *med*, which is used in several places in this paper. The median function takes any odd number of values as input and returns the unique median value. Actually, in this paper, any use of the median function takes an odd number of values as input.

4 Facility Location in a Static Model

In this section, we first formally define our static model of facility location, by specifying the notation introduced in the previous section. Let $\mathcal{I} := [0, 1]$ be an interval. We then define $\Theta_i := \mathcal{I}$ for every $i \in \mathcal{N}$, $R(\theta_i) := \Theta_i$ for every $\theta_i \in \Theta_i$, $\mathcal{O} := \mathcal{I}$, and $c(\theta_i, o) := |\theta_i - o|$. In other words, in this static model, each agent is located at a point θ_i on the interval $[0, 1]$, and her cost incurred from a social alternative o is defined as the distance between o and θ_i.

Moulin (1980) provided a necessary and sufficient condition for a social choice function defined for a fixed N to satisfy truthfulness, anonymity, and Pareto efficiency. This result can be straightforwardly extended to our static model with variable populations as the following corollary.

Corollary 1. *A social choice function f is truthful, anonymous, and Pareto efficient if and only if it has, for any positive integer $n \in \mathbb{N}$, $n-1$ parameters $p^n = (p^n_1, \ldots, p^n_{n-1}) \in \mathcal{I}^{n-1}$ such that $p^n_1 \leq \ldots \leq p^n_{n-1}$ and for any set N such that $|N| = n$ and any $\theta \in \Theta_N$,*

$$f_N(\theta) = \textit{med}(\theta_1, \ldots, \theta_n, p^n_1, \ldots, p^n_{n-1}). \qquad (1)$$

Such a social choice function is known as a median voter scheme, where those $n - 1$ parameters are called "phantom voters" and it chooses the median point out of $2n - 1$ points in total. For example, if we set $p^n = (0, \ldots, 0)$, it chooses the location of the leftmost agent among n agents. Similarly, if we set $p^n = (0, \ldots, 0, 1)$, it chooses the location of the second agent from the left, and so on. The following example shows that participation does not always hold when there is no restriction on those parameters.

Example 1. Consider the social choice function described in Eq. 1 that has the following parameters: $p^n = (0, \ldots, 0)$ when $n < 5$ and $p^n = (1, \ldots, 1)$ when $n \geq 5$. In other words, this returns the leftmost location among agents' reports when less than five agents participate and the rightmost location otherwise. Assume there are five agents such that the profile of their true types is $\theta = (0, 0, 1, 1, 1)$, i.e., two are located at 0 and the other three at 1. When all five agents participate and truthfully report their types, the outcome is 1. On the other hand, if an agent at 0 does not participate, the outcome is 0. Thus the agent at 0 has an incentive to not participate, which violates the participation property.

The following theorem is our main result for the static model, which clarifies a necessary and sufficient condition for a social choice function to satisfy participation as well as truthfulness, anonymity, and Pareto efficiency.

Theorem 1. *A social choice function f is truthful, anonymous, Pareto efficient, and satisfies participation if and only if it is described in Eq. 1 and its parameters are such that for any positive integer $n \in \{2, \ldots, |N|\}$ and any $m \in \{1, \ldots, n-2\}$,*

$$p^n_m \leq p^{n-1}_m \leq p^n_{m+1}. \qquad (2)$$

Proof. We first show the if part, i.e., we confirm that any such f satisfies all the desiderata. From Corollary 1, f is obviously truthful, anonymous, and Pareto efficient. Now assume for the sake of contradiction that there exist N, $\theta \in \Theta_N$, and $i \in N$ such that i benefits by not participating. Letting $q := f(\theta)$, we assume without loss of generality that $\theta_i < q$. Here, let L be the number of points strictly smaller than q in the profile of $2n - 1$ points (θ, p^n), L' be the number of types out of L, and L'' be the number of parameters out of L. By definition, $L = L' + L''$, $L' \geq 1$, and

$1 \leq L \leq n-1$. We also analogously define K, K', and K'' for the case where i does not participate. It is then obvious that $K' = L' - 1$, since i's type is removed in θ_{-i}. Furthermore, Eq. 2 guarantees that $L'' - 1 \leq K'' \leq L''$ holds for the number of parameters. Therefore, $K \leq L - 1 \leq n - 2$ holds, which means that when i does not participate, the median (the $n-1$-th smallest) of the $2n-3$ points (θ_{-i}, p^{n-1}) must be at least as large as q. This violates the assumption that i is better off by not participating.

