On Voting and Facility Location

MICHAL FELDMAN, Tel Aviv University
AMOS FIAT, Tel Aviv University
IDDAN GOLOMB, Tel Aviv University

"We would all like to vote for the best man, but he is never a candidate"

— Kin Hubbard

We study mechanisms for candidate selection that seek to minimize the social cost, where voters and candidates are associated with points in some underlying metric space. The social cost of a candidate is the sum of its distances to each voter. Some of our work assumes that these points can be modeled on the real line, but other results of ours are more general.

A question closely related to candidate selection is that of minimizing the sum of distances for facility location. The difference is that in our setting there is a fixed set of candidates, whereas the large body of work on facility location considers every point in the metric space to be a possible candidate. This setting gives rise to three types of candidate selection mechanisms which differ in the granularity of their input space (single candidate, ranking and location mechanisms). We study the relationships between these three classes of mechanisms.

While it may seem that Black's 1948 median algorithm is optimal for candidate selection on the line, this is not the case. We give matching upper and lower bounds for a variety of settings. In particular, when candidates and voters are on the line, our universally truthful *spike* mechanism gives a [tight] approximation of two. When assessing candidate selection mechanisms, we seek several desirable properties: (a) efficiency (minimizing the social cost) (b) truthfulness (dominant strategy incentive compatibility) and (c) simplicity (a smaller input space). We quantify the effect that truthfulness and simplicity impose on the efficiency.

CCS Concepts: •**Theory of computation**→ **Algorithmic game theory and mechanism design; Algorithmic mechanism design;** *Facility location and clustering;*•**Computing methodologies**→ Multi-agent systems;•**Applied computing**→ Economics; Voting / election technologies;

Additional Key Words and Phrases: Voting, Facility location, Algorithmic Mechanism Design, Social Choice, Approximate Mechanism Design Without Money

1. INTRODUCTION

The Hotelling-Downs model ([Downs 1957], [Hotelling 1990]) used to study political strategies, assumes that individual voters occupy some point along the real line. Non-principled political parties (or ice cream vendors) strategically position themselves at a point along the left-right axis (or along a beach) so as to garner the greatest number of supporters (clients). Implicitly, voters will vote for the closest candidate.

We consider an analogous problem to the Hotelling-Downs model, where candidates are principled (*i.e.*, non-strategic) whereas the voters have preferences but may mis-

The work of M. Feldman and I. Golomb was partially supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement number 337122. This work was done in part while the authors visited the Simons Institute for the Theory of Computing.
Authors addresses: M. Feldman, A. Fiat and I. Golomb Computer Science Department, Tel Aviv University; email: michal.feldman@cs.tau.ac.il;fiat@tau.ac.il;igolomb@gmail.com.

EC'16, July 24–28, 2016, Maastricht, The Netherlands. Copyright © 2016 ACM 978-1-4503-3936-0/16/07 ...$15.00.

http://dx.doi.org/10.1145/2940716.2940725

represent them in order to achieve what is a better outcome from their perspective. In this model, in which both voters and candidates are represented by points in the metric space, a closer candidate is preferable to one further away.

Examples for candidate selection:

— A municipality plans to erect a public library on a street, and every resident seeks to be as close as possible to the proposed library. However, the new library can only be built on suitable locations (the candidates).
— Social choice issues in which the distance is not physical: there is a set of policies ranging from left to right, and several political candidates stand for election, each one advocating a different policy. Every voter is associated with a point along the real line. An example of a collective decision problem which does not revolve around the political sphere yet may also fit this setting is the task of determining the temperature of an air conditioner in a room, where each individual has a different ideal point along the scale of temperatures (a line). There are many additional settings of relevant candidate selection problems, e.g., in the realms of recommendation systems and computational economics. While our results do not necessarily apply to all social choice settings, there are many such problems for which they do apply (whether in entirety or partially).

Assuming quasi-linear utilities, and allowing payments — then the well known Vickrey-Clarke-Groves (VCG) mechanism is truthful and can achieve the optimal social cost (see, e.g., [Nisan 2007]). However, in many real-life situations we restrict the use of money due to ethical, legal or other considerations, e.g, in democratic elections and in examples previously mentioned.

We study deterministic truthful mechanisms with no payments for the candidate selection problem. In such mechanisms, no agent can benefit from misreporting her location, regardless of the reported locations of the other agents. Such mechanisms are also known as dominant strategy incentive compatible mechanisms. We also consider randomized truthful mechanisms, both universally truthful (ex-post Nash) and truthful in expectation.

Given a set of candidate and voter locations, it is polytime to find the candidate that minimizes the social cost.

When restricted to deterministic truthful mechanisms, we show that the optimal candidate cannot be selected in the general case. Moreover, we show that the cost may be as bad as three times the optimal cost (matching lower and upper bounds). When considering randomized mechanisms on the line, the approximation factor drops to two (matching lower and upper bounds).

There are other reasons that an optimal candidate may not be chosen. In particular, this depends on the amount of information the agents supply to the mechanism. We formulate three different types of mechanisms, based on the information each agent submits to the mechanism. We remark that all three mechanism types are *candidate selection mechanisms*, that is – their output is a single candidate.

— Single Candidate [vote] mechanisms, in which each agent casts a vote for one of the candidates.
— Ranking [vote] mechanisms, in which each agent states ordinal preferences over the candidates (a permutation).
— Location [vote] mechanisms, in which each agent sends a position.

Clearly, knowing the true location of an agent allows one to infer the ranking preferred by that agent, which in turn allows one to infer the favorite candidate of the agent (up to tie-breaking).

In almost all previous work on the facility location problem every point in the metric space was considered to be a candidate. Therefore there was no difference between these three mechanism types.

The social choice literature mostly considers social choice functions (which are ranking mechanisms that are not necessarily truthful). Note that Arrow's impossibility theorem does not hold when assuming the preferences are single-peaked.

The more information an agent transmits, the more tools the mechanism has to devise an accurate solution. Albeit, this information comes at a cost, since it might disclose more private information which the agents wish to keep confidential. Furthermore, behavioral economists have long argued that the agents cannot fully acquire their utility function, or that obtaining this information requires a high cognitive cost. Additionally, sending more information also casts a higher burden on the mechanism itself. For all of these reasons deploying a simple mechanism[1] which requires less information from agents is beneficial. Indeed, in many practical scenarios, single candidate mechanisms are used rather than ranking mechanisms. Generally there is a trade-off between the accuracy of a mechanism and its simplicity.

1.1. Our Contributions

In the paper, we show the following:

— In Section 3 we formulate a framework of reductions that compare the various mechanism types. We utilize this framework to show the relations (equivalence or strict containment) between the three classes of *truthful* mechanisms – single candidate, ranking and location (see Figure 1). Furthermore, we show that for the case of two candidates, the set of truthful in expectation location mechanisms is equivalent to the set of truthful in expectation single candidate mechanisms. These results provide a significant step towards a full characterization of truthful mechanisms.

— In Section 4 we define a family of universally truthful single candidate mechanisms on the line called weighted percentile single candidate (WPSC) mechanisms, which choose the ith vote with some predetermined probability p_i. We introduce the *spike* mechanism, which is a WPSC mechanism that carefully crafts the probability distribution to account for misreports by any agent - whether they are near the center or close to the extremes (see Figure 4). We then use backwards induction to show that spike achieves an approximation ratio of two (Theorem 4.6).

— In Section 5 we show additional bounds for randomized mechanisms – On the line there is a lower bound of two, even for location mechanisms, which shows that the result for spike is tight. Furthermore, when combining this understanding with the results of Section 3, it can be concluded that two is also the tight approximation ratio for truthful in expectation mechanisms (single candidate, ranking or location) and for universally truthful single candidate mechanisms.

