

Boosting Sortition via Proportional Representation

SOROUSH EBADIAN and EVI MICHA

Sortition is an ancient age democratic paradigm, which has been revitalized the last decades, and is based on the idea of choosing randomly selected representatives for decision making. The main properties that make sortition particularly appealing are *fairness* – all the citizens can be selected with the same probability– and *proportional representation* – a randomly selected panel probably reflects the composition of the whole population. Taking a closer look to the second property, in high level, it requires that if a group consists $x\%$ of the population, then $x\%$ of the panel should consist of individuals from this group. We define this intuitive property formally, when a population lies on a representation metric, by using a notion called *core*. A panel is in the core if no group of individuals is underrepresented proportional to its size. We show that uniformly random selection of a decision panel satisfies almost *ex ante core*. In practice, however, it is often asked core to be satisfied with certainty. Can we design a selection algorithm that satisfies fairness and *ex post core* simultaneously? We answer this question affirmatively and present an efficient selection algorithm that is fair and provides asymptotically optimal ex post core. Furthermore, we provide an efficient algorithm for auditing the core of a panel, i.e. for a given panel it measures how much it violates the core. We complement our theoretical results by conducting experiments with real data.

1 INTRODUCTION

In the last centuries, representative democracy has become synonymous with elections. However, this has not been the case throughout history. Since ancient Athens, the random selection of representatives from a given population has been proposed as a means of promoting democracy and equality [Van Reybrouck, 2016]. The core idea of randomly selecting representatives, a paradigm known as *sortition*, is that equal individuals have equal chance of participation [Stone, 2011], something that is often queried under representative democracy.

Sortition has gained significant popularity in recent years, mainly because of its use for forming *citizens assemblies*, where a randomly selected panel of individuals deliberate on issues and make recommendations. Currently, citizens assemblies are being implemented by more than 40 organizations in over 25 countries [Flanigan et al., 2021a]. There are signs that the use of sortition as a form of democracy even in a national level may be just a matter of time. For example, in Belgium, permanent sortition bodies have been established within two regional parliaments to give citizens a steady voice in decision making, and more local authorities in other countries are considering adopting this practice.¹ It is also notable, the very recent use of sortition in Greece for forming an audience panel that contributed to the selection of the actor that will represent the county in the Eurovision Song Contest 2023.² This increased usage of sortition indicates that it is becoming widely accepted and promotes it to a broader audience.

The last years, there has been a growing interest in the selection of representative panels in a fair and transparent way within the computer scientific research community [Ebadian et al., 2022, Flanigan et al., 2021a, 2020, 2021b]. However, the ideal method of selecting a representative panel of size k from a given population of size n remains the same: select k individuals uniformly at random [Engelstad, 1989]. We call this simple procedure as *uniform selection*. As stated by Flanigan et al. [2020], two main reasons that make this method particularly appealing are the following:

- (1) *Fairness*: Each citizen is included in the panel with the same probability, thereby the requirement of equality of participation is satisfied. In particular, each citizen is selected with probability equal to k/n .
- (2) *Proportional Representation*: The selected panel is likely to mirror the structure of the population in the sense that if $x\%$ of the population has some particular characteristics then in expectation $x\%$ of the panel will consist of individuals with these characteristics. For example, if the female share of the population is 48%, then in expectation 48% of the panel will be females.

At a high level, uniform selection seems to achieve *ex ante* proportional representation as in expectation the selected panel reflects the composition of the population. Especially, when the size of the panel is very large, then the panel that uniform selection outputs can be *ex post* proportionally representative with high probability, as a result of the Law of Large Numbers. However, practitioners are often not satisfied with the possibility of having a panel that does not reflect the demographic characteristics of the population [Benadè et al., 2019]. As a result, instead of choosing k individuals uniformly at random, practitioners may sample individuals in a different way to ensure proportional representation. For example, in the recent application related to the Eurovision Song Contest that we mentioned above, where the goal was to form a panel of size 70, instead of sampling 70 individuals uniformly at random, the organizers sampled 25 individuals in the 18-24 age range, 20 individuals in the 25-34 age range, 15 members in the 35-44 age range and 10 members in the 45+ age range, in an attempt to reflect the age representativeness of the Eurovision audience in the panel.³

¹<https://www.buergerrat.de/en/news/permanent-climate-assemblies-in-brussels-and-milan/>

²<https://eurovoix.com/2022/12/19/greece-public-jury-tdecember-28/>

³<https://esctoday.com/188611/greece-ert-kicks-off-the-procedure-for-eurovision-2023/>

Intuitively, proportional representation requires that any demographic group is fairly represented in the decision panel. Defining proportional representation in a more rigorous way for a population that is characterized by a single feature, such as age, is relatively straightforward. In contrast, for the more general case, it presents a greater challenge. For example, consider the case that in a population that consists of 6 individuals, a, b, c, d, e, f , each individual is characterized by three binary features, A, B, C , as given in the following table.

	a	b	c	d	e	f
A	1	1	0	0	0	0
B	1	0	0	1	0	0
C	1	0	1	0	0	0

If the goal is to select a panel of size 3, a way to enforce proportional representation, could be to set proportional constraints over each feature. For example, since $2/3$ of the population has $A = 0$ and $1/3$ of the population has $A = 1$, we could require that the panel contains 2 individuals with $A = 0$ and 1 individual with $A = 1$. Similarly, we could require that the panel contains 2 individuals with $B = 0$ and 1 individual with $B = 1$, and 2 individuals with $C = 0$ and 1 individual with $C = 1$. Note that the panel $\{b, c, d\}$, satisfies all these requirements. However, while $1/3$ of the population has $A = B = C = 0$, there is no individual in the panel that represents this group. This example indicates that imposing proportional constraints over individual features may not ensure proportional representation of every group of individuals.

In this work, we aim to address the following questions:

- (1) *What is a formal definition of proportional representation of a population?*
- (2) *To what extent does uniform selection satisfy proportional representation?*
- (3) *Is it possible to design selection procedures that enhance representation guarantees while maintaining fairness?*

1.1 Proportional Representation via Core

Going back to the one-feature example, suppose that in a given population 48% are females, 50% are males and 2% self-report as non-binary. Intuitively, a panel is considered proportionally representative if it consists of $0.48 \cdot k$ females, $0.50 \cdot k$ males and $0.02 \cdot k$ non-binary individuals. In other words, any group of size s deserves $s/n \cdot k$ representatives in the panel. Motivated by this example, we borrow a notion of proportional representation by recent works on multiwinner elections and fair allocation of public goods [Aziz et al., 2017, Cheng et al., 2020, Conitzer et al., 2019, Fain et al., 2018], called *core*, which captures exactly this intuition: *Every $S \subseteq [n]$ is entitled to choose up to $|S|/n \cdot k$ representatives.* Note that this notion is not defined over predefined groups using particular features, but it provides fair representation in the panel to any subset of the population.

A panel P is called proportionally representative, or is said to be in the core, if there does not exist a subset of the population that could choose a representative panel among themselves with size proportional to the size of the group under which all of them feel more represented. In other words, given a panel P , if there is $S \subseteq [n]$ that can find a panel $P' \subseteq S$ with $|P'| \leq |S| \cdot k/n$ such that every individual in S is better represented by P' than P , then these individuals would have a justified complaint against P . A panel is in the core if no subset of the population has a justified complaint against it. However, a conceptual challenge is to qualify how much a panel represents an individual.

1.2 Representation Metric

To overcome this challenge, we use the same approach as taken by Ebadian et al. [2022] where it is assumed that the individuals lie in an underlying *representation metric space* with distance d . The distance between the individuals i and j is denoted by $d(i, j)$. We assume that the distances are symmetric, i.e. $d(i, j) = d(j, i)$, and satisfy the triangle inequality, i.e. $d(i, j) \leq d(i, \ell) + d(\ell, j)$. The smaller the distance between two individuals is the better they represent each other. The representation metric space can be constructed as a function of features that are of particular interest for an application at hand, such as gender, age, ethnicity and education. Significantly, our theoretical results depend only on the existence of such a metric space without further assumptions.

It still remains to measure how much an individual is represented by a panel which consists of a number of representatives. Again, we take the same approach as was taken by Ebadian et al. [2022], which was also taken by Caragiannis et al. [2022] in a multiwinner elections context. The cost of an individual i for a panel P with parameter $q \in [k]$, denoted by $c_q(i, P; d)$ and called q -cost (or simply cost, when q is implied from the context), is equal to the distance of i from her q -th closest member in the panel. We omit d from the notation when it is clear from the context. For $q = 1$, the cost of an individual is equal to her distance from her closest representative in the panel, while for $q = k$, the cost is equal to her distance from her furthest representative in the panel. The smaller the cost is, the more represented the individual feels from the panel.

1.3 Distributions over Panels

As mentioned above, we are particularly interested in ways of sampling panels that satisfy fairness and proportional representation concurrently. Therefore, we focus on selection procedures that return distributions over panels of size k , since otherwise fairness, which is the pivotal property of sortition, cannot be satisfied.

We denote with $\mathcal{A}_{k,q}$ a *selection algorithm* parameterized by k and q that takes as input the metric d and outputs a distribution over all panels of size k . We say that a panel is in the support of $\mathcal{A}_{k,q}$, if it is implemented with positive probability under the distribution that $\mathcal{A}_{k,q}$ outputs. We pay special attention to the *uniform selection* algorithm, denoted by \mathcal{U}_k , that always outputs a uniform distribution over all the subsets of the population of size k , independently of q .

Fairness over distributions is defined as following.

Definition 1.1 (Fairness). A selection algorithm satisfies fairness if each individual is included in the panel with probability exactly equal to k/n , i.e.

$$\forall i \in [n], \Pr[i \in \mathcal{A}_{k,q}] = k/n.$$

While above we defined core over panels, we should define core over distributions which is the output of a selection algorithm.

1.4 Distributions in the Core and Approximate Core

Here, we define core over distributions of panels with size k . A demanding extension of core over distributions is to require every panel in the support of a distribution to be in the core. Then, core is satisfied ex post, something that is often desirable in practice, as we discussed above. However, Chen et al. [2019] and Micha and Shah [2020] show that even when $q = 1$, a panel in the core is not guaranteed to exist. Therefore, as they did, we consider a multiplicative approximation of the core over panels, i.e. with respect to a multiplicative factor of the cost reduce that a subset of the population can have by choosing a panel among themselves. Then, the ex post core is defined by requiring each panel in the support of the distribution to be in the approximate multiplicative core.

Before we formally define this demanding definition of the core over distributions, we introduce the notion of α -pairwise score.

Definition 1.2 (α -Pairwise Score). Given two panels P and P' with $|P| \geq q$ and $|P'| \geq q$, the α -pairwise score, with $\alpha \geq 1$, of P' over P , is the number of individuals whose q -cost under P is larger than α times their q -cost under P' :

$$V_q(P, P', \alpha) = |\{i \in [n] : c_q(i, P) > \alpha \cdot c_q(i, P')\}|.$$

The requirement that $|P|$ and $|P'|$ are at least equal to q comes from the fact that when the cost is parameterized by q , the cost of an individual for a panel with size less than q is not even well-defined. So from now on, whenever the cost is parameterized by q , we will refer only to panels that have a size of at least q .

Now, we are ready to define ex post α -core over distributions which requires that any panel in the support of a distribution is in the α -core.

Definition 1.3 (Ex post α -Core). A panel P is in the α -core if for any panel P' with $|P'| \geq q$

$$V_q(P, P', \alpha) < |P'| \cdot \frac{n}{k}.$$

A selection algorithm $\mathcal{A}_{k,q}$ is in the ex post α -core if every P in the support of $\mathcal{A}_{k,q}$ is in the α -core. When $\alpha = 1$, we say that the selection algorithm is in the ex post core.

In other words, if for any P in the support of the selection algorithm and any other panel P' , the group of individuals, that reduce their cost by a factor of at least α by picking P' , is not sufficiently large to be eligible to choose P' , then the selection algorithm is in the ex post α -core. Above, we mentioned that a subset of individuals may choose a panel among themselves that may represent them better, and not any arbitrary panel. This restriction comes from the conceptual interpretation that a panel P' can be formed if all its representatives agree to be part of it and this happens if each of them believes that is represented better from P' than a given panel. However, this restriction is not reflected in the above definition. Note, that without this requirement it is easier to find a violation of the ex post α -core. All our negative results hold with respect to this requirement while all our positive results hold without it, and therefore this requirement does not play an important rule. Therefore, we omit it for clarity of the definition.