We next prove the only if part by showing that the participation property, combined with the other three, implies Eq. 2. For the sake of contradiction, assume that a social choice function f, described in Eq. 1, satisfies participation and for some n and some $m \in \{1, \ldots, n-2\}$, either $p_m^{n-1} < p_m^n$ or $p_{m+1}^n < p_m^{n-1}$ holds for its parameters. When $p_m^{n-1} < p_m^n$ holds, consider a profile $\theta \in \Theta_N$ of n locations such that $n-m$ agents are located at 0 and m agents are located at 1. By the definition of the median function *med*, the outcome when all the agents participate in and report truthfully is $f(\theta) = p_m^n$. When an agent i at 0 does not participate, the outcome becomes $f(\theta_{-i}) = \mathit{med}(\theta_{-i}, p^{n-1}) = p_m^{n-1}$. Since $p_m^{n-1} < p_m^n$ holds, $c(0, p_m^{n-1}) < c(0, p_m^n)$ also holds, which violates participation. The same argument holds for the case of $p_{m+1}^n < p_m^{n-1}$ from symmetry, which completes the proof. □

To confirm the independence of these four conditions, we show four social choice functions, each of which violates one of the four properties and still satisfies the remaining three. The center rule, choosing the average of the leftmost and rightmost locations, is not truthful, but it is anonymous, Pareto efficient, and satisfies participation. The dictatorship rule, which chooses the participating agent with the alphabetically youngest identity as a dictator, violates anonymity, but it is truthful, Pareto efficient, and satisfies participation. A social choice function that chooses a pre-determined fixed point as the outcome regardless of the input is not Pareto efficient, but it is truthful, anonymous, and satisfies participation. Finally, the social choice function presented in Example 1 violates participation, but it still satisfies the remaining three properties.

The condition can also be treated as a necessary and sufficient condition for a median voter scheme to satisfy participation. It is weaker than the necessary and sufficient condition for them to satisfy false-name-proofness, which requires the existence of a parameter $p^* \in \mathcal{I}$ such that for any n, $p^n = (p^*, \ldots, p^*)$ holds. This implies that under the assumption of truthfulness, anonymity, and Pareto efficiency, false-name-proofness is stronger than participation. However, these two properties themselves have no inclusion relation. Consider a social choice function that performs the above rule with parameter $p^* = 1$ when $n \leq 5$ and returns 1 regardless of the input when $n > 5$. This rule is false-name-proof and anonymous, but not Pareto efficient. Furthermore, it violates participation, since an agent has incentive to avoid participating when $n > 5$.

The above theorem also works as a characterization of a subclass of median voter schemes that is 'monotonic' on the market population. The monotonicity here means that, assuming an outcome $o \in \mathcal{I}$ for the current population, the new arrival of an agent i with $\theta_i \leq o$ does not push the outcome to the right. Indeed, any median voter scheme that violates participation does not preserve this monotonicity property. For instance, in the above example, adding a new agent at 0 moves the outcome from 0 to 1.

5 Facility Location in a Dynamic Model

In this section, we formally define our dynamic model of facility location. Let $\mathcal{I} := [0,1]$ be an interval again, and let $\mathcal{T} := \{1, \ldots, t, \ldots, T\}$ be the sequence of T periods in the market. We then define $\Theta_i \subset \mathcal{I} \times \mathcal{T} \times \mathcal{T}$ for every $i \in \mathcal{N}$. An agent's type θ_i is represented as a triple (x_i, a_i, d_i), where $x_i \in \mathcal{I}$ is her location, $a_i \in \mathcal{T}$ is her arrival period, and $d_i \in \mathcal{T}$ is her departure period. Furthermore, let $\mathcal{O} := \mathcal{I}^T$, and let $o := (o_1, \ldots, o_t, \ldots, o_T) \in \mathcal{O}$ be an outcome, where $o_t \in \mathcal{I}$ indicates the location of the facility at period t under the outcome o. For given $\theta_i = (x_i, a_i, d_i)$ and $o = (o_t)_{1 \leq t \leq T}$, the cost of an agent i with type θ under outcome o is given as $c(\theta_i, o) := \sum_{t=a_i}^{d_i} |x_i - o_t|$. In other words, in this dynamic model, each agent i enters the market at the beginning of period a_i, reports her location x_i on the interval $[0,1]$ to a social choice function, and leaves the market at the end of period d_i. The agent's cost incurred from an outcome o is defined as the sum of the distances between o_t and x_i for all the periods where it stays in the market. Also, let $\alpha \in \mathcal{I}$ be the social state at the very beginning period 0, e.g., the original location of the facility at the beginning of the facility location problem.

