

Testing Preferential Domains Using Sampling

Palash Dey
Indian Institute of Technology
Kharagpur
palash.dey@cse.iitkgp.ac.in

Swaprava Nath
Indian Institute of Technology
Kanpur
swaprava@cse.iitk.ac.in

Garima Shakya
Indian Institute of Technology
Kanpur
garima@cse.iitk.ac.in

ABSTRACT

A preferential domain is a collection of sets of preferences which are linear orders over a set of alternatives. These domains have been studied extensively in social choice theory due to both its practical importance and theoretical elegance. Examples of some extensively studied preferential domains include single peaked, single crossing, Euclidean, etc. In this paper, we study the sample complexity of testing whether a given preference profile is close to some specific domain. We consider two notions of closeness: (a) closeness via preferences, and (b) closeness via alternatives. We further explore the effect of assuming that the *outlier* preferences/alternatives to be random (instead of arbitrary) on the sample complexity of the testing problem. In most cases, we show that the above testing problem can be solved with high probability for all commonly used domains by observing only a small number of samples (independent of the number of preferences, n , and often the number of alternatives, m). In the remaining few cases, we prove either impossibility results or $\Omega(n)$ lower bound on the sample complexity. We complement our theoretical findings with extensive simulations to figure out the actual constant factors of our asymptotic sample complexity bounds.

KEYWORDS

Computational social choice; preferential domain; sampling; algorithms; testing; almost single peak; almost single crossing

ACM Reference Format:

Palash Dey, Swaprava Nath, and Garima Shakya. 2019. Testing Preferential Domains Using Sampling. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019*, IFAAMAS, 9 pages.

1 INTRODUCTION

Learning users' preferences is useful in the contexts of social choice, recommender systems, product development, and many more applications. It is often observed that preferences are never completely arbitrary, rather they possess correlated structures [32]. For example, preferences of citizens for a facility location have a *single peaked* structure [30, Section 1], i.e., a citizen has highest preference for the facility at her location and it monotonically decreases with the distance from her. This kind of preferences are also prevalent in political opinions based on the voters' bias to the conservative or liberal views [36]. Intuitively, in a single peaked preference profile,

we assume that there exists a societal axis where the alternatives have been ordered and every preference “respects” that ordering in the following sense. Every preference has an implicit most preferred point t on the societal axis and if an alternative x lies between t and another alternative y , then x is preferred over y . The advantage of preferences with such structures is that they can efficiently bypass the classic impossibility results of social choice theory [2, 33, 47].

Similarly, in the design of recommender systems, it has often been observed that users' preferences (and hence their recommendations) have patterns that are (a) demography-based, (b) knowledge-based, (c) feature-based, or (d) content based [44]. While designing a product, an enterprise may wish to look for structures in the end users' preferences, and design their product such that a collectively ‘efficient’ choice is made to cater a large number of users.

While it is difficult to predict the users' preferences apriori, data on the preferences, obtained through users' purchase and browsing patterns, or through surveys, are plentiful which are classified into demography, knowledge, affinity towards a feature or content. It remains to discover whether the preferences come from a specific class that we call *preferential domains* or simply, *domains*.

A domain is a collection of sets of preferences over a set of alternatives. A preference profile, i.e., the tuple of preferences of all the agents/users, is said to belong to a domain if, for some set in the domain, every preference in the profile belongs to that set.

Example 1.1 (Single peaked domain). Consider three alternatives a, b, c . The *single peaked domain* with these alternatives is denoted by $\mathcal{D} = \{\mathcal{B}_1, \dots, \mathcal{B}_6\}$, where $\mathcal{B}_1 = \{(abc), (bac), (bca), (cba)\}$ when the societal order over the alternatives is $a < b < c$, and similarly, $\mathcal{B}_2, \dots, \mathcal{B}_6$ are the sets of preferences over the same alternatives for different societal orders of a, b , and c .

Some prominent examples of domains are *single peaked*, *single crossing*, *Euclidean*, [32] etc. The benefit of the discovery of such domains (even as a partial population) is that a much refined plan or protocol can be designed for such domains which satisfy several desirable axioms. For example, the *median voting rule* in the single peaked domain ensures that no voter can gain by misreporting her preference [43]. Another reason to study various domains concerns computational considerations. Indeed, some of the most fundamental problems in computational social choice, for example, computing winners for many important voting rules such as Kemeny, Dodgson, and Young are computationally intractable [9]. It turns out that most of these problems become efficiently solvable in many domains, single peaked for example [8].

Our work in this paper contributes to uncovering whether a given preference profile is “close” to some domain, through sampling a small number of preferences and/or alternatives. The guarantees we provide are probabilistic that converges to unity as more

Dey is funded by DST INSPIRE grant no. 04/2016/001479 and IIT Kharagpur grant no. IIT/SRIC/CS/VTS/2018-19/247

Nath is funded by IIT Kanpur grant no. IITK/CS/2017198.

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada. © 2019 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

Input profile		Sample complexity
Possibility 1	Possibility 2	
$\varepsilon_v n$ random preferences away	random	$O(\frac{1}{(1-\varepsilon_v)^2} \log \frac{1}{\delta})$ [Theorem 3.1]
$\varepsilon_v n$ arbitrary preferences away		$O(\frac{1}{(1-3\varepsilon_v)^2} \ln \frac{1}{\delta})$ for $\varepsilon_v < 1/3^*$ [Theorem 3.4]
$\varepsilon_a m$ alternatives away		$O(\log \frac{\log_{1/\varepsilon_a} 1/\delta}{\delta} \log_{1/\varepsilon_a} \frac{1}{\delta} \log \log_{1/\varepsilon_a} 1/\delta)$ [Theorem 3.10]
$\varepsilon_v n$ arbitrary preferences away	$\varepsilon'_v n$ arbitrary preferences away	$O(\frac{1}{(\varepsilon'_v - \varepsilon_v)^2} (2^m m^2 \log^2 m + \log 1/\delta))$ [Corollary 3.9]
$\varepsilon_a m$ alternatives away	$\varepsilon'_a m$ alternatives away	$\Omega(n \log 1/\delta)$ even for $\varepsilon_a = 0$ and for every $0 < \varepsilon'_a \leq 1$ and $0 < \delta < 1/2$ [Theorem 3.13]

Table 1: Summary of results for distinguishing profiles in the first column from the profiles in the second column; all the distances are from the single peaked domain. \star : For any $0 \leq \varepsilon_v < 1$, we refer to Theorem 3.6. Refer to Section 2.2 for our sampling model.

preferences/alternatives are investigated – the cost of such an investigation is often proportional to the number of samples drawn, known as *sample complexity*. Hence our goal is to minimize the sample complexity of our algorithms. For example, our algorithms could be used to predict whether there exist at least, say 95%, of the preferences in a profile which are single peaked. If we know the societal order of the single peaked preferences (which constitute at least 95% of the profile), using median voting rule on the single peaked sub-profile would yield all the desirable properties of the median voting rule, e.g., truthfulness for those 95% of the population. These kind of truthfulness of a fraction of voters is referred to as “approximate truthfulness.” In many applications like public good provisioning, it is highly beneficial to uncover truthful opinions from the vast majority of the population.

To put our work in perspective, we revisit a question that is often asked in computational social choice for any domain. This is about the existence of an efficient recognition algorithm: given a profile \mathcal{P} , does there exist a polynomial time algorithm to decide whether \mathcal{P} belongs to the domain? There exist efficient recognition algorithms for many popular domains, for example, single peaked [6], single crossing [23], etc. [24, 25, 37, 41]. One notable exception is the Euclidean domain of dimension two where the recognition problem is NP-hard [45].

There are two main limitations of the recognition problem. First, the problem formulation is “exact.” Real world profiles are almost never perfect and thus they can only be at most “close” to some domain. More specifically, there may be few preferences or alternatives (treated as outliers) whom we need to ignore to obtain the required structure. Unfortunately, outliers’ consideration often makes the related recognition problem intractable, (e.g., the voter deletion for single peaked domain [28]). Second, the recognition problem needs access to the entire preference profile. In many situations, e.g., pre-election polls, surveys, etc., we only have access to samples. In other cases, the number of preferences may be too large and, depending on the application at hand, a sub-linear time (possibly approximation) algorithm may be more useful. We address both these issues by defining a related testing problem. As a concrete use case, a social planner could use our testing algorithms to know whether it is possible to remove, say 5% of the preferences to obtain a single peaked structure by observing a small number of samples.

