False-name-proof Mechanism Design without Money

Taiki Todo, Atsushi Iwasaki, and Makoto Yokoo
Department of Informatics, Kyushu University
Motooka 744, Fukuoka, Japan
{todo@agent., iwasaki@, yokoo@}is.kyushu-u.ac.jp

ABSTRACT

Mechanism design studies how to design mechanisms that result in good outcomes even when agents strategically report their preferences. In traditional settings, it is assumed that a mechanism can enforce payments to give an incentive for agents to act honestly. However, in many Internet application domains, introducing monetary transfers is impossible or undesirable. Also, in such highly anonymous settings as the Internet, declaring preferences dishonestly is not the only way to manipulate the mechanism. Often, it is possible for an agent to pretend to be multiple agents and submit multiple reports under different identifiers, e.g., by creating different e-mail addresses. The effect of such false-name manipulations can be more serious in a mechanism without monetary transfers, since submitting multiple reports would have no risk.

In this paper, we present a case study in *false-name-proof mechanism design without money*. In our basic setting, agents are located on a real line, and the mechanism must select the location of a public facility; the cost of an agent is its distance to the facility. This setting is called the *facility location problem* and can represent various situations where an agent's preference is *single-peaked*. First, we fully characterize the deterministic false-name-proof facility location mechanisms in this basic setting. By utilizing this characterization, we show the tight bounds of the approximation ratios for two objective functions: social cost and maximum cost. We then extend the results in two natural directions: a domain where a mechanism can be randomized and a domain where agents are located in a tree. Furthermore, we clarify the connections between false-name-proofness and other related properties.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multi-agent systems*; J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Algorithms, Economics, Theory

Cite as: False-name-proof Mechanism Design without Money, Taiki Todo, Atsushi Iwasaki, and Makoto Yokoo, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Tumer, Yolum, Sonenberg and Stone (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. 651-658.

Keywords

Auction and mechanism design, social choice theory, facility location problems

1. INTRODUCTION

1.1 Background

Mechanism design has become an integral part of electronic commerce and a promising field for applying AI and agent technologies. In particular, the celebrated Vickrey-Clarke-Groves (VCG) mechanism for combinatorial auctions, which is considered one crucial contribution of mechanism design, has been applied to several domains. One of its advantages is that it satisfies a property called *strategy-proofness*; no agent ever benefits from misreporting her preference, regardless of the other agents' strategies. The VCG mechanism achieves this property by collecting an appropriate amount of payment from each winner of the auction.

In several domains such as the Internet, however, implementing payments is sometimes impossible mainly due to security/banking issues. Moreover, there are several application fields in which monetary transfers should not be introduced due to ethical/legal considerations, including political decision making or kidney exchanges. Thus, mechanisms must be developed that satisfy strategy-proofness without involving monetary transfers. Such *mechanism design without money* is quite challenging and has attracted considerable attention among computer scientists (see [5, 10]).

Meanwhile, in such highly anonymous settings as the Internet, reporting preference insincerely is not the only way to manipulate a mechanism. Often, it is possible for an agent to pretend to be multiple agents and participate in a mechanism multiple times by using different identifiers, e.g., by creating different e-mail accounts. Since many Web applications require a valid e-mail address only, an agent can create multiple e-mail address at practically no cost. Such strategic behaviors called *false-name manipulations* have been discussed so far in the mechanism design field.

In environments in which payments can be made securely, authenticating each identifier and collecting a participation fee might discourage agents from using multiple identifiers. Furthermore, in mechanisms with monetary transfers, adding false identifiers is risky. For example, in an auction, the manipulator might have to pay a lot of money or buy unnecessary items by such false-name manipulation.

In contrast, such manipulations are more likely to occur in a mechanism without monetary transfers, since submitting multiple reports is less risky. For example, in voting, casting additional votes is unlikely to create disadvantages for the manipulator. To the best of our knowledge, there exist very few works on false-name manipulations in mechanism design without money. One notable exception is a work by Conitzer [4], which characterized anonymity-proof voting rules (i.e., rules that satisfy false-name-proofness and voluntary participation). The obtained result is rather negative. In essence, an anonymity-proof voting rule can take into account voters' preferences only when voters unanimously prefers one candidate within two candidates that are chosen at random. Furthermore, even if we only require strategy-proofness, the Gibbard-Satterthwaite theorem states that it is impossible to make a mechanism strategy-proof when agents' preferences are general (see [7]).

1.2 Our Results

In this paper, we present a case study in false-name-proof mechanism design without money. Assuming that agents' preferences are highly structured, we avoid falling into the negative result in Conitzer [4], or in more general sense, the Gibbard-Satterthwaite Theorem. We focus on *facility location problems* and discuss how difficult it is to incentivize agents to behave sincerely, even though they can use false identifiers. This is the first work to deal with false-name manipulations in facility location problems.

We discuss a facility location problem on a real line as a basic setting and characterize deterministic false-name-proof facility location mechanisms. Our characterization is inspired by Moulin's characterization of strategy-proofness [6]. To simplify expositions and notations, we define the cost of an agent as the distance between her location and a facility. It is straightforward to extend this characterization to a domain with general single-peaked preferences. Additionally, we establish the tight bounds of the approximation ratios achieved by deterministic false-name-proof mechanisms for two objective functions: social cost and maximum cost.

We then extend the results of the basic case in two further directions. One is a domain of randomized mechanisms, and the other is a facility location problem on a tree. For randomized mechanisms, we show in Section 4.1 that the left-right-middle mechanism, which was originally proposed in Procaccia and Tennenholtz [8], satisfies false-name-proofness. Furthermore, we show a lower bound of the approximation ratio for the social cost. On the other hand, for the facility location problem on a tree, we characterize deterministic false-name-proof mechanisms in Section 4.2. Our characterization can be considered a refinement of the result by Schummer and Vohra [9], in which they characterized deterministic strategy-proof mechanisms on a tree.

Furthermore, in Section 5, we clarified the connections between false-name-proofness and other related properties in a facility location problem on a tree. We focused on *population monotonicity*, *group-strategyproofness*, and *anonymity-proofness*, which have been discussed in the literature of social choice and mechanism design. By utilizing our characterization, we show that both population monotonicity and anonymity-proofness are equivalent to false-name-proofness. We also show that there exists a group-strategyproof mechanism which is not false-name-proof.