In our dynamic model, we assume that a social choice function must be *online*, meaning that a decision at period t only depends on the information available during the first t periods. For notation simplicity, we sometimes refer to the leftmost location at period t when the agents' type profile is θ as $l(t, \theta)$, and the rightmost location as $r(t, \theta)$. Note that in this dynamic model, an outcome is Pareto efficient if and only if it returns a Pareto efficient location at every period. Formally, $o = (o_t)_{1 \leq t \leq T}$ is Pareto efficient under profile θ if and only if $l(t, \theta) \leq o_t \leq r(t, \theta)$ for every $t \in \mathcal{T}$.

We also introduce a reasonable assumption on the power of misreports by agents. We assume that each agent cannot arrive earlier than its true arrival time, which is a quite popular assumption in online mechanism design (Hajiaghayi, Kleinberg, and Parkes 2004; Parkes 2007).

Assumption 1 (No-Early-Arrival)**.** An agent i whose true arrival time is a_i cannot arrive earlier than a_i. Formally, for any $i \in \mathcal{N}$ and any $\theta_i = (x_i, a_i, d_i) \in \Theta_i$, $R(\theta_i)$ is such that $\forall \theta'_i = (x'_i, a'_i, d'_i) \in R(\theta)$, it holds that $a_i \leq a'_i$.

If an online social choice function independently runs a truthful social choice function (in the static sense) for each period, it is truthful even though agents can misreport their arrivals. Actually, any combination of the above social choice rules in Eq. 2 satisfies all these properties. In this section, however, we propose a new online social choice function that is not just such a combination. Instead, it utilizes the current location of the facility to decide the location for the next period.

Mechanism 1 (Dynamic target rule)**.** The *dynamic target rule* τ is a social choice function such that for any $N \subseteq \mathcal{N}$ and any $\theta \in \Theta_N$,

$$\tau^t(\theta) = med(\tau^{t-1}(\theta), l(t,\theta), r(t,\theta)),$$

where τ^t represents a function that returns the location of the facility for period $t \in \mathcal{T}$ and $\tau^0(\theta) := \alpha$ for any θ.

Under the dynamic target rule, if the location chosen in the previous period is still Pareto efficient in the current period, the location does not change. Otherwise, the closest Pareto efficient location from the previous location is chosen. The following example demonstrates how the dynamic target rule works for a given profile of types.

Example 2. Let $\alpha = 0.4$ and $T = 6$, and there are seven agents, whose true types are given as $\theta_1 = (0.1, 3, 5)$, $\theta_2 = (0.2, 6, 6)$, $\theta_3 = (0.3, 3, 4)$, $\theta_4 = (0.4, 1, 1)$, $\theta_5 = (0.5, 1, 2)$, $\theta_6 = (0.6, 1, 3)$, and $\theta_7 = (0.7, 2, 2)$.

In period 1, there are three agents, each of which is located at 0.4, 0.5, and 0.6. The dynamic target rule therefore locates the facility at $\tau^1 = med(\alpha, 0.4, 0.6) = 0.4$ for period 1. In period 2, agents are located at 0.5, 0.6, and 0.7. Thus the facility is located at $\tau^2 = med(\tau^1, 0.5, 0.7) = med(0.4, 0.5, 0.7) = 0.5$. We can analogously calculate the sequence of locations, which is $(0.4, 0.5, 0.5, 0.3, 0.1, 0.2)$.

The following theorem shows that the dynamic target satisfies all four desiderata. The most surprising part is its truthfulness and participation; although it utilizes the current location to decide the future location, it does not provide any chance for agents to improve by any possible manipulation.

