

Temporal Panel Selection in Ongoing Citizens’ Assemblies

Yusuf Hakan Kalaycı and Evi Micha
University of Southern California
{kalayci, evi.micha}@usc.edu,

Permanent citizens’ assemblies are ongoing deliberative bodies composed of randomly selected citizens, organized into panels that rotate over time. Unlike one-off panels, which represent the population in a single snapshot, permanent assemblies enable shifting participation across multiple rounds. This structure offers a powerful framework for ensuring that different groups of individuals are represented over time across successive panels. In particular, it allows smaller groups of individuals that may not warrant representation in every individual panel to be represented across a sequence of them. We formalize this temporal sortition framework by requiring proportional representation both within each individual panel and across the sequence of panels.

Building on the work of Ebadian and Micha (2025), we consider a setting in which the population lies in a metric space, and the goal is to achieve both proportional representation, ensuring that every group of citizens receives adequate representation, and individual fairness, ensuring that each individual has an equal probability of being selected. We extend the notion of representation to a temporal setting by requiring that every initial segment of the panel sequence, viewed as a cumulative whole, proportionally reflects the structure of the population. We present algorithms that provide varying guarantees of proportional representation, both within individual panels and across any sequence of panels, while also maintaining individual fairness over time.

1 Introduction

In recent years, *citizens’ assemblies* have emerged as a promising model for community-driven governance Stone [2011], Van Reybrouck [2016]. The core idea is to randomly select a panel from the population, provide them with balanced information on an issue, allow time for deliberation, and then collect their recommendations, which are taken into account by the relevant authorities. This approach has been adopted in a variety of local contexts, with many city governments and local authorities implementing citizens’ assemblies to address policy issues, as well as in several prominent recent examples at the national level—such as in Ireland¹—and at the continental level through the European Union². At the heart of citizens’ assemblies is *sortition*: a randomized selection process creating a panel that reflects the general population’s diverse perspectives, not just technical expertise. To support this, the process aims for representativeness across key features (e.g., geography, education level) while ensuring every individual has a meaningful selection chance. In practice, organizers often implement stratified sampling with quotas on individual features—for instance, reserving a fixed percentage of seats for a particular age group or requiring that at least a certain proportion of representatives have a specific occupation Flanigan et al. [2021]. However, this method fails to account for more complex combinations of features—for example, individuals who are both from a specific region and work in a particular occupation. On the other hand, forcing representation across all such combinations is typically infeasible; with six features each taking just three values, there are already over 700 disjoint groups, far more than a typical panel in practice.

As a middle ground, recent works Ebadian et al. [2022], Ebadian and Micha [2025] propose leveraging an underlying representation metric space for defining when a panel represents an underlying population in a more rigorous way. This space measures how well an individual represents another individual or group, with smaller distances indicating stronger representation and can be constructed based on features relevant to the application. Given such a metric space, Ebadian and Micha [2025] adapt a notion of proportional representation previously used in multi-winner elections Lackner and Skowron [2022] and clustering Aziz et al. [2017], Chen et al. [2019], Micha and Shah [2020] to define when a panel is

¹<https://www.citizensinformation.ie/en/government-in-ireland/irish-constitution-1/citizens-assembly/>

²https://citizens.ec.europa.eu/european-citizens-panels_en

proportionally representative of an underlying population. At a high level, the notion requires that, given a population of size n and a panel of size k , any subset of the population of size at least $q \cdot n/k$ is entitled to up to q representatives.

As citizens’ assemblies gain increasing acceptance, the next step in this progression is the growing establishment of *permanent citizens’ assemblies*, an institutional innovation that enables ongoing participation rather than one-off deliberation. Unlike traditional, ad-hoc citizens’ assemblies that are convened for specific issues and disbanded afterward, these permanent bodies provide structured, recurring opportunities for citizens to participate directly in policymaking over extended periods. While the institution itself is continuous, its members are periodically rotated (typically every six months or annually) through sortition. The Ostbelgien Model in the German-speaking community of Belgium is one of the most prominent examples Niessen and Reuchamps [2019]. Established in 2019, it includes a standing Citizens’ Council that provides citizens a permanent voice in the process of decision making by rotating Citizens’ Panels. European Union has also recently experimented with a similar format Alemanno [2022]. This new innovation motivates the main question in this work:

In ongoing citizens’ assemblies, can we ensure that smaller groups receive representation over time, even if they do not qualify in every individual panel, while still guaranteeing that every sufficiently large group is proportionally represented in each panel?

Temporal Sortition. Permanent citizens’ assemblies create new opportunities to achieve broader representation across groups of individuals who share similar features. Smaller groups that may not qualify for representation on every individual panel can now be represented across a sequence of panels. To illustrate this, consider a toy example in which the population consists of four disjoint groups: the first makes up half of the population, the second a quarter, and the remaining two one-eighth each. Suppose we are selecting a panel of size 4. Proportionality would assign 2 seats to the first group, 1 to the second, and the last to either the third or fourth group, but not both. However, if we construct two consecutive panels, we can alternate this final seat—assigning it to the third group in the first panel and to the fourth in the second. In this way, both smaller groups receive representation over time, even if not in every individual panel.

This simple example demonstrates how permanent citizens’ assemblies can extend proportional representation beyond the limits of any single panel. To capture this formally, we introduce the *temporal sortition framework*, which models how representation can be extended across a sequence of panels. At a high level, the goal is to construct a sequence of randomly selected panels such that each panel proportionally represents every group that is large enough within that panel’s population; groups that are not eligible in a single panel still receive representation once enough seats accumulate across the sequence; and each individual has an equal probability of being selected throughout the process. Formally, given a population V of size n in a representative metric space, our objective is to construct a panel sequence $P_1, P_2, \dots, P_\ell$, where panel P_i has size k_i . Leveraging notions of proportional representation similar to those utilized by Ebadian and Micha [2025] for static panels, and assuming the existence of such a metric space, our goal is to define a distribution over these ℓ consecutive panels that satisfies the following properties:

1. *Representation per Panel:* Every subset of the population of size at least n/k_i (or $q \cdot n/k_i$) receives at least 1 (or q) representatives in panel P_i .
2. *Representation Over Time:* Every subset of the population of size at least $n/\sum_{i=1}^t k_i$ (or $q \cdot n/\sum_{i=1}^t k_i$) receives at least 1 (or q) representatives among the first t panels, i.e., in $\cup_{i=1}^t P_i$.
3. *Individual Fairness:* Each individual is included in the union of all panels $\cup_{i=1}^\ell P_i$ with equal probability, equal to $\sum_{i=1}^\ell k_i/n$.

To define representation rigorously in the presence of a metric space, we draw on two established notions from the literature: α -Proportionally Fair Clustering (α -PFC) by Chen et al. [2019], and β -Proportionally Representative Fairness (β -PRF) by Aziz et al. [2024]. Roughly speaking, α -PFC guarantees that any group of at least n/k individuals in the metric space has at least one representative whose distance from some member of the group is at most α times the group’s diameter. In contrast, β -PRF ensures that any group of size at least $q \cdot n/k$ has at least q representatives whose distance from some member of the group is at most β times the group’s diameter.

Technical Challenge and Our Contributions. In Section 3, as a warm-up, we consider the simpler setting where the objective is to ensure proportional representation for each individual panel as well as for the overall collection of

panels, which we refer to as the *global panel* $\bigcup_{i \in [\ell]} P_i$. We begin with the case of creating ℓ panels of size k . A natural approach is to proceed in two steps. First, we select $\ell \cdot k$ individuals from the population in a way that provides a constant-factor approximation to PRF, using known algorithms from the literature. Next, we partition this set of $\ell \cdot k$ individuals into ℓ panels of size k , again applying known algorithms so that each panel fairly represents this smaller group. Together, these two steps ensure that both the individual panels and the global panel satisfy a constant-factor approximation to PRF. More generally, we investigate what happens when a smaller panel is constructed from a larger one, i.e., we first form a large panel P_1 of size k_1 that satisfies α -PRF with respect to the whole population, and then extract a smaller panel P_2 of size k_2 from P_1 that satisfies α -PRF with respect to P_1 . We show that if k_1 is a multiple of k_2 , then P_2 also represents the original population within a $(2\alpha\beta + \beta)$ -approximation to PRF. Surprisingly, however, when k_1 is *not* divisible by k_2 , there exist instances where the smaller panel P_2 fails to satisfy *any* meaningful approximation of proportional representation for the population. This highlights that the problem is technically far more subtle than it might appear at first glance.

In Section 4, we strengthen the requirements by demanding that proportional representation hold not only for each individual panel and for the global panel, but also for every prefix of panels $P_{\leq t} = \bigcup_{i \in [t]} P_i$. This significantly increases the difficulty of the problem, since the goal is now to construct a global panel that is proportionally representative, decompose it into individual panels that are themselves proportionally representative, and order these panels so that every prefix $P_{\leq t}$ is proportionally representative, as if the task were to build a panel of size equal to the total size of that prefix, all while also ensuring individual fairness. While this may even appear impossible, we nevertheless design an intricate algorithm that, when the goal is to form ℓ panels of size k each, is able to guarantee individual representation and provide an approximation of PFC, with respect to both each individual panel and each prefix $P_{\leq t}$, that scales exponentially with the number of panels. The main property that the algorithm exploits is that under proportionality, any subset of individuals that deserves representation in smaller panels should also be represented in larger panels. Therefore, the algorithm constructs groups of individuals that must be represented at different panel sizes, i.e., for each prefix, and enforces that the groups for larger panels are nested within the groups for smaller panels. The algorithm then carefully leverages this nested structure to guarantee proportional representation for each panel and for every prefix. This hierarchical approach, however, comes at the cost of an exponential blow-up in the approximation factor with respect to the number of panels ℓ .

