

Designing Committees for Mitigating Biases*

Michal Feldman,^{1,2} Yishay Mansour,^{1,3} Noam Nisan,⁴
Sigal Oren,⁵ Moshe Tennenholtz⁶

¹Tel-Aviv University, Israel, ²Microsoft Research, Israel, ³Google Research, Israel, ⁴Hebrew University, Israel

⁵Ben-Gurion University of the Negev, Israel, ⁶Technion — Israel Institute of Technology, Israel

michal.feldman@cs.tau.ac.il, mansour.yishay@gmail.com, noam@cs.huji.ac.il, orensi@cs.bgu.ac.il, moshet@ie.technion.ac.il

Abstract

It is widely observed that individuals prefer to interact with others who are more similar to them (this phenomenon is termed homophily). This similarity manifests itself in various ways such as beliefs, values and education. Thus, it should not come as a surprise that when people make hiring choices, for example, their similarity to the candidate plays a role in their choice. In this paper, we suggest that putting the decision in the hands of a committee instead of a single person can reduce this bias.

We study a novel model of voting in which a committee of experts is constructed to reduce the biases of its members. We first present voting rules that optimally reduce the biases of a given committee. Our main results include the design of committees, for several settings, that are able to reach a nearly optimal (unbiased) choice. We also provide a thorough analysis of the trade-offs between the committee size and the obtained error. Our model is inherently different from the well-studied models of voting that focus on aggregation of preferences or on aggregation of information due to the introduction of similarity biases.

1 Introduction

Homophily – the tendency of people to prefer others who are more similar to them – was long studied in several disciplines including sociology and economics (see (McPherson, Smith-Lovin, and Cook 2001) and references within as well as (Currarini, Jackson, and Pin 2009)). The similarity can be over different dimensions varying from values and beliefs to social status and gender. Usually homophily is discussed in the context of structure of social networks – individuals choose friends or spouses based on similarity. However, it was observed that homophily also plays a role in other settings, such as hiring (Rivera 2012). When the decision is of a personal nature (e.g., choosing a spouse) this bias towards similar individuals is natural and expected. However, when

the decision has a more public nature (e.g., hiring for a firm or for an academic department), mitigating this bias can be valuable. As we will see in this paper, entrusting decisions such as who to hire in the hands of a committee may help in reducing this bias.

Committees are chosen to make decisions for various reasons. Sometimes individuals have different preferences and the committee is formed to make a decision that will maximize the individuals’ “happiness” (“aggregation of preferences”). Other times there is a true state of the world which individuals have noisy signals about and a planner wishes to discover what this true state is (“aggregation of information”). Here we suggest that there are some settings in which there is an additional benefit for using a committee instead of a single decision maker – that is to reduce individual biases. To further study the advantages of a committee when the individuals are biased, we devise a simple model in which a planner wishes to learn the true state of the world. To this end he organizes a committee that consists of experts that have biased views of the true state of the world. By utilizing some information on the experts’ biases the planner is able to extract a less biased estimate of the true state.

A concrete example of such a setting, from which we will also borrow our terminology, is academic hiring. Consider a dean that needs to form a hiring committee to choose among several different candidates, each working in a different sub-area of some academic field. In our setting, the dean, who is not from the academic field in question, has no information about the “quality” of the candidates. All faculty members in the academic area – the “experts” – have estimates of the candidates qualities that are shaped by the “true” qualities of the candidates and their own cognitive biases. Specifically, faculty members prefer candidates that are more similar to them and as a result underestimate the qualities of candidates in sub-fields that are far from their own sub-field. The dean, who is aware of these biases, will thus naturally attempt to construct a committee that is “diverse” in terms of its biases as to enable him to negate some of these biases and determine who the highest quality candidate is.

To model the similarities between the experts and candidates we assume that the candidates and experts reside in known locations in some metric space¹. The similar-

*This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 740282 and 740435), the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) (grant agreement No 337122), and by the Israel Science Foundation (grant numbers 317/17, 993/17 and 2167/19).
Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹This is in line with to the way similarity between voters and

ity between a candidate and an expert is determined by the distance between them. The two main metric spaces we consider are the unit interval and the more general d -dimensional unit hyper-cube. The former can model, for example, how theoretical their research is, while the latter can capture several aspects of their research simultaneously.

We assume that each candidate j has an absolute quality $q_j \in \mathbb{R}$. To model the cognitive process by which experts decide which candidate to vote for we assume that an expert i that resides at location e_i in the metric space assigns a grade $g_i(j) = q_j - d(e_i, c_j)$ to candidate j that resides in location c_j (where $d(e_i, c_j)$ is the distance between the two locations). Next, each expert votes for the candidate with the maximal grade: $vote_i = \arg \max_{j \in C} g_i(j)$, where C is the set of candidates. We stress that the grade an expert gives a candidate is not necessarily a part of some conscious computation but rather a result of biases that the expert might be unaware of. Given a committee E of experts, the dean that is aware of the locations of the experts and their voting process will aggregate the votes $vote_i$ of all committee members, and will need to pick some candidate j^* . Due to the experts' biases the dean will not always be able to find the candidate that truly maximizes q_{j^*} among all q_j , so we will aim to minimize the dean's *regret*²: $\max_{j \in C} q_j - q_{j^*}$.