Theorem 2. *The dynamic target rule is truthful, anonymous, and Pareto efficient, and satisfies participation.*

Proof Sketch. Both anonymity and Pareto efficiency obviously hold. To show participation and truthfulness, we first observe two key properties of the static functions used in each period, which is actually the target rule for the static model and thus its parameters satisfy Eq. 2.

The first property is that, any possible manipulation by a single agent in the static model, i.e., misreporting its location and not participating, moves the outcome to the opposite side. More precisely, letting i be such a manipulator, $u \in \mathcal{I}$ be the original outcome under its truth-telling, and $v \in \mathcal{I}$ be the manipulated outcome, we can show that either $x_i \le u \le v$ or $x_i \ge u \ge v$ holds. For misreporting, this is a well-known property commonly preserved by any median voter scheme. For not participating, it is essentially a special case of what we shown in the proof of Theorem 1. The second property is a kind of monotonicity on the parameter. For any two different points $\beta, \beta'(< \beta) \in \mathcal{I}$, and any profile of reported locations $x = (x_i)_{i \in N}$, $med(\beta, \min_i x_i, \max_i x_i) \ge med(\beta', \min_i x_i, \max_i x_i)$ holds.

To show both the participation and truthfulness of the dynamic target rule τ, it obviously suffices to show that no agent can benefit by any manipulation, even if it can, at each period, report any location (or even choose not to participate) regardless of its reports in the previous periods. At period 1, the target rule τ uses $\tau^0 = \alpha$ as the parameter. Then, from the above first property, the outcome by truth-telling is

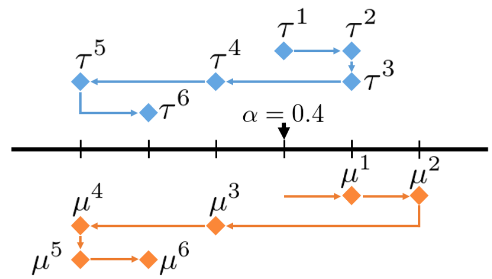

Figure 1: The outcomes of the dynamic target rule τ, indicated by blue, and the dynamic median rule μ that always chooses the median location among participating agents, indicated by orange, for the profile given in Example 2.

closer to x_i than that by any manipulation. Let us then assume that, in period $t \in \mathcal{T}$, the location $u \in \mathcal{I}$ of the facility under truth-telling and the location $v \in \mathcal{I}$ under such a manipulation satisfies either $x_i \le u \le v$ or $x_i \ge u \ge v$. Then let u' be the location returned by τ^{t+1} with parameter u and v' be the location returned by τ^{t+1} with parameter v, under the truth-telling of the manipulator i. From the second property, either $x_i \le u' \le v'$ or $x_i \ge u' \ge v'$. Furthermore, from the first property again, for both u' and v', no manipulation draws them closer to x_i. By recursively applying this argument, we show that there is no beneficial manipulation. □

There might still be some other way to utilize the current location to decide the future location. However, the following results show the optimality of the proposed rule in terms of the replacement cost among all (possibly not online) Pareto efficient ones. Note that the restriction to Pareto efficient rules is essential; otherwise choosing α for every period obviously minimizes the replacement cost.

Definition 5 (Replacement cost)**.** For given θ, the *replacement cost* of a social choice function f, denoted as $RC(f, \theta)$, is defined by

$$\sum_{t \in T} |f^t(\theta) - f^{t-1}(\theta)|,$$

where $f^0(\theta) := \alpha$.

Intuitively, the replacement cost is a quantitative measure that reflects the cost of changing the social state, such as making a new political decision in an election or rebuilding a train station like the facility location. In practice, choosing an social alternative that is quite different from the current one is more costly for society, for instance, due to an enforcement of new rules by the government.

