
Can a Few Decide for Many? The Metric Distortion of Sortition

Ioannis Caragiannis¹ Evi Micha² Jannik Peters³

Abstract

Recent works have studied the design of algorithms for selecting representative sortition panels. However, the most central question remains unaddressed: Do these panels reflect the entire population’s opinion? We present a positive answer by adopting the concept of metric distortion from computational social choice, which aims to quantify how much a panel’s decision aligns with the ideal decision of the population when preferences and agents lie on a metric space. We show that uniform selection needs only logarithmically many agents in terms of the number of alternatives to achieve almost optimal distortion. We also show that Fair Greedy Capture, a selection algorithm introduced recently by Ebadian & Micha (2024), matches uniform selection’s guarantees of almost optimal distortion and also achieves constant ex-post distortion, ensuring a “best of both worlds” performance.

1. Introduction

In recent years, *sortition* or the randomized selection of so-called *citizens’ assemblies* has gained traction around the globe. In such citizens’ assemblies, a randomly selected subset of the population discusses and deliberates on societal issues, such as challenges from climate change and AI threats. Notable examples include the citizens’ assemblies on climate change held in Belgium, France, or the United Kingdom or the Irish citizens’ assembly on abortion which lead to a referendum on the topic and finally the legalization of abortion. This surge in the usage of citizens’ assemblies has been supported by work in the AI and ML

community, which have focused on the design of the underlying sampling methods (Flanigan et al., 2021a) and the formulation of normative properties of sortition and related analysis (Ebadian et al., 2022; Meir et al., 2021; Ebadian & Micha, 2024). As a prominent example, the sampling method of Flanigan et al. (2021a) was recently deployed for the selection of the German citizens’ assembly on nutrition.¹

As outlined by the sociologist Engelstad (1989), two of the main supposed advantages of sortition are the *representation* of the underlying population and *fairness*. The interpretation of fairness by the aforementioned works is quite straightforward: *a selection procedure is fair if every member of the population is selected for the panel with equal probability*. Fairness can be achieved under *uniform selection* where a panel of given size is sampled uniformly at random. Representation on the other hand is not that easily defined. Both Ebadian et al. (2022) and Ebadian & Micha (2024) assumed that there is an underlying *representation metric*, measuring the closeness between different members of the population, for instance by measuring the similarities of certain attributes. Using this metric, Ebadian et al. (2022) measured the distance of each agent to their q -th closest member on the panel, and showed that if q is more than half the panel size, uniform selection of the panel achieves a constant factor approximation to summing these distances and thus being in some sense representative of the underlying population. Ebadian & Micha (2024) on the other hand, interpreted representation as *proportional representation* and imported the notion of the core from social choice (Gillies, 1953), which intuitively requires that each sufficiently large group has sufficiently many representatives in the panel. They show that uniform selection satisfies the core in expectation, and introduce a fair selection algorithm that always returns a panel in the core.

However, while previous work has focused on the formation of panels that are representative of the population, it is still quite unexplored whether such panels reach decisions that reflect the aggregate preference of the underlying population. So, our goal in the current paper is

(a) to explore whether panels chosen using sorti-

¹Department of Computer Science, Aarhus University, Aarhus, Denmark ²Computer Science, Harvard University, Cambridge, USA ³ Research Group Efficient Algorithms, Faculty IV – Electrical Engineering and Computer Science, TU Berlin, Berlin, Germany. Correspondence to: Ioannis Caragiannis <iannis@cs.au.dk>, Evi Micha <emicha@seas.harvard.edu>, Jannik Peters <jannik.peters@tu-berlin.de>.

Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024. Copyright 2024 by the author(s).

¹See https://www.bundestag.de/en/parliament/citizens_assemblies.

tion are able to represent the interest of the whole population for different issues, and (b) to provide bounds on the panel size that is sufficient and necessary for achieving this.

Our approach and results. To measure the extent to which a panel’s decision is close to the decision the whole population would collectively take, we utilize the concept of *distortion* from the area of computational social choice (Procaccia & Rosenschein, 2006; Boutilier et al., 2015). Specifically, first, we extend the definition of the representation metric of Ebadian et al. (2022) and Ebadian & Micha (2024) by assuming that both agents and alternatives are embedded in the same metric space. Then, we compare the (expected) social cost of the best alternative (defined as its total distance from all agents) for the panel to the social cost of the best alternative for the overall population. If the ratio between the two is low, then the will of the panel is close to the will of the electorate and the preferences of the underlying population are well represented.

In Section 4, we start by presenting the benefits of fair selection for achieving low distortion. In scenarios with arbitrarily many alternatives, deterministic selection cannot have a distortion considerably better than 5 with small panels, while, in contrast, fair selection algorithms always achieve a distortion of 3. Uniform selection, i.e., selecting each panel of a given size uniformly at random, is the most natural fair selection algorithm. For uniform selection, we specifically ask the question: for parameters m and ε , how small the panel should be so that the selected alternative has an expected social cost that is at most a factor of $1 + \varepsilon$ far from the optimal social cost for all instances with m alternatives? In other words, we seek for low *ex-ante* distortion. In Section 5, we prove an upper bound of $\mathcal{O}(\varepsilon^{-2} \ln \frac{m}{\varepsilon})$ on the panel size. This can be thought of as a “dimensionality reduction” result. No matter how large the population is, the panel size depends only on the number of alternatives and the accuracy parameter. Our proof uses a concentration inequality due to Serfling (1974) for sums of random variables with limited correlation. We also prove a matching lower bound for every selection algorithm. Our main tool here is a new anti-concentration inequality for the hypergeometric probability distribution.

One of the criticisms of sortition outlined by Engelstad (1989) is that, in extreme cases, the results will be socially disastrous. Indeed, uniform selection selects the worst possible panel with positive probability. In our terminology, the *ex-post* distortion can be as bad as $\Omega(n)$, where n is the population size. In Section 6, we revisit the Fair Greedy Capture algorithm, introduced recently by Ebadian & Micha (2024). Even though it was originally designed with a different objective in mind (i.e., forming panels that are representative of a population under the definition of the core), we show

that it achieves a “best of both worlds” performance guarantee. First, panels of size $\mathcal{O}(\varepsilon^{-3} \ln \frac{m}{\varepsilon})$ yield an *ex-ante* distortion of $1 + \varepsilon$ for all instances with m alternatives. Second, we show that the *ex-post* distortion is only constant. The proof of our bound on the panel size exploits structural properties of Fair Greedy Capture that allow for the use of Hoeffding’s inequality.

Lastly, in Section 7, we empirically compare uniform selection and Fair Greedy Capture over real data. Among other interesting observations, we notice that Fair Greedy Capture consistently achieves better distortion than uniform selection, which converges quickly to 1.

2. Related Work

We highlight three streams of research closely related to us.

Sortition. The first one consists of computer science work on formalizing and analyzing sortition. Here, starting with the work of Benade et al. (2019), there has been a focus on devising sampling methods for sortition (Flanigan et al., 2021a;b; 2020). Besides the aforementioned papers by Ebadian et al. (2022) and Ebadian & Micha (2024), the work closest to ours is by Meir et al. (2021). They also analyze the distortion of sortition. However, they focus on a series of binary issues (i.e., two issues with utility either 0 or 1 for the agents) deriving almost tight bounds for this special case, showing that for such simple options the distortion of uniformly selecting a panel of size k is $1 + \mathcal{O}(k^{-1/2})$. Our work can be seen as a generalization of theirs, in the sense that we allow many alternatives and general metric utilities, showing that sortition maintains a very low distortion even in a more general setting. Finally, Anagnostides et al. (2022) studied (among other things) the metric distortion of first forming a panel of agents and then running a voting rule. They were able to show that selecting a panel of size $\mathcal{O}(\varepsilon^{-2} \ln m)$ uniformly at random and then running Copeland’s rule is enough to achieve distortion $5 + \varepsilon$ with high probability, while uniform selection with size $\mathcal{O}(\frac{m}{\varepsilon^2})$ yields distortion $3 + \varepsilon$, using the distortion-3 rule of Gkatzelis et al. (2020). Our work additionally improves upon theirs and shows that to achieve distortion $3 + \varepsilon$ (or even better if we use the recent randomized voting rule of Charikar et al. (2024)), a panel size of $\mathcal{O}(\varepsilon^{-2} \ln \frac{m}{\varepsilon})$ is also sufficient.

Distortion. Secondly, the field of distortion in computational social choice is apparently relevant. In particular, in recent years, the study of *metric distortion*, in which the preferences of the agents come from a metric space, has gained traction, with several voting rules being shown to have good distortion. For instance, random dictatorship (Anshelevich & Postl, 2016), Copeland (Anshelevich

et al., 2015), a variant of the vote-by-veto rule (Kızılkaya & Kempe, 2022), or maximal lotteries (Charikar et al., 2024) achieve (close to) optimal metric distortion. We refer to Anshelevich et al. (2021) for a survey and to Amanatidis et al. (2022), Gonczarowski et al. (2023), Ebadian et al. (2023), and Caragiannis & Fehrs (2023) for some recent works on (non-metric) distortion. The work on distortion most related to ours is most likely the one by Jaworski & Skowron (2022). In their model, just like in the work of Meir et al. (2021), they measure distortion with regard to a series of binary decisions. However, their model differs in two aspects: (i) the voter set and the set of candidates for the committee are different, and (ii) the voters just provide a ranking over the candidates for the committee. In this setting, Jaworski & Skowron (2022) show that no (deterministic) committee selection method can achieve a bounded distortion, even when restricted to alternatives which are top-choices of some committee.

Clustering. Finally, through our usage of the Fair Greedy Capture algorithm, we are inherently related to the proportional clustering literature and the Greedy Capture mechanism introduced by Chen et al. (2019). Since Fair Greedy Capture achieves constant factor approximations to both proportional fairness as studied by Chen et al. (2019) or Li et al. (2021) and individual fairness as studied by Jung et al. (2020) (see Kellerhals & Peters (2023) for a proof of this and how these notions are related), this also shows that representation in the sense of selecting a panel with good distortion is also compatible with representation in the sense of proportionally or individually fair clustering as devised in the aforementioned works.

3. Model and Motivation

For $t \in \mathbb{N}$, define $[t] = \{1, \dots, t\}$. Throughout the paper, we assume that we are given a set of agents $N = [n]$ and a set of alternatives $C = \{c_1, \dots, c_m\}$. Together $N \cup C$ lie in an underlying (pseudo-) metric space, with a distance function $d: (N \cup C) \times (N \cup C) \rightarrow \mathbb{R}_{\geq 0}$. We denote with $d(i, j)$ the distance between each $i, j \in N \cup C$. We assume that the following properties are satisfied by d : (a) $d(i, j) = d(j, i)$ (symmetry) and (b) for $i, j, \ell \in N \cup C$, $d(i, j) \leq d(i, \ell) + d(\ell, j)$ (triangle inequality).