Finally, in Section 5 we consider a relaxed but still highly demanding requirement. Instead of enforcing proportional representation for each individual panel we require it only for every prefix $P_{\leq t} = \bigcup_{i \in [t]} P_i$ with $t \leq \ell$. While this abandons per-panel guarantees, it still captures a strong notion of representation by ensuring that cumulative representation is preserved over time. We show that this relaxation makes the problem more tractable and we present an algorithm that achieves a constant-factor approximation of PFC with respect to every prefix, while still guaranteeing that each individual is included in the global panel with equal probability. The algorithm is particularly intriguing and the key idea is to replace the hierarchical structure with a more flexible construction that carefully links together the groups that should be represented at different panel sizes, that is, across prefixes. This design ensures that representation achieved for earlier prefixes automatically extends to later ones, preventing error from accumulating and ultimately yielding a constant-factor bound.

Taken together, these results show that temporal sortition algorithms can achieve both per-panel and cumulative (over-time) representation guarantees, along with individual fairness. At the same time, our findings uncover subtle structural challenges that make the problem far from straightforward. Throughout the paper we highlight several intriguing open questions that we hope will inspire further research.

Related Work. The design of citizens’ assemblies that are representative of the broader population has received significant attention in the computer science literature Flanigan et al. [2021, 2020, 2024], Baharav and Flanigan [2024], Ebadian and Micha [2025], Ebadian et al. [2022], Caragiannis et al. [2024]. However, to the best of our knowledge, no prior work has addressed the problem of designing a sequence of panels that are representative of the population both individually and collectively over time. The only conceptually related work is the recent proposal by Halpern et al. [2025], who introduce federated assemblies: a hierarchical model in which assemblies are connected through a directed acyclic graph, and members of higher-level assemblies are drawn from lower-level ones. Their algorithms ensure individual fairness (each person has an equal probability of selection), ex ante fairness (each subassembly is expected to receive representation proportional to its size), and ex post fairness (the realized allocation closely approximates the expected proportions). In contrast, our work focuses on a sequence of stand-alone panels that, taken together over time, provide comprehensive representation, shifting the focus from structural to longitudinal representational equity.

Metric proportional representation has emerged as a central theme across computational social choice, clustering, and data summarization. Early clustering research introduced proportionally fair clustering [Chen et al., 2019] and individual fairness [Jung et al., 2020] establishing idea of ensuring representation for sufficiently large point sets. These concepts subsequently influenced committee selection, spawning frameworks for proportionally representative committees [Kalayci et al., 2024] and proportionally representative fairness [Aziz et al., 2024].

In sortition, Ebadian et al. [2022] first examined metric representation, asking whether sortition can simultaneously achieve representation and fairness objectives. Ebadian and Micha [2025] strengthened these results by introducing “core” proportionality and designing the Fair Greedy Capture algorithm that maintains individual fairness while ensuring constant-factor approximation to the ex post core. These single-shot formulations ($\ell = 1$) serve as foundation for our temporal framework extending to sequential multi-panel settings. We refer readers to [Kellerhals and Peters, 2025] for the landscape of metric proportional representation and their relationships.

Fairness across time horizons has gained momentum in perpetual voting [Lackner, 2020, Lackner and Maly, 2023, Bulteau et al., 2021], temporal committee selection [Elkind et al., 2024, 2025, Do et al., 2022], temporal clustering [Dey et al., 2017a,b], repeated matching [Gollapudi et al., 2020, Caragiannis and Narang, 2023, Trabelsi et al., 2023], and sequential allocation [Bampis et al., 2018, Igarashi et al., 2024]. Motivated by these works and the need for time-robust citizens’ assemblies, we extend single-shot metric sortition to multi-step settings where groups underrepresented in individual panels achieve fair representation across time.

2 Preliminaries

For $m \in \mathbb{N}$, let $[m] = \{1, \dots, m\}$. We denote the population by $V = [n]$. The individuals are embedded in an underlying *representation metric space* equipped with a distance function d , where the distance between any two individuals i and j is denoted by $d(i, j)$. We assume that d is a pseudo-metric³. An instance of our problem is fully specified by the set of individuals and their pairwise distances, which for simplicity denote by d . For any individual $v \in V$ and radius $r \geq 0$, we define the *ball* $B(v, r) = \{u \in V \mid d(v, u) \leq r\}$, i.e. the set of all individuals within distance r in the metric space from a given individual v .

A *temporal selection algorithm* $\mathcal{A}$ takes as input the population, the metric, the number of panels ℓ and panel sizes $k_1, k_2 \dots k_\ell$ as input and outputs a probability distribution over a sequence of ℓ disjoint panels $P_1, \dots, P_\ell$ where P_i has size equal to k_i . A sequence of panels, $P_1, \dots, P_\ell$, is in the support of $\mathcal{A}$, if the algorithm returns this sequence with positive probability.

Representation Axioms. We start by defining two notions of proportional representation that have been proposed in the literature for individual panels.

Definition 2.1 (α -Proportionally Fair Clustering (α -PFC) Chen et al. [2019]). *A panel $P \subseteq V$ of size k is α -proportionally fair for $\alpha \geq 1$ if, for any subset $S \subseteq V$ of size at least n/k , there exists an individual $v \in S$ and a panel member $p \in P$ such that*

$$d(v, p) \leq \alpha \cdot \min_{y \in V} \max_{u \in S} d(u, y).$$

In words, proportionally fair clustering ensures that no group of individuals of size at least n/k can identify an alternative representative not in the current panel that all its members would strictly prefer over their closest representative in the selected panel.

Definition 2.2 (β -Proportionally Representative Fairness (β -PRF) Aziz et al. [2024]). *A panel $P \subseteq V$ of size k is β -proportionally representative if for any set of individuals $S \subseteq V$ of size at least $q \cdot n/k$ where the maximum pairwise distance within S is r , we have:*

$$\left| P \cap \bigcup_{v \in S} B(v, \beta \cdot r) \right| \geq q.$$

³A *metric* is a non-negative function d on pairs satisfying: (i) $d(x, x) = 0$ for all x , (ii) *symmetry*: $d(x, y) = d(y, x)$ for all x, y , (iii) *triangle inequality*: $d(x, z) \leq d(x, y) + d(y, z)$ for all x, y, z , and (iv) *positivity*: $d(x, y) > 0$ whenever $x \neq y$. A *pseudo-metric* relaxes the positivity requirement, allowing distinct points to have zero distance.

In words, proportionally representative fairness ensures that for every group of individuals of size at least $q \cdot n/k$ with maximum pairwise distance r , there exist at least q representatives in the panel, each within distance r of some individual in the group.

Intuitively, proportionally fair clustering ensures that coalitions of size at least n/k receive representation by at least one panel member, while proportionally representative fairness requires that larger coalitions (of size at least $q \cdot n/k$) receive proportionally many representatives (specifically, q representatives). Recent work has established theoretical relationships between these concepts: Aziz et al. [2024] and Kellerhals and Peters [2025] showed that 1-PRF implies $(1 + \sqrt{2})$ -PFC, and this can be generalized to demonstrate that β -PRF implies $(1 + \sqrt{2}) \cdot \beta$ -PFC.

Axioms for Temporal Sortition. We evaluate the fairness and representation guarantees of a distribution across series of panels $P_1, \dots, P_\ell$ using the following axioms :

1. **Individual Fairness:** Each individual should have equal selection probability across the entire process. We say that a selection algorithm satisfies *individual fairness* if for all $v \in V$, $\Pr_{P_1, \dots, P_\ell \sim \mathcal{A}} [v \in P] = \sum_{i=1}^{\ell} k_i/n$, where $P = \bigcup_{i=1}^{\ell} P_i$.
2. **Individual Panel Representation:** For any representation axiom Π , we say that a selection algorithm $\mathcal{A}$ satisfies Π at the “panel level” if the output distribution satisfies that for any $\mathcal{P} = \{P_1, \dots, P_\ell\}$ in the support of $\mathcal{A}$ each P_i satisfies Π for population V and panel size k_i , simultaneously. In this work, we focus on α -PFC and β -PRF definitions.
3. **Global Panel Representation:** The union of all panels should satisfy representation axioms, providing long-term representation guarantees. Given a representation axiom Π , we say that a selection algorithm $\mathcal{A}$ satisfies Π at the *global level* if for every $\mathcal{P} = \{P_1, \dots, P_\ell\}$ in the support of $\mathcal{A}$, $\bigcup_{i=1}^{\ell} P_i$ satisfies axiom Π for population V and panel size $\sum_{i=1}^{\ell} k_i$.
4. **Prefix Representation:** For more stringent requirements beyond global panel representation, we demand that for every $t \in [\ell]$, the cumulative selection panel of the first t panels maintains proportional representation, ensuring representation quality at every stage. Given a proportionality axiom Π , we say that a selection algorithm $\mathcal{A}$ satisfies Π at the “prefix level” if for every $\mathcal{P} = \{P_1, \dots, P_\ell\}$ in the support of $\mathcal{A}$, the cumulative panel $P_{\leq t}$ for each $t \in [\ell]$ satisfies representation axiom Π for population V with panel size $\sum_{i=1}^t k_i$.