To familiarize the reader with our model, let us inspect a few simple examples. In all three examples depicted in Figure 1 the candidates are marked by circles, the experts by squares and there is an arrow from each expert to the candidate he votes for. First, in Figure 1(a) we observe that if we only have two candidates, an expert located at equal distances from both will always vote for the candidate of higher quality. In Figure 1(b) we see that for three candidates, having an expert at equal distances from each pair of candidates does not necessarily suffice to find the best candidate. This is due to the fact that the expert located at 0 might prefer the candidate at 0 instead of voting for the better of the candidates located at -1 and $+1$. Lastly, in Figure 1(c), we see that combining information from several experts may suffice for determining the best candidate. In the example the expert at -0.4 vote implies that the quality of the candidate at 0 is at most 0.2 less than the quality of the candidate at -1 . This is because the distance between the expert at -0.4 and the candidate at 0 is smaller by 0.2 than the distance between him and the candidate at -1 . The expert at $+0.4$ vote implies that the candidate at 1 is better by at least 0.2 from the candidate at 0. Hence, the best candidate is the one at 1.

Our Results

We start our investigation by considering given sets of candidates and experts and ask which voting method should we use to determine who the best candidate is. In the example in Figure 1(c) we have already seen that this can require combining the experts' votes in a non-trivial way. We show that

political parties is modeled in political science (Enelow and Hinich 1989; Roemer 2001).

²Our definition of regret differs from the online learning term regret, although conceptually they both address similar sub-optimality issues.




Figure 1: Examples of simple committees.

majority voting, or more generally, plurality voting, can be extremely bad for some sets of candidates. We then identify the regret-minimizing voting method, and show how it can be implemented efficiently using a set of shortest-path computations.

Next, we study the question of committee size: how large must a committee be in order to achieve a given level of regret? There are two variants of this problem: the first is the “universal committee” in which our committee must be chosen before the locations of the set of candidates are known, and the second is the “ad-hoc committee” that is chosen for a set of candidates with known locations. In our academic hiring setting, the first “universal committee” describes the scenario in which the dean chooses a hiring committee before the applications arrive, while the second “ad-hoc committee” describes a scenario in which a hiring committee is chosen after the subareas of the candidates are known. When both the candidates and the experts are located on the unit interval, we are able to give an essentially complete picture:

Theorem 1.1 *In the unit interval, for every $\varepsilon > 0$ and candidate set of size m , there exists a universal committee of size $O(m/\varepsilon)$ that achieves regret of at most ε . The bound is asymptotically tight even for ad-hoc committees.*

The universal committee is constructed by equally spacing the experts across the unit-interval. It might come as a surprise that spacing of $O(\varepsilon)$ does not suffice. In particular, the committee size must increase with the size of the candidate set. Thus, the only advantage in choosing the committee in an ad-hoc manner is when the size of the candidate set is unknown in advance. This is not the case for higher dimensional metric spaces:

Theorem 1.2 *In the d -dimensional unit hyper-cube, for every $\varepsilon > 0$ and candidate set of size m the following committees achieve regret of at most ε :*

- A universal committee of size $O(m^d/\varepsilon^d)$.
- An ad-hoc committee of size $O(m^2/\varepsilon)$.

Moreover, every universal committee for $m \geq 2d + 1$ candidates requires $\Omega(1/\varepsilon^d)$ size.

While the universal committee, again, is obtained by spreading the experts “uniformly” over the unit hypercube (in an (ε/m) -net), ad-hoc committees can be much smaller. Roughly speaking, this is done by constructing a maximal spanning tree of the candidates and constructing an ε -cover

of the tree. For $d = 2$ we can further improve this bound and construct ad-hoc committees of size $O(m^{1.5}/\varepsilon)$.

For ad-hoc committees, we also consider the computational question of efficiently constructing a committee that chooses a candidate with minimal regret. For the unit interval we show:

Theorem 1.3 *There exists a polynomial time algorithm that, given a set of candidates in the unit interval and a target committee size k , outputs a committee of size k whose regret is within a constant factor of the regret of the best committee of size k .*

Unfortunately, this is not the case for general discrete metric spaces for which we show that constant approximation is hard, even in a weak bi-criteria sense: It is NP-hard to approximate to within a factor better than $\Omega(\log m)$ the committee size needed to achieve error $\varepsilon = 0$ (“a perfect committee”).

Related Work

Collective decision making has been studied by various authors (e.g., (Young 1988; Austen-Smith and Banks 1996; Feddersen and Pesendorfer 1998)). In this setting committee members need to jointly decide which state of the world is correct. A leading example is that of a jury that is appointed to decide whether the defendant is guilty or innocent. Each committee member receives some noisy signal about the state and submits a vote accordingly. The goal of this literature is to find voting rules that minimize the probability of a mistake. Often, the search domain is the set of majority mechanisms. Of particular interest is the comparison between simple majority and unanimity. In our setting voting is for candidates; a main contribution of our paper is the consideration of a particular type of noise, namely the distance between the voter and the different candidates. A novel aspect of our work is that it provides a first formal model to deal with the optimal selection of jury while modeling the jury members’ biases. While this topic is of course of major importance (see, e.g., (Neilsona and Winterb 2000)) we are not aware of formal models dealing with this fundamental topic in economic and social systems.

Our paper is situated in an active line of work in AI that takes a utilitarian view of voting (Procaccia and Rosenschein 2006; Boutilier et al. 2015; Anshelevich et al. 2018; Caragiannis et al. 2017). In this line of work each voter has an underlined utility function. The voter then reports his preferences based on this utility function. For example, he may only report his top preference (as in our case) or give a complete ordinal ranking. Given this partial information the objective is to choose a candidate that approximately maximizes the social welfare. Note that this is different from our work as we seek to choose a candidate that is objectively the best and not to maximize the social welfare.