Theorem 3. *The dynamic target rule is optimal for replacement cost among all Pareto efficient social choice functions. Formally, for any Pareto efficient (possibly not online) social choice function f, any $N \subseteq \mathcal{N}$ and any $\theta \in \Theta_N$,*

$$RC(\tau, \theta) \le RC(f, \theta).$$

Proof. Let $N \subseteq \mathcal{N}$ and $\theta \in \Theta_N$ respectively be an arbitrarily chosen set of participating agents and arbitrarily chosen profile of types. We assume without loss of generality that the first outcome that differs from α is on the positive side from α. Formally, letting $s_0 := 0$ and $s_0^+ := \min\{t \geq s_0 \mid \tau^t \neq \alpha\} \in \mathcal{T}$, we assume that $\tau^{s_0^+} > \alpha$ holds. Also, let $s_1 \in \mathcal{T}$ be the first period where the sequence of locations changes the direction of the move from positive to negative, s_2 be the first period where the sequence of locations changes the direction of the move from negative to positive, and so on. Formally, for an odd $h \geq 1$, letting $\bar{s}_h := \min_{t'}\{t' > s_{h-1} \mid \tau^{t'+1} < \tau^{t'}\}$, we define $s_h := \min_t\{s_{h-1} \leq t \leq \bar{s}_h \mid \tau^t = \tau^{\bar{s}_h}\}$. For an even $h \geq 2$, we use $\bar{s}_h := \min_{t'}\{t' > s_{h-1} \mid \tau^{t'+1} > \tau^{t'}\}$. By this procedure, we obtain the sequence of period $s = (s_1, s_2, \dots, s_h, \dots, s_H)$. For the profile given in Example 2, the sequence is $\{s_1, s_2, s_3\} = \{2, 5, 6\}$, e.g., at period $s_1 = 2$, it stops moving in a positive direction (see Figure 1).

To complete that proof, it suffices to show that for each $h \in \{0, \dots, H-1\}$ and for each sequence of periods $\{s_h, \dots, s_{h+1}\} \subseteq \mathcal{T}$ (where s_0 is defined as zero for a technical reason), the replacement cost of τ is less than that of any Pareto efficient rule f. Formally, for any $h \in \{0, \dots, H-1\}$,

$$\sum_{t \in \{s_h \dots, s_{h+1}-1\}} |\tau^{t+1} - \tau^t| \leq \sum_{t \in \{s_h \dots, s_{h+1}-1\}} |f^{t+1} - f^t|.$$

For Example 2, it compares the total moves for each subsequence of periods: $\{0, 1, 2\}$, $\{2, 3, 4, 5\}$, and $\{5, 6\}$.

Assuming that h is even (or zero), we consider the sequence of periods ending at s_{h+1}. Since s_h is the first period where τ changes the direction of its move from negative to positive, τ^{s_h} is located at the right extreme of the Pareto efficient set at the period, i.e., $\tau^{s_h} = r(s_h, \theta)$ holds. Since f is Pareto efficient, $f^{s_h} \leq \tau^{s_h}$ holds. Analogously, $\tau^{s_{h+1}}$ is located at the left extreme of the Pareto efficient set at the period, and thus both $\tau^{s_{h+1}} = l(s_{h+1}, \theta)$ and $\tau^{s_{h+1}} \leq f^{s_{h+1}}$ hold. Moreover, since τ does not change direction during the sequence of periods, we can obtain

$$\sum_{t \in \{s_h \dots, s_{h+1}-1\}} |\tau^{t+1} - \tau^t| = |\tau^{s_{h+1}} - \tau^{s_h}|.$$

Therefore, regardless of the actual move of f during the sequence of periods,

$$\begin{aligned} \sum_{t \in \{s_h \dots, s_{h+1}-1\}} |f^{t+1} - f^t| &\geq |f^{s_{h+1}} - f^{s_h}| \\ &\geq |\tau^{s_{h+1}} - \tau^{s_h}| \end{aligned}$$

holds. It also holds for any odd h from symmetry. □

6 Balancing Two Measures in Average

Social cost is a well-known evaluation criterion of social choice functions in the literature of algorithmic mechanism design. In our dynamic model, we can also straightforwardly define the social cost of a social choice function.

Definition 6 (Social Cost)**.** For given θ, the *social cost* of a social choice function f, denoted as $SC(f, \theta)$, is defined by

$$\sum_{i \in N} c(\theta_i, f(\theta)) := \sum_{t \in T} \sum_{i \in N^t} |x_i - f^t(\theta)|,$$

where $N^t := \{i \in N \mid a_i \leq t \leq d_i\}$ indicates the set of agents present at period t.