A corresponding computational problem is: can a profile of n preferences over m alternatives belong to some domain \mathcal{D} after deleting, say at most k preferences (or alternatives), by drawing a small number of samples? However, any algorithm for this problem would need to observe $\Omega(n)$ samples which defeats the main purpose of testing (except when \mathcal{D} is empty or \mathcal{D} contains all possible profiles). To see this, let us consider a specific case of \mathcal{D} to be single peaked; the set of alternatives be $\{a, b, c\}$. Let \mathcal{P} be a profile consisting of $n/2$ (say n is an even integer) copies of $a > b > c$, $(n/2) - 1$ copies of $a > c > b$, and one $c > b > a$. We observe that \mathcal{P} is not single peaked after observing the last preference $c > b > a$. However, deletion of that preference makes it single peaked. Let us now consider another profile \mathcal{Q} consisting of $n/2$ copies of $a > b > c$, $(n/2) - 2$ copies of $a > c > b$, and two copies of $c > b > a$. Again, \mathcal{Q} is not single peaked, but deletion of the two copies of $c > b > a$ makes it single peaked. We now observe that the KL-divergence [38] between the two distributions of samples for \mathcal{P} and \mathcal{Q} is $O(1/n)$ and thus distinguishing \mathcal{P} from \mathcal{Q} (which any testing algorithm has to do) requires $\Omega(n)$ samples to succeed with any constant nonzero probability [4]. To overcome this lower bound, we introduce (as is ubiquitous in testing literature [34, 46]) a “gap” in the two possible inputs. In all our testing problems, we are given a profile as input which is guaranteed to be one of the two possible types, and we need to find which one it is. The two possibilities for the input will cover all the cases except few and thus there is a “gap.”

1.1 Our Contribution

Our specific contribution in this paper are as follows. The error probability of any algorithm below is at most $\delta \in (0, 1)$.

(i) We present a sampling based algorithm to distinguish any profile for which there exists a set \mathcal{R} of at most $\varepsilon_v n$ preferences (or $\varepsilon_a m$ alternatives) whose deletion makes the resulting profile belong to \mathcal{D} from any random profile (refer to the first three rows in Table 1). We observe that the sample complexity depends on whether we assume \mathcal{R} to be arbitrary or random. We remark that, in the testing literature [1, 35, 48], it is popular to assume the noise to be random which is equivalent to assuming the preferences in \mathcal{R} to be random in our context.

(ii) For any $0 \leq \varepsilon_v < \varepsilon'_v \leq 1$, we present a sampling based algorithm to distinguish any profile for which there exist at most $\varepsilon_v n$ preferences whose deletion makes the resulting profile belong to

\mathcal{D} from any profile where one has to delete at least $\varepsilon'_v n$ preferences to make it belong to \mathcal{D} (refer to the fourth row in Table 1).

(iii) In the case of alternatives, we prove that any algorithm for distinguishing any profile for which there exist at most $\varepsilon_a m$ alternatives whose deletion makes the resulting profile belong to \mathcal{D} from any profile where one has to delete at least $\varepsilon'_a m$ alternatives to make it belong to \mathcal{D} has sample complexity of $\Omega(n \log 1/\delta)$ for every $0 \leq \varepsilon_a < \varepsilon'_a \leq 1$ even when $\varepsilon_a = 0$ (refer to the fifth row in Table 1). This shows that detecting arbitrary *outlier alternatives* is much harder than detecting arbitrary *outlier preferences* from a sample complexity viewpoint.

We remark that all our results in Table 1 for the single peaked domain actually extend to any domain as described in Section 3. From a technical point of view, to tackle preferences which are outliers, we define and exploit a notion called *content of a domain* which, informally, is the maximum number of distinct preferences that any profile in the domain can contain as a function of the number of alternatives. On the other hand, we blend with it the ideas from the classical coupon collector problem to handle alternatives which are outliers. To develop an algorithm for the case when the outliers can be arbitrary, we prove a key structural result (in Lemma 3.7) for arbitrary domain which may be of independent interest also.

1.2 Related Work

The computational problem of recognizing whether a given profile belongs to a domain has been studied extensively in computational social choice. Trick [6] shows that the recognition problem is polynomial time solvable for single peaked profiles. Escoffier et al. [29] improve the efficiency of the recognition algorithm for the single peaked profiles. Elkind et al. [27] present a polynomial time algorithm for recognizing single crossing profiles. Barberà and Moreno [5] discover a property called top monotonicity which simultaneously generalizes both single peakedness and single crossingness. Magiera and Faliszewski [41] present polynomial time recognition algorithm for top monotonic profiles. Doignon and Falmagne [23] show that the recognition problem for the one dimensional Euclidean domain is polynomial time solvable. Knoblauch [37] and Elkind and Faliszewski [24] present alternative algorithms for recognizing one dimensional Euclidean profiles. Peters [45] shows that recognizing Euclidean profiles of dimension at least two is NP-hard.

Lackner [39] shows that the computational problem of finding if it is possible to extend a given incomplete profile to a single peaked profile is NP-complete. However, if we restrict ourselves to only weak orders, then the computational problem of recognizing incomplete single peaked profiles is polynomial time solvable [31]. The above problem is polynomial time solvable for single crossing profiles too [25]. Erdélyi [28] studies complexity of the computational problem of deciding whether a given profile can be “made” single peaked by deleting few preferences or alternatives; Bredereck et al. [11] study complexity of this problem for single peaked, single-caved, single-crossing, etc. profiles. Ballester and Haeringer [3] present characterization of single peaked profiles through succinct forbidden configurations. Bredereck et al. [10] show forbidden configurations for the single crossing profiles. Elkind et al. [26] present

forbidden configurations for profiles which are simultaneously single peaked and single crossing. A related literature studies the likelihood of a random profile being single peaked [12, 13, 40].

There is an active line of work in computational social choice which aims at sub-linear algorithms [7, 14–18, 20, 22] and our work also falls in this category.

2 PRELIMINARIES AND PROBLEM FORMULATION

For any two positive integers k and ℓ with $k \leq \ell$, we denote the set $\{j \in \mathbb{N} : 1 \leq j \leq k\}$ by $[k]$ and the set $\{j \in \mathbb{N} : k \leq j \leq \ell\}$ by $[k, \ell]$. For a set X , we denote its power set by 2^X . Let \mathcal{A} be a finite set of alternatives of cardinality m . Preferences are linear orders over \mathcal{A} . We denote the set of all linear orders over \mathcal{A} by $\mathcal{L}(\mathcal{A})$. For any positive integer n , a tuple $(\succ_i)_{i \in [n]}$ of n preferences is called a profile. If not mentioned otherwise, we use m, n , and \mathcal{A} to denote the number of alternatives, the number of preferences in a profile, and the set of alternatives, respectively. For a subset $X \subseteq \mathcal{A}$ and a preference $\succ \in \mathcal{L}(\mathcal{A})$, we denote the restriction of \succ to X by $\succ(X)$. A preferential domain or simply domain is a collection of subsets of $\cup_{|\mathcal{A}|>0} \mathcal{L}(\mathcal{A})$. We call a domain \mathcal{D} nontrivial if $\mathcal{D} \neq \emptyset$ and $\mathcal{D} \neq \cup_{|\mathcal{A}|>0} 2^{\mathcal{L}(\mathcal{A})}$. Given a domain \mathcal{D} and a profile $\mathcal{P} = (\succ_i)_{i \in [n]}$ over \mathcal{A} , we say (with slight abuse of notation) that $\mathcal{P} \in \mathcal{D}$ if there exists a $\mathcal{B} \in \mathcal{D}$ such that $\mathcal{P} \in \mathcal{B}^n$. We call a domain \mathcal{D} *neutral* if whenever $(\succ_i)_{i \in [n]} \in \mathcal{D}$, we have $(\sigma(\succ_i))_{i \in [n]} \in \mathcal{D}$ for every permutation σ of $[m]$; if \succ_i is defined as $a_1 \succ a_2 \succ \dots \succ a_m$, then $\sigma(\succ_i)$ is defined as $a_{\sigma(1)} \succ a_{\sigma(2)} \succ \dots \succ a_{\sigma(m)}$. We call a domain \mathcal{D} *normal* if whenever $(\succ_i)_{i \in [n]} \in \mathcal{D}$, we have $(\succ_i(X))_{i \in [n]} \in \mathcal{D}$ for every $X \subseteq \mathcal{A}$. In this work, we consider only neutral and normal domains. We remark that many popular domains including single peaked, single caved, single crossing, top restricted, bottom restricted, etc. satisfy these two properties (the only notable exception is the domain of top monotonic [5] profiles).