We move on to show bounds for randomized mechanisms for more general metric spaces[2] (see Figure 5). An easy observation is that the random dictator mechanism achieves an upper bound of three for any metric space. Theorem 5.3 shows a lower bound of $3 - \frac{2}{d}$ for any single candidate mechanism in $\Re^d$ (by using a counterexample based on a regular simplex). This is enough to conclude that on an arbitrary metric space, the bound of three is tight for single candidate mechanisms. Theorem 5.7 displays a lower bound of $7/3$ for any ranking mechanism in $\Re^2$ (which also holds in any higher dimension Euclidean space $\Re^d$).

[1]We use the term "simplicity" from the perspective of the voters, who have a smaller action space, i.e, less options to choose from. The mechanism itself can act in an arbitrarily complex fashion.

[2]We do not present results for deterministic mechanisms in general metric spaces, since in these cases the incentive compatibility constraints take a significant toll on the approximation ratio – according to

Fig. 1. Relationships between classes of mechanisms(Theorem 3.1): For deterministic truthful mechanisms, the set of ranking mechanisms strictly contains the set of single candidate (SC) mechanisms, yet the set of location mechanisms is equivalent to the set of ranking mechanisms.
In the randomized case, there is a hierarchy of strict containment in the following order - truthful in expectation (TIE) location mechanisms, TIE ranking mechanisms, TIE single candidate (SC) mechanisms and universally truthful (UT) single candidate (SC) mechanisms. Refer to Section 3 for formal definitions of equivalence and strict containment in our setting.

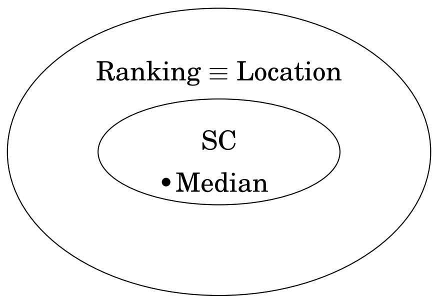

Fig. 2. **Deterministic truthful** mechanisms

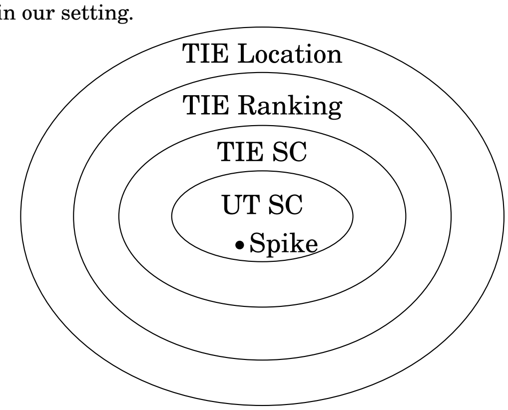

Fig. 3. **Randomized** truthful mechanisms

Fig. 4. The density function of the spike mechanism, which gives rise to the mechanism's name (the cumulative distribution function is given explicitly in Definition 4.2). In this example there are 10000 agents and 4 candidates. The candidates, when ordered from left to right, receive 2000, 2000, 3000 and 3000 votes respectively. The graph depicts probability of choosing each vote – votes are chosen with higher probability when they are closer to the 50th percentile. The area beneath the graph represents the probability that each candidate will be elected, e.g., the probability of choosing the second candidate (p_2) is the integral of the function between 2000 and 4000.

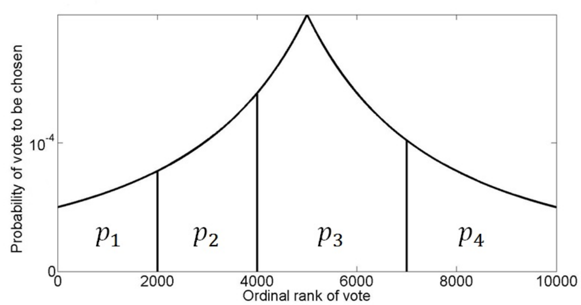

— In Section 6 we present deterministic bounds on the line – a lower bound of three is met by a matching upper bound due to the median mechanism. All the results on the line, deterministic or randomized, are displayed in the table in Figure 6.

[Anshelevich et al. 2014] in the non-strategic setting it is possible to reach a constant ratio in any metric space, while due to the characterization of [Schummer and Vohra 2002] there exist metric spaces (e.g., cycles) in which the approximation ratio is $\Omega(n)$ even in the continuous model.

Fig. 5. Summary of our results for randomized mechanisms in $\Re^d$. The columns correspond to the truthfulness constraints, while the rows show the information constraints and are further divided into lower and upper bounds. Note that for non-strategic location mechanisms the result is always optimal by definition, since there are neither information nor strategic constraints.
Most of the results here are rather straightforward, except for the upper bounds of $3 - \frac{2}{d+1}$ and of $7/3$, which are more involved.

		Strategic		Non-Strategic	
Single Candidate (low information)	LB	$3 - \frac{2}{d+1}$	(Thm. 5.3)	2	(Lemma 5.6)
	UB	3	(Lemma 5.8)	3	(Lemma 5.8)
Ranking	LB	7/3	(Thm. 5.7)	2	(Lemma 5.6)
	UB	3	(Lemma 5.8)	3	(Lemma 5.8)
Location (high information)	LB	2	(Obs. 5.4)	1	
	UB	3	(Lemma 5.8)	1	

Fig. 6. The approximation ratios of mechanisms on the line ($\Re$) in various settings. All the results in the table are tight. In particular, the randomized upper bound of two in the strategic case holds due to the spike mechanism.

	Deterministic		Randomized	
	Strategic	Non-Strategic	Strategic	Non-Strategic
Single Candidate	3^*	$3^\dagger$	$2^\ddagger$	$2^\S$
Ranking	3^*	$3^\dagger$	$2^\ddagger$	2^{**}
Location	3^*	1	$2^\ddagger$	1

*LB and UB: Thm. 6.1
†LB: Thm. 3 from [Anshelevich et al. 2014], UB: Thm. 6.1
‡LB: Obs. 5.4, UB: Thm.4.6
§LB: Lm. 5.6, UB: Thm.4.6
**Lm. 5.6, UB: Thm. 6.1

We highlight the following surprising phenomenon apparent in Figure 6. In both deterministic and randomized cases, imposing any constraint in either information or truthfulness, yields the same ratio as taking the both of these constraints simultaneously — when insisting on truthful mechanisms (in the strategic case), there is no trade-off between high and low information settings, and one can enjoy the benefits of minimal information mechanisms (single candidate mechanisms) without incurring any additional cost to the approximation ratio. Similarly, when deciding to reduce the information requirements to anything less than location mechanisms, it is possible to devise a truthful [single candidate] mechanism, without increasing the approximation ratio.

1.2. Related Work

[Procaccia and Tennenholtz 2009] introduced game theoretical aspects to the facility location problem. As mentioned before, their setting is similar to the one in this paper, except that the location of the facility is not restricted to a set of candidates, but instead can be located at any point on the line. This model was extended by these authors and by others in many different ways. The metric space was extended from the line to cycles ([Alon et al. 2009], [Alon et al. 2010]), trees ([Alon et al. 2009], [Feldman and

Wilf 2013]) and general graphs ([Alon et al. 2009]). Many papers consider building several facilities (or electing a committee of candidates), e.g. [Fotakis and Tzamos 2013], [Fotakis and Tzamos 2014], [Lu et al. 2010]. The goal of the majority of these papers is to optimize over some global target function. The most popular target functions are the social cost and the maximal cost of an agent, but additional target functions like the L_2 norm were also considered ([Feldman and Wilf 2013]). A different approach to minimize the social cost for facility location while maintaining truthfulness is [selective] verification of agents' reports [Fotakis et al. 2015]. Some papers deal with "obnoxious facility location" — a setting in which agents want to be as far away as possible from the facility, e.g., when selecting a location for a central garbage dump (e.g., [Cheng et al. 2013]).