3 Warm-up: Proportional Representation Per Panel and Global Panel

As a warm-up, we first consider the case where the goal is to achieve representation with respect to each individual panel and the global panel, without prefix representation concerns. This setup can be interpreted as a federated sortition problem where ℓ panels operate simultaneously, and we seek to ensure representation both within each individual panel and across the union of all panels Halpern et al. [2025].

To design algorithms for this case, we begin by examining how representation guarantees behave when rules satisfying PRF axioms are applied in sequence. In particular, consider a population V and suppose we apply a selection algorithm satisfying the α -PRF axiom to obtain a panel P_1 of size k_1 . We then apply another selection algorithm satisfying the β -PRF axiom to the selected panel P_1 , yielding a smaller panel $P_2 \subseteq P_1$ of size k_2 . This raises a natural question: what proportionality guarantees does P_2 achieve with respect to the original population V and panel size k_2 ? In the next theorem, we show that this approach, provides an approximation guarantee that depends on α and β when k_1 is divisible by k_2 . But surprisingly, when k_1 is not divisible by k_2 , then P_2 may not provide any finite approximation with respect to V .

Theorem 3.1. *Given a population V and panel sizes k_1 and k_2 , let P_1 be a panel of size k_1 that satisfies α -PRF for population V , and let P_2 be a smaller panel of size k_2 that satisfies β -PRF for population P_1 . Then,*

- *when k_1 is divisible by k_2 , P_2 satisfies $(2\alpha \cdot \beta + \beta)$ -PRF for population V .*
- *when k_1 is not divisible by k_2 , there exists an instance where P_2 does not satisfy any finite approximation of PRF (or PFC) for population V .*

Proof. We start by the case where k_1 is divisible by k_2 . Consider an arbitrary coalition $S \subseteq V$ of size at least $\ell \cdot n/k_2 = \ell \cdot n/k_1 \cdot k_1/k_2$ and let $r = \max d(u, v) : u, v \in S$ be the diameter.

Since P_1 satisfies α -PRF with respect to V and panel size k_1 , we have

$$\left| P_1 \cap \bigcup_{v \in S} B(v, \alpha \cdot r) \right| \geq \ell \cdot \frac{k_1}{k_2}.$$

Let $T = P_1 \cap \bigcup_{v \in S} B(v, \alpha \cdot r)$. To bound the diameter of T , consider any two individuals $u, v \in T$. By definition of T , there exist individuals $u', v' \in S$ such that $u \in B(u', \alpha \cdot r)$ and $v \in B(v', \alpha \cdot r)$. Applying the triangle inequality and using $d(u', v') \leq r$, we obtain:

$$d(u, v) \leq d(u, u') + d(u', v') + d(v', v) \leq \alpha \cdot r + r + \alpha \cdot r = (2\alpha + 1) \cdot r.$$

Thus, the diameter of T is at most $(2\alpha + 1) \cdot r$. Since P_2 satisfies β -PRF with respect to population P_1 , we have

$$\left| P_2 \cap \bigcup_{v \in T} B(v, \beta \cdot (2\alpha + 1) \cdot r) \right| \geq \ell,$$

which establishes that P_2 satisfies $(2\alpha \cdot \beta + \beta)$ -PRF for population V .

For the case, where k_1 is not divisible by k_2 consider a population V of size 9 on the real axis with locations $V = \{x_1 = 0, x_2 = 0, x_3 = 1, x_4 = 10, x_5 = 10, x_6 = 10, x_7 = 11, x_8 = 11, x_9 = 11\}$. Let $k_1 = 4$ and $k_2 = 3$.

The panel $P_1 = \{x_1, x_3, x_4, x_7\}$ satisfies 1-PRF for population V . The panel $P_2 = \{x_1, x_3, x_4\}$ satisfies 1-PRF for population P_1 . However, P_2 does not satisfy any approximate PRF for V since coalition $\{x_7, x_8, x_9\}$ requires a representative at location 11, which P_2 lacks. □

Based on the above theorem, for the simple case where all panels have the same size k across all ℓ panels, we can utilize algorithms from the literature for achieving approximately PRF for each panel and the global panel and individual fairness properties. In particular, we can first apply the algorithm called, Fair Greedy Capture by Ebadian and Micha [2025], for choosing a panel P with size $\ell \cdot k$ which satisfies 6-PRF and ensures that each individual is selected with the same probability⁴. Then we can apply, the algorithm called, Metric Expanding Approval by Aziz et al. [2024] with parameters P and k for partitioning the $\ell \cdot k$ individuals in P into k groups of size ℓ each. At a high level, when panel size evenly divides population size, Expanding Approval Rule proceeds by simultaneously expanding balls around each representative in P . Once a ball captures ℓ individuals, these are grouped together, and the process continues with the remaining representatives until k such groups are formed. Aziz et al. [2024] show that when Metric Expanding Approval is applied to an underlying population, selecting one representative from each resulting group yields a panel that satisfies 2-PRF with respect to the population. Building on this, we assign to each individual panel P_i one representative drawn uniformly at random from each of the k groups. This guarantees that every panel receives exactly one representative from each group, while each representative is assigned to a panel with probability $1/\ell$. Combined with the fact that each individual in the population is selected into P with probability $\ell \cdot k/n$, it follows that every individual is assigned into a panel P_i with probability k/n . By Theorem 3.1, we then obtain the following corollary.

Corollary 3.2. *Given population V , there exists a polynomial time algorithm that returns ℓ panels, $P_1, \dots, P_\ell$ of size k each, such that the global panel, i.e. $\bigcup_{i \in [\ell]} P_i$, is 6-PRF, each panel P_i is 26-PRF, and each $v \in V$ is included in panel P_i with probability k/n .*

However, despite significant effort, we were not able to generalize this result for the case where the panels have different sizes, and the following question remains open.

Open Question 1: Does there exist a distribution over a sequence of panels, $P_1, \dots, P_\ell$, where P_i has size k_i such that each individual panel P_i satisfies $O(1)$ -PRF and the global panel $\bigcup_{i \in [t]} P_i$ satisfies $O(1)$ -PRF while also each individual is selected with the same probability in the process?

⁴Technically, Ebadian and Micha [2025] establish this approximation under a slightly different notion of representation; however, by adapting their arguments, the guarantee carries over to PRF as well.

4 Prefix and Panel Level Representation

In this section, we turn our attention to the setting where the goal is to ensure not only global representation but also *prefix representation*—that is, for every time step $t \in [\ell]$, we require that the cumulative panel $P_{\leq t} = \bigcup_{j=1}^t P_j$ provides appropriate representation. An ideal answer would be that there is a distribution over a sequence of panels, $P_1, \dots, P_\ell$, where P_j has size k (or more generally k_j), such that each individual panel P_j and each prefix panel $P_{\leq t}$ satisfy a constant-factor approximation of PRF or at least PFC.

Despite considerable effort, we are unable to prove whether such a distribution always exists, leaving a tantalizing question open.

Open Question 2: Does there exist a distribution over a sequence of panels, $P_1, \dots, P_\ell$ such that each individual panel P_j and each prefix panel $P_{\leq t}$ satisfies $O(1)$ -PRF or $O(1)$ -PFC?

Instead, here we present a novel algorithm that provides weaker approximation guarantees with respect to individual and prefix representation for the case where each panel has size k . In particular, we show that there exists a polynomial-time algorithm that achieves an approximation to PFC within each individual panel and each prefix panel, which grows exponentially with the number of panels ℓ , while also ensuring that each individual is selected to participate in one of the panels with equal probability.

Theorem 4.1. *There exists a polynomial time algorithm that returns a distribution over a sequence of ℓ panels $P_1, \dots, P_\ell$ of size k each such that:*

- *For each $t \in [\ell]$, panel P_t satisfies $O(4^\ell)$ -PFC;*
- *For each $t \in [\ell]$, prefix panel $P_{\leq t}$ satisfies $O(4^\ell)$ -PFC;*
- *Each individual is selected in $\bigcup_{t \in [\ell]} P_t$ with probability $\ell \cdot k/n$.*

At a high level, our main algorithm, `NESTEDBASEDREPRESENTATION`, operates in two phases. In the first phase, it constructs collections $\mathcal{G}^t$ of disjoint groups of individuals, each of size at least $n/t \cdot k$ for every $t \in [\ell]$. Each group in $\mathcal{G}^t$ represents a subset of the population that should be represented within the first t panels. The algorithm proceeds in reverse order, starting from $t = \ell$ and forming the collection of groups $\mathcal{G}^t$. As t decreases, it continues forming new groups while ensuring that previously constructed groups are preserved and hierarchically nested within the new groupings.

In the second phase, the algorithm traverses this hierarchy (visualized as a tree) starting from groups corresponding to smaller values of t (i.e., larger groups) and proceeding toward groups corresponding to larger values of t (i.e., smaller groups) that are nested within the previous ones. For each such path, the algorithm identifies the terminal group, i.e. a group corresponding to the larger value of t that does not have any other group nested and selects one individual from this terminal group as the path's representative. The algorithm then assigns these selected representatives to specific panels in sequential order, ensuring that the assignment respects both per-panel representation and prefix-level representation.