1.3 Related Works

Facility location problems have also been considered an important famework of social choice due to the highly structured preferences of agents in the setting: *single-peaked preferences*. There exist many application domains with such single-peaked preferences. For example, in political decision making, an agent's peak is her most preferred alternative. Moulin [6] characterized strategy-proof, Pareto efficient, and anonymous facility location mechanisms on a real line. Schummer and Vohra [9] extended Moulin's results to facility location problems on graphs.

Procaccia and Tennenholtz [8] presented a case study in approximate mechanism design without money and established tight bounds for the approximation ratio achieved by strategy-proof facility location mechanisms on a real line. They also proposed two extensions of facility location problems: a domain where two facilities must be located and a domain where each agent owns multiple locations. Alon et al. [1] discussed the maximum cost of strategy-proof facility location mechanisms on several network topologies. Guo and Conitzer [5] is one of the most recent development of approximate mechanism design without money for strategy-proof resource allocations.

False-name manipulations have also been widely studied in combinatorial auctions. Yokoo et al. [12] proposed a condition where VCG becomes false-name-proof. Todo et al. [11] characterized false-name-proof combinatorial auction mechanisms. Besides combinatorial auctions, false-name-proofness and its relatives have been discussed in other mechanism design fields, such as voting [4] and coalitional games [2]. In particular, Conitzer [4] proposed an extended property called *anonymity-proofness* in voting and characterized anonymity-proof voting rules.

2. PRELIMINARIES

2.1 Basic Model

In this paper, we deal with facility location problems in which a mechanism locates one facility. Let n denote the number of agents (identifiers) joining a mechanism and N ($|N| = n$) the set of agents. Note that the number of agents n is defined to be variable in $\mathbb{N}$ to discuss the change of the number of agents joining a mechanism. Each agent $i \in N$ has a true location (or the most preferred location) x_i on a graph G. In this paper, we restrict our attention to *peak-only* mechanisms, i.e., each agent reports only her most preferred location. In a more general setting (e.g., voting), this can be a quite strong restriction. However, in our setting, we can assume any strategy-proof mechanism is peak-only, since the peak location of each agent is her only private information. The cost of an agent i is defined by the distance $d(\cdot, \cdot)$ between her true location and the location of a facility: if the facility is located at y, the cost of agent i with location x_i is $\text{cost}(x_i, y) = d(x_i, y)$. If a graph G is a real line, the distance is defined as $|x_i - y|$.

A (direct revelation, deterministic) *facility location mechanism* (or simply *mechanism*) is a function that maps a reported location profile $x = (x_1, \dots, x_n)$ by the set of agents to a location of a facility y on a graph G. A mechanism must locate a facility with respect to any number of agents n, since we consider an environment where each agent may use multiple identifiers (formally defined in Section 2.2). For this reason, we define a mechanism f as a set of functions $(f^n)_{n\in\mathbb{N}}$, where each function f^n is a mapping from a set of location profiles reported by n identifiers to the graph. For simplicity, we assume that a mechanism is *anonymous*, meaning that the obtained results are invariant under the permutation of identifiers.

DEFINITION 1 (FACILITY LOCATION MECHANISM). *For any natural number $n \in \mathbb{N}$, a* facility location mechanism f *assigns an outcome $f^n(x)$ to any reported location profile $x = (x_1, \ldots, x_n) \in G^n$:*

$$f = (f^n)_{n \in \mathbb{N}}, f^n : G^n \to G.$$

In facility location problems, each agent reports her location x'_i, which is not necessarily her true location x_i, to the mechanism. However, in a *strategy-proof* mechanism, it is guaranteed that each agent reports her true location x_i to the mechanism if she behaves to minimize her cost.

DEFINITION 2 (STRATEGY-PROOFNESS). *A mechanism f is* strategy-proof *if $\forall n \in \mathbb{N}$, $\forall i \in N$, $\forall x_{-i}$, $\forall x_i$, $\forall x'_i$, $\mathrm{cost}(x_i, f(x_i, x_{-i})) \leq \mathrm{cost}(x_i, f(x'_i, x_{-i}))$.*

Here x_{-i} denotes the reported location profile by agents except i. That is, $f(x'_i, x_{-i})$ is the location of a facility when agent i reports x'_i and other agents report x_{-i}. Definition 2 means that a mechanism is strategy-proof if for each agent, reporting her true location is a dominant strategy; it minimizes her cost regardless of the strategies of other agents.

Several strategy-proof mechanisms have been developed for facility location problems. For a real line, one well-known strategy-proof mechanism is the *median* mechanism, which chooses the median location among the reported locations (if the number of agents n is even, locates at the $n/2$-th smallest location). To simplify the exposition and the notations, we define a function $\mathrm{med}(\cdot)$ that returns the median point for a given profile of real numbers.

For the facility location problem on a real line, Moulin [6] characterized strategy-proof mechanisms.

THEOREM 1 (MOULIN, 1980). *A mechanism f is strategy-proof, Pareto efficient, and anonymous if and only if for all $n \in \mathbb{N}$, there exist $n-1$ real numbers $\alpha^n_1, \alpha^n_2, \ldots, \alpha^n_{n-1}$ such that for all reported location profile $x = (x_1, \ldots, x_n) \in \mathbb{R}^n$,*

$$f(x) = \mathrm{med}(x_1, \ldots, x_n, \alpha^n_1, \ldots, \alpha^n_{n-1}). \qquad (1)$$

In the case of a real line, *Pareto efficiency* requires that a facility be located at a point between the leftmost and rightmost locations among the reported locations. Theorem 1 means that any Pareto efficient, anonymous, and strategy-proof mechanism can be represented by appropriately setting the parameters in Eq. (1). Indeed, the median mechanism is represented by setting these parameters as follows:

$$\forall n \in \mathbb{N}, \forall m \in \{1, \ldots, n-1\}, \alpha^n_m = \begin{cases} -\infty & \text{if } m \text{ is odd} \\ \infty & \text{if } m \text{ is even.} \end{cases}$$

Also, the leftmost mechanism, which locates a facility at the smallest location among the reported locations, is represented by setting all parameters to $-\infty$.