[Sui et al. 2013] introduced deterministic percentile [location] mechanisms for locating multiple facilities in $\Re^d$. [Sui and Boutilier 2015] showed that even when the location of the facilities are constrained to a set of candidates (as in our paper), these percentile mechanisms are group-strategyproof on $\Re$.

[Dokow et al. 2012] characterize deterministic truthful mechanisms for locating a facility on the discrete line and the discrete cycle. In their model agents must be located precisely on candidates, and the distance between neighbor candidates is constant.

When constraining the outcome to a set of candidates, the facility location setting resembles social choice problems. The seminal Gibbard-Satterthwaite theorem ([Gibbard 1973], [Satterthwaite 1975]) shows that in a general setting the only onto truthful deterministic mechanisms are dictatorships. However, if there are limitations on the rankings, then the impossibility theorem of Gibbard-Satterthwaite does not hold. In many cases the rankings can be limited to single-peaked preferences, a notion used as early as 1948 by [Black 1948], and was later fully characterized by [Moulin 1980] and [Schummer and Vohra 2002].

There has been extensive work describing numerous candidate selection schemes (e.g, [Brandt et al. 2012]). These schemes typically have no assumptions on the preferences of the agents, and according to Gibbard-Sattethwaite they are not truthful. [Gibbard 1973] further characterizes truthful randomized mechanisms under arbitrary preferences. Some work on social choice makes use of randomized schemes in order to elicit truthfulness (e.g., [Bogomolnaia and Moulin 2001] and [Aziz et al. 2014]).

The advantages of simple mechanisms (in which each voter has less options to choose from) have been widely acknowledged. For instance, aiming for simplicity is a major reason for which the vast majority of the candidate selection schemes above accept ordinal rankings rather than cardinal rankings as input. Truthfulness, privacy and low communication are also very common and desired traits of mechanisms at large. There have been several works describing the tradeoffs between these various desired properties in mechanism design and in facility location specifically: [Feigenbaum et al. 2010] proposed a framework to analyze the tradeoff between privacy and communication, and showed examples for this using mechanisms for second price auctions. [Sui and Boutilier 2011] show the tradeoffs between efficiency and privacy in mechanism design, and analyze this for facility location problems. Nonetheless, we do not know of any work formally describing the three types of mechanisms (single candidate, ranking and location), or any similar framework for uncovering relationships between these mechanism types, as in Theorem 3.1.

Since in the lack of cardinal costs no global objective functions can be measured (e.g, the social cost), the focus of many of the aforementioned schemes is on achieving some desirable axiomatic properties. However, the use of utilitarianism in the realm of social choice has firm and ancient roots (e.g, [Fleming 1952] and [Harsanyi 1955]).

Moreover, in recent years a line of work in computational social choice regarded *distortion*, a measure for assessing social choice functions (i.e, ranking mechanisms) which also refers to the utilitarian goal of minimizing social cost. The term distortion was coined in [Procaccia and Rosenschein 2006] and was followed by several other papers (e.g., [Caragiannis and Procaccia 2011], [Boutilier et al. 2012], [Anshelevich et al. 2014] and [Anshelevich and Postl 2015]). Roughly speaking, the distortion is the worst case ratio between the social cost of the candidate elected and the social cost of the optimal candidate. Note that while the distortion stems from an information deficiency (access only to ordinal rankings of the agents), the approximation ratio in this paper is greater than one *both* because of this information deficiency (for ranking and single candidate mechanisms), *and* because of incentive compatibility constraints. Computing the approximation ratio and the distortion can quantify the affect of these two deficiencies in various settings.

[Caragiannis and Procaccia 2011] deal with a setting in which the utility functions of agents are more general than ours, which leads to a higher deterministic lower bound on the distortion.

[Anshelevich et al. 2014] provide a deterministic lower bound of 3 on the distortion in a general metric space, and show that both Copeland and uncovered set reach a distortion of 5.

Spike is a truthful mechanism which achieves an approximation ratio of 2 on the line. Independently from us, [Anshelevich and Postl 2015] recently showed a mechanism which also achieves a distortion of 2 on $\Re$, albeit it is not truthful. The results in our paper for non-strategic agents (see Figures 5 and 6) also hold in their model.

2. MODEL

Let $N = \{1, \dots n\}$ be a set of agents, where each agent $i \in N$ is located at some point x_i. We refer to the location of agent i as agent i's *type*. Let $\mathbf{x} = (x_1, \dots, x_n)$ be the *location profile* of the agents. There exist m candidates located at publicly known points $\mathbf{y} = (y_1, \dots, y_m)$ (we refer to y_i as the ith candidate and as the location of the ith candidate interchangeably). The agents and candidates are located on some metric space. Significant parts of the paper deal with specific metric spaces, and these parts will be noted. When the metric space is $\Re$, we assume that the agents and candidates are both numbered in ascending order based on their locations (since they could be renamed).

A *deterministic mechanism* M, is a function which maps an *action profile* $\mathbf{a} = (a_1, \dots, a_n) \in \mathcal{A}^n$ to a candidate, that is: $M : \mathcal{A}^n \to \mathbf{y}$. We consider three classes of mechanisms that differ in their input, *i.e.*, in the action space $\mathcal{A}$ of the agents:

— *Single candidate mechanisms*, in which each agent casts a vote for a candidate, that is: $a_i \in \mathbf{y}$.
— *Ranking mechanisms*, in which every agent reports ordinal preferences over all the m candidates. The notation $y_j \succeq y_k$ indicates a preference of candidate y_j over candidate y_k (or indifference between the two). In ranking mechanisms $a_i \in \Pi$, where Π is the set of all permutations of the set of candidates $\mathbf{y}$. These mechanisms are sometimes referred to in the literature as social choice functions.
— *Location mechanisms*, in which every agent reports a location, that is a_i is some point in the metric space.

Given a joint action profile $\mathbf{a}$, the *cost* of point x is its distance to the facility, that is: $\text{cost}_x(M, \mathbf{a}) = |x - M(\mathbf{a})|$. For agent $i \in N$ located at point x_i, we refer to $\text{cost}_{x_i}(M, \mathbf{a})$ as the *cost of agent* i. The goal of each agent is to minimize their cost.

Truthful mechanisms are usually defined in the context of direct revelation mechanisms. Since in ranking and single candidate mechanisms the action space does not coincide with the type space, we extend this notion in the following trivial manner. For an agent in location x_i and for any mechanism (location, ranking or single candidate), let $\mathcal{A}(x_i)$ be the set of *true actions* of this agent — the actions which convey the real preferences of this agent. For instance, in single candidate mechanisms $\mathcal{A}(x_i)$ is the set of candidates closest to x_i, which we refer to as the *favorite candidates* of x_i (this might be a set since there may be ties). An agent reporting $a_i \in \mathcal{A}(x_i)$ is said to be reporting *truthfully*, and an action profile a in which all agents report truthfully is called a *truthful profile*. The set of truthful profiles is denoted $\mathcal{A}(\mathbf{x})$. A *truthful mechanism* M is one in which no agent can suffer from reporting truthfully, regardless of the actions of the other agents:

$$\forall i \in N, \forall x_i, \forall a_i \in \mathcal{A}(x_i), \forall \mathbf{a}_{-i} \in \mathcal{A}^{n-1}, \forall a'_i \in \mathcal{A} : \mathrm{cost}_{x_i}\left(M, (a_i, \mathbf{a}_{-i})\right) \le \mathrm{cost}_{x_i}\left(M, (a'_i, \mathbf{a}_{-i})\right)$$

A *randomized mechanism* is a mapping from an action profile to a distribution over the candidates, that is: $M : \mathcal{A}^n \to \Delta(\mathbf{y})$. The cost of agent i is the expected cost of this agent according to the probability distribution returned by the mechanism, that is: $\mathrm{cost}_{x_i}(M, \mathbf{a}) = \mathbb{E}_{y_j \sim M(\mathbf{a})}|x_i - y_j|$.