Given how demanding this definition is, it is no surprising that uniform selection fails to achieve any bounded approximation of ex post α -core. In a high level, consider the case that in a population, the individuals are assigned into two groups A and B , such that any two individuals in the same group have distance equal to 0, while any two individuals in separate groups have distance equal to 1. Under uniform selection, while unlikely, it is possible that the selected panel contains only individuals from group A . Then, the individuals in group B could choose a panel among themselves and reduce their distance from 1 to 0, meaning that α is unbounded.

A different natural extension of core over distributions is to require the core-like property to be satisfied with respect to the expected cost that an individual has over a distribution.

Definition 1.4 (α -Core over Expected Cost). A selection algorithm $\mathcal{A}_{k,q}$ is in the α -core over expected cost (or in the core over expected cost, for $\alpha = 1$) if there is no $S \subseteq [n]$ and a panel P' with $|P'| \leq |S|/n \cdot k$ such that

$$\forall i \in S, \mathbb{E}_{P \sim \mathcal{A}_{k,q}} [c_q(i, P)] > \alpha \cdot c_q(i, P').$$

In the appendix, we show that ex post α -core and α -core over expected cost are incomparable.

Going back to the previous example, it is not hard to see that uniform selection fails to achieve any bounded approximation with respect to the α -core over expected cost, as well. Indeed, the

expected cost over uniform selection of the individuals that are in group B is a positive value, while they are eligible to choose a panel among themselves under which their cost becomes 0. However, this negative result still solely depends on the fact that a “bad” panel is returned with a positive probability. So, the fact that this is an unlike event is not reflected in this definition either.

A relaxation of the demanding ex post core, inspired by Cheng et al. [2020], is to require the core-like property to be satisfied in expectation with respect to the size of the pairwise score of panels in the support of a distribution and any other panel.

Definition 1.5 (Ex ante α -Core). A selection algorithm $\mathcal{A}_{k,q}$ is in the ex ante α -core (or in the ex ante core, for $\alpha = 1$) if for all $P' \subseteq [n]$:

$$\mathbb{E}_{P \sim \mathcal{A}_{k,q}} [V_q(P, P', \alpha)] < |P'| \cdot \frac{n}{k}.$$

The above definition says that for any panel P' , if for any realized panel P , we count the number of individuals that reduce their cost by a multiplicative factor of at least α under P' , in expectation this number is less than $|P'| \cdot n/k$, and so in expectation they are not eligible to choose it.

It is easy to see that ex post α -core implies ex ante α -core, since if for each P in the support of a distribution that $\mathcal{A}_{k,q}$ returns and each P' , it holds that $V_q(P, P', \alpha) < |P'| \cdot n/k$, then $\mathbb{E}_{P \sim \mathcal{D}_k} [V_q(P, P', \alpha)] < |P'| \cdot n/k$.

1.5 Our Results

In Section 2, we focus on the method of uniform selection. We show that for $q = k$, uniform selection is in the ex post 2-core and 2-core over expected cost, while for any other value of q , it fails to provide any reasonable approximation to any of these two definitions. On the positive side, for $q = k$, uniform selection is in the ex ante core, and for any other value of q , it is in the ex ante 4-core. We also show that no fair selection algorithm is in ex ante α -core, for $\alpha < 2$, which indicates that uniform selection is asymptotically optimal among all the fair algorithms.

In Section 3, we turn our attention to the question of the existence of a selection algorithm that is fair and also ensures that any panel that it returns is almost in the core. Since, ex post α -core implies ex ante α -core, from Section 2 we know that no fair selection algorithm is in ex post α -core, for $\alpha < 2$. However, we show that there exists an efficient selection algorithm, called Fair Greedy Capture, that, as its name indicates, is fair and is in the ex post $\frac{5+\sqrt{41}}{2}$ -core, for any value of q . For $q = 1$, the guarantee of the algorithm improves to $\frac{3+\sqrt{17}}{2}$. In addition, we show that Fair Greedy Capture is in the 6-core over expected cost.

In Section 4, we show that given a panel P , we can check in polynomial time how much it violates the core, i.e. what is the larger value of α for which $V_q(P, P', \alpha - \epsilon) \geq |P'| \cdot n/k$ for some panel P' and arbitrary small ϵ .

Finally, in Section 5, we empirically evaluate the approximation of uniform selection and Fair Greedy Capture to the ex post core on constructed metrics derived from two demographic datasets. We notice that for large values of q , uniform selection achieves an approximation to the ex post similar to the one that Fair Greedy Capture achieves. For smaller values of q , when the individuals form cohesive parts, uniform selection has unbounded approximation very often. However, when the individuals are well spread in the space, uniform selection is almost in the ex post core. Thus, the decision of using uniform selection depends on the value of q and the structure of the population.

1.6 Prior Work on Representation of Selection Algorithms

Ebadian et al. [2022] recently considered a very similar question with one that we ask in this work: How do we measure the representation that a panel or a selection algorithm achieves in a rigorous way? As we mentioned above, they also assume the existence of a representative metric space

and use the q -cost to measure to what degree a panel represents an individual. However, they use the social cost (sum of q -costs) to measure how much a panel represents a given population. In particular, under their definition, the most representative panel of a population is the one that minimizes the social cost. In the example below, we show that this measure of representation may fail to achieve proportional representation, though.

Example 1.6. Let n is odd, $k = 3$ and $q = 1$. Assume that there are four group of individuals, A , B , C and D . There are exactly one individual in group A , and exactly one individual in group B , while there are $\frac{n-1}{2}$ individuals in group C and $\frac{n-1}{2}$ individuals in group D . The distances between individuals in different groups is specified in the following table.

	A	B	C	D
A	0	∞	∞	∞
B	∞	0	∞	∞
C	∞	∞	0	10
D	∞	∞	10	0

It is not difficult to see that any panel with minimum social cost contains the single individuals in groups A and B and one individual from either group B or group C , as otherwise the social cost would be unbounded. This means that while the individuals in group B form almost 50% of the population, and similarly do the individuals in group C , in any panel with optimal social cost, either group B or C is not represented at all. On the other hand, the two eccentric individuals are always part of the panel.

1.7 Related Work

Our work is most closely related to that of Ebadian et al. [2022]. The main difference is that we use a different notion of representation over panels or distributions. As we discussed above, their notion fails to capture the intuition of proportional representation. Moreover, while their notion of representation is in some cases incompatible with fairness, in this work we show that there are selection algorithms that achieve good proportional representation and fairness simultaneously. Another related work is that of Benadè et al. [2019]. They also focused on the fact that practitioners in many cases do not simply choose k individuals uniformly at random, but they sample in ways ensuring that any realized panel is proportionally representative. However, they asked a very different question which is how this stratifying sampling may affect the variance of the representation of unknown groups. Moreover, they only consider the simple case where the population can be partitioned into a number of disjoint groups which implies that the individuals are characterized by just one feature. Nevertheless, their positive results are encouraging for our proposed sampling procedure as well.

Another way that representation is forced in practice is by setting quotas on individual features. For example, it might be asked in a panel of size 100, at least 40 representatives to be women and at least 30 representatives to be college educated. However, a problem that appears in practice is that a few people volunteer to participate in a decision panel. As a result, the representatives are selected from this pool of volunteers which usually does not reflect the composition of the population, since for example highly educated people are usually more willing to participate in a decision panel than less educated people. Flanigan et al. [2021a] proposed selection algorithms that, given a biased pool of volunteers, find distributions that maximize the minimum selection probability of any volunteer over panels that satisfy the desired quotas. However, the vision of a sortition based democracy relies on sampling the representatives directly from the underlying

population [Gastil and Wright, 2019]. In this work, as Ebadian et al. [2022] and Benadè et al. [2019], we focus on this pivotal idea of sortition. More importantly, as we discussed in the introduction, the satisfaction of particular quotas does not necessarily ensures proportional representation, and it is in general challenging to find a panel that satisfies a set of quotas as much as possible [Celis et al., 2018, Lang and Skowron, 2018].

The idea of using core as a means of measuring the proportional representation that a panel of size k provides to a population of size n , that lies in a metric space, was first introduced by [Chen et al., 2019], and then revisited by [Micha and Shah, 2020]. The main differences of these works with this one is that the cost of an individual for a panel is defined as her distance from her closest representative in the panel, i.e. by setting $q = 1$ in our q -cost function, and there is no any fairness constraint. Chen et al. [2019] show that in the general case a panel in the α -core with $\alpha < 1.5$ is not guaranteed to exist when only individuals of the population can serve as representatives and provide a greedy algorithm that always returns a panel in the $(1 + \sqrt{2})$ -core. The negative result carry over when $q = 1$, but the greedy algorithm as defined in their paper does not satisfy any fairness constraints which is one of the main desired properties of sortition. However, we show that by modifying the greedy algorithm appropriately, we can find distributions that are fair and return panels that are almost in the core.

Proportional representation through core has been extensively studied in multi winner elections [Aziz et al., 2017, Fain et al., 2018, Faliszewski et al., 2017, Lackner and Skowron, 2023]. The selection of a representative panel is a special case of committee elections where the set of candidates is the same as the set of voters. While in these works, the voters and the candidates do not lie in a metric space, but instead the voters hold rankings over candidates, in our model, the rankings could derive from the underlying metric space. Due to impossibility results [Cheng et al., 2020], relaxations of the core have been studied. In this work, we consider a relaxation of the core over distributions similar to the one that was introduced by Cheng et al. [2020]. They show that a distribution over committees in the ex ante core always exists. However, their results hold without the fairness constraint and, here we show that by imposing this constraint, a distribution over panels in the ex ante α -core with $\alpha < 2$ is not guaranteed to exist.

The representation of individuals as having an ideal point in a metric space has its roots to the spatial model of voting [Arrow, 1990, Enelow and Hinich, 1984]. This approach has been extensively used for comparing different voting rules from a more quantitative point of view [Anshelevich et al., 2021, Procaccia and Rosenschein, 2006]. Lastly, the idea of using q -cost as a measure of how much a panel represents an individual was first proposed by Caragiannis et al. [2022]. However, they, as Ebadian et al. [2022] do, use the social cost as a measurement of aggregation.

2 UNIFORM SELECTION AND CORE

In this section, we focus on the uniform selection method for selecting a panel. First, we show that for $q = k$, uniform selection is nearly in the ex post core and core over expected cost, while for any other case, it fails to provide any good approximation to any of them. On the positive side, we show that uniform selection is in the ex ante 4-core, for any q , and is asymptotically optimal among all fair selection algorithms.

2.1 Uniform Selection and Ex Post Core

We start by showing that when $q = k$, uniform selection is almost in the ex post core. This positive result derives from the fact that when $q = k$, a given panel is not in the core if and only if all individuals agree that there exists another panel of size k that is more representative for all of them.

THEOREM 2.1. *For $q = k$, uniform selection is in the ex post 2-core and in the 2-core over expected cost. The former bound is tight.*

All the missed proofs can be found in the appendix.

The above theorem shows that any selection algorithm is in the ex post 2-core and in the 2-core over expected cost when $q = k$. This is because, the only property of uniform selection that is utilized is that it returns a panel of size k .

Next, we show that for any value of q other than k , uniform selection does not satisfy any bounded approximation of the ex post core or core over expected value.

THEOREM 2.2. *For any $q \in [k - 1]$ and $\lfloor n/k \rfloor \geq k$, there exists an instance such that uniform selection is not in the ex post α -core for any bounded α , and also not in the α -core over expected cost, for any bounded α .*

2.2 Uniform selection achieves almost Optimal Ex Ante Core

As we have shown above, except when $q = k$, uniform selection fails to provide any reasonable guarantee to the ex post core and the core over expected cost. However, as previously noted in the introduction, both of these definitions do not take into account the likelihood of appearance of certain panels that are not very representative.

In this section, we turn our attention to ex ante core. First, we observe that for $q = k$, uniform selection satisfies the ex ante core criteria.

PROPOSITION 2.3. *For $q = k$, uniform selection is in the ex ante core.*

In the next theorem, we show that for any other value of q , no selection algorithm that is fair, is guaranteed to achieve ex ante α -core with $\alpha < 2$.

THEOREM 2.4. *For any $q \in [k - 1]$, when $n \geq 2k^2/(k - q)$, there exists an instance such that no selection algorithm that is fair, is in ex ante α -core with $\alpha < 2$.*

Next, we show that uniform selection achieves ex ante 4-core, and, as a result, it achieves almost optimal ex ante core among all fair selection algorithms.