We focus on a worst case analysis to consider the performance of the mechanisms. This analysis is commonly used in the literature of (algorithmic) mechanism design, especially by computer scientists. We introduce two objective functions: *social cost* and *maximum cost*. The social cost is the sum of the costs of all agents. A solution minimizing the social cost is also called a *minisum* solution. On the other hand, the maximum cost is defined by the cost of the agent whose cost is the highest among all agents. A solution minimizing the maximum cost is also called a *minimax* solution, which achieves an equitable location. We now define the *approximation ratios* of a mechanism.

DEFINITION 3 (APPROXIMATION RATIO). *The* approximation ratios *of a mechanism f for the social cost and the maximum cost are defined as follows:*

$$\max_x \frac{\sum_{i \in N} \mathrm{cost}(x_i, f(x))}{\min_{y \in G} \sum_{i \in N} \mathrm{cost}(x_i, y)},$$

$$\max_x \frac{\max_{i \in N} \mathrm{cost}(x_i, f(x))}{\min_{y \in G} \max_{i \in N} \mathrm{cost}(x_i, y)}.$$

2.2 False-name-proofness

In this subsection, we formalize false-name-proofness in facility location problems. First, we introduce some notations for discussing false-name manipulations.

Let ϕ_i denote the set of identifiers used by agent i. This is also the private information of agent i. Let x_{ϕ_i} denote a location profile reported by a set of identifiers ϕ_i, and let $x_{-\phi_i}$ denote a location profile reported by identifiers except for ϕ_i. In this definition, x_{ϕ_i} is considered a false-name manipulation by i.

DEFINITION 4 (FALSE-NAME-PROOFNESS). *A mechanism f is* false-name-proof *if $\forall n \in \mathbb{N}$, $\forall i \in N$, $\forall x_{-\phi_i}$, $\forall x_i$, $\forall \phi_i$, $\forall x_{\phi_i}$, $\mathrm{cost}(x_i, f(x_i, x_{-\phi_i})) \leq \mathrm{cost}(x_i, f(x_{\phi_i}, x_{-\phi_i}))$.*

In other words, a mechanism is false-name-proof if for each agent, reporting her true location by using a single identifier is a dominant strategy, even though she can use multiple identifiers. The following example shows that the median mechanism on a real line is not false-name-proof: an agent can reduce her cost by using multiple identifiers.

Example 1. Consider the median mechanism on a real line and $N = \{1, 2, 3\}$. Assume that $x_1 = 1$, $x_2 = 2$, and $x_3 = 3$. If they report their locations truthfully, the mechanism locates a facility at 2. However, if agent 1 adds two false identifiers and reports $x_{\phi_1} = (1, 1, 1)$, the mechanism locates a facility at 1. By this false-name manipulation, agent 1 can strictly reduce her cost.

3. BASIC RESULTS

3.1 Characterization Theorem

Now we are ready to show our characterization theorem of false-name-proof mechanisms on a real line. More precisely, we provide a necessary and sufficient condition for a mechanism to be false-name-proof, Pareto efficient, and anonymous. Lemmas 1 and 2 prove the theorem.

THEOREM 2. *A mechanism f is false-name-proof, Pareto efficient, and anonymous if and only if there exists a real number α such that for all $n \in \mathbb{N}$ and for all reported location profiles $x = (x_1, \ldots, x_n) \in \mathbb{R}^n$,*

$$f(x) = \mathrm{med}(x_1, \ldots, x_n, \underbrace{\alpha, \ldots, \alpha}_{n-1}). \qquad (2)$$

LEMMA 1. *If a mechanism f satisfies Eq. (2), f is false-name-proof, Pareto efficient, and anonymous.*

PROOF. If f satisfies Eq. (2), it also satisfies Eq. (1). Thus, f is Pareto efficient and anonymous. Therefore, we now show that f is false-name-proof if it satisfies Eq. (2).

Let us discuss false-name manipulations by agent i and show that no false-name manipulation reduces her cost. Let $\mathrm{lt}(x_{-i})$ denote the leftmost location in a location profile x_{-i} reported by agents except i, and let $\mathrm{rt}(x_{-i})$ denote the rightmost location. If $\mathrm{lt}(x_{-i}) \leq \alpha \leq \mathrm{rt}(x_{-i})$ holds, f always locates a facility at α regardless of i's strategy.

We prove that f is false-name-proof if $\alpha < \mathrm{lt}(x_{-i})$. The same argument can be applied if $\mathrm{rt}(x_{-i}) < \alpha$ from the symmetry. We show that agent i cannot reduce her cost by false-name manipulations in each of the following three cases: (i) $x_i \leq \alpha$, (ii) $\alpha < x_i \leq \mathrm{lt}(x_{-i})$, and (iii) $\mathrm{lt}(x_{-i}) < x_i$.

Case (i) If i's true location x_i satisfies $x_i \leq \alpha$, f locates a facility at α if i reports truthfully. In this situation, reporting $x_{i'} > \alpha$ by all false identifiers $i' \in \phi_i$ are the only false-name manipulations that affect the outcome. However, by these manipulations, the outcome becomes strictly further away from x_i.

Case (ii) If x_i satisfies $\alpha < x_i \leq \mathrm{lt}(x_{-i})$, f locates a facility at x_i when agent i reports her true location. In this case, agent i has no incentive to use false identifiers.

Case (iii) If x_i satisfies $\mathrm{lt}(x_{-i}) < x_i$, f locates a facility at $\mathrm{lt}(x_{-i})$ if i reports truthfully. In this situation, reporting $x_{i'} < \mathrm{lt}(x_{-i})$ by all false identifiers $i' \in \phi_i$ are the only false-name manipulations that affect the outcome. However, by these manipulations, the outcome moves further away from x_i. $\square$

LEMMA 2. *If a mechanism f is false-name-proof, Pareto efficient, and anonymous, f satisfies Eq. (2).*

PROOF. Since false-name-proofness is a generalization of strategy-proofness, if f is false-name-proof, Pareto efficient, and anonymous, then for all $n \in \mathbb{N}$, f has $n-1$ parameters $\alpha_1^n, \alpha_2^n, \dots, \alpha_{n-1}^n$ satisfying $\alpha_1^n \leq \dots \leq \alpha_{n-1}^n$ and locates a facility at the median point defined by Eq. (1) (from Theorem 1). To prove this lemma, it suffices to show that there exists $\alpha \in \mathbb{R}$ such that for all $n \geq 2$,

$$\alpha_1^n = \dots = \alpha_{n-1}^n = \alpha. \tag{3}$$

We prove this lemma by induction on n. For $n = 2$, Eq. (3) obviously holds since there exists only one parameter α_1^2.