Treating both voters and candidates as points in space, where distances between them is the basis for voting is quite standard (see, e.g., (Enelow and Hinich 1989)). An exciting recent line of work initiated by Anshelevich et al. (Anshelevich et al. 2018; Anshelevich and Postl 2017) is particularly relevant: voters and candidates are points in a metric space

and each voter prefers candidates that are closer to him. This line of work takes a utilitarian point of view and explores the distortion of known voting rules and develops new simple rules.

Moreover, some models where such type of representation of candidates and voters as points in a metric space is augmented with some objective quality for the candidates have been considered (see, e.g., (Maravall-Rodriguez 2005)). This is however studied in the context of shaping policies by candidates rather than in the context of removing experts’ biases. A similar setting was also considered in (Krishna and Morgan 2011) in which a large population should vote for one of two candidates, each has two properties: ideology and competence. Each voter might prefer a candidate of different ideology but they all agree on the competence of the candidates. The paper is interested in a different set of questions than the ones we consider, in particular it considers mechanisms for voting for the whole population (and not just a chosen committee) resulting in the more competent candidate being elected.

A different body of literature addresses the problem of strategic experts which would like to influence the outcome. The setting for two experts was studied in (Krishna and Morgan 2001) as an extension of Crawford and Sobel celebrated paper (Crawford and Sobel 1982). In general, assuming that the agents are rational, we can use mechanism design to address this concern. A major tool in this context is virtual implementation (Hitoshi 1993): ask the experts for their reports and then bias the decision (e.g., based on reported preference (Gerardi, McLean, and Postlewaite 2009)) in a way that favors experts which were more consistent with the rest. The approach we take here has some of that flavor, as we bias the decision based on information we have about agents’ biases. Another related work is the one by (Alpern and Gal 2009), which conducts an equilibrium analysis of selection committees under sequential voting with veto power.

2 Model and Preliminaries

In our setting, there is a set $C = \{1, \dots, m\}$ of m candidates, located along a metric space Ω . The location of candidate $j \in C$ is $c_j \in \Omega$ and is known to all and its quality is $q_j \in \mathbb{R}$. The quality vector of the candidates is $\mathbf{q} = (q_1, \dots, q_m) \in \mathbb{R}^m$ and the location vector is $\mathbf{c} = (c_1, \dots, c_m) \in \Omega^m$. Along the same metric space Ω , there is also a set \mathbb{E} of experts, whose locations are known to all.

A *planner* that does not know the qualities of the candidates needs to choose the best one. To make this decision the planner chooses a *committee* $E \subset \mathbb{E}$ of size k among the set of experts; we identify the members of the committee with $\{1, \dots, k\}$. The experts location vector is given by $\mathbf{e} = (e_1, \dots, e_k) \in \Omega^k$.

Each expert $i \in E$ gives candidate j a grade $g_i(j) = q_j - d(e_i, c_j)$. The expert then votes for the candidate with the maximal grade: $vote_i(\mathbf{q}) = \arg \max_j g_i(j)$. When \mathbf{q} is clear from the context we omit it and simply write $vote_i$.

A *voting rule* is a function $V : C^k \rightarrow C$, which receives the votes $\mathbf{vote}(\mathbf{q}) = (vote_1(\mathbf{q}), \dots, vote_k(\mathbf{q}))$ of the committee members, and returns a winning candidate.

Our goal is to choose the candidate with the highest quality. However, as discussed in the introduction, we are often unable to choose the candidate with the highest quality. Instead, we would like to minimize the regret. The regret of selecting candidate j with respect to a quality vector \mathbf{q} is $R(j, \mathbf{q}) = \max_{j'} \{q_{j'} - q_j\}$. The regret of a voting rule V is $R(V) = \max_{\mathbf{q}} R(V(\text{vote}(\mathbf{q})), \mathbf{q})$.

We note that for simplicity of presentation we do not deal with tie breaking explicitly throughout the paper, rather we assume that ties are broken in favor of “our needs”. It is easy to verify that this can be formally addressed by adding an arbitrarily small ε to the quality of the chosen candidate.

3 Minimal Regret Voting Rule

The problem of choosing a committee to minimize the regret is composed of two problems. First, a committee should be chosen. Second, a voting rule should be chosen; that is: an aggregation rule that, given the committee members’ votes, chooses a candidate that minimizes the regret. In this section, we deal with the second problem of choosing a voting rule: given the locations of the candidates $\mathbf{c} = (c_1, \dots, c_m)$, the locations of the experts $\mathbf{e} = (e_1, \dots, e_k)$, and the votes of the experts $\text{vote} = (\text{vote}_1, \dots, \text{vote}_k)$. Which candidate should be chosen in order to minimize the regret over all quality vectors? We first observe that natural voting rules, such as majority (or plurality), may perform very poorly:

Observation 3.1 *For every committee, there exist instances for which applying the majority voting rule will result in regret of $\Omega(1)$.*

Proof: Consider the case of 3 candidates on the unit interval located at $0, 1/2, 1$. We will show that for any possible locations of a set of experts (of any size), there is a vector of qualities where the majority rule will result in regret of $\Omega(1)$. Assume without loss of generality that the majority of the experts lie to left or exactly at $1/2$. Now consider the quality vector $q_1 = 0, q_2 = 1, q_3 = 1.4$. Notice that all the experts left of $1/2$ vote for candidate 2, and so the majority vote has regret 0.4 . \square

We next give a polynomial-time voting rule that minimizes the worst-case regret over any quality vector \mathbf{q} . Moreover, we will see that by using the optimal voting rule we can make the regret as small as desired using sufficiently large committees.