Before, we prove this result we introduce some extra notation. We denote with $\text{top}_q(i, P)$ the set of the q closest representatives of i in a panel P (ties are broken arbitrarily). Moreover, we denote with $B(x, y)$, the set of individuals that are captured from a ball centered at x with radius y , i.e. $B(x, y) = \{i \in [n] : d(x, y) \leq r\}$.

In the proof of this result, we use the following form of Chu–Vandermonde identity which we prove in the appendix for completeness.

Definition 2.5 (Chu–Vandermonde identity). Let n, k , and r be integers such that $0 \leq r \leq k \leq n$. Then, the following identity holds

$$\sum_{j=0}^n \binom{j}{r} \cdot \binom{n-j}{k-r} = \binom{n+1}{k+1}.$$

Moreover, we use the following necessary lemma.

LEMMA 2.6. *Let $S \subseteq [n]$ and let P' be a panel, with $|P'| \geq q$.*

- (1) *There exists a partition T_1, \dots, T_m of S with respect to P' , with $m \leq \lfloor |P'|/q \rfloor$, such that for each $\ell \in [m]$, there is $i_\ell^* \in T_\ell$ for which it holds that $c_q(i_\ell^*, P') \leq c_q(i, P')$ and $\text{top}_q(i, P') \cap \text{top}_q(i_\ell^*, P') \neq \emptyset$, for any $i \in T_\ell$.*




Fig. 1. Diagram for Proof of Theorem 2.7

- (2) *There exists a partition T_1, \dots, T_m of S with respect to P' , with $m \leq \lfloor |P'|/q \rfloor$, such that for each $\ell \in [m]$, there is $i_\ell^* \in T_\ell$ for which it holds that $c_q(i, P') \leq c_q(i_\ell^*, P')$ and $\text{top}_q(i, P') \cap \text{top}_q(i_\ell^*, P') \neq \emptyset$, for any $i \in T_\ell$.*

PROOF. We start by showing the first part. We partition all individuals into $m \leq \lfloor |P'|/q \rfloor$ groups, denoted by T_1, \dots, T_m iteratively as follows.

Suppose i_1^* is the individual with the smallest q -cost over P' (ties are broken arbitrary), i.e. $i_1^* = \arg \min_{i \in [n]} c_q(i, P')$. Then, T_1 is the set of all the individuals whose q closest representatives from P' includes at least one member of $\text{top}_q(i_1^*, P')$, i.e.

$$T_1 = \{i \in [n] : \text{top}_q(i, P') \cap \text{top}_q(i_1^*, P') \neq \emptyset\}.$$

Next, from the remaining individuals, suppose i_2^* is the one with the smallest q -cost over P' , i.e. $i_2^* = \arg \min_{i \in [n] \setminus T_1} c_q(i, P)$. Construct T_2 from $[n] \setminus T_1$ similarly by taking all the individuals whose at least one of their q closest representatives in P' is included in $\text{top}_q(i_2^*, P')$. We repeat this procedure, and in round ℓ , we find $i_\ell^* \in [n] \setminus (\cup_{\ell'=1}^{\ell-1} T_{\ell'})$ that has the smallest cost over P' , and construct T_ℓ by assigning any individual in $[n] \setminus (\cup_{\ell'=1}^{\ell-1} T_{\ell'})$ whose at least one of the q closest representatives belongs in $\text{top}_q(i_\ell^*, P')$. Note that for any $\ell_1, \ell_2 \in [m]$ with $\ell_1 < \ell_2$, $\text{top}_q(i_{\ell_1}^*, P') \cap \text{top}_q(i_{\ell_2}^*, P') = \emptyset$, as if at least one of the q closest representatives of $i_{\ell_2}^*$ in P is included in $\text{top}_q(i_{\ell_1}^*, P')$, then $i_{\ell_2}^*$ would have been assigned to T_{ℓ_1} and would not belong in $[n] \setminus (\cup_{\ell'=1}^{\ell_2-1} T_{\ell'})$. This means that in each round, we consider q representatives that have not been considered before, and hence after $\lfloor |P'|/q \rfloor$ rounds, less than q representatives in P' may remain unconsidered. As a result, after at most $\lfloor |P'|/q \rfloor$ rounds, all the individuals will have been assigned to some group, since at least one of their q closest representatives has been considered.

The second part follows by simply setting i_ℓ^* to be equal to the individual in $[n] \setminus (\cup_{\ell'=1}^{\ell-1} T_{\ell'})$ that has the largest cost over P' , i.e. $i_\ell^* = \arg \max_{i \in [n] \setminus (\cup_{\ell'=1}^{\ell-1} T_{\ell'})} c_q(i, P)$. All the remaining arguments remain the same. \square

Now, we are ready to prove the following theorem.

THEOREM 2.7. For any q , uniform selection is in ex ante 4-core, i.e. for any panel P' with $|P'| \geq q$

$$\mathbb{E}_{P \sim \mathcal{U}_k} [V_q(P, P', 4)] < |P'| \cdot \frac{n}{k}.$$

PROOF. Let P' be a panel, with $|P'| \geq q$. Then, by linearity of expectation, we have that

$$\mathbb{E}_{P \sim \mathcal{U}_k} [V_q(P, P', 4)] = \sum_{i \in [n]} \Pr_{P \sim \mathcal{U}_k} [c_q(i, P) > 4 \cdot c_q(i, P')].$$

Let T_1, \dots, T_m be a partition of $[n]$ with respect to P' , as given in the first part of Lemma 2.6. For each $\ell \in [m]$, we reorder the individuals in T_ℓ in an increasing order based on their distance from i_ℓ^* , and relabel them as $i_1^\ell, \dots, i_{|T_\ell|}^\ell$. This way, i_1^ℓ and $i_{|T_\ell|}^\ell$ are the individuals in T_ℓ that have the smallest and the largest distance from i_ℓ^* , respectively. Then, we get that

$$\mathbb{E}_{P \sim \mathcal{U}_k} [V_q(P, P', 4)] = \sum_{i \in [n]} \Pr_{P \sim \mathcal{U}_k} [c_q(i, P) > 4 \cdot c_q(i, P')] = \sum_{\ell=1}^m \sum_{j=1}^{|T_\ell|} \Pr_{P \sim \mathcal{U}_k} [c_q(i_j^\ell, P) > 4 \cdot c_q(i_j^\ell, P')]. \quad (1)$$

Now, we will bound $\Pr_{P \sim \mathcal{U}_k} [c_q(i_j^\ell, P) > 4 \cdot c_q(i_j^\ell, P')]$ for each i_j^ℓ . For each i_j^ℓ , let r_j^ℓ be an arbitrary representative in $\text{top}_q(i_j^\ell, P') \cap \text{top}_q(i_\ell^*, P')$. Then, we get that

$$d(i_j^\ell, i_\ell^*) \leq d(i_j^\ell, r_j^\ell) + d(r_j^\ell, i_\ell^*) \leq c_q(i_j^\ell, P') + c_q(i_\ell^*, P') \leq 2 \cdot c_q(i_j^\ell, P'), \quad (2)$$

where the last inequality follows from the fact that i_ℓ^* has the smallest cost over P' among all the individuals in T_ℓ . Now, consider the ball that is centered at i_j^ℓ and has radius $4 \cdot c_q(i_j^\ell, P')$. Note that this ball contains any individual $i_{j'}^\ell$, with $j' < j$. Indeed, for each i_ℓ^* and $i_{j'}^\ell$, with $j' < j$, we have that

$$d(i_j^\ell, i_{j'}^\ell) \leq d(i_j^\ell, i_\ell^*) + d(i_\ell^*, i_{j'}^\ell) \leq 2 \cdot d(i_j^\ell, i_\ell^*) \leq 4 \cdot c_q(i_j^\ell, P'),$$

where the second inequality follows from the fact that for each $j', j \in [|T_\ell|]$ with $j' < j$, $d(i_j^\ell, i_\ell^*) \leq d(i_{j'}^\ell, i_\ell^*)$ and the last inequality follows from Equation (2). This argument is drawn in Figure 1.

When $c_q(i_j^\ell, P) > 4 \cdot c_q(i_j^\ell, P')$, then we get that $|P \cap \{i_1^\ell, \dots, i_j^\ell\}| < q$, as otherwise there would exist at least q individuals in $B(i_j^\ell, 4 \cdot c_q(i_j^\ell, P'))$, and $c_q(i_j^\ell, P)$ would be at most $4 \cdot c_q(i_j^\ell, P')$. Hence, we have that

$$\begin{aligned} \Pr_{P \sim \mathcal{U}_k} [c_q(i_j^\ell, P) > 4 \cdot c_q(i_j^\ell, P')] &\leq \Pr_{P \sim \mathcal{U}_k} [|P \cap \{i_1^\ell, \dots, i_j^\ell\}| < q] \\ &= \Pr_{P \sim \mathcal{U}_k} \left[\bigcup_{r=0}^{q-1} |P \cap \{i_1^\ell, \dots, i_j^\ell\}| = r \right] \\ &\leq \sum_{r=0}^{q-1} \Pr_{P \sim \mathcal{U}_k} [|P \cap \{i_1^\ell, \dots, i_j^\ell\}| = r] = \sum_{r=0}^{q-1} \frac{\binom{j}{r} \binom{n-j}{k-r}}{\binom{n}{k}} \end{aligned}$$

where the second inequality follows from the Union Bound and the last equality follows from the fact that uniform selection chooses k out of n individuals uniformly at random. Then, by returning to Equation (1) we get that,

$$\begin{aligned} \mathbb{E}_{P \sim \mathcal{U}_k} [V_q(P, P', 4)] &= \sum_{\ell=1}^m \sum_{j=1}^{|T_\ell|} \Pr_{P \sim \mathcal{U}_k} [c_q(i_j^\ell, P) > 4 \cdot c_q(i_j^\ell, P')] \\ &\leq \sum_{\ell=1}^m \sum_{j=1}^{|T_\ell|} \sum_{r=0}^{q-1} \frac{\binom{j}{r} \binom{n-j}{k-r}}{\binom{n}{k}} \end{aligned}$$

ALGORITHM 1: Fair Greedy Capture

Input: Individuals $[n]$, metric d , k , q
Output: Panel P
 $R \leftarrow [n]; \delta \leftarrow 0; P \leftarrow \emptyset;$
while $|R| \geq \lceil q \cdot n/k \rceil$ **do**

 Smoothly increase δ ;

while $\exists j \in R$ such that $|B(j, \delta) \cap R| \geq \lceil q \cdot n/k \rceil$ **do**
 $S \leftarrow \lceil q \cdot n/k \rceil$ individuals arbitrary chosen from $B(j, \delta)$;

 $\hat{P} \leftarrow$ pick q individuals from S uniformly at random;

 $P \leftarrow P \cup \hat{P}$;

 $R \leftarrow R \setminus S$;

end
end
if $|P| < k$ **then**
 $\hat{P} \leftarrow k - |P|$ individuals from $[n] \setminus P$ by picking $i \in R$ with probability k/n and $i \in [n] \setminus (P \cup R)$ with probability $\frac{k - |P| - |R| \cdot k/n}{n - |P| - |R|}$;

 $P \leftarrow P \cup \hat{P}$;

end

$$\begin{aligned}
 &= \sum_{\ell=1}^m \sum_{r=0}^{q-1} \sum_{j=1}^{|T_\ell|} \frac{\binom{j}{r} \binom{n-j}{k-r}}{\binom{n}{k}} \\
 &\leq \sum_{\ell=1}^m \sum_{r=0}^{q-1} \sum_{j=0}^n \frac{\binom{j}{r} \binom{n-j}{k-r}}{\binom{n}{k}} \\
 &\stackrel{(1)}{=} \sum_{\ell=1}^m \sum_{r=0}^{q-1} \frac{\binom{n+1}{k+1}}{\binom{n}{k}} = \sum_{\ell=1}^m \sum_{r=0}^{q-1} \frac{n+1}{k+1} = m \cdot q \cdot \frac{n+1}{k+1} \leq |P'| \cdot \frac{n}{k},
 \end{aligned}$$

where (1) follows from Chu–Vandermonde identity and the last inequality follows from the facts that $m \leq \lfloor |P'|/q \rfloor$ and $\frac{n+1}{k+1} \leq \frac{n}{k}$. \square

3 FAIRNESS AND ALMOST EX POST CORE

In the previous section, we show that while uniform selection is in the ex ante 4-core, it does not guarantee to return panels that have a good approximation to the core.