Two different notions of randomized truthful mechanisms have been studied in the literature, and we extend them naturally based on our definitions of truthful reports:

— *Truthful in expectation (TIE)* mechanisms — where the expected cost of an agent reporting truthfully is never higher than any other action. That is: $\forall i \in N$, $\forall a_i \in \mathcal{A}(x_i)$, $\forall \mathbf{a}_{-i} \in \mathcal{A}^{n-1}$, $\forall a'_i \in \mathcal{A}$: $\mathrm{cost}_{x_i}\left(M, (a_i, \mathbf{a}_{-i})\right) \le \mathrm{cost}_{x_i}\left(M, (a'_i, \mathbf{a}_{-i})\right)$. In these mechanisms the agent may regret her action ex-post for some of the instances.
— *Universally truthful* mechanisms are mechanisms which can be expressed as a probability distribution over deterministic truthful mechanisms. In these mechanisms an agent never regrets reporting truthfully, even after the random outcome is unraveled.

Clearly, every universally truthful mechanism is truthful in expectation, but not necessarily vice versa. Throughout the paper, in the randomized setting we use the term "truthful" to refer to truthful in expectation mechanisms, unless otherwise stated.

The *social cost* of a mechanism is the sum of the agents' costs. For a location profile x and an action profile a the social cost is: $\mathrm{SC}(M, \mathbf{x}, \mathbf{a}) = \sum_i \mathrm{cost}_{x_i}(M, \mathbf{a})$. The cost of a candidate is the cost of the mechanism which locates the facility on that candidate, that is: $\mathrm{SC}(y_j, \mathbf{x}) = \sum_{i \in N} |y_j - x_i|$. Given a location profile x, the *optimal mechanism*, denoted $\mathrm{OPT}(\mathbf{x})$, is one which chooses a candidate that minimizes the social cost (y_{opt}). When there are when there are several optimal candidates, we break ties consistently (e.g., when the metric space is $\Re$, we refer to the leftmost among them as y_{opt}). For any truthful in expectation mechanism M (including universally truthful mechanisms), the *social cost of M given a location profile* x is the maximal social cost it yields by any truthful action profile a, that is: $\mathrm{SC}(M, \mathbf{x}) = \max_{\mathbf{a} \in \mathcal{A}(\mathbf{x})} \mathrm{SC}(M, \mathbf{x}, \mathbf{a})$. The *approximation ratio* of a truthful in expectation mechanism M is the maximal ratio for any location profile x, between social cost of M given x and the optimal social cost given x: $\max_{\mathbf{x}} \frac{\mathrm{SC}(M,\mathbf{x})}{\mathrm{SC}(\mathrm{OPT},\mathbf{x})}$.

For single candidate mechanisms when the metric space is $\Re$:

— Let τ be a permutation on indices $1, \dots, n$ such that

$$a_{\tau(1)} \le a_{\tau(2)} \le \cdots \le a_{\tau(n)}$$

Note that there are many permutations satisfying the above, each of which represents a different version of breaking ties amongst votes for the same candidate. τ is an arbitrary such permutation. Let $z_j = a_{\tau(j)}$ for $j = 1, \dots, n$. *I.e.*, z_j is the [location of the] reported ideal candidate for voter $\tau(j)$.
— A *percentile mechanism* is a mechanism specified by an index $1 \le j \le n$, which chooses candidate z_j.

Fig. 7. Illustration of candidates (white circles), agents (black circles), candidate borders and candidate zones when the metric space is $\Re$.
For example, in this case the favorite candidate of x_3 is y_2: $y(x_3) = \{y_2\}$. The candidate borders divide the distance between two consecutive candidates exactly in half, for example: $|b_1 - y_1| = |y_2 - b_1|$.

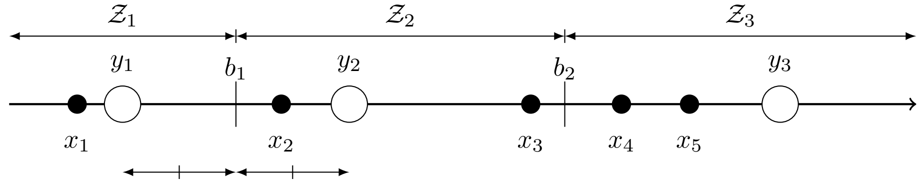

— A *weighted percentile single candidate (WPSC)* mechanism M is specified by a vector of probabilities $p_1, \ldots, p_n$, such that $\sum_j p_j = 1$. M chooses y_i with probability $\sum_{j:z_j=y_i} p_j$.
This can be interpreted as follows: a mechanism is WPSC if and only if there exists some τ as described above such that for every profile a, voter $\tau(j)$ determines the winning candidate with probability p_j.

In single candidate mechanisms, the set of candidates y induces a partition of the metric space in the following manner — the *candidate zone* of candidate y_i, denoted $\mathcal{Z}_i$, is the set of points whose favorite candidate is y_i: $\mathcal{Z}_i = \{x : \forall y_j : |x - y_i| \leq |x - y_j|\}$. The candidate zones are bounded by *candidate borders*. For example, when the metric space is $\Re$, there are $n - 1$ borders, which are the midpoints between two consecutive candidates: $b_i = \frac{y_i + y_{i+1}}{2}$ (see Figure 7). In the $\Re^d$ metric space, the candidate zones form a Voronoi diagram. A candidate who receives at least one vote is called *active*.

In ranking mechanisms, y induces a partition which divides the metric space into *ranking zones*. All points in some ranking zone $\mathcal{R}_i$ share some ranking π_i. In this case, we say that the ranking π_i is consistent with ranking zone $\mathcal{R}_i$. The ranking zones bounded by *ranking borders*. For example, when the metric space is $\Re$ the ranking borders are the midpoints between any two candidates: $b_{i,j} = \frac{y_i + y_j}{2}$.

3. CLASSES OF MECHANISMS

In this section we go over the containment hierarchy of various classes of truthful mechanisms (e.g., Figure 1). We start with some intuition, then define some necessary terms, and finally present the main theorem of this section.

Intuitively, for any mechanism M, there exists a mechanism M' which receives a "richer" input than M, and acts identically to M. For instance, for some arbitrary single candidate mechanism M, there obviously exists a ranking mechanism M' which disregards all of the preferences except the top choice of each agent, and behaves essentially just like M does.

We generalize this notion in the following informal definition — a mechanism M (whether location/ranking/or single candidate) is said to be *reducible* to a mechanism M' (location/ranking/or single candidate) if for every location profile x and true reports, the output of M is identical to the output of M' (a formal definition, which is based on M simulating M', is deferred to the full version.

As pointed out, it is clear that every single candidate mechanism M is reducible to some ranking mechanism M' (or some location mechanism M'). In these cases, if M is truthful then so is M', since M' only uses the information which is inputted to M, so any misreports to M' which would not change the input of M do not affect the outcome at all. Note that the same reasoning also shows that every ranking mechanism is reducible to some location mechanism, and that any single candidate mechanism is reducible to some location mechanism.