We suppose that Eq. (3) holds for all $n \leq k$ and show that it also holds for $n = k+1$. Assuming $\alpha_1^k = \dots = \alpha_{k-1}^k = \alpha$ holds, we prove $\alpha_1^{k+1} = \dots = \alpha_k^{k+1} = \alpha$ also holds.

First, assume that $\alpha < \alpha_1^{k+1}$ holds and derive a contradiction. Now consider that a location profile $x = (x_1, \dots, x_k)$ such that $\alpha < x_1 < \dots < x_k < \alpha_1^{k+1}$ holds. In this case, the outcome of the mechanism f is $f(x) = x_1$. If an agent k whose location is the largest among all k agents adds another identifier k' and reports $x_{\phi_k} = (x_k, x_k)$, then the outcome changes to $f(x_{\phi_k}, x_{-k}) = x_k$. By this manipulation, k's cost decreases from $x_k - x_1$ to 0. This contradicts the assumption of false-name-proofness. From symmetry, the same argument can be applied to $\alpha_k^{k+1} < \alpha$.

Next, assume that there exists $j \in \{1, \dots, k-1\}$ such that $\alpha_j^{k+1} \leq \alpha < \alpha_{j+1}^{k+1}$ holds and derive a contradiction. Consider a location profile $x = (x_1, \dots, x_k)$ such that $\alpha < x_1 < \dots < x_k < \alpha_{j+1}^{k+1}$. In this case, we have $f(x) = x_1$. If agent $l = k - j$ (whose location is the $(k-j)$-th smallest among the k agents) reports $x_{\phi_l} = (x_l, x_l)$, the outcome becomes $f(x_{\phi_l}, x_{-l}) = x_l$. By this manipulation, l's cost decreases from $x_l - x_1$ to 0. This contradicts false-name-proofness. From symmetry, we can apply the same argument to $\alpha_j^{k+1} < \alpha \leq \alpha_{j+1}^{k+1}$. $\square$

Theorem 2 means that f is false-name-proof if and only if it has a fixed parameter α regardless of the number of agents and locates a facility based on the following rule. Given a reported location profile x, f locates a facility at α if α is between the smallest and largest locations among the reported locations; otherwise, it locates at the closest location to α among the reported locations.

Since we can obtain the leftmost and the rightmost mechanism by setting the parameter α to $-\infty$ and ∞, respectively, both mechanisms satisfy false-name-proofness. However, since the median mechanism cannot be represented in this form, it does not satisfy false-name-proofness. In this way, Theorem 2 allows us to easily verify if a mechanism satisfies false-name-proofness.

One might think that a mechanism, which always locates a facility at a pre-defined point regardless of the agents' reports, satisfies false-name-proofness. This is true; there is no incentive for agents to participate at all. Even though it is false-name-proof, we cannot represent it in the form of Theorem 2 because it is not Pareto efficient. However, it is straightforward to obtain the following corollary that can deal with such non-efficient mechanisms.

COROLLARY 1. *A mechanism f is false-name-proof and anonymous if and only if there exist three real numbers $\alpha_L, \alpha, \alpha_R (\alpha_L \leq \alpha \leq \alpha_R)$ such that for all $n \in \mathbb{N}$ and for all reported location profiles $x = (x_1, \dots, x_n) \in \mathbb{R}^n$,*

$$f(x) = \mathrm{med}(x_1, \dots, x_n, \alpha_L, \overbrace{\alpha, \dots, \alpha}^{n-1}, \alpha_R). \tag{4}$$

The additional parameters α_L and α_R define the range of the mechanism; the mechanism described in Corollary 1 always locates a facility in the range $[\alpha_L, \alpha_R]$. Indeed, Eq. (4) can describe the above mechanism by defining the parameters as $\alpha_L = \alpha_R = \alpha$. Clearly, if we set these two parameters as $-\infty$ and ∞, respectively, we obtain Theorem 2.

Procaccia and Tennenholtz [8] extended the facility location problem on a real line to a domain where each agent i owns ω_i locations $\mathbf{x}_i = (x_{i,1}, \dots, x_{i,\omega_i})$. The domain is still single-peaked; when an agent hopes to minimize the sum of distances to her locations, her peak is the median point of her locations. As stated in Section 1.2, we can easily apply Theorem 2 to general single-peaked domains. Thus, we obtain the following corollary.

COROLLARY 2. *A mechanism f for the multiple locations setting is false-name-proof, Pareto efficient, and anonymous if and only if there exists a real number α such that for all $n \in \mathbb{N}$, for all $\omega_i|_{i \in N}$, and for all reported location profiles $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n) \in \mathbb{R}^n$,*

$$f(\mathbf{x}) = \mathrm{med}(\mathrm{med}(\mathbf{x}_1), \dots, \mathrm{med}(\mathbf{x}_n), \overbrace{\alpha, \dots, \alpha}^{n-1}). \tag{5}$$

Procaccia and Tennenholtz [8] developed a valuable mechanism, which creates ω_i copies of the median of each agent i and returns the median point among all copies. In their setting, the number of locations ω_i owned by agent i is public.

Table 1: Summary of the approximation ratios achieved by deterministic strategy-proof/false-name-proof mechanisms on a real line. UB and LB indicate upper and lower bounds. SP and FNP indicate strategy-proof and false-name-proof.

	SP	FNP
Social Cost	UB: 1 LB: 1	UB: $O(n)$ (Thm. 4) LB: $\Omega(n)$ (Thm. 3)
Maximum Cost	UB: 2 ([8]) LB: 2 ([8])	UB: 2 LB: 2 (Thm. 5)

This means that the domain is not anonymous; a change of agents' location profiles affects the outcome. In contrast, we deal with anonymous mechanisms in the above corollary by assuming that ω_i is the private information of i.

3.2 Approximation Ratios

In this subsection, we analyze the performance of false-name-proof mechanisms from the viewpoint of *approximate mechanism design without money* [8]. Table 1 summarizes the results of this section. Our results, shown in the right-most column, provide tight bounds of the approximation ratios for both social and maximum costs.