Let \mathcal{D} be any domain and $\mathcal{P} = (\succ_i)_{i \in [n]} \in \mathcal{L}(\mathcal{A})^n$ be a profile. If it satisfies the following conditions:

- (i) there exists a subset $J \subset [n]$ such that $|J| = \ell$ and $(\succ_i)_{i \in J} \in \mathcal{D}$, and
- (ii) for every subset $K \subset [n]$ such that $|K| > \ell$, we have $(\succ_i)_{i \in K} \notin \mathcal{D}$,

then we say that the *preference-distance* of \mathcal{P} from \mathcal{D} is $(n - \ell)$, and we call the preferences which need to be deleted to bring the profile back to \mathcal{D} to be *preference outliers*. Similarly, we can define the notion of *alternative-distance* (where only alternatives need to be deleted) and alternative outliers.

Our first problem is to distinguish a profile which is, informally speaking, $\varepsilon_a m$ alternatives and random $\varepsilon_v n$ preferences away from some domain \mathcal{D} vs a random profile. We call this problem $(\varepsilon_v, \varepsilon_a, \delta, \mathcal{D})$ – RANDOM OUTLIERS VS RANDOM PROFILE TEST which is formally defined as follows.

Problem 1 $(\varepsilon_v, \varepsilon_a, \delta, \mathcal{D})$ – RANDOM OUTLIERS VS RANDOM PROFILE TEST. Let $(\succ_i)_{i \in [n]}$ be a profile over a set \mathcal{A} of alternatives which is either one of the following kind:

- (i) There exists $\mathcal{W} \subseteq [n]$ and $\mathcal{X} \subseteq \mathcal{A}$ with $|\mathcal{W}| \geq (1 - \varepsilon_v)n$ and $|\mathcal{X}| \geq (1 - \varepsilon_a)m$ such that the profile

$(\succ_i(\mathcal{X}))_{i \in \mathcal{W}}$ belongs to the domain \mathcal{D} and $\succ_j(\mathcal{X})$ is distributed uniformly in $\mathcal{L}(\mathcal{X})$ for every $j \in [n] \setminus \mathcal{W}$.

(ii) The preference \succ_i is distributed uniformly randomly in $\mathcal{L}(\mathcal{A})$ for every $i \in [n]$.

Output 1 if the input profile is of the first kind and 0 if it is of the second kind; the probability of error can be at most δ .

Problem 1 assumes that the preference outliers are distributed uniformly randomly which can be a strong assumption depending on the application at hand. The $(\varepsilon_v, \varepsilon_a, \delta, \mathcal{D})$ - ARBITRARY OUTLIERS VS RANDOM PROFILE TEST problem in Problem 2 removes this assumption.

Problem 2 $((\varepsilon_v, \varepsilon_a, \delta, \mathcal{D})$ - ARBITRARY OUTLIERS VS RANDOM PROFILE TEST). Let $(\succ_i)_{i \in [n]}$ be a profile over a set \mathcal{A} of alternatives which is either one of the following:

(i) There exists $\mathcal{W} \subseteq [n]$ and $\mathcal{X} \subseteq \mathcal{A}$ with $|\mathcal{W}| \geq (1 - \varepsilon_v)n$ and $|\mathcal{X}| \geq (1 - \varepsilon_a)m$ such that the profile $(\succ_i(\mathcal{X}))_{i \in \mathcal{W}}$ belongs to the domain \mathcal{D} .

(ii) The preference \succ_i is distributed uniformly randomly in $\mathcal{L}(\mathcal{A})$ for every $i \in [n]$.

Output 1 if the input profile is of the first kind and 0 if it is of the second kind; the probability of error can be at most δ .

Problem 2 still retains the assumption from Problem 1 that the second possibility for the input profile is random. The $(\varepsilon_v, \varepsilon_a, \varepsilon'_v, \varepsilon'_a, \delta, \mathcal{D})$ - ARBITRARY OUTLIERS VS ARBITRARY PROFILE TEST problem in Problem 3 is the most general problem in our paper which removes all these structural assumptions from Problems 1 and 2.

Problem 3 $((\varepsilon_v, \varepsilon_a, \varepsilon'_v, \varepsilon'_a, \delta, \mathcal{D})$ - ARBITRARY OUTLIERS VS ARBITRARY PROFILE TEST). Let $(\succ_i)_{i \in [n]}$ be a profile over a set \mathcal{A} of alternatives which is either one of the following kind where $0 \leq \varepsilon_v < \varepsilon'_v \leq 1$ and $0 \leq \varepsilon_a < \varepsilon'_a \leq 1$:

(i) There exists $\mathcal{W} \subseteq [n]$ and $\mathcal{X} \subseteq \mathcal{A}$ with $|\mathcal{W}| \geq (1 - \varepsilon_v)n$ and $|\mathcal{X}| \geq (1 - \varepsilon_a)m$ such that the profile $(\succ_i(\mathcal{X}))_{i \in \mathcal{W}}$ belongs to the domain \mathcal{D} .

(ii) For every $\mathcal{W} \subseteq [n]$ and $\mathcal{X} \subseteq \mathcal{A}$ with $|\mathcal{W}| > (1 - \varepsilon'_v)n$ and $|\mathcal{X}| > (1 - \varepsilon'_a)m$, the profile $(\succ_i(\mathcal{X}))_{i \in \mathcal{W}}$ does not belong to the domain \mathcal{D} .

Output 1 if the input profile is of the first kind and 0 if it is of the second kind; the probability of error can be at most δ .

In Problems 1 to 3, the error probability is taken over the randomness used in generating the instances in (ii) and the randomness used by the algorithm.

2.1 Content and Residue of Domain

We now define the *content* and *residue* of any domain which will make the many of our results simpler to state. Let \mathcal{D} be any domain. We define the *content* of \mathcal{D} as a function $\text{con}_{\mathcal{D}} : \mathbb{N} \rightarrow [0, 1]$ such that any profile with $m \in \mathbb{N}$ alternatives in \mathcal{D} can have at most $\text{con}_{\mathcal{D}}(m)m!$ distinct preferences; we call the function $\text{res}_{\mathcal{D}} :$

$\mathbb{N} \rightarrow [0, 1]$ defined as $\text{res}_{\mathcal{D}}(m) = 1 - \text{con}_{\mathcal{D}}(m)$ the *residue* of a domain. For example, $\text{con}_{\text{single peaked}}(2) = 1$, $\text{con}_{\text{single peaked}}(3) = 2/3$, $\text{con}_{\text{single crossing}}(m) = \binom{m}{2} + 1 / m!$ [19]. For technical reason, let us assume that $\text{con}_{\mathcal{D}}(1) = 1$ for every \mathcal{D} . We observe that, for normal domains, the function $\text{con}_{\mathcal{D}}(\cdot)$ is non-increasing (and thus $\text{res}_{\mathcal{D}}(\cdot)$ is a non-decreasing function). Whenever the domain \mathcal{D} is immediate from the context, we omit \mathcal{D} from subscript of con and res .

2.2 Sampling Model and Sample Complexity

In our model, there is an oracle which, when queried, returns an agent v picked uniformly randomly with replacement from the set of all agents. Now the algorithm can ask the agent v an arbitrary number of comparison queries – in a comparison query, two alternatives x and y are presented to the agent v and it replies whether it prefers x over y or y over x . The *sample complexity* of an algorithm is defined to be the total number of comparison queries it makes during its execution. We remark that defining sample complexity (instead of the number of agents sampled) as the number of comparison queries enables us to perform more fine grained analysis of the complexity of our problems.

2.3 Chernoff Bound

We repeatedly use the following concentration inequality:

THEOREM 2.1. Let X_1, \dots, X_ℓ be a sequence of ℓ independent random variables in $[0, 1]$ (not necessarily identical). Let $S = \sum_i X_i$ and let $\mu = \mathbb{E}[S]$. Then, for any $0 \leq \delta \leq 1$:

$$\Pr[|S - \mu| \geq \delta \ell] < 2 \exp(-2\ell\delta^2), \quad (1)$$

and

$$\Pr[|S - \mu| \geq \delta \mu] < 2 \exp(-\delta^2 \mu / 3). \quad (2)$$

Equations (1) and (2) are called *additive* and *multiplicative* versions of the bound respectively.

3 RESULTS

We now present our main results. Our general approach would be to explain our algorithms for the special case of the single peaked domain first and then generalize to arbitrary domain; we make an exception for few cases where presenting the general case directly better reveals the key idea. In the interest of space, we omit some of our proofs, which can be found in the full version [21]. For ease of exposition and interest of space, we have deferred our more involved algorithms for the cases when both preferences and alternatives could simultaneously be outliers to the supplemental material.