Let $v_i = \text{vote}_i(\mathbf{q})$ and $\mathbf{v} = (v_1, \dots, v_k)$. The key for defining the minimal regret rule is observing that different qualities vector \mathbf{q}' may result in the same vector of votes \mathbf{v} . We refer to the set of such quality vectors as *indistinguishable qualities* and denote it by $Q(\mathbf{v}) = \{\mathbf{q}' : \mathbf{v} = \text{vote}(\mathbf{q}')\}$. To minimize the regret we should choose the candidate with minimal regret over the set of the indistinguishable qualities. This is a candidate $j \in C$ such that $j = \arg \min_j \max_{\mathbf{q}' \in Q(\mathbf{v})} R(j, \mathbf{q}')$ and we refer to it as the *minimal regret candidate*.

We show a set of linear programs that are used to compute the minimal regret candidate. The first step is to observe that the set of indistinguishable qualities $Q(\mathbf{v})$ can be written using linear constraints. Specifically, for each expert $i \in E$ and each candidate $j \in C$ we have the constraint:

$q_{v_i} - d(e_i, c_{v_i}) \geq q_j - d(e_i, c_j)$, where the variables are the unknown qualities q_j . For each candidate h such that there exists at least a single expert that voted for him and every $j \in C$, we can compute the regret of selecting h when j is a better candidate as

$$\max_{\mathbf{q}} R_{h,j} = q_j - q_h$$

$$\forall i \in E, \quad q_{v_i} - d(e_i, c_{v_i}) \geq q_j - d(e_i, c_j).$$

We can compute the minimal regret candidate by solving at most m^2 linear programs, one per pair of candidates. For each candidate h we compute $R_h = \max_j R_{h,j}$ and select the candidate h with the minimal R_h .

In the next theorem we show how to compute the minimal regret candidate more efficiently. This is done by reducing the problem to computing shortest paths on a directed graph.

Theorem 3.2 *Given candidate locations \mathbf{c} , experts locations \mathbf{e} and expert votes \mathbf{v} , the minimal regret candidate can be computed in time $O(mk + m^3)$.*

Proof: We construct a directed graph where every node is associated with a candidate in C . For each expert $i \in E$ we add a directed edge from the candidate that i voted for (v_i) to any other candidate h with weight $w_{v_i, h} = d(e_i, c_h) - d(e_i, c_{v_i})$. Recall that we have $q_{v_i} - d(e_i, c_{v_i}) \geq q_h - d(e_i, c_h)$ or equivalently

$$q_{v_i} \geq q_h - w_{v_i, h} \quad (1)$$

We claim that the resulting graph does not have cycles with negative weights. To show this consider a cycle

$$v_{i_1} \rightarrow v_{i_2} \rightarrow \dots \rightarrow v_{i_n} \rightarrow v_{i_1}$$

Summing the corresponding Equations (1) we get

$$\sum_{r=1}^n q_{v_{i_r}} \geq \sum_{r=1}^n q_{v_{i_r}} - \sum_{r=1}^n w_{v_{i_r}, v_{i_{r+1}}},$$

which implies that $\sum_{r=1}^n w_{v_{i_r}, v_{i_{r+1}}} \geq 0$, hence the sum of the weights along any cycle is non-negative.

Therefore, we can find the minimal regret candidate as follows: (1) Build the directed weighted graph (in time $O(mk)$), (2) Solve all-pair shortest paths (either naively in time $O(m^3)$ or use one of the more sophisticated matrix multiplication algorithms), and compute a matrix A where A_{j_1, j_2} is the regret for choosing candidate j_1 compared to candidate j_2 . Then, (3) for each candidate h compute R_h (in time $O(m)$ by computing maximum in row h of A , for a total time of $O(m^2)$) and finally, (4) compute the minimal regret candidate h^* (in time $O(m)$ by computing $\min R_h$). \square

4 Universal Committees

We are now ready to discuss our main question, how should we form a committee that will be able to choose the best candidate even when the locations of the candidates are *unknown in advance* (i.e., a universal committee). More formally, our goal is to form a committee that would be able to guarantee a regret of at most ε for any set of m candidates that may appear. We do assume that the number of candidates m is known in advance. As we shall see, this is unavoidable since the committee size must grow with the number of candidates.

4.1 General Upper bound

We begin by constructing a universal committee that achieves a regret of at most ε . We assume here that we can choose an expert at any point in the metric space (i.e., $E = \Omega$). We will need the following definition:

Definition 4.1 (δ -cover) A set $E \subset \Omega$ is an δ -cover of a metric space Ω if for each point $x \in \Omega$ there exists $y \in E$ such that $d(x, y) \leq \delta$.

We construct the universal committee by constructing an $(\varepsilon/4m)$ -cover of Ω :

Lemma 4.2 Let Ω be a connected set. Any $(\varepsilon/4m)$ -cover of Ω is a universal committee for m candidates that guarantees a regret of at most ε .