But what if we want to ensure that any realized panel is almost in the core? Is there a way to achieve fairness and a good approximation to the ex post core concurrently? From the fact that ex post α -core implies ex ante α -core and from Theorem 2.4, we get that no fair selection algorithm can return a distribution in the ex post α -core, with $\alpha < 2$. Here, we show that there exists a fair selection algorithm that returns a distribution in the ex post $\frac{5+\sqrt{41}}{2}$ -core.

We revisit the Greedy Capture algorithm introduced by Chen et al. [2019]. Briefly, their algorithm starts with an empty panel and grows a ball around any individual at the same rate. When a ball captures at least $\lceil n/k \rceil$ individuals for the first time, the center of the ball is included in the panel and all the captured individuals are disregarded. As balls continue growing, when new individuals are captured by balls that are centered to individuals that are previously added to the panel, they are immediately removed. This algorithm returns a panel in the $(1 + \sqrt{2})$ -core, when $q = 1$.

From the description above, it is clear that the Greedy Capture algorithm, as defined by Chen et al. [2019], returns a deterministic panel. Here, we define a modification of the algorithm, called

Fair Greedy Capture, which outputs a fair distribution over panels of size k that is also in the ex post $\frac{5+\sqrt{41}}{2}$ -core, for any value of q . The algorithm starts with an empty panel P and grows a ball around every individual in $[n]$ at the same rate. When a ball captures $\lceil q \cdot n/k \rceil$ individuals (if more than $\lceil q \cdot n/k \rceil$ individuals have been captured, it chooses exactly $\lceil q \cdot n/k \rceil$ by arbitrarily excluding some points on the boundary), the algorithm picks q of them uniformly at random and includes them in the panel P , and disregards all the $\lceil q \cdot n/k \rceil$ individuals. When this happens, we say that the algorithm detects this ball. Balls continue to grow only around individuals that have not been disregarded yet, i.e. detected balls are frozen. When less than $\lceil q \cdot n/k \rceil$ individuals are left, but less than k representatives have been chosen, the algorithm chooses the remaining representatives among the individuals that have not be included in the panel as following: selects each individual that has not been disregarded with probability equal to k/n and allocates the remaining probability uniformly among the individuals that have been disregarded but not selected. This can be done using systematic sampling [Yates, 1948]. Our algorithm, Fair Greedy Capture, is given in Algorithm 1.

One basic difference between Fair Greedy Capture and Greedy Capture by Chen et al. [2019] is that in the latter, one individual is included in the panel the first time that a ball captures $\lceil n/k \rceil$ individuals, but in the former q individuals are included in the panel the first time that a ball captures $\lceil q \cdot n/k \rceil$ individuals. While this modification is quite straightforward, even if we consider the original Greedy Capture with this modification, it is significantly more challenging to prove that this algorithm returns a panel that is almost in the core, for any q . The main difficulty comes from the fact that when we ask if a subset of population is eligible to choose a panel P' different than a given panel P under which all of them reduce their distance by a factor of at least α , for $q = 1$, it suffices to consider panels of size 1, while for $q > 1$, we cannot restrict our attention to panels of size q . To see that, suppose that when $q = 1$, there is a panel P' , with $|P'| > 1$ and a subset $S \subseteq [n]$, with $|S| \geq |P'| \cdot n/k$ such that each $i \in S$ reduces her cost by at least a factor of α if she is represented by P' rather than a given panel P . Then, if we partition S into $|P'|$ groups by assigning each individual to their closest representative from P' , at least one of these groups should have size at least n/k . For $q > 1$, given a panel P' and $S \subseteq [n]$ of sufficient size, it does not necessarily hold that there are q representatives in P' that are the q -closest representatives of at least $q \cdot n/k$ individuals together in S . To put it in another way, it is possible that there exists a subset of population that is eligible to choose a panel P' with $|P'| > q$ under which all of them reduce their distance by a factor of at least α , but there does not exist a subset of population that is eligible to choose a panel P'' with $|P''| = q$ under which all of them reduce their distance by a factor of at least α . To overcome this challenge, we use Lemma 2.6.

Before we prove that Fair Greedy Capture returns a distribution that is almost in the ex post core and is fair, we prove the next technical lemma which shows that a form of triangle inequality is satisfied over costs of panels.

LEMMA 3.1. *For any panel P and any $i, i' \in [n]$, it holds that $c_q(i, P) \leq d(i, i') + c_q(i', P)$*

THEOREM 3.2. *Fair Greedy Capture returns a distribution that is fair and is in the ex post $\frac{5+\sqrt{41}}{2}$ -core.*

PROOF. We start by showing that the algorithm returns a fair distribution over panels of size k .

Fairness Guarantee. Suppose that $q \cdot n/k$ is an integer. Then, each individual that is disergarded in the while loop of the algorithm is included in the panel with probability exactly k/n . Now, suppose that after the algorithm has detected t balls, less than $q \cdot n/k$ are left. Then, when the algorithm exits the while loop we have that $|R| = n - t \cdot q \cdot n/k$ and $k - |P| = k - t \cdot q$. But since,

$$|R| \cdot k/n = k - t \cdot q,$$

we conclude that the remaining $k - t \cdot q$ representatives are chosen uniformly among the individuals in R . Thus, the algorithm returns a panel of size k and each $i \in [n]$ is chosen with probability k/n .

Now, we focus on the case that $q \cdot n/k$ is not an integer. In this case, note that in the while loop of the algorithm, less than k individuals are included in the panel, since q individuals are included in the panel every time that $\lceil q \cdot n/k \rceil$ non-disregarded individuals are captured from a ball. Moreover, each individual that is disregarded is chosen with probability strictly less than k/n . Now suppose that after exiting the while loop, there are individuals that have not been disregarded, i.e. $|R| > 0$. First, we show that the algorithm correctly chooses another $k - |P|$ representatives and outputs a panel of size k . The algorithm would select each individual in R with probability k/n and allocates the remaining probability – which is equal to $k - |P| - |R| \cdot k/n$ – uniformly among the $n - |P| - |R|$ individuals that have been disregarded but not selected in P . To satisfy fairness for people in R , it suffices to show that $|R| \cdot k/n < k - |P|$. Since for each individual $i \in [n] \setminus R$ we have $\Pr[i \in P] = q/\lceil q \cdot n/k \rceil < k/n$, then we have $|P| = \mathbb{E}[|P|] = \sum_{i \in [n] \setminus R} \Pr[i \in P] < (n - |R|) \cdot k/n$. Thus, $k - |P| > k - (n - |R|) \cdot k/n = |R| \cdot k/n$. Hence, the algorithm outputs panels of size k .

It remains to show that each individual in $[n] \setminus R$, which is disregarded in the while loop, is included in the panel with probability k/n . First, note that all of them are included in the panel with the same probability. This holds, since each is selected with probability $q/\lceil q \cdot n/k \rceil$ from the ball that captured them in the while loop, and, when not selected in the while loop, they get an equal chance of selection of $\frac{k - |P| - |R| \cdot k/n}{n - |R| - |P|}$. Since the size of the final panel returned by the algorithm is always k , and by linearity of expectation, we have $k = |R| \cdot k/n + \sum_{i \in [n] \setminus R} \Pr[i]$. By equality of $\Pr[i]$'s, we conclude that all must be equal to k/n and each individual in $[n]$ is included in the panel with probability k/n .

Ex Post Core Guarantee. Next, we show that the distribution that the algorithm returns is in the ex post $\frac{5 + \sqrt{41}}{2}$ -core. Let P be any panel that the algorithm may return. Suppose for contradiction that there exists $S, P' \subseteq [n]$ with $|S| \geq |P'| \cdot n/k$, such that $V_q(P, P', (5 + \sqrt{41})/2) \geq |P'| \cdot n/k$, i.e.

$$\forall i \in S, \quad c_q(i, P) > \frac{5 + \sqrt{41}}{2} \cdot c_q(i, P').$$

Let T_1, \dots, T_m be a partition of S with respect to P' , as given in the second part of Lemma 2.6. Since $m \leq \lfloor |P'|/q \rfloor$ and $|S| \geq |P'| \cdot n/k$, we conclude that there exists a part, say T_ℓ , that has size at least $q \cdot n/k$. Since, for some $i_\ell^* \in T_\ell$, it holds that $c_q(i, P') \leq c_q(i_\ell^*, P')$ and $\text{top}_q(i, P') \cap \text{top}_q(i_\ell^*, P') \neq \emptyset$ for each $i \in T_\ell$, we can conclude that $d(i_\ell^*, i) \leq 2 \cdot c_q(i_\ell^*, P')$ for each $i \in T_\ell$, as following: Pick an arbitrary representative in $\text{top}_q(i, P') \cap \text{top}_q(i_\ell^*, P')$ and denote it as r_i . Then,

$$d(i, i_\ell^*) \leq d(i, r_i) + d(r_i, i_\ell^*) \leq c_q(i, P') + c_q(i_\ell^*, P') \leq 2 \cdot c_q(i_\ell^*, P').$$

This means that there exists a ball centered at i_ℓ^* that has radius $2 \cdot c_q(i_\ell^*, P')$ and captures all the individuals in T_ℓ . Since there are sufficiently many individuals in this ball, it is possible that the algorithm detects it during its execution. If this happened, this means that there are q representatives in P that are located within this ball. Then, we get that the distance of i_ℓ^* from her q -th closest representative in P is at less than or equal to the radius of the ball, which is at most $2 \cdot c_q(i_\ell^*, P')$. Therefore, i_ℓ^* cannot reduce her distance by a multiplicative factor larger than 2 by choosing P' , and we reach a contradiction. On the other hand, if the algorithm did not detect this ball during its execution, this means that some of the individuals in T_ℓ have been disregarded before the ball centered at i_ℓ^* captures sufficiently many of them. Hence, some individuals in T_ℓ have been captured from a different ball with radius at most $2 \cdot c_q(i_\ell^*, P')$. Suppose that i' is the first individual in T_ℓ that was captured from such a ball. Then, we have that, within this ball, q representatives are selected in P . Hence $c_q(i', P) \leq 4 \cdot c_q(i_\ell^*, P')$, since the distance of i' from any other individual in this ball is at most equal to the diameter of the ball. We consider the minimum multiplicative improvement of

ALGORITHM 2: Auditing Algorithm**Input:** P , Individuals $[n]$, metric d , k , q ,**Output:** $\hat{\alpha}$ **for** $j \in [n]$ **do**

$$\begin{aligned} & \hat{P}_j \leftarrow \{j\} \cup q - 1 \text{ closest neighbors of } j; \\ & \hat{\alpha}_j \leftarrow \text{the } \lceil q \cdot n/k \rceil \text{ largest value among } c_q(i, P)/c_q(i, \hat{P}_j); \end{aligned}$$
end $\hat{\alpha} \leftarrow \arg \max_{j \in [n]} \hat{\alpha}_j$ both i_ℓ^* and i' :

$$\begin{aligned} & \min \left(\frac{c_q(i', P)}{c_q(i', P')}, \frac{c_q(i_\ell^*, P)}{c_q(i_\ell^*, P')} \right) \\ & \leq \min \left(\frac{c_q(i', P)}{c_q(i', P')}, \frac{d(i_\ell^*, i') + c_q(i', P)}{c_q(i_\ell^*, P')} \right) && \text{(by Lemma 3.1)} \\ & \leq \min \left(\frac{c_q(i', P)}{c_q(i', P')}, \frac{d(i_\ell^*, r_{i'}) + d(i', r_{i'}) + c_q(i', P)}{c_q(i_\ell^*, P')} \right) \\ & \leq \min \left(\frac{c_q(i', P)}{c_q(i', P')}, \frac{d(i_\ell^*, r_{i'}) + c_q(i', P') + c_q(i', P)}{c_q(i_\ell^*, P')} \right) && \text{(as } r_{i'} \in \text{top}_q(i', P')) \\ & \leq \min \left(\frac{4 \cdot c_q(i_\ell^*, P')}{c_q(i', P')}, 5 + \frac{c_q(i', P')}{c_q(i_\ell^*, P')} \right) \\ & \leq \max_{z \geq 0} \min(4 \cdot z, 5 + 1/z) = \frac{5 + \sqrt{41}}{2} \end{aligned}$$

where the fourth inequality follows from the fact that $d(i_\ell^*, r_{i'}) \leq c_q(i_\ell^*, P')$, since $r_{i'} \in \text{top}_q(i', P')$, and the fact that $c_q(i', P) \leq 4 \cdot c_q(i_\ell^*, P')$. \square

The above theorem shows that for any value of q , Fair Greedy Capture returns a fair distribution in the ex post $\frac{5+\sqrt{41}}{2}$ -core. In the next theorem, we show that when $q = 1$, we can have a better a guarantee. In particular, in this case Fair Greedy Capture returns a distribution that is in the ex post $\frac{3+\sqrt{17}}{2}$ -core and this is tight. This proof is quite similar with the proof of [Chen et al., 2019] and the main difference derives from the fact that when a ball captures $\lceil n/k \rceil$ individuals, we don't necessarily include in the panel the center of the ball, but we may include any individual within it. The proof of the theorem is deferred to the appendix.