3.1 Only Preferences as Outliers

In this subsection, we focus on the case when only preferences are considered as outliers. We begin with presenting our $(\varepsilon_v, 0, \delta, \mathcal{D})$ - RANDOM OUTLIERS VS RANDOM PROFILE TESTER for the single peaked domain. Our algorithm first fixes any three alternatives, say a, b , and c . Then it samples few preferences restricted to these three alternatives only. If all the six possible permutations of a, b , and c appear nearly same number of times, then the algorithm predicts the profile to be a random profile; otherwise it predicts it to be

close to single peaked. We now formally present our algorithm in Theorem 3.1.

Algorithm 1 ($\varepsilon_v, 0, \delta$, single peak) – RANDOM OUTLIERS VS RANDOM PROFILE TESTER

Input: Oracle access to a profile \mathcal{P}

Output: 1 if there exists $\varepsilon_v n$ preferences whose deletion makes the resulting profile single peaked and 0 if \mathcal{P} has been generated randomly

- 1: Let $a, b, c \in \mathcal{A}$ be any 3 arbitrary alternatives
- 2: Sample $\ell = \frac{72}{(1-\varepsilon_v)^2} \ln \frac{6}{\delta}$ preferences restricted to $\{a, b, c\}$ uniformly at random from the input profile with replacement. Let $\mathcal{B} \in \mathcal{L}(\{a, b, c\})^\ell$ be the profile of sampled preferences
- 3: Let t be the minimum number of times any preference in $\mathcal{L}(\{a, b, c\})$ appear in \mathcal{B}
- 4: **if** $t < \frac{\ell}{12}(1 + \varepsilon_v)$ **then**
- 5: **return** 1
- 6: **else**
- 7: **return** 0
- 8: **end if**

THEOREM 3.1. *For at least 3 alternatives, there exists a ($\varepsilon_v, 0, \delta$, single peak) – RANDOM OUTLIERS VS RANDOM PROFILE TESTER with sample complexity $O(\frac{1}{(1-\varepsilon_v)^2} \log \frac{1}{\delta})$ for every $0 \leq \varepsilon_v < 1$ and $0 < \delta < 1/2$. If there are only 2 alternatives, then there does not exist any such tester.*

PROOF. For $m = 2$, the result follows from the observation that a profile where every preference is distributed uniformly in the set of all possible preferences is single peaked and thus the two cases are statistically indistinguishable. So let us assume $m \geq 3$ and a, b , and c be any three alternatives. We pick $\ell = \frac{72}{(1-\varepsilon_v)^2} \ln \frac{6}{\delta}$ preferences uniformly at random with replacement and query oracle to know how a, b , and c are ordered in these preferences. Let $p_i, i \in [6]$, be all possible permutations of $\{a, b, c\}$ and X_i be the random variable denoting the number of sampled preferences where the permutation p_i appears for $i \in [6]$. We output 1 if $\min_{i \in [6]} X_i < \frac{\ell}{12}(1 + \varepsilon_v)$ and output 0 otherwise. We observe that the sample complexity of our algorithm is $6\ell = O(\frac{1}{(1-\varepsilon_v)^2} \ln \frac{1}{\delta})$. We now turn to the correctness of our algorithm. For that we show that irrespective of the input profile, the probability of making an error is at most δ .

▷ **Case I - the input profile is single peaked after deleting at most $\varepsilon_v n$ preferences which are distributed uniformly:** Let \mathcal{P} be the input profile and \mathcal{Q} be a sub-profile of \mathcal{P} which is single peaked and contains at least $(1 - \varepsilon_v)n$ preferences. Hence, there exists an $\eta \in [6]$ such that the preference p_η does not appear in \mathcal{Q} . Since the preferences in $\mathcal{P} \setminus \mathcal{Q}$ is uniformly distributed and $|\mathcal{P} \setminus \mathcal{Q}| \leq \varepsilon_v n/6$, we have $\mathbb{E}[X_\eta] \leq \varepsilon_v \ell/6$. Using Chernoff bound (additive form), we now have the following:

$$\Pr[\text{error}] \leq \Pr[X_\eta \geq \frac{\ell}{12}(1 + \varepsilon_v)] \leq \exp\{-\frac{\ell(1-\varepsilon_v)^2}{72}\} \leq \delta$$

▷ **Case II - the input profile is distributed uniformly:** Since every preference in profile \mathcal{P} is uniformly distributed, for every $i \in [6]$, we have $\mathbb{E}[X_i] = \ell/6$. Using Chernoff bound (multiplicative

form) followed by union bound, we have the following:

$$\begin{aligned} \Pr[\text{error}] &= \Pr[\exists i \in [6], X_i \leq \frac{\ell}{12}(1 + \varepsilon_v)] \\ &\leq 6 \exp\{-\frac{\ell(1-\varepsilon_v)^2}{48}\} \leq \delta \quad \square \end{aligned}$$

The main idea in Theorem 3.1 can be easily extended to arbitrary domains.

COROLLARY 3.2. *Let \mathcal{D} be any normal and neutral domain and $m_0 = \min\{m : \text{con}_{\mathcal{D}}(m) < 1\}$. For at least m_0 alternatives, there exists a ($\varepsilon_v, 0, \delta, \mathcal{D}$) – RANDOM OUTLIERS VS RANDOM PROFILE TESTER with sample complexity $O(\frac{1}{(1-\varepsilon_v)^2} \ln \frac{1}{\delta})$ for every $0 \leq \varepsilon_v < 1$ and $0 < \delta < 1/2$. If the number of alternatives is at most $m_0 - 1$, then there does not exist any such tester.*

We now turn our attention to the ($\varepsilon_v, 0, \delta, \mathcal{D}$) – ARBITRARY OUTLIERS VS RANDOM PROFILE TESTER problem; that is when the outliers can be arbitrary (need not be randomly generated). We begin with presenting a general impossibility result in this case. Its proof follows from the observation that, in this case, one can carefully construct the set of outliers so that the distribution of samples in both the possibilities are statistically indistinguishable.

PROPOSITION 3.3. *For every domain \mathcal{D} , there does not exist any ($\varepsilon_v, 0, \delta, \mathcal{D}$) – ARBITRARY OUTLIERS VS RANDOM PROFILE TESTER for any $\varepsilon_v \geq \text{res}_{\mathcal{D}}(m)$ where m is the number of alternatives in the input profile.*

We now present our ($\varepsilon_v, 0, \delta$, single peak) – ARBITRARY OUTLIERS VS RANDOM PROFILE TESTER for $\varepsilon_v < 1/3$ in Theorem 3.4. We defer our general ($\varepsilon_v, 0, \delta, \mathcal{D}$) – ARBITRARY OUTLIERS VS RANDOM PROFILE TESTER till Theorem 3.6 which not only handles every $\varepsilon_v < 1$ but also takes care of arbitrary domain (but the sample complexity will be worse than that of Theorem 3.4). The main idea of the algorithm in Theorem 3.4 is exactly the same as the algorithm in Theorem 3.1 – it samples some preferences restricted to any 3 alternatives and outputs that the profile is random if all the 6 possible permutations appear nearly equal number of times; otherwise it says that the profile is close to single peaked.

THEOREM 3.4. *There exists a ($\varepsilon_v, 0, \delta$, single peak) – ARBITRARY OUTLIERS VS RANDOM PROFILE TESTER with sample complexity $O(\frac{1}{(1-3\varepsilon_v)^2} \ln \frac{1}{\delta})$ for every $0 \leq \varepsilon_v < 1/3$.*

PROOF. As in Theorem 3.1, we choose any 3 alternatives a, b , and c , pick $\ell = \frac{72}{(1-3\varepsilon_v)^2} \ln \frac{6}{\delta}$ preferences uniformly at random with replacement, and query oracle to know how a, b , and c are ordered in these preferences. We output 1 if $\min_{i \in [6]} X_i < \frac{\ell}{12}(1 + 3\varepsilon_v)$ and output 0 otherwise (with notation as defined in the proof of Theorem 3.1). The proof of correctness and the analysis of the sample complexity of our algorithm is similar to Theorem 3.1 using the observation that, when the input profile can be made single peaked by deleting at most $\varepsilon_v n$ preferences, there exists an $\eta \in [6]$ such that $X_\eta \leq \varepsilon_v \ell/2$ since $\text{res}_{\text{single peak}}(3) = 1/3$. \square

From the proof of Theorem 3.4, the following generalization to arbitrary domain is immediate.