Since agents have no time discount on their costs in our dynamic model, we can change the order of summations (see the right-hand side). Here, the inner summation corresponds to the social cost for one particular period. This means that, for any input θ, the social cost is minimized by choosing the median location at every period. Let μ refer to such an online social choice function, called the dynamic median rule. It is easy to see that the social cost of the dynamic target rule can be arbitrarily worse than that of the dynamic median rule. On the other hand, the replacement cost of the dynamic median rule can be arbitrarily worse than that of the dynamic target rule. Our purpose in this section is therefore to find an online social choice function that balances the performance in terms of these two measures.

The following is a class of online social choice functions, which contains the dynamic median rule μ and the dynamic target rule τ as two extremes.

Mechanism 2 (k-shifted median)**.** For any $k \in \mathbb{N}_{\geq 0}$, the *$k$-shifted median* rule σ^k is an online social choice function such that for any $N \subseteq \mathcal{N}$, any $\theta \in \Theta_N$, and any $t \in \mathcal{T}$, $\sigma^{k,t}(\theta)$ is a social choice function defined in Eq. 1, where its parameters p^n, for each n, are defined as follows:

$$\underbrace{\sigma^{k,t-1}(\theta), \dots, \sigma^{k,t-1}(\theta)}_{k}, \underbrace{0, \dots, 0}_{\lceil (n-1-k)/2 \rceil}, \underbrace{1, \dots, 1}_{\lfloor (n-1-k)/2 \rfloor}$$

where $p^n = (\sigma^{k,t-1}(\theta), \dots, \sigma^{k,t-1}(\theta))$ if $n \leq k$, and $\sigma^{k,0}(\theta) := \alpha$ for any θ.

The basic idea of k-shifted median is as follows. In a median voter scheme, we set the location of the previous period as the location of k phantom voters; the location becomes less likely to move by increasing k. We can easily observe that setting $k = 0$ corresponds to the dynamic median rule and choosing a sufficiently large k corresponds to the dynamic target rule. Note that any choice of the parameters, as long as it follows the above manner, satisfies Eq. 2, and thus it satisfies the truthfulness and participation properties, as well as anonymity and Pareto efficiency, even in the dynamic model. Here we can also clarify the relationships of any pair of the k-shifted median rules in terms of replacement cost, which shows that a k-shifted median rule with a smaller k dominates that with any larger parameter $k'(> k)$.

Theorem 4. *For any pair $k \in \mathbb{N}_{\geq 0}$, the k-shifted median rule σ^k is truthful, anonymous, Pareto efficient, and satisfies participation. Furthermore, for any $k \in \mathbb{N}_{\geq 0}$, any $N \subseteq \mathcal{N}$, and any $\theta \in \Theta_N$, $RC(\sigma^{k+1}, \theta) \leq RC(\sigma^k, \theta)$ holds.*

Proof Sketch. For all four desiderata, almost the same argument with the proof for the dynamic target rule holds for the following two reasons: (i) any of these online social choice functions uses the static social choice function described in Eq. 2 for each period, and (ii) any such function has a property that any possible manipulation by an agent, i.e., misreporting location and not participating, always moves the outcome to the opposite side, as we already observed in the proof of Theorem 2. For dominance on replacement costs, a

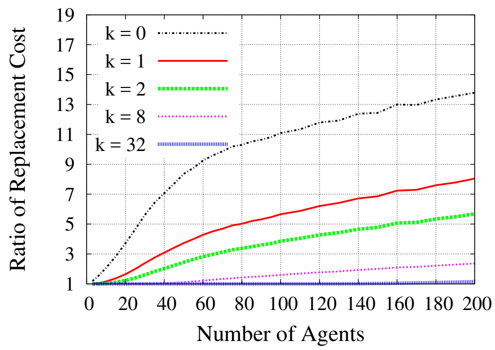

Figure 2: The simulation result for ratio of replacement cost. Each curve shows ratio of replacement cost of k-shifted median, where $k = 0$ corresponds to the dynamic median.