4.1 Modified Greedy Capture

Before describing our algorithm in detail, we introduce `MODIFIEDGREEDYCAPTURE`, a variation of Chen et al. [2019]'s algorithm that serves as a key subroutine. `MODIFIEDGREEDYCAPTURE` takes input population V , metric d , target panel size K , and a partition $\mathcal{G}$ of disjoint groups over V . The algorithm initializes an empty set $\mathcal{G}' = \emptyset$ to store the resulting groups.

The algorithm expands balls around each individual, initially marked incomplete. A ball captures a group in $\mathcal{G}$ when it encapsulates *all* group members. Crucially, partial group capture does not count: when a ball captures some but not all members of a group, neither the partially captured individuals nor the group itself are considered captured by the ball. Only when every member of a group lies within the ball's radius is the entire group deemed captured. When an incomplete ball captures at least n/K individuals, it becomes *complete*, its captured groups are consolidated as a single group and added to $\mathcal{G}'$, then disregarded from further processing. The algorithm continues expanding all balls, disregarding any groups in $\mathcal{G}$ that are captured by complete balls. New groups are formed when incomplete balls

capture total of at least n/K uncovered individuals. The process terminates when all groups are disregarded, returning $\mathcal{G}'$.

Note that by construction, each group $G \in \mathcal{G}$ is either entirely contained within some group $G' \in \mathcal{G}'$ or entirely excluded from $\mathcal{G}'$. For each group G created during the execution of **MODIFIEDGREEDYCAPTURE**, we denote by c_G its *center*, i.e., the point around which the ball was grown when the group was formed, and by r_G its *radius*, defined as the radius of the ball at the moment G was created.

For nested groups, we define the recursive function $I(\cdot)$ returning all individuals in group G and its descendants. Here, groups are understood as sets that may contain either individuals or other groups, forming a hierarchical structure:

$$I(G) = \begin{cases} \{G\} & \text{if } G \in V, \\ \bigcup_{H \in G} I(H) & \text{otherwise.} \end{cases}$$

Notice that for a single individual $v \in V$, $I(v)$ returns a singleton set containing v , and for hierarchical groups, it returns the set containing all individuals by recursively traversing the group and all its descendants.

Algorithm 1 **MODIFIEDGREEDYCAPTURE**

Input: Population V , Metric d , Parameter K , Groups $\mathcal{G}$

Output: Collection of groups $\mathcal{G}'$

```

1: Initialize  $\delta \leftarrow 0$ 
2: Let  $U \leftarrow \mathcal{G}$  be the collection of uncovered groups
3: Let  $\mathcal{G}' \leftarrow \emptyset$ .
4: Define  $B^{\mathcal{G}}(x, \delta) := \{G \in U : I(G) \subseteq B(x, \delta)\}$ 
5: while  $U \neq \emptyset$  do
6:   Increase  $\delta$  continuously
7:   for each group  $G \in \mathcal{G}'$  do
8:     Let  $S = B^{\mathcal{G}}(c_G, \delta) \cap U$  where  $c_G$  is the center of  $G$ .
9:     Disregard groups in  $S$  from the uncovered groups collection, i.e.,  $U \leftarrow U \setminus S$ .
10:  end for
11:  while there exists  $x \in V$  such that  $\left| \bigcup_{H \in B^{\mathcal{G}}(x, \delta) \cap U} I(H) \right| \geq \frac{n}{K}$  do
12:    Let  $N = B^{\mathcal{G}}(x, \delta) \cap U$ 
13:    Capture groups in  $N$ , i.e.,  $\mathcal{G}' \leftarrow \mathcal{G}' \cup \{N\}$ 
14:    Disregard groups in  $N$  from the uncovered groups collection, i.e.,  $U \leftarrow U \setminus N$ 
15:  end while
16: end while
17: return  $\mathcal{G}'$ 

```

4.2 Nested-Based Temporal Sortition Algorithm

At a high level, our main algorithm, **NESTEDBASEDREPRESENTATION**, operates in two phases. In the first phase, it constructs collections $\mathcal{G}^t$ of disjoint groups of individuals, each of size at least $n/t \cdot k$ for every $t \in [\ell]$. Each group in $\mathcal{G}^t$ represents a subset of the population that should be represented within the first t panels. The algorithm proceeds in reverse order, starting from $t = \ell$ and forming the collection of groups $\mathcal{G}^t$. As t decreases, it continues forming new groups while ensuring that previously constructed groups are preserved and hierarchically nested within the new groupings.

In the second phase, the algorithm traverses this hierarchy (visualized as a tree) starting from groups corresponding to smaller values of t (i.e., larger groups) and proceeding toward groups corresponding to larger values of t (i.e., smaller groups) that are nested within the previous ones. For each such path, the algorithm identifies the terminal group, i.e. a group corresponding to the larger value of t that does not have any other group nested and selects one individual from this terminal group as the path's representative. The algorithm then assigns these selected representatives to specific panels in sequential order, ensuring that the assignment respects both per-panel representation and prefix-level representation.

Below, we describe these two main phases separately in detail.




Figure 1: This figure demonstrates the hierarchical group structure produced by the first phase of the NESTEDBASEDREPRESENTATION algorithm through two side-by-side visualizations. The left panel shows 8 individuals positioned along a line: v_1, v_2 at location 1, v_3, v_4 at location 2, and v_5, v_6, v_7, v_8 at location 4. Given parameters $\ell = 2$ and $k = 2$, the NESTEDBASEDREPRESENTATION algorithm proceeds in two stages: it first invokes MODIFIEDGREEDYCAPTURE to construct groups $\mathcal{G}^2 = \{G_1^2, G_2^2, G_3^2\}$, then applies MODIFIEDGREEDYCAPTURE again to these formed groups and remaining individuals to create groups $\mathcal{G}^1 = \{G_1^1, G_2^1\}$. Notice that while constructing $\mathcal{G}^2$, MODIFIEDGREEDYCAPTURE disregards v_7 and v_8 without capturing them. The right panel depicts the resulting hierarchical tree structure, where internal nodes represent group formations and leaf nodes correspond to individual participants.

Tree Construction. NESTEDBASEDREPRESENTATION starts by constructing a tree rooted at $\mathcal{R}$, initially containing singleton groups $\{\{v\} : v \in V\}$ as children. Then, from $t = \ell$ down to 1, NESTEDBASEDREPRESENTATION calls MODIFIEDGREEDYCAPTURE, which at iteration t takes $t \cdot k$ as parameter K and uses the current children of $\mathcal{R}$ as its group partition, outputting a collection $\mathcal{G}^t$ in which each group contains at least $n/t \cdot k$ individuals. After each execution, NESTEDBASEDREPRESENTATION updates the tree structure as following. For each group $H \in \mathcal{G}^t$, all groups contained in H are removed from $\mathcal{R}$'s children, and H itself becomes a new child of $\mathcal{R}$. An example is illustrated at Figure 1. This tree structure ensures that if leaf node v is selected (also referring to an individual), any group on the path from $\mathcal{R}$ to v is approximately represented, as shown in the following lemma.

The tree $\mathcal{R}$ ensures that if leaf node v is selected (also referring to an individual), any group on the path from $\mathcal{R}$ to v is approximately represented, as shown in the following lemma.

Lemma 4.2. *Let P be a panel containing a representative from each group $G \in \mathcal{G}^j$. Then P satisfies $O(4^{\ell-j})$ -PFC for population V and panel size $j \cdot k$.*

Panel Assignment. The algorithm ensures representation for both individual panels and prefix panels by utilizing the constructed tree structure $\mathcal{R}$ and the above lemma by enforcing the following two properties:

- (a) Each individual panel contains one representative from each group $G \in \mathcal{G}^1$, ensuring approximate PFC representation for every panel.
- (b) For each group $G \in \mathcal{G}^t$, there exists a representative from G among the first t panels, ensuring approximate PFC representation for every prefix up to time t .

To establish these properties, we first introduce the FINDREPRESENTATIVE algorithm, which the algorithm employs in the second phase. FINDREPRESENTATIVE takes as input the tree $\mathcal{R}$ and a target group G , corresponding to a node in the tree. Starting from G , it recursively selects a subgroup contained in G that belongs to some $\mathcal{G}^j$ with the smallest possible index j , continuing this process until it reaches a group consisting only of leaf nodes. Intuitively, FINDREPRESENTATIVE traces a path from the target group down to a leaf node which we call *trajectory*, always moving through subgroups that are required to be represented earlier. Upon reaching the terminal node, it arbitrarily selects $n/(\ell \cdot k)$ leaves (equivalently saying individuals) as the sampling group Q . Notice that by construction any formed group either contains at least $n/(\ell \cdot k)$ individuals or another group, implying that such a trajectory always exists.

Algorithm 2 FINDREPRESENTATIVE

```
1: Input: Target group  $G$ 
2: Output: Modified group  $G'$  and sample group  $Q$ 
3: if  $G$  contains a group  $H \in \mathcal{G}^i$  for some  $i \leq \ell$  then
4:   Let  $H^* \in G$  be the group in  $\mathcal{G}^i$  with minimum index  $i$ .
5:   Execute FINDREPRESENTATIVE,  $H', Q \leftarrow \text{FINDREPRESENTATIVE}(H^*)$ .
6:   Define modified target group  $G' \leftarrow G \setminus \{H^*\} \cup H'$ .
7:   Return  $G', Q$ .
8: else
9:   Let  $Q \subseteq G$  be an arbitrarily selected subset of size  $\frac{n}{\ell \cdot k}$ .
10:  Return  $G \setminus Q, Q$ .
11: end if
```

Then, in two sub-phases, NESTEDBASEDREPRESENTATION enforces the two properties respectively, as follows.