COROLLARY 3.5. *Let $m_0 = \min\{m \in \mathbb{N} : \text{res}_{\mathcal{D}}(m) < 1\}$ and the number of alternatives is at least m_0 . Then there exists a ($\varepsilon_v, 0, \delta, \mathcal{D}$) – ARBITRARY OUTLIERS VS RANDOM PROFILE TESTER with sample*

complexity $O\left(\frac{1}{(1-\varepsilon_v/\text{res}_{\mathcal{D}}(m_0))^2} \ln \frac{1}{\delta}\right)$ for every ε_v with $0 \leq \varepsilon_v < \text{res}_{\mathcal{D}}(m_0)$ (the O notation in the sample complexity hides constant which depends on m_0).

We now present our $(\varepsilon_v, 0, \delta, \mathcal{D})$ – ARBITRARY OUTLIERS VS RANDOM PROFILE TESTER for any $\varepsilon_v < 1$ generalizing Theorem 3.4. Of course we need the number of alternatives to be at least $m(\varepsilon_v)$ where $m(\varepsilon_v) = \min\{m \in \mathbb{N} : \text{res}_{\mathcal{D}}(m) > \varepsilon_v\}$ due to Proposition 3.3.

THEOREM 3.6. *Given a domain \mathcal{D} , any ε_v with $0 \leq \varepsilon_v < 1$ with $\text{res}_{\mathcal{D}}(m) > \varepsilon_v$, there exists a $(\varepsilon_v, 0, \delta, \mathcal{D})$ – ARBITRARY OUTLIERS VS RANDOM PROFILE TESTER with sample complexity $O\left(\frac{m(\varepsilon_v)!m(\varepsilon_v)^2 \log^2 m(\varepsilon_v)}{(1-\varepsilon_v/\text{res}_{\mathcal{D}}(m(\varepsilon_v)))^2} \ln \frac{1}{\delta}\right)$ where $m(\varepsilon_v) = \min\{\ell \in \mathbb{N} : \text{res}_{\mathcal{D}}(\ell) > \varepsilon_v\}$.*

PROOF. Let $\mathcal{A}' \subseteq \mathcal{A}$ be any subset of alternatives with $|\mathcal{A}'| = m(\varepsilon_v)$. We pick $\ell = \frac{16m(\varepsilon_v)!m(\varepsilon_v) \log m(\varepsilon_v)}{(1-\varepsilon_v/\text{res}_{\mathcal{D}}(m(\varepsilon_v)))^2} \ln \frac{1}{\delta}$ preferences uniformly at random and elicit these preferences restricted to \mathcal{A}' . For $\succ \in \mathcal{L}(\mathcal{A}')$, let X_{\succ} be the random variable denoting the number of sampled preferences which are the same as \succ . We output 1 if $\min_{\succ \in \mathcal{L}(\mathcal{A}')} X_{\succ} < \frac{\ell}{2m(\varepsilon_v)!} (1 + (\varepsilon_v/\text{res}_{\mathcal{D}}(m(\varepsilon_v))))$ and output 0 otherwise. The sample complexity complexity of the algorithm is $O\left(\frac{m(\varepsilon_v)!m(\varepsilon_v)^2 \log^2 m(\varepsilon_v)}{(1-\varepsilon_v/\text{res}_{\mathcal{D}}(m(\varepsilon_v)))^2} \ln \frac{1}{\delta}\right)$. The proof of correctness of our algorithm is similar to that of Theorem 3.1 using the observation that, when the input profile can be made single peaked by deleting at most $\varepsilon_v n$ preferences, there exists an $\succ \in \mathcal{L}(\mathcal{A}')$ such that $X_{\succ} \leq \frac{\varepsilon_v \ell}{\text{res}_{\mathcal{D}}(m(\varepsilon_v))m(\varepsilon_v)!}$ (follows from the definition of $\text{res}_{\mathcal{D}}(m(\varepsilon_v))$). \square

We now present our result for the $(\varepsilon_v, 0, \varepsilon'_v, 0, \delta, \mathcal{D})$ – ARBITRARY OUTLIERS VS ARBITRARY PROFILE TEST problem. The following structural result provides the key building block of our algorithm. Intuitively the lemma proves that, given a profile \mathcal{P} , if we sample preferences from \mathcal{P} uniformly at random with replacement to construct another profile \mathcal{Q} (of certain size), then the “relative” distance of \mathcal{Q} from any domain \mathcal{D} is approximately same as the relative distance of \mathcal{P} from \mathcal{D} .

LEMMA 3.7. *Let \mathcal{D} be any normal and neutral domain and $(\succ_i)_{i \in [n]} \in \mathcal{L}(\mathcal{A})^n$ be a profile with preference-distance being $\varepsilon_v n$ from \mathcal{D} . Let $0 < \Delta < \min\{\varepsilon_v, 1 - \varepsilon_v\}$, $\ell = \frac{4}{\Delta^2} (\text{con}_{\mathcal{D}}(m)m \ln m + \ln 1/\delta)$, and $\succ'_i = (\succ'_i)_{i \in [\ell]}$ be a profile where \succ'_i has been picked uniformly at random with replacement from the n preferences of \succ . Then the preference-distance of \succ' from \mathcal{D} is at least $(\varepsilon_v - \Delta)\ell$ and at most $(\varepsilon_v + \Delta)\ell$ with probability at least $1 - \delta$ for every $0 < \delta < 1$.*

We now present our $(\varepsilon_v, 0, \varepsilon'_v, 0, \delta, \mathcal{D})$ – ARBITRARY OUTLIERS VS ARBITRARY PROFILE TESTER. The high level idea is to sample some number ℓ of preferences, compute the distance $\varepsilon'' \ell$ of the resulting profile from the single peaked domain, and output the distance of the original profile to be εn if and only if ε'' is closer to ε than ε' .

THEOREM 3.8. *For every domain \mathcal{D} , there exists a $(\varepsilon_v, 0, \varepsilon'_v, 0, \delta, \mathcal{D})$ – ARBITRARY OUTLIERS VS ARBITRARY PROFILE TESTER with sample complexity $O\left(\frac{1}{(\varepsilon'_v - \varepsilon_v)^2} (\text{con}_{\mathcal{D}}(m)m!m^2 \log^2 m + \log 1/\delta)\right)$ for every $0 \leq \varepsilon_v < \varepsilon'_v < 1$ and $0 < \delta < 1/2$.*

Algorithm 2 $(\varepsilon_v, 0, \varepsilon'_v, 0, \delta, \text{single peak})$ – ARBITRARY OUTLIERS VS ARBITRARY PROFILE TESTER

Input: Oracle access to a profile \mathcal{P}

Output: 1 if there exists $\varepsilon_v n$ preferences whose deletion makes the resulting profile single peaked and 0 if deleting any $\varepsilon'_v n$ preferences from \mathcal{P} does not make the resulting profile single peaked

- 1: Sample $\ell = \frac{64}{(\varepsilon'_v - \varepsilon_v)^2} (2^m m \ln m + \ln 1/\delta)$ preferences uniformly at random from the input profile with replacement. Let $\mathcal{B} \in \mathcal{L}(\mathcal{A})^\ell$ be the profile of sampled preferences
 - 2: Let t be the minimum number of times any preference in $\mathcal{L}(\{a, b, c\})$ appear in \mathcal{B}
 - 3: **if** \mathcal{B} can be made single peaked by deleting at most $(\varepsilon_v + \varepsilon'_v)\ell/2$ preferences **then**
 - 4: **return** 1
 - 5: **else**
 - 6: **return** 0
 - 7: **end if**
-

We observe that $\text{con}_{\text{single peaked}}(m) = 2^{m-1}/m!$ [29, Lemma 2] and $\text{con}_{\text{single crossing}}(m) = \binom{m}{2} + 1/m!$. Hence, from Theorem 3.8, we obtain the following result for the single peaked and single crossing domains.

COROLLARY 3.9. *There exists a $(\varepsilon_v, 0, \varepsilon'_v, 0, \delta, \mathcal{D})$ – ARBITRARY OUTLIERS VS ARBITRARY PROFILE TESTER with sample complexity $O\left(\frac{1}{(\varepsilon'_v - \varepsilon_v)^2} (2^m m^2 \log^2 m + \log 1/\delta)\right)$ for the single peaked domain and with sample complexity $O\left(\frac{1}{(\varepsilon'_v - \varepsilon_v)^2} (m^4 \log^2 m + \log 1/\delta)\right)$ for the single crossing domain for every $0 \leq \varepsilon_v < \varepsilon'_v < 1$ and $0 < \delta < 1/2$.*

3.2 Only Alternatives as Outliers

In this subsection, we now focus on the case when only alternatives are considered as outliers. We observe that when only alternatives act as outliers, the $(0, \varepsilon_a, \delta, \mathcal{D})$ – RANDOM OUTLIERS VS RANDOM PROFILE TEST and $(0, \varepsilon_a, \delta, \mathcal{D})$ – ARBITRARY OUTLIERS VS RANDOM PROFILE TEST are the same problem. We begin with presenting our $(0, \varepsilon_a, \delta, \text{single peak})$ – RANDOM OUTLIERS VS RANDOM PROFILE TESTER in Theorem 3.10 below. On a high level, our algorithm in Theorem 3.10 samples some number t of preferences restricted to some number ℓ of alternatives. If for every 3 alternatives among those ℓ alternatives, all the 6 possible permutations appear in the sampled preferences, then the algorithm outputs the profile to be random, otherwise it says that the profile is close to being single peaked.