THEOREM 3.3. *For $q = 1$, Fair Greedy Capture returns a fair distribution in the ex post $\frac{3+\sqrt{17}}{2}$ -core and there exists an instance for which this bound is tight.*

From the above theorem, we know that the analysis of Fair Greedy Capture is tight for $q = 1$. Finding tight bounds for the general case is an open question.

Moreover, we show that Fair Greedy Capture is in the 6-core over expected cost.

THEOREM 3.4. *Fair Greedy Capture returns a distribution that is in the 6-core over expected cost.*

4 AUDITING EX POST CORE

In this section, we turn our attention to the following question: Given a panel P , how much does it violate core, i.e. what is the maximum value of α such that there exists a panel P' with $V_q(P, P', \alpha - \epsilon) \geq |P'| \cdot n/k$, for arbitrary small ϵ ? This auditing question can be very useful in

practice for measuring the proportional representation of a panel that has been formed using a method that does not guarantee any panel to be almost in the core, such as uniform selection.

Chen et al. [2019] ask the same question for the case that the cost of an individual for a panel is equal to her distance from her closest representative in the panel, i.e. when $q = 1$. As we described in the previous section, in this case it suffices to restrict our attention to panels of size 1 that a subset of the population may prefer to be represented by. In other words, given a panel P , we can simply consider every individual as a potential representative and check if a sufficiently large subset of the population prefers significantly this individual than P to be their representative. Thus, we can find the maximum α such that there exists P' , with $V_q(P, P', \alpha - \epsilon) \geq n/k$ as following: For each $j \in [n]$, calculate α_j which is equal to the $\lceil n/k \rceil$ largest values among $\{c_q(i, P)/c_q(i, \{j\})\}_{i \in [n]}$. Then, α is equal to the maximum value among all α_j 's.

For $q > 1$, this question is more challenging. We show that we can approximate the value of the maximum α , by generalizing the above procedure as following: For each $j \in [n]$, calculate $\hat{\alpha}_j$ which is equal to the $\lceil n/k \rceil$ largest value among $\{c_q(i, P)/c_q(i, \hat{P}_j)\}_{i \in [n]}$, where \hat{P}_j is the panel that contains j and its $q - 1$ closest neighbors. Then, we return the maximum value among all $\hat{\alpha}_j$'s as $\hat{\alpha}$. Algorithm 2 executes this procedure. We can show that the maximum α such that there exists a panel P' with $V_q(P, P', \alpha - \epsilon) \geq |P'| \cdot n/k$, for arbitrary small ϵ , is at most $3 \cdot \hat{\alpha} + 2$.

THEOREM 4.1. *For any panel P and $q \in [k]$, if Algorithm 2 outputs $\hat{\alpha}$, then P is in the α -core with $\alpha = 3 \cdot \hat{\alpha} + 2$.*

Algorithm 2 requires at least quadratic time with respect to n , since it considers all the individuals along with their $q - 1$ closest neighbors. In Appendix K, we show that we can find a slightly worse approximation, by considering at most k individuals and their $q - 1$ closest centers, where these k individuals can be found using a modification of Fair Greedy Capture.

5 EXPERIMENTS

In the worst case, we show that uniform selection fails to guarantee that any panel that it returns is almost in the core. What happens though in the average case? Is uniform selection almost in the ex post core in the average? How much better is Fair Greedy Capture than uniform selection with respect to their approximations to the ex posts core, in the average case? In this section, we aim to answer these questions. To do that, we empirically evaluate the two algorithms, uniform selection (UNIFORM) and Fair Greedy Capture (FAIRGREEDY), using real databases.

5.1 Datasets

Following the methodology proposed by Ebadian et al. [2022], we use the same two datasets, that they also used, as a proxy for constructing the underlying metric space. These datasets capture characteristics of populations along multiple observable features. It is reasonable to assume that an individual feels closer to individuals that have similar characteristics. Thus, we construct random metric space using these data.

Adult Census Income (Adult). The first dataset we use is the Adult dataset, which is extracted from the 1994 Current Population Survey conducted by the US Census Bureau and is made accessible by the UCI Machine Learning Repository [Dua and Graff, 2017, Kohavi and Becker, 1996]. For our analysis, we use the following demographic features: sex, race, workclass, marital. status, and education. num. This dataset consists of $n = 32,561$ data points, each with a sample weight attribute (fnlwgt). By considering these five features, we identify 1513 unique data points and treat the sum of the weights associated with each unique point as a distribution across them.

European Social Survey (ESS) Our second dataset, European Social Survey (ESS), is a multi-country survey conducted every two years since 2001 in Europe.⁴ The survey collects data on a wide range of topics, including attitudes towards politics and society, social values, and well-being. We use the ESS Round 9 (2018) dataset, which covers 28 countries and contains 46,276 data points and 1451 features. Many of the features are specific to the country, resulting in an average of approximately 250 features per country (after removing non-demographic and country-unrelated data). The number country-specific data points range from 781 to 2745. Additionally, each data point has been assigned a post-stratification weight (*pspwght*), which we utilize as a representation of the distribution of the data points. We concentrated on the United Kingdom data (ESS-UK) which comprises 2204 data points.

5.2 Representation Metric Construction

In line with the work of Ebadian et al. [2022], we apply the same approach to generate synthetic metric preferences, which are used to measure the dissimilarity between individuals based on their feature values. Our datasets consist of two types of features: *categorical* features (such as sex, race, and marital status) and *continuous* features (such as income and years of education). We define the distance between individuals i and j with respect to feature f as follows:

$$d(i, j; f) := \begin{cases} \mathbb{1}[f(i) \neq f(j)], & \text{if } f \text{ is a } \textit{categorical} \text{ feature;} \\ \frac{1}{\max_{i', j'} |f(i') - f(j')|} \cdot |f(i) - f(j)|, & \text{if } f \text{ is a } \textit{continuous} \text{ feature,} \end{cases}$$

where the normalization factor for continuous features ensures that $d(i, j; f) \in [0, 1]$ for all i, j , and f , and that the distances in different features are comparable. Next, we define the distance between two individuals as the weighted sum of the distances over different features, i.e.

$$d(i, j) := \sum_{f \in F} w_f \cdot d(i, j; f),$$

where the weights w_f 's are randomly generated. Each unique set of randomly generated feature weights results in a new representation metric.

Experiment Setup. We generate 100 sets of randomly-assigned feature weights per dataset, calculate a representation metric for each set, and report the performance metrics averaged over 1000 instances. Furthermore, given that our datasets are samples of a large population (i.e. millions) and represented through a relatively small number of unique data points (i.e. few thousands), we assume that each data point represents a group of at least k people, which takes a maximum value of 40 in our study. For each experiment, we also report the approximation to the optimal social cost of the different algorithms which is the measure of representation used by Ebadian et al. [2022].

5.3 Evaluation Criteria

Ex Post Core Violation: Our main objective is to empirically measure the extent to which the selection algorithms are close to the ex post core. Similar to the question we addressed in Section 4, we aim to find the maximum value of α that the core is violated under panels that derive from the selection algorithms. To empirically measure ex post core violation, for each of the 100 instances, we sample one panel from an algorithm and compute the core violation using Algorithm 2. We note that this is not exactly equal to the worst-case core violation, but a very good approximation of it. It is infeasible to check all panels of size k (which can be as large as n^k possibilities) and find the one that maximizes the amount of core violation.

⁴ESS Round 9: European Social Survey Round 9 Data (2018). Data file edition 3.1. Sikt - Norwegian Agency for Shared Services in Education and Research, Norway – Data Archive and distributor of ESS data for ESS ERIC. doi:10.21338/NSD-ESS9-2018.

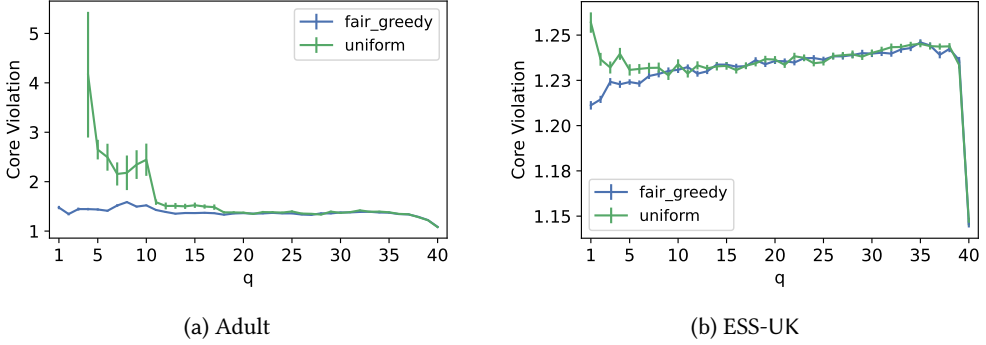


Fig. 2. Ex post core violation of FAIRGREEDY and UNIFORM with $k = 40$




Fig. 3. Approximation to the optimal social cost of FAIRGREEDY and UNIFORM with $k = 40$

Approximation to Optimal Social Cost: Ebadian et al. [2022] use a different approach to measure the representativeness of a panel by considering the social cost (sum of q -costs) over a panel. They define the representativeness of an algorithm as the worst-case ratio between the social cost of the optimal panel that minimizes the objective and the (expected) social cost obtained by the algorithm. In their empirical analysis, they measure the average approximation to the optimal social cost of an algorithm \mathcal{A} over a set of instances \mathcal{I} , defined as

$$\frac{1}{|\mathcal{I}|} \sum_{I \in \mathcal{I}} \frac{\min_P \text{sc}(P; I)}{\text{sc}(\mathcal{A}(I); I)},$$

where $\text{sc}(P; I) = \sum_{i \in [n]} c_q(i, P)$.

Since finding the optimal panel is a hard problem and the dataset and panel sizes are large, the authors use a proxy for the minimum social cost, specifically, an implementation of the algorithm of Kumar and Raichel [2013] for the fault-tolerant k -median problem that achieves a constant factor approximation of the optimal objective – which is equivalent to minimizing the q -social cost. We use the same approach and report the average approximation to the optimal social cost.

5.4 Experiment Results

Next, we present the results of our experiments for the two metrics described above. We report the mean and the standard error of the mean for each metric, computed over 100 instances generated based on both the Adult and ESS-UK datasets illustrated in figures 2 and 3.

Ex post Core Violation. In Adult dataset, we observe unbounded ex post core violation for UNIFORM when $q \leq 4$. Specifically, for $q \in \{1, 2, 3\}$, we observed unbounded core violation in 84%, 9%, and

36% of the instances respectively. This happens since $\sim 8.3\%$ of the population is mapped to a single data point and that UNIFORM fails to select q individuals from this group. When $q \leq 3$, we have $q/k \leq 8.4\%$, and this cohesive group is entitled to select at least q members of the panel from themselves, which results in q -cost of 0 for them and an unbounded violation of the core. However, FAIRGREEDY captures this cohesive group and selects at least q representatives from them. Furthermore, we see significantly higher ex post core violation for UNIFORM compared to FAIRGREEDY for smaller values of q (less than 12) and similar performance for larger values of q . This is expected as FAIRGREEDY tends to behave more similarly to UNIFORM as q increases because it selects from fewer yet larger groups ($\lfloor k/q \rfloor + 1$ groups of size qn/k to be exact).