Satisfying Property (a) - Panel Representation: This property requires that every group $G \in \mathcal{G}^1$ has a distinct representative in each of the ℓ panels. The algorithm addresses this requirement by processing groups individually through a sequential panel assignment procedure. For each group $G \in \mathcal{G}^1$, the algorithm iterates through panels $P_1, P_2, \dots, P_\ell$ in order, seeking to assign one representative from G to each panel. This systematic approach is feasible because group G initially contains at least n/k individuals, while each execution of FINDREPRESENTATIVE consumes only $n/(\ell \cdot k)$ individuals through the sampling group. Consequently, the algorithm can perform exactly ℓ iterations to populate all required panels.

However, a complication arises when G contains subgroups that demand earlier representation. Specifically, when processing panel P_t , the algorithm first examines whether any subgroup $H \in G$ belongs to $\mathcal{G}^j$ for some $j < t$. Such a subgroup might violate the prefix representation guarantee when it is represented in panel P_t , since it requires representation in the first j panels. To resolve this conflict, the algorithm extracts this problematic subgroup from G and relocates it as direct children of the root $\mathcal{R}$, setting it aside for representation during the prefix sub-phase. Crucially, H has not been represented yet and we are simply deferring its assignment to ensure it gets placed in an appropriately early panel.

The key is that removing H and its entire sub-tree still preserves sufficient population in G to support ℓ executions, as we establish in our analysis below. After removing problematic subgroups, all remaining groups in G lives in some $\mathcal{G}^i$ with $i \geq t$, allowing for the safe execution of FINDREPRESENTATIVE to obtain sampling group Q . The algorithm then samples an individual $v \in Q$, assigns v to panel P_t (where v represents all groups along execution trajectory of FINDREPRESENTATIVE), and updates the tree structure.

The tree update process permanently removes all sampled individuals in Q from future consideration, eliminates all intermediate groups along the trajectory except the root group G , and flattens the structure by having G directly contain the remaining children groups of these intermediate groups. After completing all ℓ panels for group G , the algorithm removes G from $\mathcal{R}$ and promotes any remaining children to become direct children of the root.

Satisfying Property (b) - Prefix Representation: Groups requiring prefix representation (i.e., those set aside in the previous step) plus any other children of $\mathcal{R}$, accumulate as children of $\mathcal{R}$. While the root contains groups ($\mathcal{R} \neq \emptyset$), the algorithm executes FINDREPRESENTATIVE starting from $\mathcal{R}$ to obtain sampling group Q . It samples an individual $v \in Q$ and assigns v to panel P_t where t is the minimum index such that $|P_t| < k$, and then it applies the same tree structure updates as before. As we show in the analysis, if a group on execution trajectory of FINDREPRESENTATIVE belongs to $\mathcal{G}^t$ for some minimum value t , then there must be available capacity in the first t panels to accommodate the selected representative. The algorithm terminates when $\mathcal{R} = \emptyset$, at which point all groups have received appropriate representation in accordance with their prefix requirements.

These two properties together with Lemma 4.2 imply representation guarantees of Theorem 4.1.

4.3 Proof of Theorem 4.1

In the remainder of this section, we prove Theorem 4.1.

Proof of Lemma 4.2. Fix an arbitrary index $j \in [\ell]$. Consider an arbitrary coalition S with $|S| \geq n/(j \cdot k)$, and suppose

Algorithm 3 NESTEDBASEDREPRESENTATION

```
1: Input:  $V, d$ , Groups  $\mathcal{G}$ , Parameter  $\ell$ , Panel size  $k$ .
2: Output: Panels  $P_1, \dots, P_\ell$ .
3: —Phase 1: Tree Formation—
4: Initialize tree  $\mathcal{R}$  with children  $V$ , i.e.,  $\mathcal{R} \leftarrow V$ 
5: for  $t = \ell$  down to 1 do
6:   Let  $\mathcal{G}^t \leftarrow \text{MODIFIEDGREEDYCAPTURE}(V, d, t \cdot k, \mathcal{R})$ 
7:   for  $H \in \mathcal{G}^t$  do
8:     Remove groups in  $H$  from the root  $\mathcal{R}$ 
9:     Add  $H$  as a child of the root  $\mathcal{R}$ 
10:  end for
11: end for
12: —Phase 2: Panel Assignment—
13: —Phase 2.1: Panel Representation—
14: for  $G \in \mathcal{G}^1$  do
15:   for  $t = 1$  to  $\ell$  do
16:    while there is a group  $H \in G \cap \mathcal{G}^i$  for some  $i < t$  do
17:       $G \leftarrow G \setminus \{H\}$ .
18:       $\mathcal{R} \leftarrow \mathcal{R} \cup \{H\}$ .
19:    end while
20:    Execute  $\text{FINDREPRESENTATIVE}$ ,  $G', Q \leftarrow \text{FINDREPRESENTATIVE}(G)$ .
21:    Sample  $v \in Q$  uniformly, assign to  $P_t$ 
22:    Update  $G \leftarrow G'$ 
23:  end for
24:  Remove  $G$  from  $\mathcal{R}$ , add its children to  $\mathcal{R}$ 
25: end for
26: —Phase 2.2: Prefix Representation—
27: while  $\mathcal{R} \neq \emptyset$  do
28:   Execute  $\text{FINDREPRESENTATIVE}$ ,  $\mathcal{R}', Q \leftarrow \text{FINDREPRESENTATIVE}(\mathcal{R})$ .
29:   Sample  $v \in Q$ , assign to  $P_t$  with min  $t$  s.t.  $|P_t| < k$ 
30:   Update the root  $\mathcal{R} \leftarrow \mathcal{R}'$ 
31: end while
32: return  $P_1, \dots, P_\ell$ 
```

there exists an individual $x \in S$ such that $d(x, y) \leq r$ for all $y \in S$ and some $r \geq 0$. We will prove that there exists a panel member $p \in P$ such that $\min_{y \in S} d(y, p) \leq r \cdot 2 \cdot 4^{\ell-j+1}$.

For each index $t \in \{j, j+1, \dots, \ell\}$, let $\mathcal{G}^t$ denote the collection of groups formed during the execution of $\text{MODIFIEDGREEDYCAPTURE}$ at iteration t in the $\text{NESTEDBASEDREPRESENTATION}$ algorithm. For any individual $v \in S$ and index t , define $G_v^t \in \mathcal{G}^t$ as the unique group in $\mathcal{G}^t$ that disregards individual v during the execution of $\text{MODIFIEDGREEDYCAPTURE}$. We say that individual v is *processed* by group G_v^t if v is captured while the ball centered at $c_{G_v^t}$ was still incomplete and included in G_v^t , or if v is captured after the ball has opened and is simply disregarded.

We prove by backward induction on t (from $t = \ell$ down to $t = j$) that for every individual $v \in S$, the following properties hold:

- (a) **Proximity:** $d(c_{G_v^t}, v) \leq r \cdot (4^{\ell-t+1} - 1)$
- (b) **Bounded radius:** $\max_{y \in I(G_v^t)} d(c_{G_v^t}, y) \leq r \cdot (4^{\ell-t+1} - 1)$

where c_G denotes the center of group G and $I(G)$ denotes the set of individuals in group G .

Base Case ($t = \ell$): In this case, $\text{MODIFIEDGREEDYCAPTURE}$ operates with parameter $K = \ell \cdot k$, so groups become complete when they contain at least $n/(\ell \cdot k)$ individuals. We claim that there exists an individual $v^* \in S$ such that both $d(c_{G_{v^*}^\ell}, v^*) \leq r$ and $\max_{y \in I(G_{v^*}^\ell)} d(c_{G_{v^*}^\ell}, y) \leq r$.

Suppose for contradiction that no such individual exists. Then when $\text{MODIFIEDGREEDYCAPTURE}$ reaches radius $\delta = r$, no member of S has been disregarded by any complete group. However, since $|S| \geq n/(j \cdot k) \geq n/(\ell \cdot k)$ and all

members of S lie within distance r of point x , the ball $B(x, r)$ contains at least $n/(\ell \cdot k)$ individuals from S . By the algorithm's design, this ball would become complete and capture some of these individuals, disregarding all others in S , which contradicts our assumption. Therefore, such a v^* exists.

For this v^* , properties (a) and (b) hold with bound $r = r \cdot (4^{\ell-\ell+1} - 1) = r \cdot (4^1 - 1) \leq 3r$. For any other individual $v \in S$, the triangle inequality yields:

$$d(c_{G_{v^*}^\ell}, v) \leq d(c_{G_{v^*}^\ell}, v^*) + d(v^*, x) + d(x, v) \leq r + r + r = 3r.$$

Since every individual $v \in S$ must be processed by some group (either $G_{v^*}^\ell$ or another group with appropriately bounded center and radius), the base case is established.

Inductive Step ($t < \ell$): Assume the inductive hypothesis holds for all indices $t' > t$. We prove the claim for index t . At this step, MODIFIEDGREEDYCAPTURE operates with parameter $K = t \cdot k$, so groups become complete when they contain at least $n/(t \cdot k)$ individuals.