The *cost* an agent $i \in N$ has for an alternative $c \in C$ is equal to $d(i, c)$, with a lower cost being “better” for the agent. The social cost of an alternative c is equal to $SC(c) := \sum_{i \in N} d(i, c)$, i.e., the sum of the costs over all agents. A panel P is a subset of the agents, i.e., $P \subseteq N$. For such a panel P and alternative $c \in C$, we let $SC(P, c) := \sum_{i \in P} d(i, c)$ denote that social cost of that alternative for the panel. Using this, we define $c(P) := \arg \min_{c \in C} SC(P, c)$ to be the alternative with the

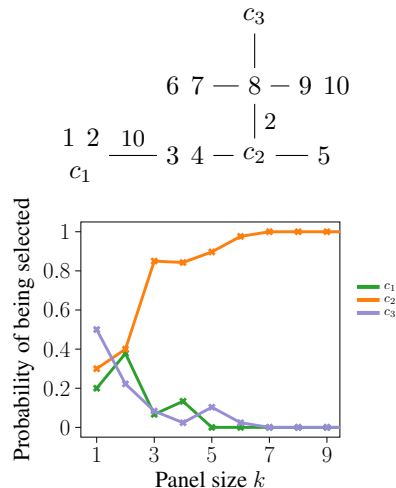


Figure 1. Metric spaces for Example 1. Edges without labels have length 1. The graph indicates the probability of the different alternatives being the top choice of a uniformly selected panel, for different panel sizes.

lowest social cost for this panel, with ties among possible alternatives being broken arbitrarily.

While the locations of alternatives in the metric space may be unknown, the position of the agents is typically known. For example, each agent is characterized by a set of features crucial for a specific application, and its position in the metric space is a function of these features. Therefore, a *selection algorithm* \mathcal{A}_k takes as input only the locations of the agents and outputs a distribution over panels of size k . A panel is in the support of \mathcal{A}_k if it is returned with non-zero probability. We call a selection algorithm fair, if the probability of any agent being on the panel is equal to $\frac{k}{n}$. We pay special attention to *uniform selection*, \mathcal{U}_k , which selects each panel of size k with equal probability. Uniform selection is fair by definition.

We consider the decision of panels selected by an algorithm to be “good” if its social cost is not much larger than the social cost of the optimum $c(N)$. More formally, we define the distortion of a selection algorithm \mathcal{A}_k as the worst-case ratio of the expected social cost of the alternatives selected by panels in the support of \mathcal{A}_k and the minimum possible social cost over any instance, i.e.,

$$\text{dist}(\mathcal{A}_k) = \sup_d \frac{\mathbb{E}_{P \sim \mathcal{A}_k} [SC(c(P))]}{SC(c(N))}.$$

We will call this as ex-ante distortion, as well.

Example 1. As an example consider the graph depicted in the upper part of Figure 1 together with its shortest distance metric, e.g. the distance from agent 1 to alternative c_2 is $10 + 1 = 11$.

Now for the social costs, one can calculate that the total

distance to alternative c_1 is 101 to alternative c_2 is 39 and to alternative c_3 is 49. If a mechanism would select $\{1, 2\}$ as the panel, their best alternative would be c_1 leading to a distortion of $\frac{101}{39} \sim 2.59$, while for instance $\{5, 8\}$ would have c_2 as their best alternative with a distortion of 1. In the lower part of Figure 1 we plot the probability for each alternative of being $c(P)$ for a uniformly at random selected panel P of size k , with tie-breaking in favor of c_1 . For instance, for $k = 2$, the probability for alternative c_1 is $\frac{17}{45}$, for c_2 is $\frac{2}{5}$, and for c_3 it is $\frac{2}{9}$. Thus the distortion of uniform selection for $k = 2$ in this instance is

$$\frac{\frac{17}{45}101 + \frac{2}{5}39 + \frac{2}{9}49}{39} = \frac{64.6\bar{4}}{39} \sim 1.66.$$

4. The Benefit of Fair Selection

Before we turn to randomized mechanisms, we start the investigation of distortion with a general lower bound on the distortion of *any* deterministic mechanism. We show that for any panel size k there is an instance in which no deterministic mechanism can reach a distortion of better than 5.

Theorem 1. *Let k, n be positive integers with $k \leq n$ and \mathcal{A}_k a deterministic algorithm that selects a panel of size k among n agents. For any $\varepsilon > 0$, there exists an instance with n agents and two alternatives in which \mathcal{A}_k has distortion of at least $5 - \frac{12k}{n+2k} - \varepsilon$.*

Proof. Consider an instance with two alternatives a and b and a set N of agents at distance 2 from each other. Alternative a has distance $3 - \varepsilon$ from each agent in \mathcal{A}_k and distance $5 - \varepsilon$ from each agent in set $N \setminus \mathcal{A}_k$. Alternative b has distance 3 from each agent in set \mathcal{A}_k and distance 1 from each agent in set $N \setminus \mathcal{A}_k$. Finally, we can set $d(a, b) = 4 - \varepsilon$. It can be easily seen that the distances among agents and alternatives form a metric. Clearly, the panel \mathcal{A}_k will select alternative a , who has social cost $\text{SC}(a) = (5 - \varepsilon)(n - k) + (3 - \varepsilon)k = (5 - \varepsilon)n - 2k$. The social cost of alternative b is $\text{SC}(b) = n - k + 3k = n + 2k$. Thus, the distortion is $\frac{(5 - \varepsilon)n - 2k}{n + 2k} \geq 5 - \frac{12k}{n + 2k} - \varepsilon$. \square

So, deterministic algorithms that select panels of size much smaller than the agent population cannot have a distortion considerably better than 5. In contrast, we can show that *any* fair selection algorithm achieve a strictly better ex-ante distortion.

Theorem 2. *Let k, n be positive integers with $k \leq n$ and \mathcal{A}_k a fair algorithm that selects an assembly of size k among n agents. Then, the ex-ante distortion of \mathcal{A}_k is at most $3 - \frac{2k}{n}$.*

We note that for the case of $k = 1$ this is the bound for the famous random dictatorship mechanism (Anshelevich

& Postl, 2016). Comparing the exact bound derived in Theorem 1, this shows that to even theoretically being able to match the ex-ante distortion of fair mechanisms for $k = 1$, a deterministic mechanism would need to have at least a panel size of $k = \frac{n}{4}$.

We can further show that the bound in Theorem 2 is tight.

Theorem 3. *For every $\varepsilon > 0$ and positive integers k, n with $k \leq n$, there exists an instance in which any selection algorithm \mathcal{A}_k has an ex ante distortion of at least $3 - \frac{2k}{n} - \varepsilon$.*

Note that the instance constructed in the proof of Theorem 3 requires an exponential number of alternatives in terms of the number of agents. However, in the next sections, we see that when the number of alternatives does not depend on the number of agents, we can achieve almost optimal distortion, by sampling only logarithmically many number of agents in terms on the number of alternatives m .

5. Uniform Selection

We now consider uniform selection and show that panels of size $\mathcal{O}(\varepsilon^{-2} \ln \frac{m}{\varepsilon})$ yield ex-ante distortion of $1 + \varepsilon$ for every instance with m alternatives.² To prove the next theorem, we need to bound the contribution to the expected social cost of the alternative that can be returned by the panel, distinguishing between three classes of alternatives. Alternatives that can incur (ex-post) distortion at most $1 + \varepsilon/2$ are easy to handle. To take care of alternatives of distortion between $1 + \varepsilon/2$ and 73 , we use Serfling’s inequality for the concentration of sampling without replacement (Serfling, 1974). However, applying this to alternatives incurring distortion higher than 73 would introduce a (logarithmic) dependency on the number of agents, since the distortion of such alternatives can be $\Omega(n)$. So, to bound the contribution of these alternatives, we construct an explicit hypergeometric random variable underlying the selection process and use it instead to bound the selection probability appropriately.

Theorem 4. *For every $\varepsilon \in (0, 1]$ and integer $m > 0$, there exists integer $k \in \mathcal{O}(\varepsilon^{-2} \ln \frac{m}{\varepsilon})$, so that a uniformly random k -sized panel of agents, has ex ante distortion at most $1 + \varepsilon$, for every set of m alternatives.*

Proof. Let $\varepsilon \in (0, 1]$ and k be the smallest integer multiple of 3 that is at least $\max\{\frac{25}{2\varepsilon^2} \ln \frac{144m}{\varepsilon}, 3 + 3 \log \frac{72m}{\varepsilon}\}$. Clearly, $k \in \mathcal{O}(\varepsilon^{-2} \ln \frac{m}{\varepsilon})$.³ We will abbreviate $c(N)$ as c' . We denote by W the set of alternatives of social cost at most $(1 + \frac{\varepsilon}{2}) \cdot \text{SC}(c')$. We also denote by P the random panel selected. Unless otherwise specified, expectations and

²Throughout this paper we assume that $\varepsilon \rightarrow 0$ for the \mathcal{O} -notation.

³Here \ln is the logarithm with base e and \log is the logarithm with base 2.

probabilities are defined over the selection of set P among all k -sized subsets of agents uniformly at random. Then,

$$\begin{aligned}
 & \mathbb{E}[\text{SC}(c(P))] \\
 & \leq \left(1 + \frac{\varepsilon}{2}\right) \cdot \Pr[c(P) \in W] \cdot \text{SC}(c') \\
 & \quad + \sum_{a \in C \setminus W} \Pr[c(P) = a] \cdot \text{SC}(a) \\
 & \leq \left(1 + \frac{\varepsilon}{2}\right) \cdot \left(1 - \sum_{a \in C \setminus W} \Pr[c(P) = a]\right) \text{SC}(c') \\
 & \quad + \sum_{a \in C \setminus W} \Pr[c(P) = a] \cdot \text{SC}(a) \\
 & \leq \left(1 + \frac{\varepsilon}{2}\right) \cdot \text{SC}(c') \\
 & \quad + \sum_{a \in C \setminus W} \Pr[c(P) = a] \cdot (\text{SC}(a) - \text{SC}(c')). \quad (1)
 \end{aligned}$$

We will show that, for every alternative $a \in C \setminus W$, it holds that

$$\Pr[c(P) = a] \leq \frac{\varepsilon \cdot \text{SC}(c')}{2m \cdot (\text{SC}(a) - \text{SC}(c'))}. \quad (2)$$

This, using (2), equation (1) will yield

$$\begin{aligned}
 \mathbb{E}[\text{SC}(c(P))] & \leq \left(1 + \frac{\varepsilon}{2}\right) \cdot \text{SC}(c') + \frac{\varepsilon |C \setminus W|}{2m} \cdot \text{SC}(c') \\
 & \leq (1 + \varepsilon) \cdot \text{SC}(c'),
 \end{aligned}$$

as desired.

In the following, we prove inequality (2) for every alternative $a \in C \setminus W$. We distinguish between alternatives in $C \setminus W$ of low and high social cost.