We observe a similar pattern in ESS-UK that UNIFORM obtains worse ex post core violations when q is smaller and similar performance as FAIRGREEDY for larger values of q . However, in contrast to Adult, we do not observe similar unbounded violations for UNIFORM in ESS-UK. The reason is that ESS-UK consists of 250 features (compared to the 5 we used from Adult) and any data points represent at most 0.2% of the population. Thus, no group is entitled to choose enough representatives from their own to significantly improve their cost or make it 0. The decline in core violation for $q = k$ happens as it measures the minimum improvement in cost over the whole population, which is more demanding than lower values of q . Lastly, FAIRGREEDY performs consistently for all values of q and achieves an ex post core violation less than 1.6 and 1.25 in Adult and ESS-UK respectively.

Approximation to Optimal Social Cost. For ESS-UK, we observe a similar behaviour from both UNIFORM and FAIRGREEDY, while for Adult, FAIRGREEDY outperforms UNIFORM for $q \in [3]$, which is again due to FAIRGREEDY capturing the cohesive group.

All considered, we observe that FAIRGREEDY can maintain at least the same level or even better optimal social cost approximation as UNIFORM would, while achieving significantly better empirical core guarantees in the two datasets.

6 DISCUSSION

This work introduces a novel notion of proportional representation, called core, in sortition. We show that uniform selection, under which representatives are chosen uniformly at random, achieves almost ex ante core, but it fails to guarantee that any panel that it might return is in the core, i.e. it is not in the ex post core. However, we present a selection algorithm, called Fair Greedy Capture, that is almost in the ex post core and is also fair (all the individuals are selected with the same probability). Lastly, we suggest an efficient auditing algorithm for measuring how much a given panel violates the core.

There are many directions for future work. First, while we show that Fair Greedy Capture is in the $\frac{5+\sqrt{41}}{2}$ -core, we did not find a lower bound that indicates if this analysis is tight. Moreover, we show that there is no fair algorithm that is in the ex post α -core for $\alpha < 2$. Closing the gap between this lower bound and the bound that Fair Greedy Capture guarantees is another interesting direction. Ebadian et al. [2022] show that while in some cases, fairness and a good approximation to the optimal social cost are incompatible, by relaxing a bit the fairness requirement, it is possible to find distributions that are almost fair and achieve good expected social cost. It would be interesting to consider the trade-off between fairness and approximation to the ex post core, as well. How much should we relax the fairness requirement to achieve ex post core for any value of q ? Moreover, Micha and Shah [2020] show that for $q = 1$, Greedy Capture, that was introduced by Chen et al. [2019], provides better guarantees for the Euclidean space. So, another interesting question is to see if when the metric d consists of usual distance functions such as norms L^2 , L^1 and L^∞ , Fair Greedy Capture can provide better guarantees.

REFERENCES

- Elliot Anshelevich, Aris Filos-Ratsikas, Nisarg Shah, and Alexandros A Voudouris. 2021. Distortion in social choice problems: The first 15 years and beyond. *arXiv preprint arXiv:2103.00911* (2021).
- Kenneth Arrow. 1990. *Advances in the spatial theory of voting*. Cambridge University Press.
- Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. 2017. Justified representation in approval-based committee voting. *Social Choice and Welfare* 48, 2 (2017), 461–485.
- G. Benadè, P. Gözl, and A. D. Procaccia. 2019. No Stratification Without Representation. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*. 281–314.
- I. Caragiannis, N. Shah, and A. A. Voudouris. 2022. The Metric Distortion of Multiwinner Voting. *Artificial Intelligence* 313 (2022), 103802.
- L. E. Celis, L. Huang, and N. K. Vishnoi. 2018. Multiwinner Voting with Fairness Constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*. 144–151.
- X. Chen, B. Fain, L. Lyu, and K. Munagala. 2019. Proportionally fair clustering. In *International Conference on Machine Learning*. 1032–1041.
- Y. Cheng, Z. Jiang, K. Munagala, and K. Wang. 2020. Group fairness in committee selection. *ACM Transactions on Economics and Computation (TEAC)* 8, 4 (2020), 1–18.
- V Conitzer, R Freeman, N Shah, and J. W. Vaughan. 2019. Group Fairness for the Allocation of Indivisible Goods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*. 1853–1860.
- D. Dua and C. Graff. 2017. UCI Machine Learning Repository. <http://archive.ics.uci.edu/ml>
- S. Ebadian, G. Kehne, E. Micha, A. D. Procaccia, and N. Shah. 2022. Is Sortition Both Representative and Fair?. In *Proceedings of the 36th Annual Conference on Neural Information Processing Systems (NeurIPS)*. 25720–25731.
- James M Enelow and Melvin J Hinich. 1984. *The spatial theory of voting: An introduction*. CUP Archive.
- F. Engelstad. 1989. The assignment of political office by lot. *Social Science Information* 28, 1 (1989), 23–50.
- B. Fain, K. Munagala, and N. Shah. 2018. Fair allocation of indivisible public goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*. 575–592.
- Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. 2017. Multiwinner voting: A new challenge for social choice theory. *Trends in computational social choice* 74, 2017 (2017), 27–47.
- B. Flanigan, P. Gözl, A. Gupta, B. Hennig, and A. D. Procaccia. 2021a. Fair Algorithms for Selecting Citizens’ Assemblies. *Nature* 596 (2021), 548–552.
- B. Flanigan, P. Gözl, A. Gupta, and A. D. Procaccia. 2020. Neutralizing Self-Selection Bias in Sampling for Sortition. In *Proceedings of the 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*. 6528–6539.
- B. Flanigan, G. Kehne, and A. D. Procaccia. 2021b. Fair Sortition Made Transparent. In *Proceedings of the 35th Annual Conference on Neural Information Processing Systems (NeurIPS)*. 25720–25731.
- J. Gastil and E. O. Wright (Eds.). 2019. *Legislature by Lot: Transformative Designs for Deliberative Governance*. Verso.
- Ronny Kohavi and Barry Becker. 1996. Adult data set. *UCI machine learning repository* 5 (1996), 2093.
- N. Kumar and B. Raichel. 2013. Fault Tolerant Clustering Revisited. In *Proceedings of the 25th Canadian Conference on Computational Geometry, CCCG 2013, Waterloo, Ontario, Canada, August 8-10, 2013*. Carleton University, Ottawa, Canada.
- Martin Lackner and Piotr Skowron. 2023. *Multi-winner voting with approval preferences*. Springer Nature, 9783031090158.
- J. Lang and P. Skowron. 2018. Multi-attribute proportional representation. *Artificial Intelligence* 227 (2018), 74–106.
- E. Micha and N. Shah. 2020. Proportionally fair clustering revisited. In *47th International Colloquium on Automata, Languages, and Programming (ICALP)*. Schloss Dagstuhl, Saarbrücken, Germany, 85:1–85:16.
- A. D. Procaccia and J. S. Rosenschein. 2006. The Distortion of Cardinal Preferences in Voting. In *Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA)*. 317–331.
- P. Stone. 2011. *The Luck of the Draw: The Role of Lotteries in Decision Making*. Oxford University Press.
- D. Van Reybrouck. 2016. *Against Elections: The Case for Democracy*. Random House.
- Frank Yates. 1948. Systematic sampling. *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences* 241, 834 (1948), 345–377.

A RELATIONSHIP BETWEEN EX POST α -CORE AND α -CORE OVER EXPECTED COST

PROPOSITION A.1. *For any $q \in [k]$, ex post α -core and α -core over expected cost are incomparable.*

PROOF. First, we show that there exists an instance and a selection algorithm that achieves ex post core that instance, but violates α -core over expected cost for any $\alpha \geq 1$.

Assume that n is divisible by k and q is divisible by 3. Consider an instance where there are five groups of individuals, A, B, C, D and E . The first three groups contain $(q \cdot n/k - q)/3$ individuals each, the fourth groups contains q individuals and the last group contains $n - q \cdot n/k$ individuals. The distance between individuals belonging to given groups is specified in the following table.

	A	B	C	D	E
A	0	2	2	1	∞
B	2	0	2	1	∞
C	2	2	0	1	∞
D	1	1	1	0	∞
E	∞	∞	∞	∞	0

Suppose that a selection algorithm $\mathcal{A}_{k,q}$ returns with probability $1/3$ a panel that contains q individuals from group A and the remaining representatives are from group E , with probability $1/3$ a panel that contains q individuals from group B and the remaining representatives are from group E and with probability $1/3$ a panel that contains q individuals from group C and the remaining representatives are from group E . All these panels are in the ex post core since there is no sufficiently large group such that if they choose another panel, all of them reduce their distance. Now, we see that for each i in A or B or C , it holds that

$$\mathbb{E}_{P \sim \mathcal{A}_{k,q}} [c_q(i, P)] = \frac{2}{3} \cdot 2 = 4/3$$

while for each i in D , it holds that

$$\mathbb{E}_{P \sim \mathcal{A}_{k,q}} [c_q(i, P)] = 1.$$

If all the individuals in A, B, C and D choose a panel P' that contains q individuals from D , then all of them reduce their distance by a factor larger than α .

Next, we show an instance and a selection algorithm that achieves core over expected cost. Now, suppose that $n > q^2$. Consider an instance where there are four groups of individuals, A, B, C, D . The first group contains $q \cdot n/k - q$ individuals each, the second group contains q individuals, the third group contains q individuals and the last group contains all the remaining individuals. The distance between individuals belonging to given groups is specified in the following table.

	A	B	C	D
A	0	1	2	∞
B	1	0	1	∞
C	2	1	0	∞
D	∞	∞	∞	0

Suppose that a selection algorithm $\mathcal{A}_{k,q}$ returns with probability $1/2$ a panel P_1 that contains q individuals from group C and $k - q$ individuals from group D , and with the remaining probability returns a panel P_2 that contains q individuals from group A and $k - q$ individuals from group D . Then, for each i in A , we have that

$$\mathbb{E}_{P \sim \mathcal{A}_{k,q}} [c_q(i, P)] = \frac{1}{2} \cdot 2 = 1$$

while for each i in b , we have that

$$\mathbb{E}_{P \sim \mathcal{A}_{k,q}}[c_q(i, P)] = \left(1 - \frac{1}{2}\right) \cdot 1 + \frac{1}{2} \cdot (1 - 2) = 0.$$

Hence, this algorithm is in the core over expected cost. But when the algorithm returns P_1 , all the individuals in A and B can reduce their cost by a factor of at least α by choosing q representatives in B . \square

B PROOF OF THEOREM 2.1

PROOF. Consider a panel P returned by uniform selection. According to Definition 1.3 of the ex-post core, it suffices to show that for any arbitrary panel P' of size k , the q -cost of all individuals cannot be improved by a factor of greater than $\alpha = 2$.

Let i_1 and i_2 be the two individuals in the population with the maximum distance between them. Now, consider an arbitrary representative r in panel P' . Without loss of generality, suppose that $c_k(i_1, P') \leq c_k(i_2, P')$. Then, we have

$$\begin{aligned} c_k(i_2, P) &= \max_{j \in P} d(i_2, j) \leq d(i_1, i_2) && \text{(by the choice of } i_1 \text{ and } i_2\text{)} \\ &\leq d(i_1, r) + d(r, i_2) && \text{(triangle inequality)} \\ &\leq c_k(i_1, P') + c_k(i_2, P') && \text{(as } r \in P'\text{)} \\ &\leq 2 \cdot c_k(i_2, P'). \end{aligned}$$

This implies that for any panel P in \mathcal{U}_k , $V_q(P, P', 2) < |P'| \cdot n/k = n$, since q -cost for i_2 does not improve by a factor of more than two. Thus, uniform selection is in the ex post 2-core. Furthermore, this implies that we have $\mathbb{E}_{P \sim \mathcal{U}_k}[c_k(i_2, P)] \leq \max_{P \in \mathcal{U}_k} c_k(i_2, P) \leq 2 \cdot c_k(i_2, P')$, which means uniform selection is also in the 2-core over expected cost. This is because to violate 2-core over expected cost, the q -cost of the entire population would have to improve by a factor of more than 2, which does not hold for individual i_2 .

Next, we show that there exists an instance such that uniform selection is not in the ex post α -core for $\alpha < 2$. Consider the case that the individuals are assigned into three groups, A , B and C , with $\lfloor k/2 \rfloor$, $\lceil k/2 \rceil$, and $n - k$ individuals, respectively. The distances between individuals is as specified in the following table.