3.2.1 Social Cost

We first consider social cost as an objective function and show the lower bound of the approximation ratio achieved by deterministic false-name-proof mechanisms for it in Theorem 3. Theorem 4 shows that the lower bound is tight.

THEOREM 3. *Any deterministic false-name-proof mechanism has an approximation ratio of $\Omega(n)$ for social cost.*

PROOF. First, consider a location profile $x = (0, 1)$ and assume that $f(x) = y \in \mathbb{R}$. If $y \neq 0$, then consider another location profile $x' = (0, \dots, 0, 1)$ where $|x'| = n$. From false-name-proofness, $y' = f(x')$ must satisfy $|y'| \geq |y| = \mathrm{cost}(0, f(x))$. In this case, the social cost with respect to x' becomes $(n-1)|y'| + |y' - 1| \geq (n-1)|y| = (n-1) \cdot \mathrm{cost}(0, f(x))$, which depends on the number of agents n. If $y = 0$, we can apply a similar argument by considering a location profile $x' = (0, 1, \dots, 1)$. $\square$

THEOREM 4. *The leftmost mechanism has an approximation ratio of $n - 1$ for the social cost.*

PROOF. The leftmost mechanism is a false-name-proof mechanism whose parameter α is defined as $-\infty$. For any reported location profile x, the approximation ratio of the leftmost mechanism for the social cost with respect to x is defined as $\sum_{i \neq 1} |x_i - x_1| / \sum_{i \neq \lceil n/2 \rceil} |x_i - x_{\lceil n/2 \rceil}|$. Here the denominator is the social cost of the median mechanism that has an approximation ratio of 1 for social cost. This ratio is at most

$$\frac{\sum_{i \neq 1} |x_n - x_1|}{\sum_{i=1,n} |x_i - x_{\lceil n/2 \rceil}|} = \frac{(n-1)(x_n - x_1)}{x_n - x_1} = n - 1,$$

and we have equality if x satisfies $x_1 < x_2 = \ldots = x_n$. $\square$

3.2.2 Maximum Cost

In contrast to Section 3.2.1, we consider maximum cost as an objective function in this subsection. First, we show the lower bound of an approximation ratio for maximum cost.

THEOREM 5. *Any deterministic false-name-proof mechanism has an approximation ratio of at least 2 for maximum cost.*

The proof is straightforward. It was shown by Procaccia and Tennenholtz [8] that any deterministic strategy-proof mechanism has an approximation ratio of at least 2 for maximum cost. Since false-name-proofness implies strategy-proofness, the lower bound does not decrease if we require false-name-proofness.

As stated in Section 3.1, the leftmost mechanism is false-name-proof. Furthermore, the leftmost mechanism has an approximation ratio of 2 for maximum cost by [8]. This implies that the bound obtained in Theorem 5 is tight.

4. EXTENDED RESULTS

In Section 3, we showed a basic result on false-name-proof mechanisms on a real line. We then extend the result in two directions. In Section 4.1, we discuss the bound of approximation ratios achieved by randomized false-name-proof mechanisms. In Section 4.2, we characterize deterministic false-name-proof mechanisms on a tree.

4.1 Randomized Mechanisms

Our results shown in Section 3.2 suggest the difficulty of designing a deterministic false-name-proof mechanism that achieves good approximation ratios, even if the domain is a real line. One natural approach to this problem is to use *randomized* mechanisms, which return a probability distribution over a real line for a given location profile. In this subsection, we discuss whether allowing randomization enables mechanisms to achieve better approximation ratios.

First, let us introduce a mechanism called *left-right-middle*, which was developed by Procaccia and Tennenholtz [8].

MECHANISM 1 (LEFT-RIGHT-MIDDLE). *Given a location profile $x = (x_1, \dots, x_n)$, the* left-right-middle *mechanism locates a facility at x_1 with probability 1/4, x_n with probability 1/4, and $(x_1 + x_n)/2$ with probability 1/2.*

Note that the cost of an agent is defined as the expected distance from the location. Also, approximation ratios can be redefined over a distribution. Now we confirm that the left-right-middle mechanism is false-name-proof and calculate the approximation ratio for social cost.

THEOREM 6. *The left-right-middle mechanism is false-name-proof and has an approximation ratio of $n/2$ for social cost.*

PROOF. First we show that the left-right-middle mechanism is false-name-proof. The mechanism defines the outcome depending only on the leftmost and rightmost locations. Thus, from a similar argument for strategy-proofness, no agent can be better off by any false-name manipulations.

We now turn to proving that the left-right-middle mechanism is $n/2$-approximation. For any reported location profile x, the approximation ratio of the left-right-middle mechanism for social cost with respect to x is defined as

$$\frac{\frac{1}{4}\sum_i |x_i - x_1| + \frac{1}{4}\sum_i |x_i - x_n| + \frac{1}{2}\sum_i |x_i - \frac{x_1 + x_n}{2}|}{\sum_i |x_i - x_{\lceil n/2 \rceil}|}.$$

This is at most

$$\frac{\frac{1}{4}\sum_{i \in N}(x_i - x_1) + \frac{1}{4}\sum_{i \in N}(x_n - x_i) + \frac{1}{2}\sum_{i \in N}(x_n - \frac{x_1 + x_n}{2})}{\sum_{i=1,n} |x_i - x_{\lceil n/2 \rceil}|}$$
$$= \frac{\frac{n}{2}(x_n - x_1)}{x_n - x_1} = \frac{n}{2},$$

and equality holds if x satisfies $x_1 < x_2 = \ldots = x_n$. $\square$

This shows us that with randomization, we can slightly improve the social cost than with deterministic mechanisms, e.g., the leftmost mechanism, when the number of agent n is large. However, from an algorithmic point of view, the performances of these mechanisms are essentially the same: both have an approximation ratio of $O(n)$ for social cost. Thus, we next discuss if there exist randomized false-name-proof mechanisms which have an essentially better approximation ratio for social cost. We show a lower bound of the approximation ratio for social cost and answer the question.

THEOREM 7. *Any randomized false-name-proof mechanism has an approximation ratio of $\Omega(n)$ for social cost.*

PROOF. Consider arbitrary randomized false-name-proof mechanism f. Let $x = (0, 1)$ be a location profile when there are two agents and let $P = f(x)$ be the outcome distribution over $\mathbb{R}$. Intuitively, $\text{cost}(0, P) + \text{cost}(1, P) \geq 1$ holds (formally proved in [8], Lemma 2.6). Thus, we assume $\text{cost}(1, P) \geq 1/2$ without loss of generality.