Proof: We will show that given the votes of the experts on any set locations $\mathbf{c} = (c_1, \dots, c_m)$ of the m candidates, we can determine the difference in quality between any two candidates that were voted by at least one expert up to an additive error of ε . We assume that every candidate received at least a single vote by one of the committee members.³

Consider the following directed graph. Every candidate $j \in C$ is associated with a vertex v_j . There exists an edge between candidates j and j' if there are experts i and i' such that $\text{vote}_i = j$, $\text{vote}_{i'} = j'$, and $d(i, i') \leq 2\varepsilon/(4m)$. Such an edge is assigned weight $w(j, j') = d(i, j) - d(i, j')$. The votes of the experts imply that $q_j - d(i, j) \geq q_{j'} - d(i, j')$ and $q_j - d(i', j) \leq q_{j'} - d(i', j')$.

Combining the last inequalities with the bounds on distances and the triangular inequality implies that

$$d(i, j) - d(i, j') \leq q_j - q_{j'} \leq d(i, j) - d(i, j') + \varepsilon/m.$$

Therefore, the difference in quality is $w(j, j')$ up to an additive error of ε/m .

Consider any two nodes in the graph (which correspond to candidates in C) which are connected by a path v_{j_1}, \dots, v_{j_r} . Since the difference in qualities between any two consecutive candidates j, j' is $w(j, j')$ up to an error of ε/m , the difference between q_{j_r} and q_{j_1} is $\sum_{i=1}^{r-1} w(j_i, j_{i+1})$ up to an additive error of $\pm\varepsilon$.

It remains to show that the resulting graph is strongly connected. Consider two nodes j and j' which are not connected. Consider a shortest path p in the metric space connecting j and j' . For any point on the path, there is an expert of distance at most $\varepsilon/(4m)$. Consider the set of experts $E_p \subset E$ which are at distance at most $\varepsilon/(4m)$ from the path p . We can create a path from the experts in E_p such that the distance between any two is at most $2\varepsilon/(4m)$, where the first expert votes for j and the last votes for j' . As we march along the path, if an expert votes differently than its predecessor, we add an edge to the graph, and thus we have a path in the graph connecting j and j' . \square

The following theorem is a straightforward corollary of the last lemma, using standard bounds on covers.

³The proof can be easily extended to the case that some of the candidates were not voted by any expert. This is done by noting that the quality of a candidate j that did not receive any votes is at most $q_j \leq q_{j'} + \varepsilon/4m$ where j' is another candidate that was voted for.

Theorem 4.3 For any number of candidates m , and any $d \geq 1$, the metric space $\Omega = [0, 1]^d$ admits a universal committee of size $(4m/\varepsilon)^d$ that guarantees a regret of ε . \square

We establish a better bounds for the special case of d -dimensional unit hyper-cubes with two candidates:

Theorem 4.4 For the case of $m = 2$ candidates and any $d \geq 1$, the metric space $\Omega = [0, 1]^d$ admits a universal committee of size $O(2^d/\varepsilon)$ that guarantees a regret of ε .

Proof: We show how to choose a set of experts E with the property that for any two candidates there exists a single expert in E that can distinguish between them with accuracy ε . For $\mathbf{x} \in \{0, 1\}^d$ define $\bar{\mathbf{x}} = 1 - \mathbf{x}$, component-wise, i.e., $\bar{x}_i = 1 - x_i$. We call the line connecting x to $\bar{\mathbf{x}}$ the diagonal of $(\mathbf{x}, \bar{\mathbf{x}})$. Along each diagonal of $(\mathbf{x}, \bar{\mathbf{x}})$ we place experts spaced $\varepsilon/2$ away from each other. This gives a total of $O(2^d/\varepsilon)$ experts.

Given two candidates in locations c_1 and c_2 , consider the equi-distance hyperplane from c_1 and c_2 . This hyperplane must intersect at least one of the diagonals $(\mathbf{x}, \bar{\mathbf{x}})$. By the construction, there exists an expert at distance at most $\varepsilon/2$ from the intersection point, which implies that the difference between its distance to c_1 and c_2 is at most ε . \square

4.2 Lower Bounds

We begin by considering d -dimensional unit hyper-cubes and present two lower bounds on the required size of the committee: one for achieving a regret smaller than $1/2$ and the other for achieving a regret smaller than $\varepsilon \leq 1/\sqrt{d}$:

Theorem 4.5 For any $d \geq 1$, a universal committee that guarantees a regret of at most $1/2$ in $\Omega = [0, 1]^d$ must be of size at least $m - 1$.

Proof: For contradiction, assume there is a universal committee E of size $m - 2$. For every expert in E , locate a candidate in the exact same location as the expert. Then, locate the additional two candidates in $(0, \dots, 0)$ and $(1, \dots, 1)$, respectively. The quality of a candidate is determined as follows. We have two states of the world. In state *plus* the quality of each candidate is its first coordinate, i.e., x_1 , and in state *minus* it is the complement of its first coordinate, i.e., $1 - x_1$.

Note that each expert selects the candidate which is at the same location as him, since for any other candidate the difference in quality is at most the distance between them. Therefore, the two states of the world are indistinguishable given the experts' votes, and the same candidate will be selected. The chosen candidate has quality either q or $1 - q$, while the best candidate has quality 1. Consequently, the regret is $\max\{1 - q, 1 - (1 - q)\} \geq 1/2$. \square

Theorem 4.6 For any $d \geq 1$, a universal committee that guarantees a regret of at most $\varepsilon \leq 1/\sqrt{d}$ for any set of $m \geq 2d + 1$ candidates in $[0, 1]^d$ must be of size at least $(2\varepsilon\sqrt{d})^{-d}$.