On the other hand, it is not true that every location mechanism is reducible to some ranking mechanism (as location mechanisms have a "richer" input space). Somewhat surprisingly, we will soon show that when we restrict ourselves to deterministic truthful mechanisms this does hold, that is — every *deterministic truthful* location mechanism is reducible to some deterministic truthful ranking mechanism.

Two sets of mechanisms, A and B, are said to be *equivalent* if every $a \in A$ is reducible to some $b \in B$, and every $b \in B$ is reducible to some $a \in A$.

A set of mechanisms A is said to be *strictly contained* in a set of mechanisms B if every mechanism $a \in A$ is reducible to some mechanism $b \in B$, yet not every mechanism $b \in B$ is reducible to some mechanism $a \in A$. This is a slight abuse of terminology since the sets A and B may be disjoint, as their input space may be different.

The following theorem shows several claims regarding relations (equivalence or strict containment) between sets of truthful mechanisms. Notice that not only does this theorem show the hierarchy of the different classes, but it also provides notions relevant to a full characterization of truthful mechanisms. For instance, the second claim proves that no mechanism can use any information regarding the location of the agents beyond their ranking, while maintaining truthfulness. In addition, in the claims showing strict containment, the examples in the proofs portray the expressiveness that the additional information gives the mechanism.

THEOREM 3.1. *The following claims hold in the Euclidean metric space $\Re^d$ (for any $d \in \mathbb{N}$):*

(*1*) *The set of truthful deterministic ranking mechanisms strictly contains the set of truthful deterministic single candidate mechanisms.*
(*2*) *The set of truthful deterministic location mechanisms is equivalent to the set of truthful deterministic ranking mechanisms.*
(*3*) *The set of truthful in expectation randomized ranking mechanisms strictly contains the set of truthful in expectation randomized single candidate mechanisms.*
(*4*) *The set of truthful in expectation randomized location mechanisms strictly contains the set of truthful in expectation randomized ranking mechanisms.*
(*5*) *The set of truthful in expectation randomized single candidate mechanisms strictly contains the set of universally truthful randomized single candidate mechanisms.*
(*6*) *When there are two candidates, the set of truthful in expectation randomized location mechanisms is equivalent to the set of truthful in expectation randomized single candidate mechanisms.*

For the ease of readability, we defer the proof of this theorem to the full version.

4. THE SPIKE MECHANISM

In the upcoming sections we will prove that both the median mechanism and the random dictator mechanism achieve an approximation ratio of three on $\Re$. However, the source for this ratio in these two cases is different - for median it is due to an instance which is costly for the median agent, while for random dictator it is due to a bad instance for an agent in one of the extremes. The spike mechanism was devised with the objective of being resistant to costly instances of *any* agent.

This section contains foundations needed for the introduction of the spike mechanism, the definition of spike, and the theorem showing that spike achieves an approximation ratio of 2. The reductions in Section 3 show that this positive result extends to ranking and location mechanisms as well. In the entirety of this section, the metric space is $\Re$ and the mechanisms are single candidate mechanisms.

We start by showing a basic lemma regarding WPSC mechanisms. Recall that these are single candidate mechanisms, specified by a vector of probabilities $p_1, \ldots, p_n$, which choose candidate y_i with probability $\sum_{j:z_j=y_i} p_j$.

LEMMA 4.1. *Any weighted percentile single candidate (WPSC) mechanism M on $\Re$ is universally truthful.*

PROOF. [Sui and Boutilier 2015] show that percentile mechanisms are truthful on $\Re$. WPSC mechanisms take a given distribution over truthful [deterministic] mechanisms, and are therefore universally truthful. □

Definition 4.2 (*Spike Mechanism*). Let $P(j)$ be the following function for $0 \leq j \leq n$:

$$P(j) = \begin{cases} 0 & \text{if } j = 0 \\ \frac{j}{2(n-j)} & \text{if } 0 < j \leq n/2 \\ 1.5 - \frac{n}{2j} & \text{if } j > n/2 \end{cases}$$

Let $p_j = P(j) - P(j-1)$.
The spike mechanism chooses candidate y_i with probability $\sum_{j:z_j=y_i} p_j$.

The previous sentence is equivalent to the following: the spike mechanism chooses voter $\tau(j)$ with probability p_j, and then locates the facility on $a_{\tau(j)}$.

Spike is a WPSC mechanism due to its definition and since for any value of n, $\{p_j\}_{j=1}^n$ is uniquely defined. The mechanism is named after the shape of the function that $\{p_j\}_{j=1}^n$ creates (see Figure 4). We note that the result of the mechanism depends on the number of votes that each candidate received and on the order of the candidate along the line, but not on the distances between the candidates.

OBSERVATION 4.3. *Spike induces a symmetric distribution: $\forall i : F(i) = 1 - F(n-i)$.*

PROOF. Without loss of generality let $1 \leq i \leq n/2$, then:

$$F(i) = \frac{i}{2(n-i)}$$

$$1 - F(n-i) \;=\; 1 - \left(1.5 - \frac{n}{2(n-i)}\right) = \frac{n}{2(n-i)} - \frac{1}{2} = \frac{n-(n-i)}{2(n-i)} = \frac{i}{2(n-i)}$$

□

We now define a few terms needed for the proof of the approximation ratio. Recall that y_{opt} is uniquely defined for a location profile x, since ties are broken in favor of the leftmost candidate. Denote the set of borders $\{b_i\}_{i=1}^{m-1}$ by B.

Definition 4.4 (*Tight profile of* x). Given a location profile x, the profile x′ is said to be the *tight profile of* x if it moves all agents who are not on a border as close as possible to y_{opt} within their zones, that is:

$$\forall i : x'_i = \begin{cases} x_i & \text{if } x_i \in B \\ y_{\text{opt}} & \text{if } x_i \in \mathcal{Z}_{\text{opt}} \setminus B \\ b_j & \text{if } x_i \in \mathcal{Z}_j \setminus B \text{ and } j < y_{\text{opt}} \\ b_{j-1} & \text{if } x_i \in \mathcal{Z}_j \setminus B \text{ and } j > y_{\text{opt}} \end{cases}$$

A graphical visualization of a tight profile can be seen in the transition between Figures 8 and 9, or in a figure specifically designated to this in the full version.

Definition 4.5 (*Left-compressed profile of* x). Given a tight location profile x, a left-compressed profile of x moves all the agents on the leftmost border to their neighboring border on the right, if this border is left of y_{opt}. Formally: let the location of the leftmost agent be $x_1 = b_j$, then the left-compressed profile of x is:

$$\forall i : x'_i = \begin{cases} b_{j+1} & \text{if } (x_i = b_j) \wedge (b_{j+1} < y_{\text{opt}}) \\ x_i & \text{otherwise.} \end{cases}$$

Left-compressed profiles can be seen in the transition between Figures 9 and 10, or in a figure specifically designated to this in the full version. Note that the left-compressed profile of a tight profile is also a tight profile. The right-compressed profile of x is defined in a completely symmetrical fashion.

After compressing location profiles, there are likely to be locations in which there are many agents. We therefore use the following notation: the location profile is written as $\mathrm{x} = \{(\hat{x}_1, n_1), ..., (\hat{x}_k, n_k)\}$, which means that for each $j : 1 \leq j \leq k$, there are n_j agents located at $\hat{x}_j$ (see, e.g, Figure 10).

We now use these definitions to prove the main result of this section:

THEOREM 4.6. *The spike mechanism is universally truthful, and it achieves an approximation ratio of 2 on* $\Re$.

PROOF. Spike is a WPSC mechanism, so from Lemma 4.1 it is universally truthful.