Then, we consider the case with n agents $1, \ldots, n$ and the reported location profile $x' = (0, 1, \ldots, 1)$. Let $P' = f(x')$ be the outcome distribution. Since f is false-name-proof, $\text{cost}(1, P') \geq \text{cost}(1, P)$. Thus, the social cost is at least $(n-1)/2$. On the other hand, the optimal solution with respect to the profile x' is to locate a facility at 1, in which the social cost is 1. Thus, the ratio is $(n-1)/2$. $\square$

That is, even if randomization is allowed, the approximation ratio of a false-name-proof mechanism for social cost ever depends on the number of agents n.

In contrast to social cost, we can obtain a tight bound for maximum cost from the result of Procaccia and Tennenholtz [8]. They showed that the left-right-middle mechanism achieves an optimal approximation ratio of $3/2$ for maximum cost. Furthermore, as shown in Theorem 6, the left-right-middle is false-name-proof. Thus, the tight bound of the approximation ratio for maximum cost is $3/2$.

4.2 Location on a Tree

Several application domains of facility locations have much more complicated structures, i.e., graph structure, than a simple line. Thus, facility location problems on a graph are natural extensions of the 1-dimensional case to such application domains, as discussed in Section 3. One simple structure of graphs is a tree [1, 9]. Therefore, we characterize deterministic false-name-proof mechanisms on a tree.

First, let us introduce additional notations. Let G be a *tree*, which is a finite connected graph composed of the union of a finite number of curves of finite length and contains no cycle. Let $L(G) \subset G$ be a set of leaves of G. For any two points $p, q \in G$, let $[p, q]$ denote the *path* between p, q. Note that we can define a unique path $[p, q]$ for all p, q since G contains no cycle. Also, let $d(p, q)$ denote the *distance* between two points p, q, which is defined as the path-length between the two points. When a facility is located at $y \in G$, the cost of an agent i with true location $x_i \in G$ is defined as $\text{cost}(x_i, y) = d(x_i, y)$.

Although each agent still has a peak on the tree G, this setting is no longer a single-peaked domain because we cannot order all the points on the tree G according to any linear order in which every agent has a single-peaked preference. Thus, we cannot straightforwardly apply our result obtained in Section 3.1 to this setting.

Let us introduce a well-known (group) strategy-proof mechanism, which is an generalization of the median mechanism on a real line. We refer to it as the *tree-median* mechanism.

MECHANISM 2 (TREE-MEDIAN). *A* tree-median *mechanism on a tree G has a fixed parameter (root) $\beta \in G$ and, for all n and for all reported location profiles, starts from β. As long as the current point has a subtree that contains at least $n/2$ locations, it smoothly moves down this subtree. When it reaches a point that does not have such a subtree, locates a facility at this point.*

As stated in Alon et al. [1], the tree-median mechanism achieves the optimal approximation ratio for social cost. However, obviously it is not false-name-proof, since it behaves in the same manner as the original median mechanism when all agents are on a single path.

We then characterize false-name-proof mechanisms on a tree. First, to simplify notations, let us define a *Pareto efficient set* with respect to a given location profile.

DEFINITION 5 (PARETO EFFICIENT SET). *For a tree G and for a location profile $x = (x_1, \ldots, x_n) \in G^n$, a set of points $PE(x) \subseteq G$ is said to be* Pareto-efficient *for x if $\forall y \in PE(x)$, $\forall y' \in G$, y' does not dominate y for x.*

Here, we say $y' \in G$ *dominates* $y \in G$ for a location profile x if $\forall i \in N, d(x_i, y) \leq d(x_i, y')$ and $\exists j \in N, d(x_j, y) < d(x_j, y')$. By using this notation, we define a class of mechanisms on a tree called *Pareto-improving relocation* rules.

MECHANISM 3 (PARETO-IMPROVING RELOCATION). *A mechanism f on a tree G is a* Pareto-improving relocation *rule if it has a fixed point $\beta \in G$ such that for all $n \in \mathbb{N}$ and for all reported location profiles $x = (x_1, \ldots, x_n) \in G^n$,*

$$f(x) = \arg \min_{z \in \mathrm{PE}(x)} d(z, \beta). \tag{6}$$

Now we show our characterization theorem; a class of false-name-proof, Pareto efficient, and anonymous mechanisms consists exactly of Pareto-improving relocation rules. It is shown separately in Lemmas 3 and 4.

THEOREM 8. *For any tree G, a mechanism f is false-name-proof, Pareto efficient, and anonymous if and only if it is a Pareto-improving relocation rule.*

LEMMA 3. *If a mechanism f for a tree G is false-name-proof, Pareto efficient, and anonymous, then it is a Pareto-improving relocation rule.*

PROOF. Consider a deterministic mechanism f that is false-name-proof, Pareto efficient, and anonymous. Since f is deterministic, it returns a point with probability 1 for a reported location profile. Now choose a location profile x^L that exactly contains every leaf $L(G)$ of the tree G. We then prove that f is a Pareto-improving relocation rule with a parameter $\beta = f(x^L)$. More precisely, we show that $f(x) = \arg\min_{z \in \mathrm{PE}(x)} d(z, \beta)$ holds for the above $\beta = f(x^L)$ and any location profile x.

First, we show that $f(x^L, x) = \beta$ holds for all x. Suppose not; there exists at least one location profile x such that $f(x^L, x) \neq \beta$. Here, let $f(x^L, x)$ indicate an outcome

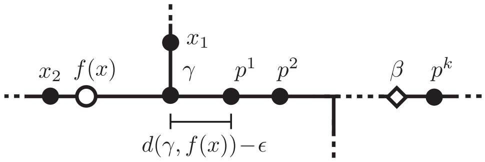

Figure 1: Sequence $p^1, p^2, \ldots, p^k$ in the proof of Lemma 3. Note that $x = (x_1, x_2)$. False-name-proofness implies $f(x, p_1, p_2, \ldots, p_k) = f(x)$. However, it contradicts Eq. (7).

location for an input location profile that is a joint of two location profiles, x^L and x. From Pareto efficiency, there exists at least one location x_i^L in the profile x^L such that $d(x_i^L, f(x^L, x)) < d(x_i^L, \beta)$. Thus, when x^L is a true location of the agents, agent i at x_i^L can strictly reduce her cost by reporting (x_i^L, x) under false identifiers. This contradicts the assumption that f is false-name-proof.