THEOREM 3.10. *There exists a $(0, \varepsilon_a, \delta, \text{single peak})$ – RANDOM OUTLIERS VS RANDOM PROFILE TESTER with sample complexity $O\left(\log \frac{\log_{1/\varepsilon_a} 1/\delta}{\delta} \log_{1/\varepsilon_a} \frac{1}{\delta} \log \log_{1/\varepsilon_a} 1/\delta\right)$. Hence, there also exists a $(0, \varepsilon_a, \delta, \text{single peak})$ – ARBITRARY OUTLIERS VS RANDOM PROFILE TESTER with the same sample complexity for every $0 < \varepsilon_a < 1$ and $0 < \delta < 1/2$ such that $\text{con}_{\text{single peak}}((1 - \varepsilon_a)m) < 1$.*

PROOF. We sample $\ell = \min\{(1 - \varepsilon_a)m, 2 \log_{1/\varepsilon_a} 1/\delta\}$ alternatives uniformly at random without replacement. Let \mathcal{B} be the set of sampled alternatives. We now sample $t = 18 \ln \frac{2 \log_{1/\varepsilon_a} 1/\delta}{\delta}$ preferences uniformly at random with replacement restricted to \mathcal{B} . Let

Algorithm 3 ($0, \varepsilon_a, \delta$, single peak) – RANDOM OUTLIERS VS RANDOM PROFILE TESTER**Input:** Oracle access to a profile \mathcal{P}

Output: 1 if there exists $\varepsilon_a m$ alternatives whose deletion makes the resulting profile single peaked and 0 if \mathcal{P} has been generated randomly

- 1: Sample $\ell = \min\{(1 - \varepsilon_a)m, 2 \log_{1/\varepsilon_a} 1/\delta\}$ alternatives uniformly at random from \mathcal{A} without replacement. Let \mathcal{B} be the set of sampled alternatives.
- 2: Sample $t = 18 \ln \frac{2 \log_{1/\varepsilon_a} 1/\delta}{\delta}$ uniformly random preferences restricted to \mathcal{B} . Let the sampled profile be $Q \in \mathcal{L}(\mathcal{B})^t$
- 3: **for** Every distinct $a, b, c \in \mathcal{B}$ **do**
- 4: **if** at least one permutation in $\mathcal{L}(\{a, b, c\})$ is not present in Q **then**
- 5: **return** 1
- 6: **end if**
- 7: **end for**
- 8: **return** 0

Q be the set of sampled preferences. We output 1 if there exist 3 alternatives $a, b, c \in \mathcal{B}$ such that at least one permutation in $\mathcal{L}(\{a, b, c\})$ is not present in Q and output 0 otherwise. We observe that the sample complexity of our algorithm is $O(t\ell \log \ell) = O(\log \frac{\log_{1/\varepsilon_a} 1/\delta}{\delta} \log_{1/\varepsilon_a} \frac{1}{\delta} \log \log_{1/\varepsilon_a} 1/\delta)$. We now turn to the correctness of our algorithm. For that we show that irrespective of the input profile, the probability of making an error is at most δ .

▷ **Case I - the input profile is single peaked after deleting at most $\varepsilon_a m$ alternatives:** Let \mathcal{A} be the set of alternatives and $\mathcal{W} \subset \mathcal{A}$ with $|\mathcal{W}| \leq \varepsilon_a m$ such that the input profile restricted to $(\mathcal{A} \setminus \mathcal{W})$ is single peaked. Then we have the following for the chosen value of ℓ :

$$\begin{aligned} \Pr[\text{error}] &\leq \Pr[|\mathcal{B} \cap \mathcal{W}| \geq \ell - 2] \\ &= \varepsilon_a^\ell + \binom{\ell}{1}(1 - \varepsilon_a)\varepsilon_a^{\ell-1} + \binom{\ell}{2}(1 - \varepsilon_a)^2\varepsilon_a^{\ell-2} \\ &\leq \varepsilon_a^\ell + \ell\varepsilon_a^{\ell-1} + \ell^2\varepsilon_a^{\ell-2} \leq \delta \end{aligned}$$

▷ **Case II - the input profile has been generated uniformly at random:** For any 3 alternatives $a, b, c \in \mathcal{A}$, we define a random variable $X_{\{a, b, c\}}$ to be 1 if all 6 possible permutations in $\mathcal{L}(\{a, b, c\})$ are present in $Q(\{a, b, c\})$ and 0 otherwise. Using folklore tail bound for the *coupon collector problem* (for example, see [42, Chap 3.6]), we obtain the following for the chosen value of t .

$$\Pr[X_{\{a, b, c\}} = 0] \leq 6^{-t/3 \ln 6} \leq e^{-t/6}$$

Now using union bound, we obtain the following for the chosen values of ℓ and t .

$$\Pr[\text{error}] \leq \Pr[\exists \{a, b, c\} \subset \mathcal{B}, X_{\{a, b, c\}} = 0] \leq \binom{\ell}{3} e^{-\frac{t}{6}} \leq \delta \quad \square$$

From the proof of Theorem 3.10, Corollary 3.11 follows.

COROLLARY 3.11. *For every domain \mathcal{D} , there exists a $(0, \varepsilon_a, \delta, \mathcal{D})$ – ARBITRARY OUTLIERS VS RANDOM PROFILE TESTER with sample complexity $O(\log \frac{\log_{1/\varepsilon_a} 1/\delta}{\delta} \log_{1/\varepsilon_a} \frac{1}{\delta} \log \log_{1/\varepsilon_a} 1/\delta)$ for every $0 < \varepsilon_a < 1$ and $0 < \delta < 1/2$ such that $\text{con}_{\mathcal{D}}((1 - \varepsilon_a)m) < 1$.*

We show below that the condition $\text{con}_{\mathcal{D}}((1 - \varepsilon_a)m) < 1$ in Theorem 3.10 and Corollary 3.11 is necessary. We prove Proposition 3.12 by carefully constructing a set of outliers such the the sample distribution in both the possibilities are statistically indistinguishable.

PROPOSITION 3.12. *For every domain \mathcal{D} , there does not exist any $(0, \varepsilon_a, \delta, \mathcal{D})$ – ARBITRARY OUTLIERS VS RANDOM PROFILE TESTER if $\text{con}_{\mathcal{D}}((1 - \varepsilon_a)m) = 1$.*

We now turn to the $(0, 0, 0, \varepsilon'_a, \delta, \text{single peak})$ – ARBITRARY OUTLIERS VS ARBITRARY PROFILE TESTER problem. The following results show that the sample complexity of this problem is $\Omega(n \log 1/\delta)$ even for the single peaked and single crossing domains.

THEOREM 3.13. *Any $(0, 0, 0, \varepsilon'_a, \delta, \text{single peak})$ – ARBITRARY OUTLIERS VS ARBITRARY PROFILE TESTER has sample complexity $\Omega(n \log 1/\delta)$ for every $0 < \varepsilon'_a \leq 1$ and $0 < \delta < 1/2$ such that $\text{con}_{\text{single peak}}((1 - \varepsilon'_a)m) < 1$.*

THEOREM 3.14. *Any $(0, 0, 0, \varepsilon'_a, \delta, \mathcal{D})$ – ARBITRARY OUTLIERS VS ARBITRARY PROFILE TESTER has sample complexity $\Omega(n \log 1/\delta)$ for single crossing domain for every $0 < \varepsilon'_a \leq 1$ and $0 < \delta < 1/2$ such that $\text{con}_{\text{single crossing}}((1 - \varepsilon'_a)m) < 1$.*

4 EMPIRICAL EVALUATION

The algorithms presented in Section 3 provide upper bounds on the sample complexities of the problems of outlier detection. These algorithms distinguish between two possibilities of profile generation with a probability of correctness of at least $(1 - \delta)$. It is interesting to find out the optimal multiplying factors of the sampling complexities inside $O(\cdot)$ in these algorithms. This is why an empirical evaluation is called for.