The analysis of the approximation ratio is more involved, and is based on backwards induction which follows these steps (see Figures 8 through 11):

(1) Figure 8: Let x an arbitrary location profile. Compute its optimal candidate, y_{opt}.
(2) Figure 9: Let $\mathrm{x}^{(1)}$ be the tight profile of x. We show that the transition from x to $\mathrm{x}^{(1)}$ cannot reduce the approximation ratio (Lemma 4.7).
(3) Figure 10: Let $\mathrm{x}^{(2)}$ be the left-compression of $\mathrm{x}^{(1)}$. We show that if the ratio of $\mathrm{x}^{(2)}$ is not higher than 2, then so is the ratio of $\mathrm{x}^{(1)}$ (Lemma 4.8).
(4) Figure 11: Repeat left and right compressions until we can no longer compress. At this stage, the profile is tight with at most 3 active candidates, and we note this profile $\mathrm{x}^{(3)}$. We show that the ratio of $\mathrm{x}^{(3)}$ is not above 2 (Lemma 4.9).

Proving these steps is sufficient to complete the proof of the theorem, since in Lemma 4.9 we show that the ratio of $\mathrm{x}^{(3)}$ is not higher than 2 (the base case). According to Lemma 4.8, this implies that the ratio of $\mathrm{x}^{(1)}$ (prior to all of the compressions) is also not higher than 2 (the induction steps). Since the ratio of the x is not higher than that of $\mathrm{x}^{(1)}$, this means that the approximation ratio of x is not above 2, as needed.

Notice that throughout this process y_{opt} remains the optimal candidate, since it was optimal in the original profile x, and in each step all agents move towards it, so the cost of any other candidate can decrease by no more than what the cost of y_{opt} decreases.

Since spike is a single candidate mechanism and agents may be on borders, truthful reports are not necessarily unique. In cases of ties, we show that the worst-case ratio always occurs when the agents vote for the candidate located farther away from y_{opt} (the formal lemma and its proof are provided in the full version).

We now present the aforementioned lemmas formally. Their proofs (deferred to the the full version) prove the backwards induction and conclude the proof of the theorem.

LEMMA 4.7. *Let* M *be an arbitrary WPSC mechanism on* $\Re$*, let* x *be an arbitrary location profile and let* x′ *be the tight profile of* x*. Then the approximation ratio of* M

Fig. 8. Initial state: A general location profile $\mathbf{x}$

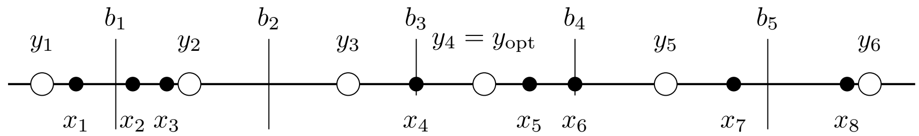

Fig. 9. $\mathbf{x}^{(1)}$: The tight profile of $\mathbf{x}$, using the notation in which there are n_i agents at point $\hat{x}_i$.

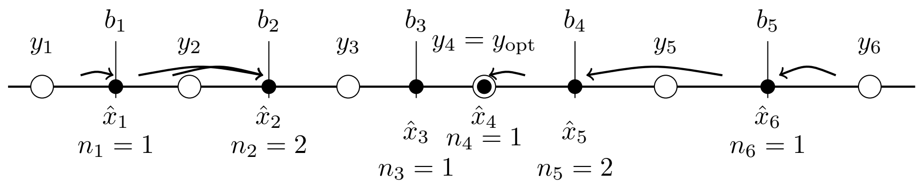

Fig. 10. $\mathbf{x}^{(2)}$: The left-compressed profile of $\mathbf{x}^{(1)}$ (moves n_1 agents from b_1 to b_2).

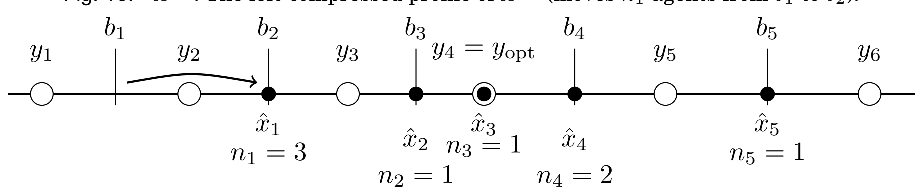

Fig. 11. $\mathbf{x}^{(3)}$: The final profile after reapplying left and right compressions on $\mathbf{x}^{(2)}$ repeatedly. The active candidates are denoted by y_L, y_C and y_R, the number of agents by L, C and R respectively, and we scale the distances such that $\frac{b_C - y_C}{2} = 1$ and $\frac{y_C - b_L}{2} = \beta$.

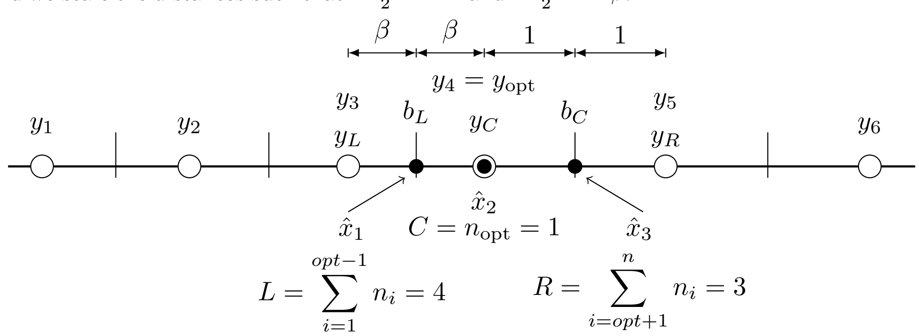

given $\mathbf{x}'$ *is not lower than that of* M *given* $\mathbf{x}$*:*

$$\frac{\mathrm{SC}(M, \mathbf{x})}{\mathrm{SC}(\mathrm{OPT}, \mathbf{x})} \leq \frac{\mathrm{SC}(M, \mathbf{x}')}{\mathrm{SC}(\mathrm{OPT}, \mathbf{x}')}.$$

The previous lemma holds for any WPSC mechanism, and in particular for spike.

LEMMA 4.8. *Let* $\mathbf{x}$ *be a tight location profile on* $\Re$*, let* $\mathbf{x}'$ *be the left-compressed profile of* $\mathbf{x}$ *and let* S *be the spike mechanism. Then if the approximation ratio of* S *given* $\mathbf{x}'$ *is not higher than 2, then so is that of* S *given* $\mathbf{x}$*:*

$$\frac{\mathrm{SC}(S,\mathbf{x}')}{\mathrm{SC}(\mathrm{OPT},\mathbf{x}')} \leq 2 \Rightarrow \frac{\mathrm{SC}(S,\mathbf{x})}{\mathrm{SC}(\mathrm{OPT},\mathbf{x})} \leq 2$$

Observation 4.3 shows that the cumulative function defining the spike mechanism is symmetrical, so the lemma can be trivially extended to right-compressions as well.

After reapplying compressions on both sides, the resulting profile has agents in three locations at most (see Figure 11). In this profile, the ratio is not higher than 2:

LEMMA 4.9. *Let* $\mathbf{x}$ *be a tight location profile on* $\Re$ *in which there are at most 3 active candidates:* $y_{\mathrm{opt}-1} < y_{\mathrm{opt}} < y_{\mathrm{opt}+1}$*. The ratio of the spike mechanism* S *given* $\mathbf{x}$ *is not higher than 2:* $\frac{\mathrm{SC}(S,\mathbf{x})}{\mathrm{SC}(\mathrm{OPT},\mathbf{x})} \leq 2$*.*

□

5. ADDITIONAL RESULTS FOR RANDOMIZED MECHANISMS

5.1. Lower Bounds

This section shows lower bounds of randomized mechanisms in different settings:

— For truthful in expectation single candidate mechanisms in $\Re^d$, we show a lower bound of $3 - \frac{2}{d+1}$. Hence, for a general metric space this lower bound approaches 3 (we will later show that the random dictator mechanism matches this bound).