Proof: Consider partitioning the space into boxes of size $(2\delta)^d$ each; there are $(2\delta)^{-d}$ such boxes. We will show that if there is a box that does not include any expert, then we

can select $2d$ candidates at the exterior of the box and one candidate in the interior of the box, such that for any expert outside the box, there exists at least one exterior candidate whose distance from the expert is smaller by at least ε than the distance between the expert and the internal candidate. This implies that if the quality of the internal candidate is higher than the other candidates by ε , no expert will vote for the internal candidate. Therefore, we can force a regret of ε in this case. Thus, any box should have an expert in its interior, and hence the number of experts is at least $(2\delta)^{-d}$.

Let $\delta = \varepsilon\sqrt{d} < 1$, then $\varepsilon = \delta/\sqrt{d}$. Consider the hyper-box, $[-\delta, +\delta]^d$ and the $2d$ unit vectors δe_i and their negation $-\delta e_i$. Consider any point $x \notin [-\delta, +\delta]^d$, and assume that $|x_1| = \max_i |x_i|$. We compare its distance to the origin versus its distance to $\pm\delta e_1$. The distance of x to the origin is $\sqrt{\sum_{j=1}^d x_j^2} = R$. The distance of x from δe_1 is $\sqrt{(|x_1| - \delta)^2 + \sum_{j=2}^d x_j^2} \geq R - \frac{2\delta x_1}{2R}$. Since $|x_1| \geq R/\sqrt{d}$ we have that the difference in the distance is at least δ/\sqrt{d} . This implies that any point outside the box has a distance smaller by $\frac{\delta}{\sqrt{d}}$ to at least one of the $\pm\delta e_i$ compared to the origin.

To conclude, if there is a box which does not have any expert, we build the bad example as follows. We locate $2d+1$ candidates, one in each $\pm\delta e_i$ with quality 1 and one in the origin with quality $1 + \frac{\delta}{\sqrt{d}}$. No expert outside the box would prefer the candidate at the origin which implies the lower bound. \square

Next, we reconcile the last two lower bounds for the case of the unit interval. Interestingly, this lower bound holds even for the case of an ad-hoc committee (the committee is chosen only after the candidates' locations are known).

Theorem 4.7 *For the unit interval, every $m \geq 3$ and $\varepsilon > 0$ there exists a set of m candidates such that every (even ad-hoc) committee that achieves regret at most ε requires size $k = \Omega(m/\varepsilon)$.*

Proof: We sketch the proof of the dual argument that for every $m \geq 3$ and k , if the set of candidates are located along the unit interval in equal distances: $c_i = \frac{i-1}{m-1}$ for $i = 1, \dots, m$, then every committee of size k achieves regret at least $\Omega(m/k)$. The complete proof can be found in the full version.

Before presenting the exact construction, we give a high-level description of the intuition behind it. We construct two different quality vectors \mathbf{q}^1 and \mathbf{q}^2 that, on the one hand, admit identical voting of all k experts (and are therefore indistinguishable from the planner's point of view), and on the other hand, admit an $\Omega(m/k)$ gap between the highest and the lowest qualities.

Both quality vectors consist of a decreasing sequence of qualities followed by an increasing sequence. The difference is that in \mathbf{q}^1 , qualities first decrease sharply and then increase mildly, while in \mathbf{q}^2 they first decrease mildly and then increase sharply. This way, the best candidate in \mathbf{q}^1 is the left-most candidate, while the best candidate in \mathbf{q}^2 is the right-most one. The actual construction is subtle: the challenge is to delicately find the right sequence of qualities that

ensure identical voting on the one hand, yet a sufficiently large gap in qualities on the other hand.

We now proceed with the exact construction. Divide the line into $10k$ equal intervals of size $\delta = 1/(10k)$ each. This way, there are $\frac{10k}{m-1}$ intervals between any two candidates. Let $e = (e_1, \dots, e_k)$ be the location of the experts such that $e_i \leq e_{i+1}$ for every expert i . Given an index $\ell \in 1, \dots, \frac{10k}{2(m-1)}$ (i.e., the index of an interval that resides to the left of the center of two consecutive candidates), two consecutive candidates are said to be ℓ -empty if there are no experts in the ℓ^{th} interval between them, nor in its "mirror" interval $\hat{\ell} := \frac{10k}{m-1} - \ell + 1$. Pick $\ell \leq \frac{10k}{2(m-1)}$ for which the number of ℓ -empty pairs is maximal and let m' denote this number. Observe that $m' > 4m/5$.⁴ Finally, further divide the set of ℓ -empty pairs into two equal-size sets — the left set and the right set.⁵

Let the pair $(j, j+1)$ be an ℓ -empty pair of candidates from the left set. We wish to build the qualities such that every expert in $[c_j, c_j + (\hat{\ell} - 1)\delta]$ votes for candidate j , while every expert in $[c_j + \hat{\ell}\delta, c_{j+1}]$ votes for candidate $j+1$. (recall, there are no experts inside the interval $[c_j + (\hat{\ell} - 1)\delta, c_j + \hat{\ell}\delta]$). All experts in $[c_j, c_j + (\hat{\ell} - 1)\delta]$ will prefer candidate j over candidate $j+1$ if and only if

$$q_{j+1} \leq q_j + \frac{1}{m-1} - 2(\hat{\ell} - 1)\delta. \quad (2)$$

Similarly, all experts in $[c_j + \hat{\ell}\delta, c_{j+1}]$ will prefer candidate $j+1$ over candidate j if and only if

$$q_{j+1} \geq q_j + \frac{1}{m-1} - 2\hat{\ell}\delta. \quad (3)$$

The combination of the two inequalities leave us a leeway of 2δ . In particular, we can decrease the quality of q_{j+1} by value $2(\hat{\ell} - 1)\delta - \frac{1}{m-1} + \Delta$ for any $\Delta \in [0, 2\delta]$. (This is indeed a decrease since $(\hat{\ell} - 1)\delta \geq \frac{1}{2(m-1)}$).