In this section, we empirically find the factors for the results of Theorems 3.1, 3.4 and 3.10, which provide constant time algorithms for the testing problem. The other two cases as shown in Table 1 either consider an exponential time (Corollary 3.9) algorithm or provide a lower bound (Theorem 3.13), which are unsuitable for an empirical study.

4.1 Approach for Theorems 3.1 and 3.4:

We generate $n = 10,000$ preferences with m alternatives uniformly at random to form a preference profile. The sampling algorithm of Theorem 3.1 (given by Algorithm 1) picks an ℓ for a given ε_v . In this experiment, we choose a sampling size l that is smaller than ℓ , and apply the same algorithm using l preferences sampled with replacement from the population of n . We generate the preference profile 100 times and for every profile, sample l preferences 100 times. We consider the fraction of correct classifications given by this modified sampling algorithm and plot it with increasing l . We fix $\delta = 0.001$ for these evaluations. We show the plot of the fraction of correct classification (denoted by ρ) for Theorem 3.1 with $m = 3$ in Figure 1.

The plot shows the growth of the empirical probability of correctness (and therefore does not need any errorbar). The x-axis shows the normalized sample size (that is l/ℓ). Notice that the growth of the curves almost overlaps for different ε_v s, and reaches $(1 - \delta)$ nearly at 0.5. This empirically shows that when other parameters are held fixed at the chosen values, the hidden constant in the upper

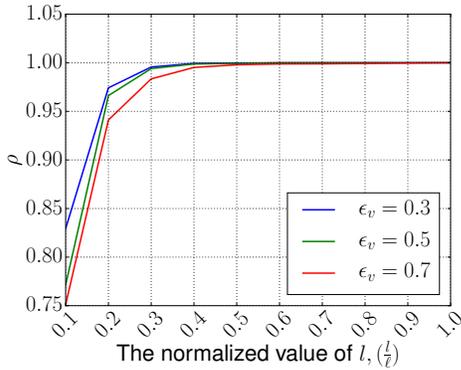


Figure 1: Fraction of correct classification (ρ) of the adaptation of Algorithm 1 when $l(\leq \ell)$ preferences have been sampled uniformly at random from a random preference profile of size $n = 10,000$, $\delta = 0.001$ (x-axis shows the normalized value, l/ϵ_v).

bound of the sample complexity in the context of random outliers can be reduced by almost 50%, and is independent of ϵ_v .

We perform a similar exercise with different sampling sizes for the algorithm in the proof of Theorem 3.4 (given by Algorithm 2) with $m = 5$ in Figure 2. Here too, the proportionality factor is independent of the ϵ_v s, and the hidden constant factor in this case can be reduced by 60%.

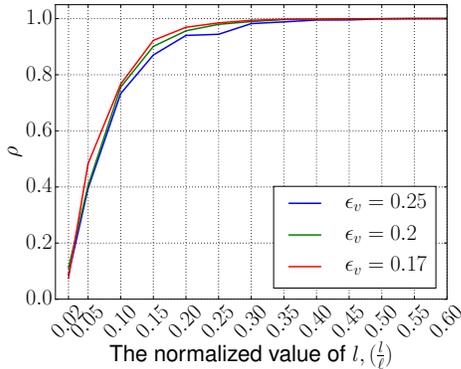


Figure 2: Fraction of correct classification (ρ) of the adaptation of Algorithm 2 when $l(\leq \ell)$ preferences have been sampled uniformly at random from a random preference profile of size $n = 10,000$, $\delta = 0.001$ (x-axis shows the normalized value, l/ϵ_v).

Why the error with a random/arbitrary outliers profile being classified as a random profile is not considered? We argue that such an error is not very likely in the algorithms of these theorems, which is also manifested in our simulations. Therefore we omit them presenting here. For Theorem 3.1, since the focus is only on the three alternatives a, b , and c , the number of random outliers will be close to $\epsilon_v n/6$ for large enough n . If l preferences are drawn uniformly at random with replacement from this profile, it is very likely that $\min_i X_i$ will be at most close to $\epsilon_v l/6$ for reasonably sized l . The algorithm classifies the profile as random outlier profile if $\min_i X_i \leq l(1+\epsilon_v)/12$ and since $\epsilon_v l/6 \leq l(1+\epsilon_v)/12$, it is unlikely that

a random outlier profile will be classified as random profile under this algorithm. Similar observation is true for Theorem 3.4.

4.2 Approach for Theorem 3.10:

Here we consider the alternatives as outliers. The algorithm in the proof of this theorem (given by Algorithm 3) samples ℓ alternatives uniformly at random and samples t preferences restricted to the sampled alternatives uniformly at random. In this case, we pick the values of δ and n as before. We fix $m = 9$, and pick $\ell = \min\{(1 - \epsilon_a)m, 2 \log_{1/\epsilon_a} 1/\delta\}$ as given in the proof of Theorem 3.10, and vary the value of $\tau(\leq t)$, which is the sampling size of the preferences restricted to the chosen ℓ alternatives. The alternatives of size ℓ are sampled 100 times. Figure 3 shows the plot of the fraction of correct classification (ρ) under this setting. It empirically shows that when other parameters are held fixed at the chosen values, the hidden constant of the upper bound of the probability in the case of random alternative outliers can be reduced by almost 75%, and is independent of ϵ_a .

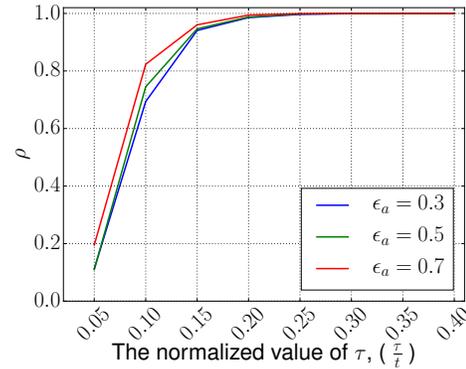


Figure 3: Fraction of correct classification (ρ) of the adaptation of Algorithm 3 when $\tau(\leq t)$ preferences have been sampled uniformly at random from a random preference profile of size $n = 10,000$ with $\ell = \min\{(1 - \epsilon_a)m, 2 \log_{1/\epsilon_a} 1/\delta\}$ as given in the proof, $m = 10$, $\delta = 0.001$ (x-axis shows the normalized value, τ/ϵ_a).

In a way similar to the previous paragraph we can argue that this algorithm also has a bias towards classifying a profile as random alternative outlier, which also is empirically manifested. Hence, we omit presenting them here.

5 DISCUSSION

In this paper, we have developed sampling based algorithms for testing if a profile is close to some specific domain. These testing problem can be quite accurately solved by observing a small number of samples for most of the cases, and the numbers are often independent to the number of preferences or alternatives. In other cases, we have proved impossibility results. Our extensive empirical study further improve the constants of the asymptotic theoretical upper bounds on the sample complexity by 50% to 75% depending on the problem. As a future work, there exist more sophisticated notion of distances, namely swap distance, footrule distance, maximum displacement distance, etc. where it will be interesting to extend our results to those fine grained measures of distance.