— For the special case of $d = 1$, the previous result implies a lower bound of 2. We show that on the line, the lower bound of 2 extends to truthful in expectation location mechanisms and ranking mechanisms as well. Additionally, we show that on the line, there is a lower bound of 2 for any ranking mechanism, even when the mechanism need not be truthful (in the non-strategic setting). Collectively this shows that the bound of spike is tight in several aspects (see Figure 6).

— Finally, for truthful in expectation ranking mechanisms, we present a lower bound of $7/3$ in $\Re^2$ (which also holds in $\Re^d$, for any $d > 2$).

We start by proving a helpful lemma. Informally, the lemma states that when there is an agent located on a border (and can therefore has several true actions in ranking and single candidate mechanisms), their cost should not change under any truthful report they submit.

Proofs of most of the results in this section are deferred to the full version.

LEMMA 5.1. *For any truthful in expectation ranking mechanism* M *in any metric space, let* $b_{i,j}$ *be the border between ranking zones* $\mathcal{R}_i, \mathcal{R}_j$*. Let* π_i, π_j *be the rankings consistent with* $\mathcal{R}_i, \mathcal{R}_j$ *respectively. Let agent* l *be located on this border, that is:* $x_l \in b_{i,j}$*. Then the cost at point* x_l *remains the same whether the agent reports* π_i *or* π_j*, that is:*

$$\mathrm{cost}_{x_l}(M, (a_l = \pi_i, a_{-l})) = \mathrm{cost}_{x_l}(M, (a_l = \pi_j, a_{-l}))$$

OBSERVATION 5.2. *Lemma 5.1 also holds for single candidate mechanisms.*

THEOREM 5.3. *In the* d *dimensional real space* $\Re^d$*, any truthful in expectation single candidate mechanism has an approximation ratio of at least* $3 - \frac{2}{d+1}$*.*

PROOF. Let there be $d+1$ candidates, located on the vertices of a regular simplex H (all $d+1$ vertices are equally distanced from one another). Let there be $d+1$ agents, and let M be an arbitrary truthful in expectation single candidate mechanism.

Let **x** be the profile in which each agent i is located precisely on candidate y_i. Therefore **a** $= (y_1, y_2 \ldots y_{d+1})$ is the only truthful single candidate profile for **x**. Denote the probability of choosing candidate i by $p_i(a)$, that is: $p_i(a) = \Pr(M(\mathbf{a}) = y_i)$. Clearly there exists some candidate which is chosen by M with probability at least $\frac{1}{d+1}$. Assume without loss of generality that this candidate is y_{d+1}, that is: $p_{d+1}(a) \geq \frac{1}{d+1}$.

We move on to define another location profile, **x**$'$, which is also consistent with the single candidate profile **a**. Let H' be the regular simplex in which candidates $y \setminus y_{d+1}$ are on the vertices. Let P be the point with equal distance to all d vertices in H' (H' is a regular simplex, so such a point necessarily exists). Denote this distance as t. However, this distance is different from the distance from P to y_{d+1}: $|P - y_{d+1}| = u \neq t$. Let **x**$'$ be the profile in which there are k agents at P and one agent at y_{d+1} (see Figure 12).

According to Observation 5.2, the cost of an agent at point P should not change under any truthful vote, that is for any vote $y_j : 1 \leq j \leq d$. In particular, this holds when any agent on point P votes for candidate y_1. We make use of this observation several times by changing the votes for each of the points at P to y_1, one at a time, such that the final single candidate profile is **a**$' = (y_1, y_1 \ldots y_1, y_{d+1})$ (d agents vote for y_1, one agent votes for y_{d+1}). Due to Observation 5.2, the cost of point P must remain the same throughout these transitions, that is: $\mathrm{cost}_P(M, a') = \mathrm{cost}_P(M, a)$. Therefore:

$$
\begin{aligned}
& \mathrm{cost}_P(M,a) = \mathrm{cost}_P(M,a') \\
\Rightarrow\ & u \cdot p_{d+1}(a) + t \cdot (1 - p_{d+1}(a)) = u \cdot p_{d+1}(a') + t \cdot (1 - p_{d+1}(a')) \\
\Rightarrow\ & t + (u - t) \cdot p_{d+1}(a) = t + (u - t) \cdot p_{d+1}(a') \\
\Rightarrow\ & p_{d+1}(a') = p_{d+1}(a) \\
\Rightarrow\ & p_{d+1}(a') \geq \frac{1}{d+1}
\end{aligned}
$$

Denote the midpoint between y_1 and y_{d+1} by Q. Without loss of generality, scale the distances such that $|y_1 - Q| = |Q - y_{d+1}| = 1$. Examine the following location profile

Fig. 12. An illustration of the proof of Theorem 5.3 for the case of $d = 2$. The three candidates are on the vertices of an equilateral triangle (a regular simplex with 3 vertices). The lines within the triangle denote the borders between a pair of candidates. The simplex H' is the line between y_1, y_2, and its midpoint is P.
These three figures, from left to right, show the dynamics of the proof (location profiles **x**,**x**$'$ and **x**$''$ respectively). The key observation is that the agents at point P are at equal distance to all candidates except y_3, therefore the transitions do not change the probability that y_3 be chosen.

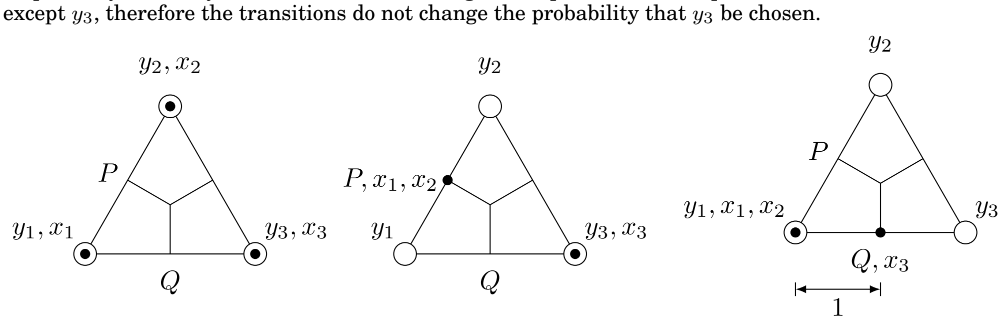

$\mathbf{x}'' = (y_1, y_1 \ldots y_1, Q)$, which is also consistent with the single candidate profile $\mathbf{a}'$. In this case the cost of y_1, which is the optimal candidate, is: $\text{SC}(y_1, \mathbf{x}'') = 1$. The cost of y_{d+1} is $\text{SC}(y_{d+1}, \mathbf{x}'') = d \cdot 2 + 1 = 2d + 1$. Ergo the approximation ratio of M is at least:

$$\begin{aligned}\frac{\text{SC}(M, \mathbf{x}'')}{\text{SC}(\text{OPT}, \mathbf{x}'')} &= p_{d+1}(a')(2d+1) + (1 - p_{d+1}(a'))(1) \\ &= 2d \cdot p_{d+1}(a') + 1 \\ &\geq (2d)\frac{1}{d+1} + 1 = 3 - \frac{2}{d+1}\end{aligned}$$

□

OBSERVATION 5.4. *Any truthful in expectation location mechanism has an approximation ratio of at least 2, even on the line.*

PROOF. Let there be two candidates on the line. According to Theorem 3.1 (6), any truthful in expectation location mechanism is equivalent to a single candidate mechanism. One can apply the same proof as in Theorem 5.3 for the case of $d = 1$ to achieve a lower bound of $3 - \frac{2}{d+1} = 2$. □

OBSERVATION 5.5. *The bound of 3 by random dictator is tight for general metric spaces.*

LEMMA 5.6. *No randomized ranking mechanism can achieve an approximation ratio strictly below 2, even if the metric is $\Re$ and even if there are no truthfulness requirements from the mechanism (non-strategic case).*

The following result shows a lower bound on truthful in expectation ranking mechanisms in $\Re^2$ (and trivially extends to higher order Euclidean metric spaces).