Case 1: $\text{SC}(a) - \text{SC}(c') \leq 72 \cdot \text{SC}(c')$. In this case we use a concentration inequality by [Serfling \(1974\)](#).

Theorem 5 (Serfling (1974)). *Let c_1, c_2, \dots, c_n be n real numbers from the interval $[\alpha, \beta]$ and consider the process of selecting k of them uniformly at random without replacement. Let X be the random variable denoting the sum of the k numbers selected. Then, $\mathbb{E}[X] = \frac{k}{n} \sum_{i=1}^n c_i$ and for every $t > 0$, it holds that*

$$\Pr[X - \mathbb{E}[X] \geq t] \leq \exp\left(-\frac{2t^2}{k(\beta - \alpha)^2}\right).$$

Recall that for any panel $P \subseteq N$ of agents,

$$\text{SC}(P, c') - \text{SC}(P, a) = \sum_{i \in P} (d(i, c') - d(i, a)).$$

Furthermore, using $D = d(a, c')$, by the triangle inequality, we have

$$-D \leq d(i, c') - d(i, a) \leq D,$$

for every agent $i \in N$. Thus, when selecting panel P uniformly at random among all k -sized sets of agents, the random variable $\text{SC}(P, c') - \text{SC}(P, a)$ follows the distribution of random variable X in [Theorem 5](#) with $c_i = d(i, c') - d(i, a)$ for $i \in N$, $\alpha = -D$, and $\beta = D$. Then, [Theorem 5](#) yields

$$\mathbb{E}[\text{SC}(P, c') - \text{SC}(P, a)] = -\frac{k}{n} (\text{SC}(a) - \text{SC}(c'))$$

and, using $t = \frac{k}{n} \cdot (\text{SC}(a) - \text{SC}(c'))$,

$$\begin{aligned}
 & \Pr[\text{SC}(P, c') \geq \text{SC}(P, a)] \\
 & \leq \exp\left(-\frac{2k(\text{SC}(a) - \text{SC}(c'))^2}{n^2 D^2}\right). \quad (3)
 \end{aligned}$$

Notice that,

$$\begin{aligned}
 \text{SC}(a) - \text{SC}(c') & \geq \frac{\varepsilon}{5} (\text{SC}(a) + \text{SC}(c')) \\
 & = \frac{\varepsilon}{5} \cdot \sum_{i \in N} (d(i, a) + d(i, c')) \\
 & \geq \frac{\varepsilon n D}{5}. \quad (4)
 \end{aligned}$$

The first inequality follows by our assumption that $a \in C \setminus W$ and since $\varepsilon \in (0, 1]$ and the second one by the triangle inequality.

Now, notice that an alternative a is selected by the panel P only if $\text{SC}(P, c') \geq \text{SC}(P, a)$. Using this observation, together with equations (4) and (3), and the facts that $k \geq \frac{25}{2\varepsilon^2} \ln \frac{144m}{\varepsilon}$ and $\text{SC}(a) - \text{SC}(c') \leq 72 \cdot \text{SC}(c')$, we obtain

$$\begin{aligned}
 \Pr[c(P) = a] & \leq \Pr[\text{SC}(P, c') \geq \text{SC}(P, a)] \\
 & \leq \exp(-2k\varepsilon^2/25) \leq \frac{\varepsilon}{144m} \\
 & \leq \frac{\varepsilon \cdot \text{SC}(c')}{2m \cdot (\text{SC}(a) - \text{SC}(c'))},
 \end{aligned}$$

completing the proof of inequality (2) in Case 1.

Case 2: $\text{SC}(a) - \text{SC}(c') \geq 72 \cdot \text{SC}(c')$. Define $L = \{i \in N : d(i, c') > \frac{D}{4}\}$ and let $\ell = |L|$. We will prove two useful lemmas.

Lemma 6. *It holds that $\ell \leq \frac{4n \cdot \text{SC}(c')}{\text{SC}(a) - \text{SC}(c')}$.*

Proof. By the definition of set L , we have $\text{SC}(c') \geq \sum_{i \in L} d(i, c') > \frac{D\ell}{4}$. Thus, we have

$$\begin{aligned}
 \text{SC}(a) - \text{SC}(c') & = \sum_{i \in N} (d(i, a) - d(i, c')) \\
 & \leq nD \leq \frac{4n \cdot \text{SC}(c')}{\ell}.
 \end{aligned}$$

□

Lemma 7. For every k -sized set P of agents, it that holds $c(P) = a$ only if $|L \cap P| \geq k/3$.

Proof. Assume otherwise that $|L \cap P| < k/3$. Using the triangle inequality and the definition of set L , it holds that $d(i, a) - d(i, c') \geq D - 2d(i, c') \geq D/2$ for $i \notin L$. Also, again by the triangle inequality, $d(i, a) - d(i, c') \geq -D$ for $i \in L$. Thus,

$$\begin{aligned} \text{SC}(P, a) - \text{SC}(P, c') &= \sum_{i \in P \cap L} (d(i, a) - d(i, c')) \\ &\quad + \sum_{i \in P \setminus L} (d(i, a) - d(i, c')) \\ &\geq |P \setminus L| \cdot D/2 - |P \cap L| \cdot D \\ &= kD/2 - |P \cap L| \cdot 3D/2 > 0, \end{aligned}$$

contradicting the assumption that $c(P) = a$. \square

Lemma 7 implies that $\Pr[c(P) = a] = 0$ if $\ell < k/3$. In the following, we bound the probability $\Pr[c(P) = a]$ assuming that $\ell \geq k/3$.

Notice that the quantity $|P \cap L|$ follows a hypergeometric distribution with a population size of n , ℓ success states, and k draws. The mode of the distribution of $|P \cap L|$ is at most $\frac{(k+1)(\ell+1)}{n+2}$ which, due to Lemma 6 and the assumption $\text{SC}(a) - \text{SC}(c') \geq 72 \cdot \text{SC}(c')$ in Case 2, is at most $\frac{(k+1) \cdot (n/18+1)}{n+2} < k/3$. Thus, $\Pr[|P \cap L| = j]$ is decreasing in j for $j \geq k/3$ and

$$\begin{aligned} \Pr[|P \cap L| \geq k/3] &= \sum_{j=k/3}^k \Pr[|P \cap L| = j] \\ &\leq \left(\frac{2k}{3} + 1\right) \Pr[|P \cap L| = k/3] \\ &= \left(\frac{2k}{3} + 1\right) \frac{\binom{\ell}{k/3} \cdot \binom{n-\ell}{2k/3}}{\binom{n}{k}} \\ &= \left(\frac{2k}{3} + 1\right) \binom{k}{k/3} \cdot \prod_{i=1}^{k/3} \frac{(\ell - k/3 + i)}{n - k/3 + i} \\ &\quad \cdot \prod_{i=1}^{2k/3} \frac{(n - \ell - 2k/3 + i)}{n - k + i} \\ &\leq \left(\frac{2k}{3} + 1\right) \binom{k}{k/3} \left(\frac{\ell}{n}\right)^{k/3} \leq \left(\frac{9\ell}{n}\right)^{k/3}. \quad (5) \end{aligned}$$

The second inequality follows since $k/3 \leq \ell \leq n$ and the third one by verifying that $\left(\frac{2k}{3} + 1\right) \cdot \binom{k}{k/3} \leq 9^{k/3}$ for every $k \geq 3$ which is a multiple of 3.

Now, we use Lemma 7, inequality (5), the fact that $k \geq 3 + 3 \log \frac{72m}{\varepsilon}$, Lemma 6, and the assumption $\text{SC}(a) - \text{SC}(c') \geq 72 \cdot \text{SC}(c')$ in Case 2 to get

$$\Pr[c(P) = a] \leq \Pr[|P \cap L| \geq k/3]$$

$$\begin{aligned} &\leq \left(\frac{9\ell}{n}\right)^{1 + \log \frac{72m}{\varepsilon}} \\ &\leq \left(\frac{36 \cdot \text{SC}(c')}{\text{SC}(a) - \text{SC}(c')}\right)^{1 + \log \frac{72m}{\varepsilon}} \\ &\leq \frac{36 \cdot \text{SC}(c')}{\text{SC}(a) - \text{SC}(c')} \cdot \left(\frac{1}{2}\right)^{\log \frac{72m}{\varepsilon}} \\ &= \frac{\varepsilon \cdot \text{SC}(c')}{2m \cdot (\text{SC}(a) - \text{SC}(c'))}. \end{aligned}$$

This completes the proof of inequality (2) in Case 2 and, subsequently, of Theorem 4. \square

We now show that Theorem 4 is nearly tight.

Theorem 8. Let $\varepsilon \in (0, 1/30]$, and n, m, k be positive integers such that n is even, $m \geq 1 + 6^9$, and $k \leq \min\left\{\frac{n}{2}, m-1, \frac{\ln(m-1)}{288\varepsilon^2}\right\}$. Then, there exists a set of n agents such that every probability distribution over panels of agents of size k has distortion more than $1 + \varepsilon$ for some instance with m alternatives.

6. Distortion Guarantees for Fair Greedy Capture

One downside of uniform selection is that it does not ensure that the decision of every realized panel is good for the whole population, as we see in the following example.

Example 2. Consider an instance with two alternatives, a and b , such that k agents have distance equal to 0 from a and distance equal to 1 from b , while the remaining $n - k$ agents have distance equal to 0 from b and distance equal to 1 from a . With non-zero probability, under uniform selection we can see a panel P that consists from the first k agents, and then $c(P) = a$. Then, $\text{SC}(c(P))/\text{SC}(c(N)) \geq (n - k)/k$. As k decreases this ratio goes to n .

Under this observation, in this section, we define the *ex-post* distortion of a selection algorithm \mathcal{A}_k as the worst-case ratio of the largest social cost of the alternatives selected by panels in the support of \mathcal{A}_k and the minimum possible social cost, i.e.,

$$\sup_d \sup_{P \in \text{supp}(\mathcal{A}_k)} \frac{\text{SC}(c(P))}{\text{SC}(c(N))}.$$

The above example indicates that uniform selection achieves bad ex-post distortion. This problem arises because uniform selection does not guarantee that any realized panel is a good representative of the whole population. The key question here is whether we can guarantee that *each* selected panel achieves good ex-post distortion, without affecting the expected distortion.