Next, let us show that

$$\forall x \text{ such that } \beta \in \mathrm{PE}(x), f(x) = \beta. \tag{7}$$

Suppose not; there exists at least one profile x such that $\beta \in \mathrm{PE}(x) \wedge f(x) \neq \beta$. From Pareto efficiency, there exists at least one location x_i in the profile x such that $d(x_i, \beta) < d(x_i, f(x))$. Thus, when x is a true location profile, agent i at x_i can strictly reduce her cost by reporting (x_i, x^L). This contradicts the assumption that f is false-name-proof.

Finally, let us show that

$$\forall x \text{ such that } \beta \notin \mathrm{PE}(x), f(x) = \gamma \tag{8}$$

where $\gamma = \arg\min_{z \in \mathrm{PE}(x)} d(z, \beta)$. Suppose not; there exists at least one profile x such that $\beta \notin \mathrm{PE}(x) \wedge f(x) \neq \gamma$. From Pareto efficiency, $f(x) \in \mathrm{PE}(x)$ holds. Then we can move from γ toward β with distance $d(\gamma, f(x)) - \epsilon$ and refer to the point as p^1. For this p^1, there exists at least one location x_i such that $\forall z \in [\gamma, p^1], d(x_i, z) < d(x_i, f(x))$. If $f(x, p^1) \in \mathrm{PE}(x) \setminus f(x)$, there exists at least one agent who strictly prefers $f(x, p^1)$ to $f(x)$. She can reduce her cost by adding a location p^1 under a false identifier. Also, if $f(x, p^1) \in (\gamma, p^1]$, agent i at x_i has an incentive to report (x_i, p^1). Note that $(\gamma, p^1]$ indicates a half-open interval. Thus, from false-name-proofness, $f(x, p^1) = f(x)$ must hold.

If $\beta \in (\gamma, p^1]$, the above equation contradicts Eq. (7). If $\beta \notin (\gamma, p^1]$, we can construct a finite sequence of points $p_1, p_2, \ldots, p_k$ by applying the same argument (see Fig. 1) and obtain $\beta \in (\gamma, p_k] \wedge f(x, p_1, \ldots, p_k) = f(x) \neq \beta$. This contradicts Eq. (7). $\square$

Lemma 4. *If a mechanism f for a tree G is a Pareto-improving relocation rule, then it is false-name-proof, Pareto efficient, and anonymous.*

Proof. Clearly, a Pareto-improving relocation rule is Pareto efficient and anonymous. To prove this lemma, it suffices to show that f is false-name-proof if it is a Pareto-improving relocation rule.

From the definition, a Pareto-improving relocation rule f has a fixed parameter $\beta \in G$. Let us fix a location profile x_{-i} and consider false-name manipulations by i. If $\beta \in \mathrm{PE}(x_{-i})$ holds, f always locates a facility at β regardless of i's strategy and satisfies false-name-proofness. Then we focus on showing that no false-name manipulation strictly reduces her cost when $\beta \notin \mathrm{PE}(x_{-i})$ holds.

Let us define a point $\gamma = \arg\min_{z \in \mathrm{PE}(x_{-i})} d(z, \beta)$. From the definition of a Pareto-improving relocation rule, we have $f(x_{\phi_i}, x_{-i}) \in [\gamma, \beta]$ for any x_{ϕ_i} and $f(x_i, x_{-i}) = \arg\min_{z \in [\gamma, x_i]} d(z, \beta)$ for any x_i. This means that $f(x_i, x_{-i})$ is the closest point to x_i in the range $[\gamma, \beta]$. Thus, we obtain $d(x_i, f(x_i, x_{-i})) \leq d(x_i, f(x_{\phi_i}, x_{-i}))$ for any x_i and x_{ϕ_i}. $\square$

Theorem 8 can be considered an extension of the result by Schummer and Vohra [9], which characterized the class of strategy-proof and Pareto efficient mechanisms on a tree. Now let us introduce the relationship between these two characterizations. We assume in this paper that mechanisms are anonymous, while [9] did not. With the assumption of anonymity, mechanisms characterized in [9] behave in the same manner as those described in Eq. (1), when all agents are on a single path. To achieve false-name-proofness when all agents are on a single path, each "partial" mechanism defined on each single path must be described in Eq. (2); for any two leaves $l_1, l_2 \in L(G)$, the path $[l_1, l_2]$ has a fixed parameter $\alpha_{l_1,l_2} \in [l_1, l_2]$. Here, as discussed in [9], these partial mechanisms must be self-consistent in some way; they must agree on the intersection of their paths. This consistency requires that each parameter of each partial mechanism must be defined as the closest point to a fixed $\beta \in G$ on the path, which is identical to Mechanism 3.

5. DISCUSSIONS

In the literature of social choice and mechanism design, several properties have been introduced. In this section, we clarify the connections between false-name-proofness and three other properties in facility location problems on a tree.

5.1 Population Monotonicity

Population monotonicity in public goods environments was originally identified in Ching and Thomson [3] Informally, population monotonicity requires that the arrival of a new agents affects all agents initially present in the same direction. However, with the assumption of Pareto efficiency, we can define the property in a more restricted way:

Definition 6 (Population Monotonicity). *A mechanism f is* population monotonic *if $\forall n \in \mathbb{N}$, $\forall x$, $\forall j \in N$, $\forall i \neq j$, $\mathrm{cost}(x_i, f(x)) \geq \mathrm{cost}(x_i, f(x_{-j}))$.*

This is somehow reminiscent of false-name-proofness. Both deal with the change of the number of agents. The following theorem shows the equivalence of these two properties.

Theorem 9. *Under the assumptions of Pareto efficiency and anonymity, a mechanism f is population monotonic if and only if it is false-name-proof.*

Proof. Ching and Thomson [3] gave a characterization of population monotonic mechanisms under the assumption of Pareto efficiency. Their characterization is identical to our characterization of false-name-proofness (Theorem 8). $\square$

Note that without the assumption of Pareto efficiency, false-name-proofness and population monotonicity do not coincide even in the case of a real line. Consider the following mechanism for a real line: if $n < 3$, then locate a facility at the point that is slightly smaller than the leftmost reported location, otherwise use the leftmost mechanism. This mechanism is population monotonic, although it is not false-name-proof (not even strategy-proof).