THEOREM 5.7. *In $\Re^2$, any truthful in expectation ranking mechanism has an approximation ratio of at least $7/3$.*

5.2. Upper Bound

We previously showed that spike achieves an approximation ratio of 2 on the line. We now show that for a general metric space, random dictator achieves a ratio of 3 (i.e, the upper bound is 3). Recall that random dictator locates the facility on vote a_i with probability $1/n$ for all $i \in N$, and achieves a ratio of 2 in the continuous model.

LEMMA 5.8. *On any metric space, random dictator yields a 3 approximation of the optimal social cost.*

We note that in contradiction to the continuous model, in our candidate model random dictator is not group-strategyproof (proof in full version).

COROLLARY 5.9. *The combination of the lower bound in Theorem 5.3 with the upper bound of random dictator show that for a general metric space the bound of 3 is tight for single candidate mechanisms.*

6. DETERMINISTIC MECHANISMS ON THE LINE

In the continuous model, choosing the location of the median agent is both truthful and optimal [Procaccia and Tennenholtz 2009]. The following theorem shows that in the candidate model, *median* results in a ratio of 3, and that this is the best one can hope for with any deterministic mechanism (even location mechanisms). The proof, as well as the formal definition of the median mechanism, are deferred to the full version.

THEOREM 6.1. *The following claims hold:*

(*1*) *No deterministic truthful mechanism (location, ranking or single candidate) can achieve an approximation ratio strictly below* 3 *for the social cost, even when the metric space is* $\Re$.
(*2*) *Median is a truthful mechanism on* $\Re$ *which results in a 3 approximation of the social cost.*

7. DISCUSSION AND OPEN PROBLEMS

We defined three types of truthful mechanisms (single candidate, ranking and location), and showed the relations between these sets of truthful mechanisms. Then we gave bounds on the approximation ratio of these mechanism types in various settings. In particular, the paper introduced the spike mechanism, a truthful single candidate mechanism which achieves a [tight] bound of 2 on $\Re$. Ideas for future work include:

— Electing a committee of multiple candidates, i.e., locating multiple facilities.
— Closing the gap in the bounds for ranking and location mechanisms in $\Re^d$.
— More generally, we studied the affects of information deficiency and truthfulness in the context of voting. These affects can be addressed for many additional problems.

ACKNOWLEDGMENTS

We would like to thank Alon Eden, Ilan Cohen, Ophir Friedler, Shai Vardi, Rachel Matichin, Sella Nevo, Guy Reiner, Lirong Xia, Elliot Anshelevich and Ariel Procaccia for interesting discussions.

REFERENCES

Noga Alon, Michal Feldman, Ariel D Procaccia, and Moshe Tennenholtz. 2009. Strategyproof approximation mechanisms for location on networks. *arXiv preprint arXiv:0907.2049* (2009).

Noga Alon, Michal Feldman, Ariel D Procaccia, and Moshe Tennenholtz. 2010. Walking in circles. *Discrete Mathematics* 310, 23 (2010), 3432–3435.

Elliot Anshelevich, Onkar Bhardwaj, and John Postl. 2014. Approximating Optimal Social Choice under Metric Preferences. (2014).

Elliot Anshelevich and John Postl. 2015. Randomized Social Choice Functions Under Metric Preferences. *arXiv preprint arXiv:1512.07590* (2015).

Haris Aziz, Florian Brandl, and Felix Brandt. 2014. On the Incompatibility of Efficiency and Strategyproofness in Randomized Social Choice.. In *AAAI*. Citeseer, 545–551.

Duncan Black. 1948. On the rationale of group decision-making. *The Journal of Political Economy* (1948), 23–34.

Anna Bogomolnaia and Hervé Moulin. 2001. A new solution to the random assignment problem. *Journal of Economic theory* 100, 2 (2001), 295–328.

Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D Procaccia, and Or Sheffet. 2012. Optimal social choice functions: A utilitarian view. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 197–214.

Felix Brandt, Vincent Conitzer, and Ulle Endriss. 2012. Computational social choice. *Multiagent systems* (2012), 213–283.

Ioannis Caragiannis and Ariel D Procaccia. 2011. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence* 175, 9 (2011), 1655–1671.

Yukun Cheng, Wei Yu, and Guochuan Zhang. 2013. Strategy-proof approximation mechanisms for an obnoxious facility game on networks. *Theoretical Computer Science* 497 (2013), 154–163.

Elad Dokow, Michal Feldman, Reshef Meir, and Ilan Nehama. 2012. Mechanism design on discrete lines and cycles. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 423–440.

Anthony Downs. 1957. An economic theory of political action in a democracy. *The journal of political economy* (1957), 135–150.

Joan Feigenbaum, Aaron D Jaggard, and Michael Schapira. 2010. Approximate privacy: foundations and quantification. In *Proceedings of the 11th ACM conference on Electronic commerce*. ACM, 167–178.

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

Marcus Fleming. 1952. A cardinal concept of welfare. *The Quarterly Journal of Economics* (1952), 366–384.

Dimitris Fotakis and Christos Tzamos. 2013. Strategyproof facility location for concave cost functions. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. ACM, 435–452.

Dimitris Fotakis and Christos Tzamos. 2014. On the power of deterministic mechanisms for facility location games. *ACM Transactions on Economics and Computation* 2, 4 (2014), 15.

Dimitris Fotakis, Christos Tzamos, and Emmanouil Zampetakis. 2015. Who to Trust for Truthfully Maximizing Welfare? *arXiv preprint arXiv:1507.02301* (2015).

Allan Gibbard. 1973. Manipulation of voting schemes: a general result. *Econometrica: journal of the Econometric Society* (1973), 587–601.

J Harsanyi. 1955. ICardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Welfare. *J. Polit. Econ* 63 (1955).

Harold Hotelling. 1990. *Stability in competition*. Springer.

Pinyan Lu, Xiaorui Sun, Yajun Wang, and Zeyuan Allen Zhu. 2010. Asymptotically optimal strategy-proof mechanisms for two-facility games. In *Proceedings of the 11th ACM conference on Electronic commerce*. ACM, 315–324.

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

Noam Nisan. 2007. Introduction to mechanism design (for computer scientists). *Algorithmic game theory* 209 (2007), 242.

Ariel D Procaccia and Jeffrey S Rosenschein. 2006. The distortion of cardinal preferences in voting. In *Cooperative Information Agents X*. Springer, 317–331.

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

Mark Allen Satterthwaite. 1975. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of economic theory* 10, 2 (1975), 187–217.

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

Xin Sui and Craig Boutilier. 2011. Efficiency and Privacy Tradeoffs in Mechanism Design.. In *AAAI*.

Xin Sui and Craig Boutilier. 2015. Approximately Strategy-proof Mechanisms for (Constrained) Facility Location. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 605–613.

Xin Sui, Craig Boutilier, and Tuomas W Sandholm. 2013. Analysis and optimization of multi-dimensional percentile mechanisms. AAAI.