Algorithm 1 Fair Greedy Capture (Ebadian & Micha, 2024)

Input: $(N, d), k$
Output: Panel P
 $R \leftarrow N, \delta \leftarrow 0, P \leftarrow \emptyset;$
while $|R| \geq \lceil \frac{n}{k} \rceil$ **do**
 Smoothly increase $\delta;$
 while there exists $i \in R$ such that $|B(i, \delta)| \geq \lceil \frac{n}{k} \rceil$ **do**
 $S \leftarrow \lceil n/k \rceil$ agents from $B(i, \delta)$ by breaking ties arbitrarily among agents at the boundaries of the ball;
 Choose j from S uniformly at random;
 $P \leftarrow P \cup \{j\};$
 $R \leftarrow R \setminus S;$
 end while
end while
if $|P| < k$ **then**
 Sample at most one agent from R by picking each agent in R with probability equal to $1/\lceil n/k \rceil;$
 Sample the last $k - |P|$ members from $N \setminus P$ uniformly at random;
end if

We provide an affirmative answer to this question by revisiting Fair Greedy Capture (Algorithm 1) introduced by Ebadian & Micha (2024). This algorithm works as follows: It starts with an empty panel P and continuously grows a ball around every agent with rate δ , until there is a ball that contains at least $\lceil \frac{n}{k} \rceil$ agents.⁴ Then, one agent within this ball (if there are more than $\lceil n/k \rceil$ agents, we arbitrarily exclude agents at the boundaries of the ball) is selected uniformly at random and added to the panel and the rest of the $\lceil \frac{n}{k} \rceil$ agents are temporarily deleted. This process continues this way, until less than $\lceil \frac{n}{k} \rceil$ agents are left. Then, at most one agent is selected among the agents that have not been disregarded yet by picking each with probability $\lceil n/k \rceil$, and the last positions in the panel are filled by picking each agent that has not been selected so far with the same probability.⁵

Here, we show that any panel in the support of this algorithm achieves constant distortion when $k \geq 3$.

Theorem 9. *For any $k \geq 3$, Fair Greedy Capture guarantees ex-post distortion of at most 127.*

⁴This can be equivalently described by calculating the distance of each agent from her $\lceil n/k \rceil$ closest agent, and pick the ball that is centered at the agent with the smallest such distance and has a radius equal to this distance.

⁵This description is slightly different than the one in the original paper. We can see that under this version again each agent is selected with probability equal to k/n , since before the last step of the algorithm each agent is selected with probability exactly equal to $1/\lceil n/k \rceil$, and at the last step the remaining probability is distributed equally among the agents that have not been selected yet.

Further, we show that for $k \leq 2$ there exists an instance, in which Fair Greedy Capture can have linear ex-post distortion in terms of the size of the population.

Proposition 10. *For $k \leq 2$ there exists an instance in which Fair Greedy Capture has $\Omega(n)$ ex-post distortion.*

Proof. When $k = 1$, Fair Greedy Capture operates as uniform selection, and therefore the statement follows by Example 2.

For $k = 2$, consider n agents on the line, with one agent located at position 1 and $n - 1$ agents located at position 0. Fair Greedy Capture could select a panel consisting of one agent on 0 and the agent on 1. Now consider an instance with two alternatives, one at location $-\delta$ (for $\delta \in (0, 1)$) and one at 1. The panel would choose the alternative at 1. However, the ex-post distortion of this alternative is $\frac{n-1}{1+n\delta}$, which approaches $n - 1$ as δ approaches 0. \square

Now we exploit the fact that the outcomes from Fair Greedy Capture are at most a constant factor away from the social cost of the optimum alternative, together with some structural properties to show that $k = \mathcal{O}(\varepsilon^{-3} \ln \frac{m}{\varepsilon})$ is sufficient to reach a distortion of $1 + \varepsilon$.

Theorem 11. *For every $\varepsilon > 0$ and integer $m > 0$, there exists integer $k \in \mathcal{O}(\varepsilon^{-3} \ln \frac{m}{\varepsilon})$, so that selecting a k -sized panel of agents with Fair Greedy Capture, has ex-ante distortion at most $1 + \varepsilon$, for every set of m alternatives.*

Given that both uniform selection and Fair Greedy Capture are fair selection algorithms and that there is a general constant upper bound for fair selection algorithms, one might wonder whether $\text{poly}(\varepsilon^{-1}) \cdot \ln \frac{m}{\varepsilon}$ samples are sufficient for any fair selection algorithm. We give a simple example that this is not the case.

Theorem 12. *There exists a fair selection algorithm that has distortion arbitrarily close to 2 on instances with an even number of agents and two alternatives.*

Proof. Let $\delta \in (0, 1)$ and $k = n/2$. Consider an instance on the line with $n/2$ agents at location 0, $n/2$ agents at location 1, an alternative at location $-1 + \delta$ and an alternative at location 1. Let \mathcal{A}_k be the selection algorithm, which returns a panel consisting, equiprobably, of all agents at location 0 or all agents at location 1. The algorithm is clearly fair as each agent is selected with probability $k/n = 1/2$. Notice that the panel of agents located at position 0 will select the alternative at location $-1 + \delta$ of social cost $(3/2 - \delta)n$, and the panel of agents located at position 1 will select the alternative at location 1 of social cost $n/2$. The expected social cost of the alternative selected by \mathcal{A}_k is thus $(1 - \delta/2)n$, which yields an ex-ante distortion $2 - \delta$, approaching 2 as δ approaches 0. \square




Figure 2. Average distortion for k from 1 to 40. The figures include the 95% confidence interval.

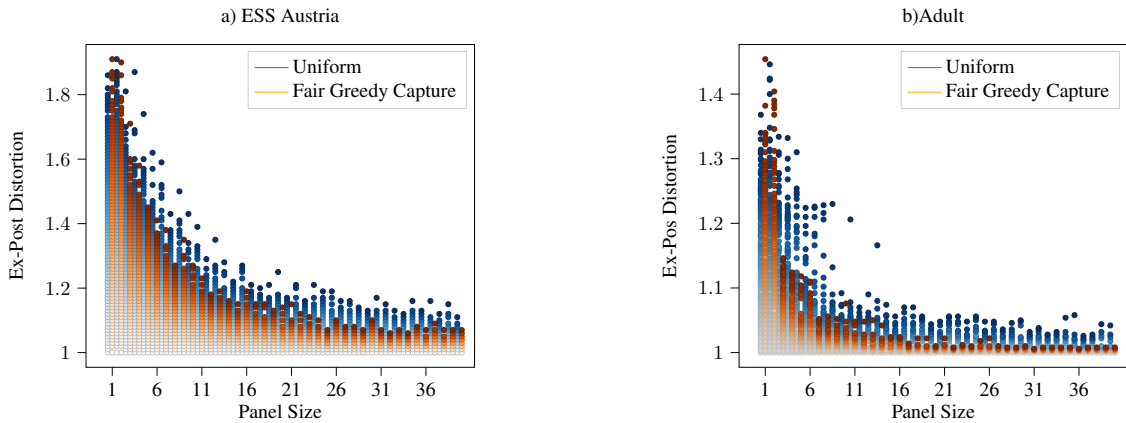


Figure 3. Distribution of ex-post distortion values for k from 1 to 49.

7. Experiments

We close off our paper with an experimental section studying the distortion achieved by both uniform selection and Fair Greedy Capture on real world data.

7.1. Experimental Setup

We use the same datasets used by Ebadian & Micha (2024) and Ebadian et al. (2022) in their papers: Firstly, the Adult dataset is probably one of the most popular datasets in research on algorithmic fairness (Fabris et al., 2022). It contains 32,561 data points from a 1994 US population survey. Following Ebadian & Micha (2024) we use the attributes sex, race, workclass, marital.status, and education.num from this dataset. Secondly, the European Social Survey (ESS) collection of datasets consists of population survey datasets

from multiple European countries. We also use the data from *ESS Round 9* from 2018, more specifically, the data from Austria ($n = 2498$).

Representation Metric Construction. We follow the method of Ebadian et al. (2022), and use the features of the datasets as a proxy for generating representation metrics. Similarly, with the aforementioned work, we make the natural assumption that the more characteristics two individuals share the “closer” they feel to each other. In more detail, for each dataset, we have a set of features $F = \{f_1, \dots, f_\ell\}$. If a feature is *categorical*, e.g., sex, the distance of two individuals with respect to this feature is equal to 1 if the values are different and 0 otherwise. For *continuous* features, the distance of two individuals with respect to this feature is equal to absolute difference of the two values, normalized by the maximum distance. We then uniformly at random

select a weight between 0 and 1 for each feature and set the weighted sum of all feature distances as the actual distances between the agents.

The original datasets contains only features of individuals, which are perceived as the agents. To construct the alternatives, we use the original agents as a proxy. That is, given a sampled panel P , we compare the agent (i.e. alternative) minimizing the distance to the panel P to the agent minimizing the distance to the whole population.

For each dataset, we average over 10 different generated metrics by randomly assigned weights to each feature and over 50 sample panels from each metric. Average results are plotted with 95% confidence intervals.

7.2. Our Results

In Figure 2, we see the average distortion of uniform selection and Fair Greedy Capture for ESS Austria (left) and Adult (right) datasets. In the ESS Austria dataset, we observe that the distortion of both algorithms converges to 1 more rapidly than in the Adult dataset. This difference is attributed to the larger number of alternatives ($m = 32561$) in the Adult dataset compared to the ESS Austria dataset ($m = 2498$). As discussed in Section 5 and Section 6, the size of panel required to achieve nearly optimal distortion depends on m .

Another interesting phenomenon is that while both uniform selection and Fair Greedy Capture behave very similarly for $k = 1$ and $k = 2$, Fair Greedy Capture seems to converge to a distortion of 1 more quickly than uniform selection and consistently achieves better distortion overall. The former similarity arises because when $k = 1$, uniform selection and Fair Greedy Capture perform identically. Moreover, as detailed in Section 6, when $k = 2$, Fair Greedy Capture can return panels that exhibit significant distortion. The latter observation is particularly surprising, indicating that despite theoretical bounds suggesting both methods require the same number of samples for convergence to almost optimal distortion, Fair Greedy Capture consistently outperforms uniform selection.

Ex-Post Distortion. In Figure 3, we can also observe the ex-post distortion of the two algorithms over any instance⁶ under the same datasets. As expected, for small values of k , both algorithms exhibit outliers, representing realized panels with high distortion. With increasing values of k , Fair Greedy Capture consistently achieves low distortion, in contrast to uniform selection, where the presence of outliers remains significant. Notably, in the Adult dataset, uniform selection shows more extreme outliers than in the ESS-

⁶Due to the large number of instances (500 in total), we rounded the values to the second digit to enhance visual clarity

Austria dataset. This difference is due to the fact that in the Adult dataset, a significant percentage of the population ($\sim 8.3\%$) is concentrated at the same point, and in some cases, uniform selection fails to choose enough representatives from this group. On the other hand, Fair Greedy Capture always ensures proportional representation of any sufficiently cohesive group.

8. Discussion

This work addresses a critical gap in understanding the ability of representative sortition panels to reflect the preferences of an entire population. There are several open questions. First, we show that Fair Greedy Capture achieves ex-post distortion equal to at most 127, and from Theorem 1, we know that no deterministic algorithm can ensure distortion better than 5. Closing this gap is an immediate open question. Another interesting direction is to show that indeed Fair Greedy Capture always achieves better ex-ante distortion than uniform selection, something that is observed in the experiments.