Consider next an ℓ -empty pair $(j, j+1)$ from the right set. We build the qualities such that every expert in the interval $[c_j, c_j + (\ell - 1)\delta]$ votes for candidate j , while every expert in the interval $[c_j + \ell\delta, c_{j+1}]$ votes for candidate $j+1$. By replacing $\hat{\ell}$ by ℓ in Equations (2) and (3) we can increase the quality of q_{j+1} by $\frac{1}{m-1} - 2(\ell - 1)\delta + \Delta$ for any $\Delta \in [0, \delta]$.

We now present the explicit quality vectors q^1 and q^2 . In both q^1 and q^2 , for every pair $(j, j+1)$ that is not ℓ -empty, we set $q_j = q_{j+1}$. In q^1 , for a pair of candidates $j, j+1$ in the left set, we set $q_{j+1}^1 = q_j^1 + \frac{1}{m-1} - 2\hat{\ell}\delta$, while for a pair in the right set, we set $q_{j+1}^1 = q_j^1 + \frac{1}{m-1} - 2\ell\delta$. In q^2 , for every ℓ -empty pair in the left set, we set $q_{j+1}^2 = q_j^2 + \frac{1}{m-1} - 2\hat{\ell}\delta + 2\delta$, and for every ℓ -empty pair in the right set, we set $q_{j+1}^2 = q_j^2 + \frac{1}{m-1} - 2\ell\delta + 2\delta$. For these qualities it

⁴Otherwise, for any $\ell \leq \frac{10k}{2(m-1)}$ there exists at least $m/5$ pairs that are not ℓ -empty. The total number of experts required for this is $\frac{10k}{2(m-1)} \cdot \frac{m}{5} = \frac{m \cdot k}{m-1} > k$.

⁵We briefly discuss the case that m' is odd later.

holds that $q_m^1 = q_1^1 - m'\delta$ while $q_m^2 = q_1^2 + m'\delta$. This implies that q^1 and q^2 are indistinguishable and since $m' \geq 4m/5$, the regret is at least $m\delta = \Omega(m/k)$, as required. \square

5 Ad-hoc Committees

We return to discussing the case of forming a committee after observing the locations of the candidates. In the full version, we provide a polynomial time approximation for the unit interval. Consider the candidates location vector $\mathbf{c} = (c_1, \dots, c_m)$. Roughly speaking, the heart of the algorithm is identifying the interval which is at the center between the leftmost and rightmost candidates and locating about half of the experts in this interval. The algorithm then, cleverly decides how many experts are needed in each of the rest of the intervals $[c_j, c_{j+1}]$ and essentially equally spaces them in each interval.

Our techniques also imply the dual approximation: given a target regret ε we can find a committee whose size is within a constant factor of the size of the smallest committee with ε regret for that set of candidates.

For general discrete metric spaces we show that constant approximation is hard, even in a weak bi-criteria sense. To this end, we consider the problem of computing a minimal size committee that always chooses the optimal candidate (i.e., *perfect committee*) and show:⁶

Theorem 5.1 *Approximating the perfect committee problem to within an $o(\log m)$ -factor is NP-hard.*

To prove the theorem we define the *pair cover problem* in which we are given a bipartite graph and need to compute a minimal subset of left vertices such that each pair of right vertices has a left vertex that is connected to both. We show that the latter problem is NP-hard to approximate with in $o(\log m)$ -factor by reducing from set cover.

Next, we contrast ad-hoc committees and universal committees with respect to the required size of a committee guaranteeing at most ε regret. As discussed in Section 4, for the unit interval there is no difference between the performance of universal and ad-hoc committees. Interestingly, for higher dimensions, observing the locations of the candidates can help reduce the committee size needed to guarantee a certain regret. This is done by first building a minimum spanning tree (MST) over the candidates, then densely populating the edges of the MST with experts. The length (or weight) of an MST T is denoted by $w(T)$.

Theorem 5.2 *Let Ω be a connected metric space and T a minimum spanning tree of the locations of the candidates $\mathbf{c} = (c_1, \dots, c_m)$ over Ω with a total weight of $w(T)$. There exists a committee of size $k = O(w(T)m/\varepsilon)$ which guarantees a regret of at most ε .*

Proof: Define a new metric space Ω_T that includes only points on the edges of T . We now create an $(\varepsilon/4m)$ -cover for Ω_T , which has size at most $k = O(w(T)m/\varepsilon)$ and locate the k experts in those locations. By Lemma 4.2 this is a universal committee for Ω_T which guarantees a regret of at most ε . Since our set of candidates, by construction, are

⁶Note that a perfect committee does not always exist.

in Ω_T , it implies that in particular the committee will have regret at most ε . \square

Let D denote the diameter of Ω . Clearly, the weight of the MST is bounded by Dm , which implies the following corollary.