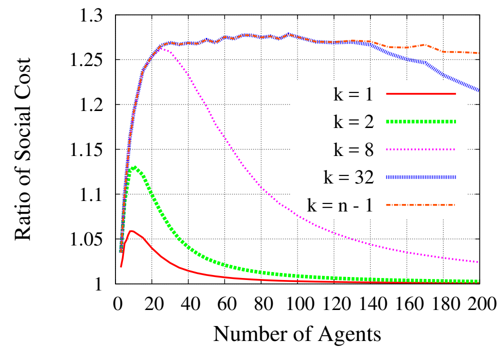

Figure 3: The simulation result for ratio of social cost. Each curve shows ratio of social cost of k-shifted median, where $k = n - 1$ corresponds to the dynamic target.

quite similar argument for the dynamic target rule also holds, by first obtaining a sequence of periods from the move of the k+1-shifted median rule and comparing its replacement cost to $k + 1$-shifted median rule for each subsequence. □

On the other hand, for any pair $k, k'(> k) \in \mathbb{N}_{\geq 0}$, the k-shifted median rule cannot dominate the k'-shifted median rule with respect to the social cost (except for $k = 0$). That is, we can find at least one profile of types under which the social cost of the k-shifted median rule is strictly larger than that of the k'-shifted median rule.

Example 3. Let $\alpha = 1 - \epsilon$ for a sufficiently small ϵ and assume there are $2k + 2$ agents, whose types are given as follows: agent 1 has type $(0, 1, T)$, agents $2, \ldots, k + 1$ have type $(1 - \epsilon, 1, 1)$, agents $k + 2, \ldots, 2k + 1$ have type $(\epsilon, 2, T)$, and agent $2k + 2$ has type $(1, 1, T)$. In period 1, agents $1, 2, \ldots, k + 1, 2k + 2$ participate. The k-shifted median σ^k then adds $k + 1$ phantom voters, k of which are at $\alpha = 1 - \epsilon$ and one is at 0. Thus, the facility is built at ϵ, and the sum of the participating agents' costs at this period is 1. At every subsequent periods, k agents located at $1 - \epsilon$ participate, while the k agents located at ϵ no longer exist. The facility is kept at ϵ, where the sum of the costs at each period is $1 + k(1 - 2\epsilon)$. Therefore, the social cost is $2 + k(T - 1)(1 - 2\epsilon)$. On the other hand, for any $k' > k$, the k'-shifted median rule $\sigma^{k'}$ builds the facility at $1 - \epsilon$ in period 1, and keeps it unchanged at every subsequent period. Therefore, the social cost is $k(1 - 2\epsilon) + (T - 1)$, which is strictly smaller than the above when $T > 2$ (and the gap expands when T grows).

In contrast, our simulation results show that such an "unfortunate" example rarely occurs when agents' types are sufficiently distributed, and these social choice functions actually performs well on average. We simulate some shifted median rules by changing the parameters, as well as the dynamic target and dynamic median rules, and compare their replacement and social costs. More specifically, we set $T = 100$ and randomly generated 10,000 instances of the dynamic facility location problem, where each agent's location x_i is independently drawn from a predefined identical distribution over $[0, 1]$, and the arrival period a_i is also independently drawn from a predefined identical distribution over $[1, T]$. For the distribution of the departure period d_i, we use both the uniform and exponential distributions over $[a_i, T]$. In both figures, the x-axis denotes the number of participating agents, and the y-axis denotes the ratio of the replacement and social cost against the optimal solution, i.e., the dynamic target and the dynamic median, respectively.

One of the most important observation from our experiments is that under any of the randomly generated 10000 instances and under both distributions for d_i, a shifted median rule with a smaller k outperforms any shifted median rule with a larger k in terms of social cost. This means that, if we can guarantee that agents' types are sufficiently distributed, the proposed k-shifted median rule is a reasonable candidate that balances the quality of the replacement and social costs. Furthermore, the figures show that the performance of the k-shifted median rules in average varies based on the choice of parameter k. The ratio of replacement cost increases when the number of agents n raises, since the (k-shifted) median location frequently changes. On the other hand, the ratio of social cost converges to one for large n, since the optimal social cost by the dynamic target rule is already very large.

7 Conclusion

We studied variable populations in the static and dynamic facility location models. For static model, we clarified a necessary and sufficient condition for a social choice function to satisfy participation, as well as truthfulness, anonymity, and Pareto efficiency. For dynamic model, we proposed a class of online social choice functions and analyzed their performance in both theoretical and experimental ways.

One obvious future work is to clarify the existence of online social choice functions that theoretically outperform the dynamic target for social cost and the dynamic median for replacement cost. Studying a different measure, such as minimizing the number of replacements, i.e., $\min \#\{t \in \mathcal{T} \mid f^t(\theta) \neq f^{t-1}(\theta)\}$ would also be a possible future direction. Considering more powerful manipulations such as renaming in the dynamic model might also be interesting.