More broadly, in this work, we assume that the optimal alternative for a population is the one that minimizes the social cost. Various alternatives, such as the one minimizing the maximum cost or the one preferred by the majority in one-to-one comparisons, could be considered. Additionally, in real-life applications, a sortition panel may converge to multiple suggestions and not to just one. The question then arises: Do these suggestions proportionally reflect the opinions of different groups within the underlying population? These intriguing questions are left for future work.

Acknowledgments

The work of Ioannis Caragiannis was partially supported by the Independent Research Fund Denmark (DFR) under grant 2032-00185B. Jannik Peters was partially funded by the Deutsche Forschungsgemeinschaft (DFG) under the grant BR 4744/2-1 and the Graduiertenkolleg ‘‘Facets of Complexity’’ (GRK 2434).

Impact Statement

This paper explores two critical questions: (a) whether representative panels selected from a broader population can reach good decisions for the entire population, and (b) the optimal size these panels should be to achieve this goal. While this work is purely theoretical, the findings can have societal implications, particularly if the studied algorithms are implemented in real-world scenarios. Throughout the paper, we discuss the immediate disadvantages of the different approaches that should be taken into account before these algorithms are implemented, despite the positive results we

provide in their favor.

References

- Amanatidis, G., Birmpas, G., Filos-Ratsikas, A., and Voudouris, A. A. Don't roll the dice, ask twice: The two-query distortion of matching problems and beyond. In *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS)*, pp. 30665–30677, 2022.
- Anagnostides, I., Fotakis, D., and Patsilinakos, P. Metric-distortion bounds under limited information. *Journal of Artificial Intelligence Research*, 74:1449–1483, 2022.
- Anshelevich, E. and Postl, J. Randomized social choice functions under metric preferences. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 46–59. AAAI Press, 2016.
- Anshelevich, E., Bhardwaj, O., and Postl, J. Approximating optimal social choice under metric preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 777–783, 2015.
- Anshelevich, E., Filos-Ratsikas, A., Shah, N., and Voudouris, A. A. Distortion in social choice problems: The first 15 years and beyond. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 4294–4301, 2021.
- Benade, G., Gözl, P., and Procaccia, A. D. No stratification without representation. In *Proceedings of the 20th ACM Conference on Economics and Computation (ACM-EC)*, pp. 281–314, 2019.
- Boutilier, C., Caragiannis, I., Haber, S., Lu, T., Procaccia, A. D., and Sheffet, O. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227:190–213, 2015.
- Caragiannis, I. and Fehrs, K. Beyond the worst case: Distortion in impartial culture electorate. *CoRR*, abs/2307.07350, 2023.
- Charikar, M., Ramakrishnan, P., Wang, K., and Wu, H. Breaking the metric voting distortion barrier. In *Proceedings of the 35th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2024. Forthcoming.
- Chen, X., Fain, B., Lyu, L., and Munagala, K. Proportionally fair clustering. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pp. 1032–1041, 2019.
- Ebadian, S. and Micha, E. Boosting sortition via proportional representation. Technical report, arXiv:2406.00913 [cs.GT], 2024.
- Ebadian, S., Kehne, G., Micha, E., Procaccia, A. D., and Shah, N. Is sortition both representative and fair? In *Proceedings of the 36th Conference on Neural Information Processing Systems (NeurIPS)*, pp. 3431–3443, 2022.
- Ebadian, S., Filos-Ratsikas, A., Latifian, M., and Shah, N. Explainable and efficient randomized voting rules. In *Proceedings of the 36th Conference on Neural Information Processing Systems (NeurIPS)*, 2023. Forthcoming.
- Engelstad, F. The assignment of political office by lot. *Social Science Information*, 28(1):23–50, 1989.
- Fabris, A., Messina, S., Silvello, G., and Susto, G. A. Algorithmic fairness datasets: the story so far. *Data Mining and Knowledge Discovery*, 36(6):2074–2152, 2022.
- Flanigan, B., Gözl, P., Gupta, A., and Procaccia, A. D. Neutralizing self-selection bias in sampling for sortition. In *Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS)*, pp. 6528–6539, 2020.
- Flanigan, B., Gözl, P., Gupta, A., Hennig, B., and Procaccia, A. D. Fair algorithms for selecting citizens' assemblies. *Nature*, 596:548–552, 2021a.
- Flanigan, B., Kehne, G., and Procaccia, A. D. Fair sortition made transparent. In *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS)*, pp. 25720–25731, 2021b.
- Gillies, D. B. *Some theorems on n-person games*. Princeton University, 1953.
- Gkatzelis, V., Halpern, D., and Shah, N. Resolving the optimal metric distortion conjecture. In *Proceedings of the 61st Symposium on Foundations of Computer Science (FOCS)*, pp. 1427–1438, 2020.
- Gonczarowski, Y. A., Kehne, G., Procaccia, A. D., Schiffer, B., and Zhang, S. The distortion of binomial voting defies expectation. In *Proceedings of the 36th Conference on Neural Information Processing Systems (NeurIPS)*, 2023. Forthcoming.
- Jaworski, M. and Skowron, P. Phragmén rules for degressive and regressive proportionality. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 328–334. IJCAI, 2022.
- Jung, C., Kannan, S., and Lutz, N. Service in your neighborhood: Fairness in center location. In *Proceedings of the 1st Symposium on Foundations of Responsible Computing (FORC)*, pp. 5:1–5:15, 2020.
- Kellerhals, L. and Peters, J. Proportional fairness in clustering: A social choice perspective. Technical report, arXiv:2310.18162 [cs.LG], 2023.

- Kızılkaya, F. E. and Kempe, D. Pluralityveto: A simple voting rule achieving optimal metric distortion. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 349 – 355, 2022.
- Li, B., Li, L., Sun, A., Wang, C., and Wang, Y. Approximate group fairness for clustering. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, pp. 6381–6391, 2021.
- Meir, R., Sandomirskiy, F., and Tennenholtz, M. Representative committees of peers. *Journal of Artificial Intelligence Research*, 71(401–429), 2021.
- Procaccia, A. D. and Rosenschein, J. S. The distortion of cardinal preferences in voting. In *International Workshop on Cooperative Information Agents (CIA)*, pp. 317—331, 2006.
- Serfling, R. J. Probability inequalities for the sum in sampling without replacement. *The Annals of Statistics*, 2(1): 39–48, 1974.

A. Missing proofs of Section 4

Theorem 2. *Let k, n be positive integers with $k \leq n$ and \mathcal{A}_k a fair algorithm that selects an assembly of size k among n agents. Then, the ex-ante distortion of \mathcal{A}_k is at most $3 - \frac{2k}{n}$.*

Proof. Consider an alternative c' ; we will show that $\mathbb{E}[\text{SC}(c(\mathcal{A}_k))] \leq (3 - 2k/n) \cdot \text{SC}(c')$. For any k -sized set P of agents, we clearly have

$$\text{SC}(c(P)) = \text{SC}(P, c(P)) + \text{SC}(N \setminus P, c(P)). \quad (6)$$

By the definition of alternative $c(P)$, we have

$$\text{SC}(P, c(P)) \leq \text{SC}(P, c'). \quad (7)$$

Also, using the triangle inequality,

$$\begin{aligned} \text{SC}(N \setminus P, c(P)) &= \sum_{i \in N \setminus P} d(i, c(P)) \\ &\leq \sum_{i \in N \setminus P} (d(i, c') + d(c', c(P))) \\ &= \text{SC}(N \setminus P, c') + (n - k) \cdot d(c', c(P)). \end{aligned} \quad (8)$$

The triangle inequality implies that $d(c', c(P)) \leq d(i, c') + d(i, c(P))$ for every alternative $i \in P$. Thus,

$$\begin{aligned} d(c', c(P)) &\leq \frac{1}{k} \sum_{i \in P} (d(i, c') + d(i, c(P))) \\ &= \frac{1}{k} \text{SC}(P, c') + \frac{1}{k} \text{SC}(P, c(P)) \leq \frac{2}{k} \text{SC}(P, c'), \end{aligned} \quad (9)$$

where the last inequality follows by equation (7). Using equation (9), equation (8) yields

$$\text{SC}(N \setminus P, c(P)) \leq \text{SC}(N \setminus P, c') + 2 \frac{n - k}{k} \text{SC}(P, c'). \quad (10)$$

Now, using equations (7) and (10), equation (6) yields

$$\begin{aligned} \text{SC}(c(P)) &\leq \text{SC}(P, c') + \text{SC}(N \setminus P, c') + 2 \frac{n - k}{k} \text{SC}(P, c') \\ &= \text{SC}(c') + 2 \frac{n - k}{k} \text{SC}(P, c'). \end{aligned} \quad (11)$$

Thus, using equation (11), we have

$$\begin{aligned} \mathbb{E}[\text{SC}(c(\mathcal{A}_k))] &= \sum_{P \in N_k} \Pr[\mathcal{A}_k = P] \cdot \text{SC}(c(P)) \\ &\leq \sum_{P \in N_k} \Pr[\mathcal{A}_k = P] \cdot \text{SC}(c') \\ &\quad + 2 \frac{n - k}{k} \sum_{P \in N_k} \Pr[\mathcal{A}_k = P] \cdot \text{SC}(P, c'). \end{aligned} \quad (12)$$

The first sum in the RHS of equation (12) becomes

$$\sum_{P \in N_k} \Pr[\mathcal{A}_k = P] \cdot \text{SC}(c')$$

$$= \text{SC}(c') \cdot \sum_{P \in \mathcal{N}_k} \Pr[\mathcal{A}_k] = \text{SC}(c'), \quad (13)$$

while, since algorithm \mathcal{A}_k is fair, the second one becomes

$$\begin{aligned} & \sum_{P \in \mathcal{N}_k} \Pr[\mathcal{A}_k = P] \text{SC}(P, c') \\ &= \sum_{P \in \mathcal{N}_k} \Pr[\mathcal{A}_k = P] \cdot \sum_{i \in P} d(i, c') \\ &= \sum_{i \in N} d(i, c') \cdot \sum_{P \in \mathcal{N}_k: i \in P} \Pr[\mathcal{A}_k = P] \\ &= \sum_{i \in N} d(i, c') \cdot \Pr[i \in \mathcal{A}_k] \\ &= \frac{k}{n} \cdot \sum_{i \in N} d(i, c') = \frac{k}{n} \cdot \text{SC}(c'). \end{aligned} \quad (14)$$

Hence, using equations (13) and (14), equation (12) yields

$$\mathbb{E}[\text{SC}(c(\mathcal{A}_k))] \leq \left(3 - \frac{2k}{n}\right) \cdot \text{SC}(c'),$$

as desired. \square

Theorem 3. *For every $\varepsilon > 0$ and positive integers k, n with $k \leq n$, there exists an instance in which any selection algorithm \mathcal{A}_k has an ex ante distortion of at least $3 - \frac{2k}{n} - \varepsilon$.*