Corollary 5.3 *For any connected metric space Ω of diameter D , there exists an ad-hoc committee of size $k = O(Dm^2/\varepsilon)$ which guarantees a regret of at most ε .*

The last corollary implies a big separation between ad-hoc and universal committees. While the required size of a universal committee is exponential in the dimension of the metric space, this exponential growth can be avoided in the case of ad-hoc committees. For the special case of $\Omega = [0, 1]^2$ we provide a better bound by establishing a better bound on the weight of the MST. We show:

Theorem 5.4 *For $\Omega = [0, 1]^2$, there exists a committee of size $k = O(m^{1.5}/\varepsilon)$ which guarantees a regret of at most ε .*

6 Discussion

In this paper we initiate the study of forming committees to reduce biases in voting. We show that eliminating biases is a difficult task and often requires very large and diverse committees. Furthermore, the required size of the committee increases as the dimension of the metric space increases. To remedy this, one should prefer forming ad-hoc committees whenever possible, as we show that for d -dimensional hyper-cubes ad-hoc committees can be substantially smaller. Alternatively, to simplify the metric space, one should seek to minimize the exposure of the experts to irrelevant information that may bias them. For hiring decisions, one can, for example, refrain from disclosing the family status of the candidates.

In the full version of the paper we explore additional aspects of committee formation. We consider *randomized voting rules* and show that for the unit interval they can perform better than deterministic voting rules by a constant gap. We also revisit the assumption that experts vote for the candidate with the maximal grade. Instead we allow the experts to vote strategically in order to get a candidate they prefer to be chosen. We show that if it is possible to locate many experts at the same location then we can transform any non-strategic voting rule to a strategyproof voting rule that obtains the same regret by tripling the size of the committee and locating three experts in each original location. If we are not allowed to locate multiple experts at the same location then even for 3 candidates on the unit interval there is no strategyproof voting rule that is guaranteed to do well.

Our work raises some new and exciting questions. On a more concrete level, we leave open the question of closing the gap between the lower bound and upper bound on the size of ad-hoc committees for d -dimensional unit hyper-cubes. More generally, in line with (Anshelevich et al. 2018), it is interesting to consider the tradeoff between the quality of the chosen candidate and the amount of information each expert should report. In our work each expert only reports his top candidate. By how much can we improve the quality of the chosen candidate (or reduce the size

of the committee) if we allow each expert to report his two (or more) favorite candidates?

References

- Alpern, S., and Gal, S. 2009. Analysis and design of selection committees: a game theoretic secretary problem. *Int. J. Game Theory* 38(3):377–394.
- Anshelevich, E., and Postl, J. 2017. Randomized social choice functions under metric preferences. *Journal of Artificial Intelligence Research* 58:797–827.
- Anshelevich, E.; Bhardwaj, O.; Elkind, E.; Postl, J.; and Skowron, P. 2018. Approximating optimal social choice under metric preferences. *Artificial Intelligence* 264:27–51.
- Austen-Smith, D., and Banks, J. S. 1996. Information aggregation, rationality, and the condorcet jury theorem. *American Political Science Review* 90:34–45.
- Boutilier, C.; Caragiannis, I.; Haber, S.; Lu, T.; Procaccia, A. D.; and Sheffet, O. 2015. Optimal social choice functions: A utilitarian view. *Artificial Intelligence* 227:190–213.
- Caragiannis, I.; Nath, S.; Procaccia, A. D.; and Shah, N. 2017. Subset selection via implicit utilitarian voting. *Journal of Artificial Intelligence Research* 58:123–152.
- Crawford, V. P., and Sobel, J. 1982. Strategic information transmission. *Econometrica: Journal of the Econometric Society* 1431–1451.
- Currarini, S.; Jackson, M. O.; and Pin, P. 2009. An economic model of friendship: Homophily, minorities, and segregation. *Econometrica* 77(4):1003–1045.
- Enelow, J. M., and Hinich, M. J. 1989. A general probabilistic spatial theory of elections. *Public Choice* 61:101–113.
- Feddersen, T., and Pesendorfer, W. 1998. Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting. *American Political Science Review* 92:23–35.
- Gerardi, D.; McLean, R. P.; and Postlewaite, A. 2009. Aggregation of expert opinions. *Games and Economic Behavior* 65:339–371.
- Hitoshi, M. 1993. Bayesian monotonicity with side payments. *Journal of Economic Theory* 59:107–121.
- Krishna, V., and Morgan, J. 2001. A model of expertise. *Quarterly Journal of Economics* 116(2).
- Krishna, V., and Morgan, J. 2011. Overcoming ideological bias in elections. *Journal of Political Economy* 119(2):183–211.
- Maravall-Rodriguez, C. 2005. A spatial election with common values. Working Paper.
- McPherson, M.; Smith-Lovin, L.; and Cook, J. M. 2001. Birds of a feather: Homophily in social networks. *Annual review of sociology* 415–444.
- Neilsona, W. S., and Winterb, H. 2000. Bias and the economics of jury selection. *International Review of Law and Economics* 20:223–250.
- Procaccia, A. D., and Rosenschein, J. S. 2006. The distortion of cardinal preferences in voting. In *International Workshop on Cooperative Information Agents*, 317–331. Springer.
- Rivera, L. A. 2012. Hiring as cultural matching the case of elite professional service firms. *American Sociological Review* 77(6):999–1022.
- Roemer, J. E. 2001. *Political Competition: Theory and Applications*.
- Young, P. 1988. Condorcet’s theory of voting. *American Political Science Review* 82:1231–1244.