	A	B	C
A	0	2	1
B	2	0	1
C	1	1	0

Panel P consisting of all the k people in groups A and B is in the support of uniform selection. Then, for $i \in A \cup B$, $c_k(i, P) = 2$ as the k th closest representative in P lies in the other group. For $i \in C$, the $c_k(i, P) = 1$. Now, consider panel P' that consists of k individuals from group C . The q -costs of all individuals improve by a factor of at least 2. Hence, \mathcal{U}_k violates ex post 2-core in this example. \square

C PROOF OF THEOREM 2.2

PROOF. Consider an instance in which there are $\lfloor n/k \rfloor$ individuals in group A and the remaining individuals in group B . Suppose that the distance between two individuals in the same group is 0, and the distance between two individuals in different groups is 1. Since, $\lfloor n/k \rfloor \geq k$, uniform selection has a non-zero probability of returning a panel where all representatives are from group A . In this scenario, for any $q \in [k - 1]$, the cost of all the individuals in group B is equal to 1. However, individuals in group B are entitled to choose up to $k - 1$ representatives among themselves, and

if they do so, their cost becomes 0, resulting in an unbounded improvement in cost. Therefore, uniform selection is not in the ex post α -core for any bounded α .

Furthermore, individuals in group B have positive expected costs under uniform selection, while if they choose a panel among themselves, they would all have a cost of 0. Thus, uniform selection is not in the α -core over expected cost for any bounded α . \square

D PROOF OF PROPOSITION 2.3

PROOF. To satisfy ex ante core when $q = k$, by Definition 1.5, for any panel $|P'|$ of size k , we should have

$$\mathbb{E}_{P \sim \mathcal{D}_k} [V_q(P, P', \alpha)] < |P'| \cdot \frac{n}{k} = n.$$

Since $V_q(P, P', \alpha) \leq n$, it suffices to show that there exists a non-zero probability that $V_q(P, P', \alpha) < n$. Since, \mathcal{U}_k chooses any panel, including P' , with non-zero probability, there is a positive probability that we realize panel $P = P'$ for which $V_q(P, P', \alpha) = 0$ – since the q -costs do not strictly improve for any individual. Thus, the expectation of the pairwise score is strictly less than n , satisfying the ex ante core. \square

E PROOF OF THEOREM 2.4

PROOF. Consider a star graph with $n - q$ leaves and an internal node. Suppose q individuals $I = \{i_1, \dots, i_q\}$ lie on the internal node, and exactly one individual lies on each of the $n - q$ leaves. Individuals in I have a distance of 0 from each other and a distance of 1 from $[n] \setminus I$; and, the distance between a pair of individuals from $[n] \setminus I$ is equal to 2. These distances satisfy the triangle inequality.

Let P be an arbitrary panel of size k that does not contain i_1 . We show that for $P' = I$ and any $\alpha < 2$, we have that $V_q(P, P', \alpha) \geq n - k$. For any $i \in I$, it holds $c_q(i, P) = 1$ and $c_q(i, P') = 0$ – which is an unbounded improvement. For any individual i in $[n] \setminus (I \cup P)$, $c_q(i, P) = 2$ since their q th closest representative in P would be on another leaf, while $c_q(i, P') = 1$ – which yields a 2 factor improvement. Therefore, $V_q(P, P', \alpha) \geq |([n] \setminus (I \cup P)) \cup I| \geq n - |P| = n - k$.

Under any fair selection algorithm, i_1 is not included in the panel with probability $1 - k/n$. Thus, we have that

$$\mathbb{E}_{P \sim \mathcal{U}_k} [V_q(P, P', \alpha)] \geq \Pr[i_1 \notin P] \cdot (n - k) = (1 - k/n) \cdot (n - k) \geq q \cdot n/k = |P'| \cdot n/k,$$

where the last inequality follows from the assumption that $n \geq 2k^2/(k - q)$. \square

F PROOF OF CHU-VANDERMONDE IDENTITY

PROOF. We give a combinatorial argument for this identity. Suppose we want to select $k + 1$ items out of a set of size $n + 1$. For $i \in [1, n + 1]$, let P_i be the number of such subsets in which the $(r + 1)$ th picked item is item i . As each subset is counted exactly once among P_i 's (at the position of its $(r + 1)$ th item), we have $\sum_{i=1}^{n+1} P_i = \binom{n+1}{k+1}$. Now, we calculate P_i . Suppose the $(r + 1)$ th item is i . Then, r items should be selected from the first $i - 1$ items and $k + 1 - (r + 1) = k - r$ items should be selected from the last $n + 1 - i$ items. Therefore, $P_i = \binom{i-1}{r} \cdot \binom{n-(i-1)}{k-r}$. Then, we have

$$\binom{n+1}{k+1} = \sum_{i=1}^{n+1} P_i = \sum_{i=1}^{n+1} \binom{i-1}{r} \cdot \binom{n-(i-1)}{k-r} = \sum_{j=0}^n \binom{j}{r} \cdot \binom{n-j}{k-r}. \quad \square$$

G PROOF OF LEMMA 3.1

PROOF. Consider a ball centered at i' with radius $c_q(i', P)$. This ball contains at least q representatives of P . Hence, $c_q(i, P)$ is less than or equal to the distance of i to one of the q representatives that are included in $B(i', c_q(i', P))$ which is at most $d(i, i') + c_q(i', P)$. \square

H PROOF OF THEOREM 3.3

PROOF. We have already shown that Fair Greedy Capture is fair. Hence, we focus on showing its approximation to the ex post $\frac{3+\sqrt{17}}{2}$ -core, when $q = 1$.

Let P be any panel in the support of the distribution that algorithm returns. Suppose for contradiction that there exists $S \subseteq [n]$ and a panel P' , with $|S| \geq |P'| \cdot n/k$, such that

$$\forall i \in S, \quad c_q(i, P) > \frac{3 + \sqrt{17}}{2} \cdot c_q(i, P').$$

As, we have described in Section 3, when $q = 1$, we may assume without loss of generality that $|P'| = 1$. Let $P' = \{i^*\}$ and i' be the individual in S that has the largest distance from i^* . Since there are sufficiently many individuals in the ball $B(i^*, d(i', i^*))$, it is possible that the algorithm detected it during its execution. If this happened, this means that there is 1 representative in P that is located within this ball. Then, we get that i' has a distance at most equal to the diameter of the ball from her closest representative in P which is at most $2 \cdot d(i', i^*)$. Hence, i' cannot reduce her distance by a multiplicative factor larger than 2 by choosing P' , and we arrive at a contradiction. On the other hand, if the algorithm did not detect this ball during its execution, this means that some of the individuals in T have been disregarded before the ball centered at i^* captures sufficiently many of them. Hence, some individuals in T have been captured from a different ball with radius at most $d(i', i^*)$. Suppose that i'' is the first individual in T that was captured from such a ball. Then, we have that within this ball there is 1 representative in P . Hence $c_q(i'', P) \leq 2 \cdot d(i', i^*)$, since the distance of i'' from any other individual in this ball is at most equal to the diameter of the ball. We consider the minimum multiplicative improvement of both i' and i'' :

$$\begin{aligned} & \min \left(\frac{c_q(i', P)}{d(i', i^*)}, \frac{c_q(i'', P)}{d(i'', i^*)} \right) \\ & \leq \min \left(\frac{d(i', i'') + c_q(i'', P)}{d(i', i^*)}, \frac{c_q(i'', P)}{d(i'', i^*)} \right) && \text{(by Lemma 3.1)} \\ & \leq \min \left(\frac{d(i', i^*) + d(i^*, i'') + c_q(i'', P)}{d(i', i^*)}, \frac{c_q(i'', P)}{d(i'', i^*)} \right) \\ & \leq \min \left(\frac{d(i', i^*) + d(i^*, i'') + 2 \cdot d(i', i^*)}{d(i', i^*)}, \frac{2 \cdot d(i', i^*)}{d(i'', i^*)} \right) && \text{(as } c_q(i'', P) \leq 2 \cdot d(i', i^*) \text{)} \\ & \leq \max_{z \geq 0} \min(3 + 1/z, 2 \cdot z) = (3 + \sqrt{17})/2. \end{aligned}$$

To show that this bound is tight consider the case that $n = 28$ and $k = 7$. Assume that the individuals form four isomorphic sets of 7 individuals each such that each set is sufficiently far from all other sets. The distances between the individuals in one set are given in the table below. Since $k = 7$ and there are four isomorphic groups, there exists a group that has at most one representative in some realized panel. Note that the algorithm first detects the balls that are centered at a_5 and have radius equal to $1 - \epsilon$. Assume that when this ball was detected in the group that has one representative in the panel, the algorithm chooses a_7 to be included in the panel. Then, in this group the individuals a_1, a_2, a_3 and a_4 are eligible to choose a_2 and all of them reduce their distance by a multiplicative factor of at least $(\sqrt{17} + 3)/2$ as ϵ goes to zero. \square

	a_1	a_2	a_3	a_4	a_5	a_6	a_7
a_1	0	1	2	$\frac{\sqrt{17}-1}{2}$	$\frac{\sqrt{17}+1}{2} - \epsilon$	$\frac{\sqrt{17}+1}{2} - \epsilon$	$\frac{\sqrt{17}+3}{2} - 2 \cdot \epsilon$
a_2	1	0	1	$\frac{\sqrt{17}-3}{2}$	$\frac{\sqrt{17}-1}{2} - \epsilon$	$\frac{\sqrt{17}-1}{2} - \epsilon$	$\frac{\sqrt{17}+1}{2} - 2 \cdot \epsilon$
a_3	2	1	0	$\frac{\sqrt{17}-1}{2}$	$\frac{\sqrt{17}+1}{2} - \epsilon$	$\frac{\sqrt{17}+1}{2} - \epsilon$	$\frac{\sqrt{17}+3}{2} - 2 \cdot \epsilon$
a_4	$\frac{\sqrt{17}-1}{2}$	$\frac{\sqrt{17}-3}{2}$	$\frac{\sqrt{17}-1}{2}$	0	$1 - \epsilon$	$1 - \epsilon$	$2 - 2\epsilon$
a_5	$\frac{\sqrt{17}+1}{2} - \epsilon$	$\frac{\sqrt{17}+1}{2} - \epsilon$	$\frac{\sqrt{17}-1}{2} - \epsilon$	$1 - \epsilon$	0	0	$1 - \epsilon$
a_6	$\frac{\sqrt{17}+1}{2} - \epsilon$	$\frac{\sqrt{17}+1}{2} - \epsilon$	$\frac{\sqrt{17}-1}{2} - \epsilon$	$1 - \epsilon$	0	0	$1 - \epsilon$
a_7	$\frac{\sqrt{17}+3}{2} - 2\epsilon$	$\frac{\sqrt{17}+3}{2} - 2\epsilon$	$\frac{\sqrt{17}+1}{2} - 2\epsilon$	$2 - 2\epsilon$	$1 - \epsilon$	$1 - \epsilon$	10

I PROOF OF THEOREM 3.4

PROOF. Let $\mathcal{D}_{k,q}$ be the distribution that Fair Greedy Capture returns. Suppose for contradiction that there exists $S \subseteq [n]$ and $P' \subseteq [n]$, with $|S| \geq |P'| \cdot n/k$, such that

$$\forall i \in S, \quad \mathbb{E}_{P \sim \mathcal{D}_{k,q}} [c_q(i, P)] > 6 \cdot c_q(i, P').$$

Let T_1, \dots, T_m be a partition of S with respect to P' , as given in the second part of Lemma 2.6. Since $m \leq \lfloor |P'|/q \rfloor$ and $|S| \geq |P'| \cdot n/k$, we conclude that there exists a part, say T_ℓ , that has size at least $q \cdot n/k$. Moreover, since, for each $i \in T_\ell$, it holds that $c_q(i, P') \leq c_q(i_\ell^*, P')$ and $\text{top}_q(i, P') \cap \text{top}_q(i_\ell^*, P') \neq \emptyset$ for some $i_\ell^* \in T_\ell$, we can conclude that for each $i \in T_\ell$, $d(i, i_\ell^*) \leq 2 \cdot c_q(i_\ell^*, P')$, by considering a representative in $\text{top}_q(i, P') \cap \text{top}_q(i_\ell^*, P')$ and applying the triangle inequality, i.e. $d(i, i_\ell^*) \leq d(i, r_i) + d(r_i, i_\ell^*) \leq 2 \cdot c_q(i_\ell^*, P')$, where $r_i \in \text{top}_q(i, P') \cap \text{top}_q(i_\ell^*, P')$.