Proof. Consider the instance with a set N of n agents and $\binom{n}{k} + 1$ alternatives defined as follows. There is an alternative c at distance $3 + \varepsilon$ from every agent in N . For every distinct set K of k agents from N , there exists an alternative c_K at distance 3 from the agents in K and distance 9 from the agents in $N \setminus K$. Notice that every panel K returned by algorithm \mathcal{A}_k with positive probability will select alternative c_K , who has social welfare $\text{SC}(c_K) = 3k + 9(n - k) = 9n - 6k$. Alternative c has minimum social welfare of $\text{SC}(c) = (3 + \varepsilon)n$. The ex ante distortion of algorithm \mathcal{A}_k is then $\frac{9n - 6k}{(1 + \varepsilon)n} = 3 - \frac{2k}{n} - \frac{\varepsilon(3n - 2k)}{(3 + \varepsilon)n} \geq 3 - \frac{2k}{n} - \varepsilon$. \square

B. Proof of Theorem 8

Theorem 8. *Let $\varepsilon \in (0, 1/30]$, and n, m, k be positive integers such that n is even, $m \geq 1 + 6^9$, and $k \leq \min\left\{\frac{n}{2}, m - 1, \frac{\ln(m-1)}{288\varepsilon^2}\right\}$. Then, there exists a set of n agents such that every probability distribution over panels of agents of size k has distortion more than $1 + \varepsilon$ for some instance with m alternatives.*

Proof. We will define a family of instances with the same set of n agents and different sets of m alternatives and a probability distribution \mathcal{I} to select a random instance among them. All instances have an alternative of social cost $(2 - 4\varepsilon)n$, and all other alternatives of social cost $2n$. Then, we will show that any k -sized panel P satisfies $\mathbb{E}_{\mathcal{I}}[\text{SC}(c(P))] \geq (2 - 2\varepsilon)n$. Notice that $\mathbb{E}_{\mathcal{I}}[\text{SC}(c(P))]$ is the expectation of the social cost of alternative $c(P)$, taken over the instances produced by \mathcal{I} . Thus, for every selection algorithm \mathcal{A}_k , the expectation, taken over the random selection of the instance according to \mathcal{I} , of the expected social cost of alternative $c(\mathcal{A}_k)$ is at least $(2 - 2\varepsilon)n$. Then, using the probabilistic method, we have that there exists a single instance for which $\mathbb{E}[\text{SC}(c(\mathcal{A}_k))] \geq (2 - 2\varepsilon)n$. The ex-ante distortion of \mathcal{A}_k , when applied to this instance, will then be $\frac{(2 - 2\varepsilon)n}{(2 - 4\varepsilon)n} > 1 + \varepsilon$; this will complete the proof.

Our construction of the family of instances is as follows. First, there is a set N of n agents, who have distance 2 to each other. Let \mathcal{S} be the collection of all subsets of $n/2$ agents from N . For every different $(m - 1)$ -tuple $(S_1, S_2, \dots, S_{m-1})$ of subsets from \mathcal{S}^{m-1} , we define the instance with alternatives c_1, c_2, \dots, c_{m-1} so that alternative c_i is at distance 1 from the agents in S_i and at distance 3 from the agents in $N \setminus S_i$. All instances have an additional alternative c_0 at distance $2 - 4\varepsilon$ from every agent. We can easily verify that the distances between agents and alternatives form a metric. Clearly, in all instances defined in this way, alternative c_0 has social cost $(2 - 4\varepsilon)n$, while all other alternatives have social cost $2n$.

Now, the probability distribution \mathcal{I} returns the instance identified by the $(m-1)$ -tuple $(S_1, S_2, \dots, S_{m-1})$ by selecting, for $i = 1, 2, \dots, m-1$, set S_i uniformly at random among all elements of \mathcal{S} , so that the selection is independent for the different i s.

We will prove the desired property of the probability distribution \mathcal{I} using the following lemma.

Lemma 13. *Consider a panel P of k agents. Then, $\Pr_{\mathcal{I}}[c(P) = c_0] \leq \frac{1}{2}$.*

The crucial property our construction and the probability distribution \mathcal{I} have is that, for $i \in [m-1]$, the quantity $|P \cap S_i|$ is a random variable following the hypergeometric distribution with n as population size, $n/2$ success states, and k draws. Furthermore, the random variables $|P \cap S_i|$ and $|P \cap S_j|$ are independent for $i \neq j$.

Let us warm up by proving the lemma assuming that $k \leq 9$. Using the inequality $(t/s)^s \leq \binom{t}{s} \leq (et/s)^s$ for positive integers s and t with $t \geq s$, we have that, for $i \in [m-1]$, the probability that all agents in P belong to set S_i is at least

$$\frac{\binom{n/2}{k}}{\binom{n}{k}} \geq \frac{\left(\frac{n}{2k}\right)^k}{\left(\frac{en}{k}\right)^k} = \frac{1}{(2e)^k} \geq \frac{1}{6^9}.$$

Notice that we would have $\text{SC}(P, c_i) = k < \text{SC}(P, c_0)$ in this case. Thus, omitting \mathcal{I} from notation and writing just \Pr instead of $\Pr_{\mathcal{I}}$, we can bound the probability that panel P selects alternative c_0 as follows:

$$\Pr[c(P) = c_0] \leq \left(1 - \frac{1}{6^9}\right)^{m-1} \leq 1/e < 1/2.$$

The second inequality follows since $m \geq 1 + 6^9$ and using the inequality $(1 - 1/t)^t \leq 1/e$ for $t > 1$.

In the following, we assume that $k \geq 10$. Observe that alternative c_0 is selected by the panel only if $\text{SC}(P, c_i) \geq \text{SC}(P, c_0)$ for every $i \in [m-1]$. Furthermore, observe that $\text{SC}(P, c_i) = |P \cap S_i| + 3(k - |P \cap S_i|) = 3k - 2|P \cap S_i|$ and recall that $\text{SC}(P, c_0) = (2 - 2\varepsilon)k$. Thus, the condition $\text{SC}(P, c_i) \geq \text{SC}(P, c_0)$ is equivalent to $|P \cap S_i| \leq \frac{k}{2} + 2\varepsilon k$. Using these observations, we have

$$\begin{aligned} \Pr[c(P) = c_0] &\leq \prod_{i=1}^{m-1} \Pr[\text{SC}(P, c_i) \geq \text{SC}(P, c_0)] \\ &= \prod_{i=1}^{m-1} \Pr\left[|P \cap S_i| \leq \frac{k}{2} + 2\varepsilon k\right] \\ &= \prod_{i=1}^{m-1} \left(1 - \Pr\left[|P \cap S_i| > \frac{k}{2} + 2\varepsilon k\right]\right). \end{aligned} \tag{15}$$

Now, let ℓ^* be the smallest integer that is strictly higher than $\frac{k}{2} + 2\varepsilon k$. Then, inequality (15) implies that

$$\Pr[c(P) = c_0] \leq \prod_{i=1}^{m-1} (1 - \Pr[|P \cap S_i| = \ell^*]). \tag{16}$$

To continue, we will need an anti-concentration bound for the random variable $|P \cap S_i|$. The proof appears in Appendix C.

Lemma 14. *Let k be an integer such that $10 \leq k \leq n/2$. For every integer ℓ with $k/2 \leq \ell \leq 2k/3$, it holds that*

$$\Pr[|P \cap S_i| = \ell] \geq \frac{1}{\sqrt{k}} \cdot \left(\frac{k}{\ell} - 1\right)^{2\ell - k}$$

Notice that $\ell^* < 1 + k/2 + 2\varepsilon k$ and, using the assumptions $k \geq 10$ and $\varepsilon \leq 1/30$, we get $\ell^* < k(3/5 + 2\varepsilon) \leq 2k/3$. By applying Lemma 14 for $\ell = \ell^*$ and substituting the value of $\Pr[|P \cap S_i| = \ell^*]$ in (16), we get

$$\Pr[c(P) = c_0] \leq \left(1 - \frac{1}{\sqrt{k}} \left(\frac{k}{\ell^*} - 1\right)^{2\ell^* - k}\right)^{m-1} \tag{17}$$

We will complete the proof of Lemma 13 by distinguishing between two cases for ℓ^* .

Case 1. $\ell^* > \frac{k}{2} + 3\epsilon k$. Then, by the definition of ℓ^* , it is $\ell^* - 1 < \frac{k}{2} + 2\epsilon k$ which implies that $\epsilon k < 1$. Hence, $2\ell^* - k < 2 + 4\epsilon k < 6$, i.e., $2\ell^* - k \leq 5$ since both ℓ^* and k are integers. Recall that $\ell^* \leq 2k/3$. Then, we have $(\frac{k}{\ell^*} - 1)^{2\ell^* - k} \geq (1/2)^{2\ell^* - k} \geq 1/32$. Now, inequality (17) yields

$$\begin{aligned} \Pr[c(P) = c_0] &\leq \left(1 - \frac{1}{32\sqrt{k}}\right)^{m-1} \leq \exp\left(-\frac{m-1}{32\sqrt{k}}\right) \\ &\leq \exp\left(-\frac{\sqrt{m-1}}{32}\right) \leq 1/e < 1/2. \end{aligned}$$

The second inequality follows since $1 - t \leq e^{-t}$, the third one follows by our assumption $k \leq m - 1$, and the fourth one is due to the assumption $m \geq 1 + 6^9$.

Case 2. $\ell^* \leq \frac{k}{2} + 3\epsilon k$. Thus, $\frac{k}{\ell^*} - 1 \geq 1 - \frac{12\epsilon k}{k+6\epsilon k} \geq 1 - 12\epsilon$ and inequality (17) yields

$$\begin{aligned} \Pr[c(P) = c_0] &\leq \left(1 - \frac{1}{\sqrt{k}}(1 - 12\epsilon)^{6\epsilon k}\right)^{m-1} \\ &\leq \left(1 - \frac{1}{\sqrt{k}} \cdot e^{-144\epsilon^2 k}\right)^{m-1} \\ &\leq \exp\left(-\frac{m-1}{\sqrt{k}} \cdot e^{-144\epsilon^2 k}\right) \\ &\leq \exp\left(-\sqrt{m-1} \cdot e^{-\frac{1}{2}\ln(m-1)}\right) \\ &= 1/e < 1/2. \end{aligned}$$

The second inequality follows since $\epsilon \in (0, 1/30]$ and using the property $1 - t \geq e^{-2t}$ for $t \in [0, 1/2]$. The third inequality follows since $1 - t \leq e^{-t}$ and the fourth one is due to our assumption $k \leq \frac{\ln(m-1)}{288\epsilon^2}$. \square

We are now ready to complete the proof of Theorem 8. Notice that Lemma 13 implies that, for any panel P of k agents,

$$\begin{aligned} \mathbb{E}_{\mathcal{I}}[\text{SC}(c(P))] &= (2 - 4\epsilon)n \cdot \Pr_{\mathcal{I}}[c(P) = c_0] \\ &\quad + 2n \cdot \Pr_{\mathcal{I}}[c(P) \neq c_0] \\ &= 2n - 4\epsilon n \cdot \Pr_{\mathcal{I}}[c(P) = c_0] \\ &\geq (2 - 2\epsilon)n. \end{aligned}$$

The proof of the desired property of distribution \mathcal{I} is complete, and the theorem follows.