REFERENCES

- [1] Donald WK Andrews, Xuemei Liu, and Werner Ploberger. 1998. Tests for white noise against alternatives with both seasonal and nonseasonal serial correlation. *Biometrika* 85, 3 (1998), 727–740.
- [2] Kenneth J Arrow. 1950. A difficulty in the concept of social welfare. *J. Polit. Econ.* (1950), 328–346.
- [3] Miguel A Ballester and Guillaume Haeringer. 2011. A characterization of the single-peaked domain. *Soc. Choice Welf.* 36, 2 (2011), 305–322.
- [4] Ziv Bar-Yossef. 2003. Sampling lower bounds via information theory. In *Proc. 35th Annual ACM Symposium on Theory of Computing (STOC)*, 335–344. <https://doi.org/10.1145/780542.780593>
- [5] Salvador Barberà and Bernardo Moreno. 2011. Top monotonicity: A common root for single peakedness, single crossing and the median voter result. *Games Econ. Behav.* 73, 2 (2011), 345–359.
- [6] John Bartholdi III and Michael A Trick. 1986. Stable matching with preferences derived from a psychological model. *Oper. Res. Lett.* 5, 4 (1986), 165–169.
- [7] Arnab Bhattacharyya, Palash Dey, and David P. Woodruff. 2016. An Optimal Algorithm for l_1 -Heavy Hitters in Insertion Streams and Related Problems. In *Proc. 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS '16)*. San Francisco, California, USA, 385–400. <https://doi.org/10.1145/2902251.2902284>
- [8] Felix Brandt, Markus Brill, Edith Hemaspaandra, and Lane A Hemaspaandra. 2015. Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. *J. Artif. Intell. Res.* (2015), 439–496.
- [9] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. 2016. *Handbook of Computational Social Choice*. Cambridge University Press.
- [10] Robert Bredereck, Jiehua Chen, and Gerhard J Woeginger. 2013. A characterization of the single-crossing domain. *Soc. Choice Welf.* 41, 4 (2013), 989–998.
- [11] Robert Bredereck, Jiehua Chen, and Gerhard J. Woeginger. 2016. Are there any nicely structured preference profiles nearby? *Math. Soc. Sci.* 79 (2016), 61–73. <https://doi.org/10.1016/j.mathsocsci.2015.11.002>
- [12] Shurojit Chatterji, Arunava Sen, and Huaxia Zeng. 2016. A characterization of single-peaked preferences via random social choice functions. *Theor. Econ.* 11, 2 (2016), 711–733.
- [13] Jiehua Chen and Ugo Paavo Finnendahl. 2018. On the number of single-peaked narcissistic or single-crossing narcissistic preference profiles. *Discrete Math.* 341, 5 (2018), 1225–1236. <https://doi.org/10.1016/j.disc.2018.01.008>
- [14] Vincent Conitzer. 2009. Eliciting Single-Peaked Preferences Using Comparison Queries. *J. Artif. Intell. Res.* 35 (2009), 161–191.
- [15] Palash Dey. 2017. Query Complexity of Tournament Solutions. In *Proc. 31st AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA, 2992-2998*. <http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14180>
- [16] Palash Dey and Arnab Bhattacharyya. 2015. Sample Complexity for Winner Prediction in Elections. In *Proc. 2015 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2015, Istanbul, Turkey, May 4-8, 2015, 1421-1430*. <http://dl.acm.org/citation.cfm?id=2773334>
- [17] Palash Dey and Neeldhara Misra. 2016. Elicitation for Preferences Single Peaked on Trees. In *Proc. 25th International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016, 215-221*. <http://www.ijcai.org/Abstract/16/038>
- [18] Palash Dey and Neeldhara Misra. 2016. Preference Elicitation for Single Crossing Domain. In *Proc. Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI)*, 222–228. <http://www.ijcai.org/Abstract/16/039>
- [19] Palash Dey and Neeldhara Misra. 2016. Preference Elicitation for Single Crossing Domain. In *Proc. 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 222–228. <http://www.ijcai.org/Abstract/16/039>
- [20] Palash Dey and Y. Narahari. 2015. Estimating the Margin of Victory of an Election Using Sampling. In *Proc. 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 1120–1126. <http://ijcai.org/Abstract/15/162>
- [21] Palash Dey, Swaprava Nath, and Garima Shakya. 2019. Testing Preferential Domains Using Sampling. *CoRR* abs/1902.08930 (2019). <http://arxiv.org/abs/1902.08930>
- [22] Palash Dey, Nimrod Talmon, and Otniel van Handel. 2017. Proportional Representation in Vote Streams. In *Proc. 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2017, São Paulo, Brazil, May 8-12, 2017, 15-23*. <http://dl.acm.org/citation.cfm?id=3091134>
- [23] Jean-Paul Doignon and Jean-Claude Falmagne. 1994. A Polynomial Time Algorithm for Unidimensional Unfolding Representations. *J. Algorithms* 16, 2 (1994), 218–233. <https://doi.org/10.1006/jagm.1994.1010>
- [24] Edith Elkind and Piotr Faliszewski. 2014. Recognizing 1-Euclidean Preferences: An Alternative Approach. In *Proc. 7th International Symposium on Algorithmic Game Theory (SAGT)*, 146–157. https://doi.org/10.1007/978-3-662-44803-8_13
- [25] Edith Elkind, Piotr Faliszewski, Martin Lackner, and Svetlana Obraztsova. 2015. The Complexity of Recognizing Incomplete Single-Crossing Preferences. In *Proc. 29th AAAI Conference on Artificial Intelligence (AAAI)*, 865–871.
- [26] Edith Elkind, Piotr Faliszewski, and Piotr Skowron. 2014. A Characterization of the Single-Peaked Single-Crossing Domain. In *Proc. 28th AAAI Conference on Artificial Intelligence (AAAI)*, 654–660.
- [27] Edith Elkind, Piotr Faliszewski, and Arkadii M. Slinko. 2012. Clone structures in voters' preferences. In *Proc. 13th ACM Conference on Electronic Commerce (EC)*, 496–513. <https://doi.org/10.1145/2229012.2229050>
- [28] Gábor Erdélyi, Martin Lackner, and Andreas Pfandler. 2017. Computational Aspects of Nearly Single-Peaked Electorates. *J. Artif. Intell. Res.* 58 (2017), 297–337. <https://doi.org/10.1613/jair.5210>
- [29] Bruno Escoffier, Jérôme Lang, and Meltem Öztürk. 2008. Single-peaked consistency and its complexity. In *Proc. 18th European Conference on Artificial Intelligence (ECAI)*, 366–370. <https://doi.org/10.3233/978-1-58603-891-5-366>
- [30] Aris Filos-Ratsikas, Mingming Li, Jie Zhang, and Qiang Zhang. 2017. Facility Location with Double-peaked Preferences. *Auton. Agents Multi Agent Syst* 31, 6 (2017), 1209–1235.
- [31] Zack Fitzsimmons. 2015. Single-Peaked Consistency for Weak Orders Is Easy. In *Proc. 15th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 127–140. <https://doi.org/10.4204/EPTCS.215.10>
- [32] Wulf Gaertner. 2001. *Domain Conditions in Social Choice Theory*. Cambridge University Press.
- [33] A. Gibbard. 1973. Manipulation of Voting Schemes: a General Result. *Econometrica* (1973), 587–601.
- [34] Oded Goldreich. 1999. Combinatorial property testing (a survey). *Randomization Methods in Algorithm Design* 43 (1999), 45–59.
- [35] Oded Goldreich, Shafi Goldwasser, and Dana Ron. 1998. Property Testing and its Connection to Learning and Approximation. *J. ACM* 45, 4 (1998), 653–750. <https://doi.org/10.1145/285055.285060>
- [36] Melvin J Hinich and Michael C Munger. 1997. *Analytical politics*. Cambridge University Press.
- [37] Vicki Knoblauch. 2010. Recognizing one-dimensional Euclidean preference profiles. *J. Math. Econ.* 46, 1 (2010), 1–5.
- [38] Solomon Kullback and Richard A Leibler. 1951. On information and sufficiency. *The annals of mathematical statistics* 22, 1 (1951), 79–86.
- [39] Martin Lackner. 2014. Incomplete Preferences in Single-Peaked Electorates. In *Proc. 28th AAAI Conference on Artificial Intelligence (AAAI)*, 742–748.
- [40] Marie-Louise Lackner and Martin Lackner. 2017. On the likelihood of single-peaked preferences. *Soc. Choice Welf.* 48, 4 (2017), 717–745.
- [41] Krzysztof Magiera and Piotr Faliszewski. 2017. Recognizing Top-Monotonic Preference Profiles in Polynomial Time. In *Proc. 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 324–330. <https://doi.org/10.24963/ijcai.2017/46>
- [42] Rajeev Motwani and Prabhakar Raghavan. 2010. *Randomized algorithms*. Chapman & Hall/CRC.
- [43] Hervé Moulin. 1991. *Axioms of cooperative decision making*. Number 15. Cambridge University Press.
- [44] David M. Pennock, Eric Horvitz, and C. Lee Giles. 2000. Social Choice Theory and Recommender Systems: Analysis of the Axiomatic Foundations of Collaborative Filtering. In *Proc. 17th National Conference on Artificial Intelligence and 12th Conference on Innovative Applications of Artificial Intelligence*, 729–734. <http://www.aaai.org/Library/AAAI/2000/aaai00-112.php>
- [45] Dominik Peters. 2017. Recognising Multidimensional Euclidean Preferences. In *Proc. 31st AAAI Conference on Artificial Intelligence (AAAI)*, 642–648. <http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14947>
- [46] Dana Ron. 2001. Property testing. *Comb. Opt.* 9, 2 (2001), 597–643.
- [47] M.A. Satterthwaite. 1975. Strategy-proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions. *J. Econ. Theory* 10, 2 (1975), 187–217.
- [48] Xiaofeng Shao. 2011. Testing for white noise under unknown dependence and its applications to diagnostic checking for time series models. *Econ. Theory* 27, 2 (2011), 312–343.