Claim 4.3. *There exists an individual $v^* \in S$ such that:*

- $d(c_{G_{v^*}^t}, v^*) \leq r \cdot (1 + 2 \cdot (4^{\ell-t} - 1))$,
- $\max_{y \in I(G_{v^*}^t)} d(c_{G_{v^*}^t}, y) \leq r \cdot (1 + 2 \cdot (4^{\ell-t} - 1))$.

Proof of Claim. Suppose for a contradiction that no such individual exists. By the inductive hypothesis, for each $v \in S$, we have:

$$d(c_{G_v^{t+1}}, v) \leq r \cdot (4^{\ell-t} - 1).$$

During the execution of MODIFIEDGREEDYCAPTURE at index t , when $\delta = r \cdot (1 + 2 \cdot (4^{\ell-t} - 1))$, consider the ball $B(x, r)$. Recall that $S \subseteq B(x, r)$.

Now note that for any $v \in S$ and any $y \in I(G_v^{t+1})$:

$$\begin{aligned} d(x, y) &\leq d(x, v) + d(v, c_{G_v^{t+1}}) + d(c_{G_v^{t+1}}, y) \\ &\leq r + r \cdot (4^{\ell-t} - 1) + r \cdot (4^{\ell-t} - 1) \\ &= r \cdot (1 + 2 \cdot (4^{\ell-t} - 1)). \end{aligned}$$

Therefore, the ball $B(x, r \cdot (1 + 2 \cdot (4^{\ell-t} - 1)))$ contains all individuals from groups $\{G_v^{t+1} : v \in S\}$, which includes at least $|S| \geq n/(j \cdot k) \geq n/(t \cdot k)$ individuals. By the algorithm's design, this ball becomes complete, contradicting our assumption. Therefore, such a v^* exists. $\square$

For all other individuals $v \in S$, the triangle inequality ensures that properties (a) and (b) hold with the bound:

$$r \cdot (1 + 2 \cdot (4^{\ell-t} - 1) + 2 + 2 \cdot (4^{\ell-t} - 1)) = r \cdot (4^{\ell-t+1} - 1).$$

Applying the inductive claim at index j , we obtain that there exists a group $G \in \mathcal{G}^j$ such that:

$$\min_{v \in S} d(c_G, v) \leq r \cdot (4^{\ell-j+1} - 1) \leq r \cdot 4^{\ell-j+1}$$

Since panel P contains a representative from each group in $\mathcal{G}^j$, there exists $p \in P$ with $d(p, c_G) \leq \max_{y \in I(G)} d(c_G, y) \leq r \cdot 4^{\ell-j+1}$. By the triangle inequality:

$$\min_{v \in S} d(v, p) \leq \min_{v \in S} d(v, c_G) + d(c_G, p) \leq r \cdot 4^{\ell-j+1} + r \cdot 4^{\ell-j+1} = r \cdot 2 \cdot 4^{\ell-j+1}$$

This establishes $O(4^{\ell-j})$ -PFC. $\square$

Lemma 4.4. *FINDREPRESENTATIVE algorithm always terminates properly.*

Proof. We prove FINDREPRESENTATIVE terminates in both Phase 2 and Phase 3 of the NESTEDBASEDREPRESENTATION algorithm.

During Phase 2, the algorithm processes each group $G \in \mathcal{G}^1$ by executing FINDREPRESENTATIVE exactly ℓ times (once for each panel). Each execution removes a trajectory of groups and a sampling group of size $n/(\ell \cdot k)$ from G .

Importantly, while the contents of G change, the groups in $\bigcup_{i=2}^{\ell} \mathcal{G}^i$ maintain their structure through out the process, as they are only removed from the hierarchy, not modified internally.

First, we establish a baseline: when `FINDREPRESENTATIVE` is executed on any group whose content has not been modified, it must terminate. This follows because every recursive path eventually reaches a terminal group containing at least $n/(\ell \cdot k)$ individuals, which is sufficient to create the required sampling group.

Now suppose, for contradiction, that `FINDREPRESENTATIVE` fails to terminate when executed on group G during Phase 2. Since `FINDREPRESENTATIVE` would succeed on any unmodified subgroup (by our baseline observation), the failure can only occur if G contains only individuals (no subgroups) and $|I(G)| < n/(\ell \cdot k)$.

Let t^* be the first iteration where, after executing `FINDREPRESENTATIVE`, we have $|I(G_{t^*})| < (\ell - t^*) \cdot n/(\ell \cdot k)$. Here, G_{t^*} denotes the content of G after completing iteration t^* . Such a t^* must exist since we assumed $|I(G)| < n/(\ell \cdot k)$ eventually. By the minimality of t^* , after iteration $t^* - 1$ we have:

$$|I(G_{t^*-1})| \geq (\ell - t^* + 1) \frac{n}{\ell \cdot k}.$$

Since each execution of `FINDREPRESENTATIVE` removes exactly $n/(\ell \cdot k)$ individuals from G , and we know that $|I(G_{t^*-1})| - n/(\ell \cdot k) < (\ell - t^*)n/(\ell \cdot k)$, the deficit at iteration t^* must trigger the while-loop condition at line 16. This while-loop removes groups from G that belong to $\mathcal{G}^i$ for $i < t^*$. Let H_{t^*} be one such group removed during this while-loop.

Let us denote by G'_{t^*} the content of G at iteration t^* right after the while-loop completes but before calling `FINDREPRESENTATIVE`. This is the critical state we need to analyze.

Consider the trajectories $L_1, \dots, L_{t^*-1}$ obtained from executing `FINDREPRESENTATIVE` in iterations 1 through $t^* - 1$, with corresponding sampling groups $Q_1, \dots, Q_{t^*-1}$. For each trajectory L_i , let H_i be the group on L_i that belongs to $\mathcal{G}^{t_i}$ where t_i is maximal subject to $t_i \leq t^* - 1$.

We establish two crucial properties:

1. For any $i \in [t^* - 1]$, the group H_i cannot contain any group from $\mathcal{G}^t$ for $t_i \leq t \leq t^* - 1$. If it did, `FINDREPRESENTATIVE` would have recursed through that group, contradicting the maximality of t_i for H_i .
2. The groups $H_1, \dots, H_{t^*-1}, H_{t^*}$ are mutually disjoint. To see why, note that if $I(H) \cap I(H') \neq \emptyset$ for two groups in our collection, then by the hierarchical structure, either $H \subseteq H'$ or $H' \subseteq H$. But property (1) prevents any H_i from containing another H_j when $i, j \leq t^* - 1$. In addition, $H_{t^*} \in \mathcal{G}^i$ for $i < t^*$ ensures it cannot be contained in any H_j and H_{t^*} cannot contain any H_j 's as they appear in some other trajectories implying that H_{t^*} cannot be their ancestor in the hierarchy.

Property (1) has an important consequence: the individuals in $\bigcup_{i=1}^{t^*-1} (I(H_i) \setminus Q_i)$ cannot be removed during the while-loop at iteration t^* . This is because these individuals either belong to groups in $\mathcal{G}^t$ for $t \geq t^*$ (and thus don't satisfy the while-loop condition) or are direct members of G not belonging to any subgroup.

Now we can lower bound $|I(G'_{t^*})|$. Using property (2) and the fact that each H_i contains at least $\frac{n}{(t^*-1) \cdot k}$ individuals:

$$\begin{aligned} |I(G'_{t^*})| &\geq \left| \bigcup_{i=1}^{t^*-1} (I(H_i) \setminus Q_i) \right| \\ &= \sum_{i=1}^{t^*-1} |I(H_i) \setminus Q_i| \quad (\text{by disjointness of the } H_i\text{'s}) \\ &= \sum_{i=1}^{t^*-1} (|I(H_i)| - |Q_i|) \\ &\geq (t^* - 1) \cdot \frac{n}{(t^* - 1) \cdot k} - (t^* - 1) \cdot \frac{n}{\ell \cdot k} \\ &= \frac{n}{k} - \frac{(t^* - 1) \cdot n}{\ell \cdot k} \\ &= (\ell - t^* + 1) \cdot \frac{n}{\ell \cdot k} \end{aligned}$$

But this means G'_{t^*} contains at least $(\ell - t^* + 1) \cdot n/(\ell \cdot k)$ individuals. Therefore, `FINDREPRESENTATIVE` at this step can terminate properly and removes $n/(\ell \cdot k)$ individuals implying that $|I(G_{t^*})| \geq (\ell - t^*) \frac{n}{\ell \cdot k}$, contradicting our assumption.

During Phase 3, a similar counting argument shows that the root $\mathcal{R}$ maintains $|I(\mathcal{R})| \geq n/(\ell \cdot k)$ until it becomes empty. Moreover, whenever $\mathcal{R}$ contains groups (not just individuals), `FINDREPRESENTATIVE` can recurse through these groups and terminate successfully. This completes the proof. $\square$

Corollary 4.5. *For any panel P_i and any group $G \in \mathcal{G}^1$, we have $P_i \cap I(G) \neq \emptyset$.*

Proof. This corollary follows from Lemma 4.4, as `FINDREPRESENTATIVE` is executed ℓ times for each group $G \in \mathcal{G}^1$. $\square$

Lemma 4.6. *For any cumulative panel $P_{\leq i}$ and any group $G \in \mathcal{G}^i$, we have $P_{\leq i} \cap I(G) \neq \emptyset$.*

Proof. For any index t , if a group $G \in \mathcal{G}^t$ is removed from the tree during the second phase of `NESTEDBASEDREPRESENTATION`, by construction an individual $v \in I(G)$ is assigned to one of the first t panels.