References

Arribillaga, R. P., and Mass, J. 2016. Comparing generalized median voter schemes according to their manipulability. *Theoretical Economics* 11(2):547–586.

Atlamaz, M., and Klaus, B. 2007. Manipulation via endowments in exchange markets with indivisible goods. *Social Choice and Welfare* 28(1):1–18.

Bochet, O., and Gordon, S. 2012. Priorities in the location of multiple public facilities. *Games and Economic Behavior* 74(1):52–67.

Brandl, F.; Brandt, F.; Geist, C.; and Hofbauer, J. 2015. Strategic abstention based on preference extensions: Positive results and computer-generated impossibilities. In *Proceedings of IJCAI-15*, 18–24.

Brandl, F.; Brandt, F.; and Hofbauer, J. 2015. Incentives for participation and abstention in probabilistic social choice. In *Proceedings of AAMAS-15*, 1411–1419.

Brandl, F.; Brandt, F.; and Hofbauer, J. 2017. Random assignment with optional participation. In *Proceedings of AAMAS-17*, 326–334.

Brandt, F.; Geist, C.; and Peters, D. 2016. Optimal bounds for the no-show paradox via SAT solving. In *Proceedings of AAMAS-16*, 314–322.

Bu, N. 2013. Unfolding the mystery of false-name-proofness. *Economics Letters* 120(3):559–561.

Ching, S., and Thomson, W. forthcoming. Population-Monotonic Solutions in Public Good Economies with Single- Peaked Preferences. *Social Choice and Welfare*.

Conitzer, V. 2008. Anonymity-proof voting rules. In *Proceedings of WINE-08*, 295–306.

Fishburn, P. C., and Brams, S. J. 1983. Paradoxes of preferential voting. *Mathematics Magazine* 56(4):207–214.

Fujita, E.; Lesca, J.; Sonoda, A.; Todo, T.; and Yokoo, M. 2015. A complexity approach for core-selecting exchange with multiple indivisible goods under lexicographic preferences. In *Proceedings of AAAI-15*, 907–913.

Guo, M., and Deligkas, A. 2013. Revenue maximization via hiding item attributes. In *Proceedings of IJCAI-13*, 157–163.

Hajiaghayi, M. T.; Kleinberg, R. D.; and Parkes, D. C. 2004. Adaptive limited-supply online auctions. In *Proceedings of EC-04*, 71–80.

Klaus, B. 2001. Target rules for public choice economies on tree networks and in euclidean spaces. *Theory and Decision* 51(1):13–29.

Lesca, J.; Todo, T.; and Yokoo, M. 2014. Coexistence of utilitarian efficiency and false-name-proofness in social choice. In *Proceedings of AAMAS-14*, 1201–1208.

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

Moulin, H. 1988. Condorcet's principle implies the no show paradox. *Journal of Economic Theory* 45(1):53–64.

Ono, T.; Todo, T.; and Yokoo, M. 2017. Facility location with public residence. To appear in *Proceedings of PRIMA-17*.

Parkes, D. C., and Procaccia, A. D. 2013. Dynamic social choice with evolving preferences. In *Proceedings of AAAI-13*, 767–773.

Parkes, D. C. 2007. Online mechanisms. In Nisan, N.; Roughgarden, T.; Éva Tardos; and Vazirani, V. V., eds., *Algorithmic Game Theory*. Cambridge University Press. chapter 16.

Procaccia, A. D., and Tennenholtz, M. 2013. Approximate mechanism design without money. *ACM Transactions on Economics and Computation* 1(4):18.

Sonoda, A.; Todo, T.; and Yokoo, M. 2016. False-name-proof locations of two facilities: Economic and algorithmic approaches. In *Proceedings of AAAI-16*, 615–621.

Tennenholtz, M. 2004. Transitive voting. In *Proceedings of EC-04*, 230–231.

Todo, T.; Iwasaki, A.; and Yokoo, M. 2011. False-name-proof mechanism design without money. In *Proceedings of AAMAS-11*, 651–658.

Todo, T.; Sun, H.; and Yokoo, M. 2014. Strategyproof exchange with multiple private endowments. In *Proceedings of AAAI-14*, 805–811.