Thus, we conclude that there exists a ball centered at i_ℓ^* that has radius $2 \cdot c_q(i_\ell^*, P')$ and captures all the individuals in T_ℓ . Since there are sufficiently many individuals in this ball, for any panel P in the support of $\mathcal{D}_{k,q}$, we conclude that either there exist q representatives in the ball $B(i_\ell^*, 2 \cdot c_q(i_\ell^*, P'))$ (meaning that this ball was detected), or there exists an i' in this ball that was captured from a different ball with radius at most $2 \cdot c_q(i_\ell^*, P')$. In the first case, we have that $c_q(i_\ell^*, P) \leq 2 \cdot c_q(i_\ell^*, P')$, while in the second case we have that

$$c_q(i_\ell^*, P) \leq d(i_\ell^*, i') + c_q(i', P) \leq 6 \cdot c_q(i_\ell^*, P')$$

where the first inequality follows from Lemma 3.1 and the second inequality follows from the fact that i' was captured from a ball of radius at most $2 \cdot c_q(i_\ell^*, P)$, and hence the distance of i' to any representative in this ball is at most $4 \cdot c_q(i_\ell^*, P)$ and the fact that $d(i_\ell^*, i') \leq 2 \cdot c_q(i_\ell^*, P')$. Hence, for any P in the support of the algorithm, we have that $c_q(i_\ell^*, P) \leq 6 \cdot c_q(i_\ell^*, P')$, and hence $\mathbb{E}_{P \sim \mathcal{D}_{k,q}} c_q(i_\ell^*, P) \leq 6 \cdot c_q(i_\ell^*, P')$ which is a contradiction. \square

J PROOF OF THEOREM 4.1

PROOF. Suppose for contradiction that while the algorithm returns $\hat{\alpha}$, there exists $S \subseteq [n]$ and $P' \subseteq [n]$, with $|S| \geq |P'| \cdot n/k$, such that

$$\forall i \in S, \quad c_q(i, P^*) > (3 \cdot \hat{\alpha} + 2) \cdot c_q(i, P').$$

First, note that if the algorithm outputs $\hat{\alpha}$, this means that for every $j \in [n]$, it holds that

$$V_q(P, \hat{P}_j, \hat{\alpha}) < |\hat{P}_j| \cdot n/k, \quad (3)$$

as otherwise the algorithm would output a value strictly larger than $\hat{\alpha}$.

Let T_1, \dots, T_m be a partition of S with respect to P' , as given in the second part of Lemma 2.6. Since $m \leq \lfloor |P'|/q \rfloor$ and $|S| \geq |P'| \cdot n/k$, we conclude that there exists a part, say T_ℓ , that has size at least $q \cdot n/k$. Moreover, since, for each $i \in T_\ell$, it holds that $c_q(i, P') \leq c_q(i_\ell^*, P')$ and

ALGORITHM 3: Auditing Greedy Capture

```

 $\delta \leftarrow 0; \hat{P} \leftarrow \emptyset; C \leftarrow \emptyset; N \leftarrow [n];$ 
while  $N \neq \emptyset$  do
  Smoothly increase  $\delta$ ;
  while  $\exists j \in C$  such that  $|B(j, \delta) \cap N| \geq 1$  do
     $W_j \leftarrow W_j \cup B(j, \delta);$ 
     $N \leftarrow N \setminus B(j, \delta);$ 
  end
  while  $\exists j \in N$  such that  $|B(j, \delta) \cap N| \geq \lceil q \cdot n/k \rceil$  do
     $\hat{P}_j \leftarrow \{j\} \cup q - 1$  closest neighbors of  $j$ ;
     $\hat{P} \leftarrow \hat{P} \cup \hat{P}_j;$ 
     $C \leftarrow C \cup \{j\};$ 
     $W_j \leftarrow B(j, \delta);$ 
     $N \leftarrow N \setminus B(j, \delta);$ 
  end
end

```

$\text{top}_q(i, P') \cap \text{top}_q(i_\ell^*, P') \neq \emptyset$ for some $i_\ell^* \in T_\ell$, we can conclude that for each $i \in T_\ell$, $d(i, i_\ell^*) \leq 2 \cdot c_q(i_\ell^*, P')$, by considering a representative in $\text{top}_q(i, P') \cap \text{top}_q(i_\ell^*, P')$ and applying the triangle inequality, i.e. $d(i, i_\ell^*) \leq d(i, r_i) + d(r_i, i_\ell^*) \leq 2 \cdot c_q(i_\ell^*, P')$, where $r_i \in \text{top}_q(i, P') \cap \text{top}_q(i_\ell^*, P')$. This means there exists a ball centered at i_ℓ^* that has radius $2 \cdot c_q(i_\ell^*, P')$ and captures all the individuals in T_ℓ . Now, note that there exists $i' \in T_\ell$ such that $\hat{\alpha} \cdot c_q(i', \hat{P}_{i'}^*) \geq c_q(i', P)$, since otherwise for each $i \in T_\ell$ would hold that $\hat{\alpha} \cdot c_q(i', \hat{P}_{i'}^*) < c_q(i', P)$ and then $V_q(P, \hat{P}_{i'}^*, \hat{\alpha}) \geq |q| \cdot n/k = |\hat{P}_{i'}^*| \cdot n/k$ which contradicts Equation (3). Hence,

$$\begin{aligned} c_q(i_\ell^*, P) &\leq d(i_\ell^*, i') + c_q(i', P) \leq d(i_\ell^*, i') + \hat{\alpha} \cdot c_q(i', \hat{P}_{i'}^*) \leq d(i_\ell^*, i') + \hat{\alpha} \cdot (d(i_\ell^*, i') + c_q(i_\ell^*, \hat{P}_{i'}^*)) \\ &\leq (3 \cdot \hat{\alpha} + 2) \cdot c_q(i_\ell^*, P'). \end{aligned}$$

where the first and the third inequalities follows from Lemma 3.1 and the last inequality follows from the facts that for each $i \in T$, $d(i, i_\ell^*) \leq 2 \cdot c_q(i_\ell^*, P')$, and $c_q(i_\ell^*, \hat{P}_{i'}^*) \leq c_q(i_\ell^*, P')$ for each P' since $\hat{P}_{i'}^*$ consists of the q closest neighbors of i_ℓ^* . \square

K AUDITING USING GREEDY CAPTURE

We first find a panel \hat{P} using Algorithm 3 which will server as the means to measure the maximum reduce to the cost that any individual in a subset of individuals can have by choosing a different panel with size proportional to the size of the subset. This algorithm is similar to Greedy Capture by Chen et al. [2019] with the difference that every time that a ball captures $\lceil q \cdot n/k \rceil$ individuals, the center along with its $q - 1$ closest representatives are included in the panel. We denote with \hat{P}_j the set of the q individuals that are included in the panel when the ball that centered at j captures $\lceil q \cdot n/k \rceil$ individuals for first time. When this happens, j is added to a set C that contains all the individuals whose ball caused to the addition of q individuals in the panel. Lastly, we denote with W_j the set that contains all the individuals that were disregarded because they were captured from the ball centered at j . In the theorem below, we show that if we cannot find a sufficient large subset of some W_j such that the reduce of the cost for each individual in the subset is at least α under \hat{P}_j , then P is in the expected $(4 \cdot \alpha + 6)$ -core.

THEOREM K.1. *If for no \hat{P}_j , there is $S' \subseteq W_j$ with $|S'| \geq q \cdot n/k$ such that for each $i \in S'$, $\alpha \cdot c_q(i, \hat{P}_j) < c_q(i, P)$, then for any P' with $|P'| \geq q$, we get that*

$$V_q(P, P', 4 \cdot \alpha + 6) \leq |P'| \cdot n/k.$$

PROOF. Suppose for contradiction that while for no \hat{P}_j , there is $S' \subseteq W_j$ with $|S'| \geq q \cdot n/k$ such that for each $i \in S'$, $\alpha \cdot c_q(i, \hat{P}_j) < c_q(i, P)$, there exists $S \subseteq [n]$ and $P' \subseteq [n]$, with $|S| \geq |P'| \cdot n/k$ and $|P'| \geq q$, such that

$$\forall i \in S, \quad c_q(i, P^*) > (4 \cdot \alpha + 6) \cdot c_q(i, P').$$

Let T_1, \dots, T_m be a partition of S with respect to P' , as given in the second part of Lemma 2.6. Since $m \leq \lfloor |P'|/q \rfloor$ and $|S| \geq |P'| \cdot n/k$, we conclude that there exists a part, say T_ℓ , that has size at least $q \cdot n/k$. Moreover, since, for each $i \in T_\ell$, it holds that $c_q(i, P') \leq c_q(i_\ell^*, P')$ and $\text{top}_q(i, P') \cap \text{top}_q(i_\ell^*, P') \neq \emptyset$ for some $i_\ell^* \in T_\ell$, we can conclude that for each $i \in T_\ell$, $d(i, i_\ell^*) \leq 2 \cdot c_q(i_\ell^*, P')$, by considering a representative in $\text{top}_q(i, P') \cap \text{top}_q(i_\ell^*, P')$ and applying the triangle inequality, i.e. $d(i, i_\ell^*) \leq d(i, r_i) + d(r_i, i_\ell^*) \leq 2 \cdot c_q(i_\ell^*, P')$, where $r_i \in \text{top}_q(i, P') \cap \text{top}_q(i_\ell^*, P')$. This means there exists a ball centered at i_ℓ^* that has radius $2 \cdot c_q(i_\ell^*, P')$ and captures all the individuals in T_ℓ . Since there are sufficiently many individuals in this ball, it is possible that Algorithm 3 detects a ball centered at i_ℓ^* with radius $r \leq 2 \cdot c_q(i_\ell^*, P')$ that contains at least $q \cdot n/k$ individuals that has not been disregarded yet. If this happened then the algorithm includes q representatives from $B(i_\ell^*, r)$ in $\hat{P}_{i_\ell^*}$ and sets $W_{i_\ell^*} = B(i_\ell^*, r)$. Since $|B(i_\ell^*, r)| \geq q \cdot n/k$, from the initial assumption that for no \hat{P}_j , there is $S' \subseteq W_j$ with $|S'| \geq q \cdot n/k$ such that for each $i \in S'$, $\alpha \cdot c_q(i, \hat{P}_j) < c_q(i, P)$, we get that there exists $i' \in B(i_\ell^*, r)$ such that $c_q(i', P) \leq \alpha \cdot c_q(i', \hat{P}_{i_\ell^*})$. Hence,

$$\begin{aligned} c_q(i_\ell^*, P) &\leq d(i_\ell^*, i') + c_q(i', P) \leq d(i_\ell^*, i') + \alpha \cdot c_q(i', \hat{P}_{i_\ell^*}) \leq d(i_\ell^*, i') + \alpha \cdot (d(i_\ell^*, i') + c_q(i_\ell^*, \hat{P}_{i_\ell^*})) \\ &\leq r + \alpha \cdot (r + r) \\ &\leq (4 \cdot \alpha + 2) \cdot c_q(i_\ell^*, P'). \end{aligned}$$

On the other hand, if the algorithm did not find this ball during its execution, this means that some of the individuals in T_ℓ have been disregarded before the ball centered at i_ℓ^* captures sufficiently many of them, which in its turn means that some individuals in T_ℓ have been captured from a different ball with radius at most $2 \cdot c_q(i_\ell^*, P')$. Suppose that i' is the first individual in T_ℓ that was captured from a different ball centered at c with radius r . Since, $|B(c, r)| \geq q \cdot n/k$, from the assumption that for no \hat{P}_j , there is $S' \subseteq W_j$ with $|S'| \geq q \cdot n/k$ such that for each $i \in S'$, $\alpha \cdot c_q(i, \hat{P}_j) < c_q(i, P)$, we know that there exists $i'' \in B(c, r)$ such that $c_q(i'', P) \leq \alpha \cdot c_q(i'', \hat{P}_\ell)$. Hence,

$$\begin{aligned} c_q(i_\ell^*, P) &\leq d(i_\ell^*, i'') + c_q(i'', P) \leq d(i_\ell^*, i'') + \alpha \cdot c_q(i'', \hat{P}_\ell) \leq d(i_\ell^*, i'') + d(i', i'') + \alpha \cdot c_q(i'', \hat{P}_\ell) \\ &\leq 2 \cdot c_q(i_\ell^*, P') + 2 \cdot r + \alpha \cdot 2 \cdot r \\ &\leq (4 \cdot \alpha + 6) \cdot c_q(i_\ell^*, P'). \end{aligned}$$

where the fourth inequality follows from the fact that both i' and i'' are captured from $B(c, r)$, and the fact that there are at least q representatives in $B(c, r)$. \square