Assume for contradiction that there exists a group $G \in \mathcal{G}^t$ not represented by the first t panels, and assume t is the minimum such index. When G is removed from the tree, G must belong to $\mathcal{R}$ and `FINDREPRESENTATIVE` will provide a sample group belonging to $I(G)$. `NESTEDBASEDREPRESENTATION` can only violate the representation guarantee if the first t panels are completed.

Note that at the end of the second phase, each panel contains exactly $|\mathcal{G}^1|$ individuals. Moreover, since the first t panels are full, there must be at least $t \cdot (k - |\mathcal{G}^1|) + 1$ groups belonging to $\mathcal{G}^t$ existing in the tree at the beginning of the third phase. Since $\mathcal{G}^t$ groups are disjoint and sample groups obtained by running `FINDREPRESENTATIVE` during the second phase are disjoint from these $\mathcal{G}^t$ groups, we obtain

$$|V| \geq \frac{n}{\ell \cdot k} \cdot \ell \cdot |\mathcal{G}^1| + (t \cdot (k - |\mathcal{G}^1|) + 1) \cdot \frac{n}{t \cdot k} \geq n + \frac{n}{t \cdot k} > n.$$

This yields a contradiction. $\square$

Lemma 4.7. *The output panels of `NESTEDBASEDREPRESENTATION` satisfy individual fairness.*

Proof. `FINDREPRESENTATIVE` is executed a total of $\ell \cdot k$ times, outputting a sample group Q of size $\frac{n}{\ell \cdot k}$ after each execution. Sample groups from these executions form a partition of V . Since the algorithm selects uniformly random individuals from each sample group, it ensures every individual appears in some panel with probability $\frac{\ell \cdot k}{n}$. $\square$

Proof of Theorem 4.1. Panel level representation guarantee follows from Lemmas 4.2 and 4.5. Prefix level representation follows from Lemmas 4.2 and 4.6. Individual fairness follows from Lemma 4.7. $\square$

5 Prefix Level Representation

In the previous section, we show that by nesting groups across panel sizes, we can ensure approximate PFC representation for both prefixes and individual panels. However, the approximation factor increases exponentially with the number of panels ℓ .

In this section, we adopt a different approach and show that a *constant*-factor approximation to PFC can be achieved with respect to every prefix panel $P_{\leq t}$, though at the cost of abandoning representation guarantees for individual panels. The key idea is to replace the nested-group construction with a distinct family of groups $\mathcal{G}^t$ for each prefix $t \in [\ell]$, capturing the groups that must be represented within the first t panels. Unlike the nested setting, this method does not require that for overlapping groups $G \in \mathcal{G}^t$ and $G' \in \mathcal{G}^{t'}$ with $t < t'$, the group from the larger prefix size (G') be fully contained in the group from the smaller prefix size (G), i.e., $G' \subseteq G$. Instead, the algorithm links groups across different prefix sizes into *chains*, i.e., sequences of the form $G^t \rightarrow G^{t+1} \rightarrow \dots \rightarrow G^\ell$, with $G^j \in \mathcal{G}^j$, where each group overlaps with some of its predecessor and has radius no larger than those that overlaps with. Representing the final group in such a chain then suffices to approximate the representation of all groups in the sequence, ensuring that coverage for the last group automatically yields coverage for every earlier one.

Note that simply linking overlapping groups with radii no larger than their predecessors is not enough, as the approximation error can accumulate at each step, leading to a bound that grows linearly with ℓ . To prevent this, the algorithm designates an anchor group as the guiding reference for building the chain, starting from a group $G^t \in \mathcal{G}^t$ with the minimum index t . A group is added to the chain only if it overlaps with the anchor and its radius is no larger than that of the anchor. The anchor is updated to the newly added group only when its radius is significantly smaller

than that of the current anchor (e.g., by at least a factor of two). Each update therefore corresponds to a substantial reduction in radius, and as a result, the approximation error remains constant instead of accumulating with the length ℓ of the chain.

Theorem 5.1. *There exists a polynomial time algorithm that returns a distribution over a sequence of ℓ panels $P_1, \dots, P_\ell$, where panel P_t has size k_t , such that:*

- *For each $t \in [\ell]$, prefix panel $P_{\leq t} = \cup_{t \in [t]} P_t$ satisfies $O(1)$ -PFC;*
- *Each individual is selected in $\cup_{t \in [\ell]} P_t$ with probability $\sum_{t=1}^{\ell} k_t/n$.*

Description of the Algorithm. Our algorithm, called CHAINBASEDREPRESENTATION consists of three main phases.

In the first phase, for each $t \in [\ell]$, the algorithm runs MODIFIEDGREEDYCAPTURE on the population V with metric d , panel size $\sum_{j=1}^t k_j$, and with each individual initially assigned to their own singleton group (i.e., there is no pre-existing grouping that must be nested). This yields a collection of groups $\mathcal{G}^t$, representing the groups that should be represented among the first t panels. For each group G created during the execution of MODIFIEDGREEDYCAPTURE, we denote by c_G its center, i.e., the point around which the ball was grown when the group was formed, and by r_G its radius, defined as the radius of the ball at the moment G was created.

In the subsequent phases, the algorithm aims to build panels $P_1, \dots, P_\ell$ so that every prefix $P_{\leq t}$ and for every group $G^t \in \mathcal{G}^t$, there should exist an individual among the first t panels who is no farther than a constant factor of the radius r_{G^t} from the center c_{G^t} . When this requirement is met for a group, it is marked as covered. Initially, all the groups are marked as uncovered. The main property that the algorithm utilizes for achieving this (proved in Lemma 5.2) is the following: for every group $G^t \in \mathcal{G}^t$ and every $j \geq t$, there exists a group $G^j \in \mathcal{G}^j$ such that (i) G^t and G^j overlap and (ii) the radius of G^j is at most equal to the radius of G^t . Essentially this means that by (approximately) representing group G^j , then G^t is also (approximately) represented.

By exploiting this property, the second phase proceeds as follows. Initially, all the groups are marked as uncovered. For each group G^t , the algorithm attempts to find a group in $\mathcal{G}^\ell$ such that once this group is represented, G^t is also approximately represented. More precisely, the algorithm starts from the set $\mathcal{G}^t$ with uncovered groups of smallest index and picks an arbitrary uncovered group G^t . It then calls the subroutine CONSTRUCTCHAIN, which tries to construct a chain of groups $G^t \rightarrow \dots \rightarrow G^\ell$, with each $G^j \in \mathcal{G}^j$, ensuring that selecting an individual from G^ℓ represents every group in the chain within a constant approximation. The subroutine begins by marking G^t as the *anchor* and uses it to guide the chain extension. For each $j \in \{t+1, \dots, \ell\}$, it finds a group $G^j \in \mathcal{G}^j$ that satisfies the conditions of Lemma 5.2 relative to the anchor. If such a group is found and is uncovered, it is added to the chain; furthermore, if the radius of G^j is at most half that of the current anchor, G^j is promoted as the new anchor. This process continues until either a group in $\mathcal{G}^\ell$ is reached, at which point the chain construction succeeds, or a covered group is encountered, in which case CONSTRUCTCHAIN returns an unsuccessful chain consisting only of G^t . CHAINBASEDREPRESENTATION marks all groups in the returned chain as covered and repeats the process, always by calling CONSTRUCTCHAIN on an uncovered group in $\mathcal{G}^j$ with the smallest index j . The phase ends once all groups are marked as covered.

In the third phase, the algorithm uses the previously constructed chains to generate a distribution over panels $P_1, \dots, P_\ell$. To ensure individual fairness, it begins by sampling $\sum_{t=1}^{\ell} k_t$ individuals from the population as following: one individual is selected uniformly at random from each group in $\mathcal{G}^\ell$, and the remaining $\sum_{t=1}^{\ell} k_t - |\mathcal{G}^\ell|$ individuals are sampled uniformly at random from the rest of the population⁵. Next, to ensure proportional representation, the algorithm assigns each sampled individual a priority label based on the chains created in the previous phase. In particular, if an individual is sampled from a group G^ℓ that belongs to some chain $G^t \rightarrow \dots \rightarrow G^\ell$, the individual is assigned the index t , corresponding to the head of the chain. If the individual is not sampled from any such group, they are assigned the priority label ℓ . Finally, the algorithm allocates individuals to panels sequentially, from P_1 to P_ℓ , in increasing order of their assigned priority labels. Intuitively, this gives higher priority to individuals representing groups that must be covered earlier, ensuring that such representatives appear sooner in the panel sequence.

Lemma 5.2. *Let V be a population in a metric space with distance function d , and consider panel sizes $k_1 \leq k_2$. Let $\mathcal{G}$ and $\mathcal{H}$ be the outputs of MODIFIEDGREEDYCAPTURE with panel sizes k_1 and k_2 , respectively. For every group $G \in \mathcal{G}$ with center c_G and radius r_G , there exists a group $H \in \mathcal{H}$ with center c_H and radius r_H such that $d(c_H, c_G) \leq 2 \cdot r_G$ and $r_H \leq r_G$.*

⁵When n is divisible by $\sum_{t=1}^{\ell} k_t$, this can be achieved by applying standard dependent rounding techniques. Otherwise, we first create groups of size exactly $n/\sum_{t=1}^{\ell} k_t$ by fractional assignment, and then apply Birkhoff's decomposition, as in Ebadian and Micha [2025].