C. Proof of Lemma 14

Lemma 14. *Let k be an integer such that $10 \leq k \leq n/2$. For every integer ℓ with $k/2 \leq \ell \leq 2k/3$, it holds that*

$$\Pr[|P \cap S_i| = \ell] \geq \frac{1}{\sqrt{k}} \cdot \left(\frac{k}{\ell} - 1\right)^{2\ell - k}$$

Proof. The random variable $|P \cap S_i|$ follows the hypergeometric distribution with a population size of n , $n/2$ success states, and k draws. Thus,

$$\Pr[|P \cap S_i| = \ell] = \frac{\binom{n/2}{\ell} \binom{n/2}{k-\ell}}{\binom{n}{k}} = \frac{(n/2)!^2 \cdot k! \cdot (n-k)!}{n! \cdot \ell! \cdot (n/2 - \ell)! \cdot (k - \ell)! \cdot (n/2 - k + \ell)!}. \quad (18)$$

We will now use Stirling's approximation which states that $t! = \alpha(t) \cdot \beta(t) \cdot \gamma(t)$ with $\alpha(t) = \sqrt{2\pi t}$, $\beta(t) = (\frac{t}{e})^t$, and $\gamma(t) \in [e^{\frac{1}{12t+1}}, e^{\frac{1}{12t}}]$. Then, equation (18) can be written as $\Pr[|P \cap S_i| = \ell] = A \cdot B \cdot \Gamma$, where

$$A = \frac{\alpha(n/2)^2 \cdot \alpha(k) \cdot \alpha(n-k)}{\alpha(n) \cdot \alpha(\ell) \cdot \alpha(n/2 - \ell) \cdot \alpha(k - \ell) \cdot \alpha(n/2 - k + \ell)},$$

$$B = \frac{\beta(n/2)^2 \cdot \beta(k) \cdot \beta(n-k)}{\beta(n) \cdot \beta(\ell) \cdot \beta(n/2 - \ell) \cdot \beta(k - \ell) \cdot \beta(n/2 - k + \ell)}, \text{ and}$$

$$\Gamma = \frac{\gamma(n/2)^2 \cdot \gamma(k) \cdot \gamma(n-k)}{\gamma(n) \cdot \gamma(\ell) \cdot \gamma(n/2 - \ell) \cdot \gamma(k - \ell) \cdot \gamma(n/2 - k + \ell)}.$$

Using the definition of $\alpha(t)$, we have

$$A = \sqrt{\frac{kn(n-k)}{2\pi\ell(n-2\ell)(k-\ell)(n-2k+2\ell)}} \geq \sqrt{\frac{k^3}{4\pi\ell^2(k-\ell)^2}} \geq \frac{2}{\sqrt{\pi k}}.$$

The first inequality follows by observing that the assumption $k \leq n/2$ implies that $\frac{n}{n-2\ell} \geq \frac{k}{k-\ell}$ and $\frac{n-k}{n-2k+2\ell} \geq \frac{k}{2\ell}$. The second inequality follows since the expression $\ell^2(k-\ell)^2$ is maximized for $\ell = k/2$ at $k^4/16$.

Using the definition of $\beta(t)$, we have

$$\begin{aligned} B &= \frac{\left(\frac{n}{2e}\right)^n \cdot \left(\frac{k}{e}\right)^k \cdot \left(\frac{n-k}{e}\right)^{n-k}}{\left(\frac{n}{e}\right)^n \cdot \left(\frac{\ell}{e}\right)^\ell \cdot \left(\frac{n-2\ell}{2e}\right)^{n/2-\ell} \cdot \left(\frac{k-\ell}{e}\right)^{k-\ell} \cdot \left(\frac{n-2k+2\ell}{2e}\right)^{n/2-k+\ell}} \\ &= \frac{k^k \cdot (n-k)^{n-k}}{2^k \cdot \ell^\ell \cdot (n-2\ell)^{n/2-\ell} \cdot (k-\ell)^{k-\ell} \cdot (n-2k+2\ell)^{n/2-k+\ell}} \\ &= \left(\frac{k(n-2k+2\ell)}{2(k-\ell)(n-k)}\right)^k \cdot \left(\frac{(k-\ell)(n-2\ell)}{\ell(n-2k+2\ell)}\right)^\ell \cdot \left(\frac{(n-k)(n-k)}{(n-2\ell)(n-2k+2\ell)}\right)^{n/2} \\ &\geq \left(\frac{\ell}{k-\ell}\right)^k \cdot \left(\frac{k-\ell}{\ell}\right)^{2\ell} = \left(\frac{k}{\ell} - 1\right)^{2\ell-k}. \end{aligned}$$

The inequality follows since the assumption $k/\ell \geq n/2$ implies that $\frac{n-2k+2\ell}{n-k} \geq \frac{2\ell}{k}$ and $\frac{n-2\ell}{n-2k+2\ell} \geq \frac{k-\ell}{\ell}$, while among all pairs of positive numbers with a given sum their product is maximized when they are equal. Thus, $\frac{(n-k)(n-k)}{(n-2\ell)(n-2k+2\ell)} \geq 1$.

Finally, using the definition of $\gamma(t)$, we have

$$\begin{aligned} \Gamma &\geq \frac{1}{\gamma(n) \cdot \gamma(\ell) \cdot \gamma(n/2 - \ell) \cdot \gamma(k - \ell) \cdot \gamma(n/2 - k + \ell)} \\ &\geq \exp\left(-\frac{1}{12n} - \frac{1}{12\ell} - \frac{1}{12(n/2 - \ell)} - \frac{1}{12(k - \ell)} - \frac{1}{12(n/2 - k + \ell)}\right) \\ &\geq \exp\left(-\frac{1}{24k} - \frac{1}{6k} - \frac{1}{4k} - \frac{1}{4k} - \frac{1}{6k}\right) = \exp\left(-\frac{7}{8k}\right) \geq \exp(-7/80) > \frac{\sqrt{\pi}}{2}. \end{aligned}$$

In the third inequality, we have used the inequalities $k \leq n/2$ and $\ell \in [k/2, 2k/3]$, while the fourth one follows since $k \geq 10$. We conclude that $\Pr[|P \cap S_i| = \ell] \geq \frac{1}{\sqrt{k}} \left(\frac{k}{\ell} - 1\right)^{2\ell-k}$, as desired. \square

D. Missing Proofs of Section 6

Theorem 9. *For any $k \geq 3$, Fair Greedy Capture guarantees ex-post distortion of at most 127.*

Proof. Let P be any panel of size k that has positive probability to be chosen by Fair Greedy Capture. Let $D = d(c(P), c(N))$ and denote by L the set of agents that have distance less than $D/12$ from $c(N)$, i.e., $L := \{i \in N : d(i, c(N)) < D/12\}$.

Note that

$$\begin{aligned} \text{SC}(c(P)) &= \sum_{i \in N} d(i, c(P)) \\ &\leq \sum_{i \in N \setminus L} (d(i, c(N)) + D) + \sum_{i \in L} (d(i, c(N)) + D) \end{aligned}$$

$$\begin{aligned} &\leq \sum_{i \in N \setminus L} 13 \cdot d(i, c(N)) + \sum_{i \in L} d(i, c(N)) + |L| \cdot D \\ &\leq 13 \cdot SC(c(N)) + |L| \cdot D. \end{aligned} \quad (19)$$

The first inequality follows from the triangle inequality. The second one is due to the fact that for each $i \in N \setminus L$, it holds that $d(i, c(N)) \geq D/12$. From the same fact, we also get that

$$SC(c(N)) \geq \sum_{i \in N \setminus L} d(i, c(N)) \geq (n - |L|) \cdot D/12. \quad (20)$$

Therefore, from Equation (19) and Equation (20), we have

$$SC(c(P)) \leq \left(13 + 12 \frac{|L|}{n - |L|}\right) \cdot SC(c(N)). \quad (21)$$

Next, we show that $|L|$ cannot be significantly large when $c(P) \neq c(N)$.

Suppose that $|L|$ satisfies $\ell \cdot \lceil n/k \rceil \leq \ell < (\ell + 1) \cdot \lceil n/k \rceil$ for some integer ℓ with $0 \leq \ell \leq k$. Note that since $|L| \geq \ell \cdot \lceil n/k \rceil$, we can partition L into at least ℓ disjoint parts such that each of them contains at least $\lceil n/k \rceil$ agents with distance less than $D/6$ from each other (since each $i \in L$ has distance less than $D/12$ from $c(N)$). This means that during the execution of Fair Greedy Capture, at least ℓ balls captured individuals from L and each such ball has diameter less than $D/6$. Therefore, there are at least ℓ representatives in P that have distance less than $D/4$ from $c(N)$ (since each such representative has distance less than $D/6$ from an agent in L and each agent in L has distance less than $D/12$ from $c(N)$). Thus, for the subset $P' := \{i \in P : d(i, c(N)) < D/4\}$ it holds that $|P'| \geq \ell$ (and therefore $|P \setminus P'| \leq k - \ell$). In other words, each panel contains at least ℓ representatives from L .

Since $c(P)$ is the alternative selected by panel P , we have

$$\begin{aligned} 0 &\geq SC(P, c(P)) - SC(P, c(N)) \\ &= \sum_{i \in P'} (d(i, c(P)) - d(i, c(N))) \\ &\quad + \sum_{i \in P \setminus P'} (d(i, c(P)) - d(i, c(N))) \\ &\geq \sum_{i \in P'} (D - 2d(i, c(N))) - \sum_{i \in P \setminus P'} D \\ &> \ell \cdot D/2 - (k - \ell) \cdot D = (3\ell/2 - k) \cdot D. \end{aligned}$$

The second inequality follows by the triangle inequality using $d(i, c(P)) \geq D - d(i, c(N))$ for $i \in P'$ and $d(i, c(P)) \geq d(i, c(N)) - D$ for $i \in P \setminus P'$. The third inequality follows by the definition of set P' . We conclude that $\ell < 2k/3$, which can be expressed as $\ell \leq \lceil \frac{2k}{3} \rceil - 1$ since k and ℓ are integers.

By the definition of ℓ , we have that $|L| < (\ell + 1) \cdot \lceil n/k \rceil$, implying that

$$|L| \leq \left\lceil \frac{2k}{3} \right\rceil \cdot \left\lceil \frac{n}{k} \right\rceil - 1. \quad (22)$$

Furthermore, panel has $k - \ell$ agents from set $N \setminus L$. Thus,

$$|L| \leq n - k + \ell \leq n - k + \left\lceil \frac{2k}{3} \right\rceil - 1. \quad (23)$$

We will use equations (22) and (23) to show that $|L| \leq 19n/21$; this is done in the next lemma.