5.2 Group-strategyproofness

Group-strategyproofness has been widely discussed in economics. A mechanism is *group-strategyproof* if for any location profile and any coalition of agents, there is no joint deviation of the coalition such that every agent in the coalition strictly reduces her cost. For the connection to false-name-proofness, we show the next theorem. For space reasons, we omit the proof.

THEOREM 10. *Under the assumptions of Pareto efficiency and anonymity, any false-name-proof mechanism f is group-strategyproof.*

It has been known that the tree-median mechanism is group-strategyproof. However, as stated in Section 4.2, it is not false-name-proof. In other words, the class of false-name-proof mechanisms is a strict subset of the class of group-strategyproof mechanisms under the assumptions of Pareto efficiency and anonymity.

5.3 Anonymity-proofness

Anonymity-proofness, which was first proposed by Conitzer [4], is an extension of false-name-proofness. First, to define anonymity-proofness, we introduce the notion of *participation*. A mechanism f satisfies participation if $\forall n \in \mathbb{N}$, $\forall i \in N$, $\forall x_{-i}$, $\forall x_i$, $\text{cost}(x_i, f(x_i, x_{-i})) \leq \text{cost}(x_i, f(x_{-i}))$. That is, for each agent, it never hurts her to join the mechanism as long as she behaves sincerely.

DEFINITION 7 (ANONYMITY-PROOFNESS). *A mechanism f is* anonymity-proof *if it is false-name-proof and satisfies participation.*

The next theorem shows the equivalence of anonymity-proofness and false-name-proofness.

THEOREM 11. *Under the assumptions of Pareto efficiency and anonymity, a mechanism f is anonymity-proof if and only if f is false-name-proof.*

PROOF. From the definition of anonymity-proofness, f is false-name-proof if it is anonymity-proof. To prove this theorem, it suffices to show that f satisfies participation if it is false-name-proof. From Theorem 8, a false-name-proof f is a Pareto-improving relocation rule; it has a parameter β. Now let us fix the location profile x_{-i} reported by agents except i and focus on i's strategy.

Let us define $\gamma = \arg\min_{z \in \text{PE}(x_{-i})} d(z, \beta)$. Note that $f(x_{-i}) = \gamma$ holds from the definition of the Pareto-improving relocation rule. If $\text{PE}(x_i, x_{-i}) \cap (\gamma, \beta] = \emptyset$ holds, f always locates a facility at γ regardless whether agent i participates. Thus, f satisfies participation.

If $\text{PE}(x_i, x_{-i}) \cap (\gamma, \beta] \neq \emptyset$ holds, we can find a point γ' such that $\gamma' = \arg\min_{z \in \text{PE}(x_i, x_{-i})} d(z, \beta) \wedge \gamma' \in (\gamma, x_i]$. Here $f(x_i, x_{-i}) = \gamma'$ holds from the definition of Pareto-improving relocation rule. Thus, we obtain $\text{cost}(x_i, f(x_{-i})) = d(x_i, \gamma) > d(x_i, \gamma') > \text{cost}(x_i, f(x_i, x_{-i}))$, and f satisfies participation. □

6. CONCLUSIONS AND FUTURE WORKS

In this paper, we presented a case study of false-name-proof mechanism design without money by dealing with facility location problems. We first characterized deterministic false-name-proof mechanisms on a real line and established the tight bounds of approximation ratios. We then discussed the approximation ratios achieved by randomized false-name-proof mechanisms. Also, we characterized deterministic false-name-proof mechanisms on a tree. Furthermore, we clarified the connections between false-name-proofness and other related properties.

We outline our future direction of false-name-proof mechanism design without money. To the best of our knowledge, there exists no work that discussed the effect of false-name manipulations in private goods environments without monetary transfers, e.g., resource allocations [5]. Intuitively, false-name manipulations must become much more powerful strategic behaviors in such environments. We would like to find a solution to prevent false-name manipulations, evaluate it using techniques of approximate mechanism design without money, and characterize the solutions.

7. ACKNOWLEDGMENTS

This research was partially supported by the Ministry of Education, Science, Sports and Culture, Grant-in-Aid for Scientific Research (A), 20240015, 2008, and Grant-in-Aid for JSPS research fellows. The authors would like to thank anonymous reviewers of AAMAS'11 and attendees of WINE'10.

8. REFERENCES

[1] N. Alon, M. Feldman, A. D. Procaccia, and M. Tennenholtz. Strategyproof approximation of the minimax on networks. *Mathematics of Operations Research*, 35(3):513–526, 2010.

[2] Y. Bachrach and E. Elkind. Divide and conquer: false-name manipulations in weighted voting games. In *AAMAS'08*.

[3] S. Ching and W. Thomson. Population-monotonic solutions in public good economies with single-peaked preferences. *Social Choice and Welfare*, forthcoming.

[4] V. Conitzer. Anonymity-proof voting rules. In *WINE'08*.

[5] M. Guo and V. Conitzer. Strategy-proof allocation of multiple items between two agents without payments or priors. In *AAMAS'10*.

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

[7] B. Peleg. Game-theoretic analysis of voting in committees. In K. J. Arrow, A. K. Sen, and K. Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 1, chapter 8, pages 395–423, 2002.

[8] A. D. Procaccia and M. Tennenholtz. Approximate mechanism design without money. In *ACM-EC'09*.

[9] J. Schummer and R. V. Vohra. Strategy-proof location on a network. *Journal of Economic Theory*, 104(2):405–428, 2004.

[10] J. Schummer and R. V. Vohra. Mechanism design without money. In N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, chapter 10, 2007.

[11] T. Todo, A. Iwasaki, M. Yokoo, and Y. Sakurai. Characterizing false-name-proof allocation rules in combinatorial auctions. In *AAMAS'09*.

[12] M. Yokoo, Y. Sakurai, and S. Matsubara. The effect of false-name bids in combinatorial auctions: New fraud in internet auctions. *Games and Economic Behavior*, 46(1):174–188, 2004.