Lemma 15.

$$\min \left\{ \frac{\lceil \frac{2k}{3} \rceil \cdot \lceil \frac{n}{k} \rceil - 1}{n}, 1 - \frac{k - \lceil \frac{2k}{3} \rceil + 1}{n} \right\} \leq 19/21.$$

Proof. Notice that if $n \leq 9k - 9 \lceil \frac{2k}{3} \rceil + 9$, then the second term in the min-expression above is at most $8/9$ and the claim is clearly true.

In the following, we assume that $n \geq 9k - 9 \lceil \frac{2k}{3} \rceil + 9$. Notice that the first term in the min-expression is equal to $\frac{\lceil \frac{2k}{3} \rceil \cdot t - 1}{n}$ for $n = (t-1)k + 1, (t-1)k + 2, \dots, tk$ and integer $t \geq 1$, and is thus decreasing for values of n in that range. We will show that the value $\frac{\lceil \frac{2k}{3} \rceil \cdot t - 1}{(t-1)k + 1}$ of the first term for $n = (t-1)k + 1$ is at most $19/21$, for all possible integers $t \geq 1$ and $k \geq 3$.

Actually, the assumption $n \geq 9k - 9 \lceil \frac{2k}{3} \rceil + 9$ implies that we have to consider values of t satisfying $(t-1)k + 1 \geq 9k - 9 \lceil \frac{2k}{3} \rceil + 9$, which yields $t \geq 10 - \frac{9}{k} \cdot \lceil \frac{2k}{3} \rceil + \frac{8}{k}$. We distinguish between four cases:

If $\lceil \frac{2k}{3} \rceil = \frac{2k}{3}$, then $t \geq 4 + \frac{8}{k}$, i.e., $t \geq 5$. Thus, $\frac{\lceil \frac{2k}{3} \rceil \cdot t - 1}{(t-1)k + 1} \leq \frac{\frac{2k}{3} \cdot 5 - 1}{4k + 1} < 5/6 < 19/21$.

If $\lceil \frac{2k}{3} \rceil = \frac{2k+1}{3}$ and $k \geq 7$, then $t \geq 4 + \frac{5}{k}$, i.e., $t \geq 5$. Thus, $\frac{\lceil \frac{2k}{3} \rceil \cdot t - 1}{(t-1)k + 1} \leq \frac{\frac{2k+1}{3} \cdot 5 - 1}{4k + 1} < 5/6 < 19/21$.

If $\lceil \frac{2k}{3} \rceil = \frac{2k+2}{3}$, then $k \geq 5$ and $t \geq 4 + \frac{2}{k}$, i.e., $t \geq 5$. Thus, $\frac{\lceil \frac{2k}{3} \rceil \cdot t - 1}{(t-1)k + 1} \leq \frac{\frac{2k+2}{3} \cdot 5 - 1}{4k + 1} = \frac{5}{6} + \frac{3}{8k+2} \leq 19/21$.

If $k = 4$ then $t \geq 6$ and $\frac{\lceil \frac{2k}{3} \rceil \cdot t - 1}{(t-1)k + 1} \leq \frac{\lceil \frac{8}{3} \rceil \cdot 6 - 1}{5 \cdot 4 + 1} = 17/21 < 19/21$. □

Using Lemma 15, Equation (21) now yields

$$\begin{aligned} \text{SC}(c(P)) &\leq \left(13 + 12 \cdot \frac{19/21}{1 - 19/21}\right) \cdot \text{SC}(c(N)) \\ &= 127 \cdot \text{SC}(c(N)), \end{aligned}$$

as desired. □

Theorem 11. For every $\varepsilon > 0$ and integer $m > 0$, there exists integer $k \in \mathcal{O}(\varepsilon^{-3} \ln \frac{m}{\varepsilon})$, so that selecting a k -sized panel of agents with Fair Greedy Capture, has ex-ante distortion at most $1 + \varepsilon$, for every set of m alternatives.

Proof. Let P be the panel selected by FGC, $c' = c(N)$, and $a \in C$ be any candidate. Following the exact same arguments as in the proof of Theorem 4, for getting distortion of at most $(1 + \varepsilon)$, we will show that $\Pr[c(P) = a] \leq \frac{\varepsilon \text{SC}(c')}{2m(\text{SC}(a) - \text{SC}(c'))}$.

When Fair Greedy Capture breaks ties in a consistent way (with respect to the agents that are located at the boundaries of a ball), it can be equivalently described as following. The voters are partitioned into $k' \leq k$ sets, $S_1, \dots, S_{k'}$, where each S_j with $j < k'$ contains $\lceil n/k \rceil$ voters and the last set contains at most $\lceil n/k \rceil$ voters. Then, the sampling procedure is executed into two stages. In Stage 1, from each S_j , we sample one voter with probability $1/\lceil n/k \rceil$ and in Stage 2, the remaining probability is allocated equally among the voters that were not selected in Stage 1. Note, that because $|S_{k'}| \leq \lceil n/k \rceil$, it is probable no agent to be selected from $S_{k'}$, and thus we denote with $K \in [k' - 1, k']$ the number of voters that are selected in Stage 1. Then, if p_1, \dots, p_k are the random variables that indicate the agents in the panel P , each p_j with $j \leq K$ is sampled from S_j independently and the last $k - K$ agents are sampled from $N \setminus \{p_1, \dots, p_K\}$ uniformly at random without replacement. Let $X_i = d(p_i, a) - d(p_i, c(N))$. A panel P returns the candidate a instead of $c(N)$, when $X = \sum_{i=1}^k X_i \leq 0$. Let $X' = \sum_{i=1}^K X_i$. Next, we will show that (a) $\Pr[X' \leq kcD/4] \leq \frac{\varepsilon \text{SC}(c')}{4m(\text{SC}(a) - \text{SC}(c'))}$ and (b) $\Pr[X \leq 0 \mid X' > kcD/4] \leq \frac{\varepsilon \text{SC}(c')}{4m(\text{SC}(a) - \text{SC}(c'))}$. Then, we have that

$$\begin{aligned} \Pr[X \leq 0] &= \Pr[X \leq 0 \mid X' \leq kcD/4] \cdot \Pr[X' \leq kcD/4] + \Pr[X \leq 0 \mid X' > kcD/4] \cdot \Pr[X' > kcD/4] \\ &\leq \frac{\varepsilon \text{SC}(c')}{2m(\text{SC}(a) - \text{SC}(c'))} \end{aligned}$$

where the last inequality follows from (a) and (b), and the theorem follows.

We start by showing (a). Using identical arguments as in Case 1 in the proof of Theorem 4, we conclude that we can choose

a $1 \geq c \geq \frac{\varepsilon}{4}$ such that $\text{SC}(a) - \text{SC}(c(N)) = cnD$. Thus, we have

$$\begin{aligned} \mathbb{E}[X'] &= \sum_{j=1}^{k'} \sum_{i \in S_j} \Pr[i \text{ is selected}] \cdot (d(i, a) - d(i, c(N))) \geq \sum_{j=1}^{k'} \sum_{i \in S_j} \frac{1}{\lceil n/k \rceil} \cdot (d(i, a) - d(i, c(N))) \\ &\geq \frac{k}{2n} \cdot cnD = \frac{kcD}{2} \end{aligned}$$

Since the random variables X_i for $i = 1, \dots, K$ are independent, we can apply Hoeffding's inequality and obtain

$$\begin{aligned} &\Pr \left[X' < \frac{kcD}{4} \right] \\ &\leq \Pr \left[\left| X' - \mathbb{E}[X'] \right| \geq \frac{kcD}{4} \right] \\ &\leq 2 \exp \left(-2 \left(\frac{kcD}{4} \right)^2 / (K4D^2) \right) \\ &\leq 2 \exp(-kc^2/32). \end{aligned}$$

where the last inequality follows since $K \leq k$. Thus, it is sufficient to choose

$$k \geq \frac{\ln(4 \frac{\text{SC}(a) - \text{SC}(c')}{\text{SC}(c')} m / \varepsilon) 64}{c^2}.$$

Now, since from Theorem 9, we know that $\frac{\text{SC}(a) - \text{SC}(c')}{\text{SC}(c')}$ is at least equal to $\varepsilon/2$ and at most 127, we conclude that there is a $k \in \mathcal{O}(\frac{\ln(m/\varepsilon)}{\varepsilon^2})$, for which (a) is true.

Next, we show (b). First, we note that we can assume that $k - K \geq \frac{kc}{4}$ since otherwise X cannot be negative anymore. We have that

$$\mathbb{E}[X - X'] = \frac{k - K}{n - K} \cdot (cnD - X').$$

Using this, we get that

$$\begin{aligned} &\Pr[X < 0 \mid X' > kcD/4] = \Pr[X - X' < -X' \mid X' > kcD/4] \\ &= \Pr[X - X' - \mathbb{E}[X - X'] < -X' - \frac{k - K}{n - K} \cdot (cnD - X') \mid X' > kcD/4] \\ &\leq \Pr[|X - X' - \mathbb{E}[X - X']| > |-X' - \frac{k - K}{n - K} \cdot (cnD - X')| \mid X' > kcD/4] \\ &= \Pr[|X - X' - \mathbb{E}[X - X']| > |X'(1 - \frac{k - K}{n - K}) + \frac{k - K}{n - K}(cnD)| \mid X' > kcD/4] \\ &\leq \Pr[|X - X' - \mathbb{E}[X - X']| > \frac{kcD}{8} \mid X' > kcD/4] \end{aligned}$$

where the last inequality follows from the fact that $k \leq 2n$.

Thus, we can also apply Serfling's inequality (Theorem 5) here and obtain that

$$\begin{aligned} &\Pr[X < 0 \mid X' > kcD/4] \\ &\leq 2 \exp\left(-\frac{2(kcD)^2}{64(k - K)D^2}\right) \\ &\leq 2 \exp\left(-\frac{2k^2c^2}{64\frac{kc}{4}}\right) \\ &= 2 \exp\left(-\frac{kc^3}{8}\right). \end{aligned}$$

Thus, to get a probability of at most $\frac{\varepsilon \text{SC}(c')}{4m(\text{SC}(a) - \text{SC}(c'))}$ of selecting a in this case, we choose $k \geq \frac{8}{\varepsilon^3} \ln(4 \frac{\text{SC}(a) - \text{SC}(c')}{\text{SC}(c')} m / \varepsilon)$. As before since from Theorem 9, we know that $\frac{\text{SC}(a) - \text{SC}(c')}{\text{SC}(c')}$ is at least equal to $\varepsilon/2$ and at most equal to a constant, this gives us the desired asymptotic bound. \square