Algorithm 4 CHAINBASEDREPRESENTATION

```
1: Input:  $V, d$ , Panel sizes  $k_1, \dots, k_\ell$ 
2: Output: A distribution over a sequence of panels  $P_1, \dots, P_\ell$ 
3: — Phase 1: Construct Groups —
4: for level  $t = 1$  to  $\ell$  do
5:    $\mathcal{G}^t \leftarrow \text{MODIFIEDGREEDYCAPTURE}(V, d, \sum_{j=1}^t k_j, \cup_{v \in V} \{v\})$ 
6: end for
7: — Phase 2: Build Chains and Assign Priorities —
8: for  $t = 1$  to  $\ell - 1$  do
9:   for each uncovered  $G^t \in \mathcal{G}^i$  do
10:     $(\text{status}, T) \leftarrow \text{CONSTRUCTCHAIN}(G^t, \{\mathcal{G}^1, \dots, \mathcal{G}^\ell\})$ 
11:    Mark all groups in chain  $T$  as covered
12:    if status = succeed then
13:      Assign to the last group priority number  $t$ 
14:    end if
15:  end for
16: end for
17: — Phase 3: Sampling and Assignment to Panels —
18: for each  $G^\ell \in \mathcal{G}^\ell$  do
19:   Sample  $v$  from  $G^\ell$  uniformly at random.
20:   Assign  $u$  the priority label of its group  $G^\ell$ .
21: end for
22: Sample  $\sum_{t=1}^\ell k_j - |\mathcal{G}^\ell|$  representatives uniformly at random from  $V \setminus \bigcup_{G^\ell \in \mathcal{G}^\ell} G^\ell$ , and assign each a priority label of  $\ell$ .
23: Assign representatives to panels by ordering them in increasing order of priority labels, and placing each into the earliest panel  $P_t$  with available capacity (i.e., the smallest such index  $t$ ).
24: return  $P_1, \dots, P_\ell$ 
```

Algorithm 5 CONSTRUCTCHAIN

```
1: Input: Target group  $G^t \in \mathcal{G}^t, \{\mathcal{G}^1, \dots, \mathcal{G}^\ell\}, V, d$ 
2: Output: status (succeed when a new chain was constructed and failed otherwise) and chain  $T$ 
3: Initialize the chain and anchor:  $T \leftarrow G^t$  and  $H \leftarrow G^t$ 
4: for  $j = t + 1$  to  $\ell$  do
5:   Let  $G^j \in \mathcal{G}^j$  be such that  $d(c_H, c_{G^j}) \leq 2 \cdot r_H$  and  $r_{G^j} \leq r_H$ 
6:   if  $G^j$  is already covered then
7:     Remove from  $T$  any group except for  $G^t$ 
8:     return fail,  $T$ 
9:   end if
10:  Add to  $G^j$  at the end of chain  $T$ 
11:  if  $r_{G^j} < r_H/2$  (radius shrinks significantly) then
12:    Update anchor:  $H \leftarrow G^j$ 
13:  end if
14: end for
15: return succeed,  $T$ 
```

Proof. Consider any group $G \in \mathcal{G}$ formed during the execution of MODIFIEDGREEDYCAPTURE with panel size k_1 . When group G is created, the radius threshold has reached r_G , and G contains at least n/k_1 individuals that were previously not disregarded. Since $k_1 \leq k_2$, we have $n/k_1 \geq n/k_2$, which implies $|G| \geq n/k_2$.

Consider the execution of MODIFIEDGREEDYCAPTURE with panel size k_2 . Let δ denote the radius threshold at any point during this execution. We claim that when $\delta = r_G$, at least one individual in G must be disregarded by some group in $\mathcal{H}$. If all individuals in G remain available at this point, then the ball centered at c_G would capture all $|G| \geq \frac{n}{k_2}$ individuals and become complete, forming a new group in $\mathcal{H}$.

Let $v \in G$ be an individual that is disregarded by some group $H \in \mathcal{H}$ when $\delta = r_G$. Since H captures v at radius threshold r_G , we have $r_H \leq r_G$. Moreover, by the triangle inequality:

$$d(c_H, c_G) \leq d(c_H, v) + d(v, c_G) \leq r_H + r_G \leq r_G + r_G = 2 \cdot r_G.$$

This proves the lemma. □

Before proving Theorem 5.1, we first prove two more auxiliary lemmas.

Lemma 5.3. *For every group $G^t \in \mathcal{G}^t$, there exists a representative v among the first t panels, such that $d(c_{G^t}, v) \leq 16 \cdot r_{G^t}$.*

Proof. We start by showing that whenever a group $G^\ell \in \mathcal{G}^\ell$ is assigned a priority number i , some representative from this group is included in one of the first i panels. To establish this, we show that for each t , at most $\sum_{j=1}^t k_j$ groups can have priority number at most t . This implies that for every t , there are at most $\sum_{j=1}^t k_j$ representatives that must be assigned to the first t panels—exactly matching the number of available seats across these panels. Hence, in the final phase, the algorithm ensures that the representative from every group in $\mathcal{G}^\ell$ is assigned to a panel no later than its priority number. We prove this by induction on t . For $t = 1$, there are at most $|\mathcal{G}^1| \leq k_1$ groups with priority number 1, satisfying the bound. Suppose the claim holds for $t - 1$. Then, at step t , there are at most $|\mathcal{G}^t| \leq k_t$ additional groups that can be assigned priority number t . By the inductive hypothesis, at most $\sum_{j=1}^{t-1} k_j$ groups have priority number at most $t - 1$. Therefore, the total number of groups with priority number at most t is at most $\sum_{j=1}^{t-1} k_j + k_t = \sum_{j=1}^t k_j$, as required.

It remains to show that for each $G^t \in \mathcal{G}^t$, with $t \in [\ell]$, there exists a representative that is at most $16 \cdot r_{G^t}$ far away from its center and is assigned a priority label at most t . Note that, during the second phase, the algorithm marks G^t as covered after either including it in a chain of groups $G^j \rightarrow \dots \rightarrow G^t \rightarrow \dots \rightarrow G^\ell$ or by marked it as covered when G^t is the head of the chain and the algorithm fails to construct a new chain. We analyze the two cases separately.

Case 1: G^t is included in a chain. When G^t is first added to a chain with current anchor H , the algorithm checks whether $r_{G^t} \leq r_H/2$. If so, G^t becomes the new anchor. Otherwise, the anchor remains H . In either case, we denote the resulting anchor by H_1 . The algorithm continues to add groups to the chain by performing a shrinking test on the radius. Each time the anchor changes, the radius decreases by at least a factor of 2. Let $H_1, H_2, \dots, H_m$ denote the sequence of successive anchors during the execution of the algorithm after adding G^t to the chain. Then, we have

$$\begin{aligned} d(c_{G^t}, c_{G^\ell}) &\leq d(c_{G^t}, c_{H_1}) + \sum_{i=1}^{\max\{1, m-1\}} d(c_{H_i}, c_{H_{i+1}}) + d(c_{H_m}, c_\ell) \\ &\leq 2 \cdot r_{G^t} + \left(\sum_{i=1}^{\max\{1, m-1\}} 2 \cdot r_{H_i} \right) + 2 \cdot r_{H_m} \\ &\leq 2 \cdot r_{G^t} + 2 \cdot r_{H_1} \cdot \left(\sum_{i=2}^{\max\{1, m-1\}} \frac{1}{2^i} \right) + 2 \cdot r_{H_1} \\ &\leq 2 \cdot r_{G^t} + 2 \cdot r_{G^t} \cdot 2 + 2 \cdot r_{G^t} \leq 8 \cdot r_{G^t}, \end{aligned} \tag{1}$$

where the first inequality follows by the triangle inequality, the second inequality follows by Lemma 5.2, the third inequality follows since the radius shrinks by at least a factor of 2 each time the anchor changes and the fourth inequality follows again by Lemma 5.2. Moreover, if G^ℓ is the end point of the chain, then G^ℓ is assigned a priority number that is at most t , as the head of the chain belongs to $\mathcal{G}^j$ with $j \leq t$. This means that when a representative from G^ℓ is chosen is assigned a priority number of at most t .

Case 2: G^t is not included in a chain. When the algorithm fails to construct the chain, this means that during the process, the algorithm failed to find a group from some set $\mathcal{G}^j$ with $j > t$ that is uncovered and satisfies the above conditions. This means that there exists a group $G^j \in \mathcal{G}^j$ that has been assigned to a previous chain with end point

some $G^\ell \in G^\ell$ and from the above case we know that $d(c_{G^j}, c_{G^\ell}) \leq 8 \cdot r_{G^j}$. Moreover with similar arguments as above, we can conclude that $d(c_{G^t}, c_{G^j}) \leq 8 \cdot r_{G^t}$ and $r_{G^j} \leq r_{G^t}$ implying that $d(c_{G^t}, c_{G^\ell}) \leq 16 \cdot r_{G^t}$. Lastly, G^ℓ is assigned a priority number at most t , because the algorithm initiates chains from groups in $\mathcal{G}^r$ with the smallest index r that still contain uncovered groups, and at the time this chain was started, G^t was still uncovered. $\square$

Lemma 5.4. *Let V be a population and k be a panel size. Consider the collection of groups $\mathcal{G}$ formed by MODIFIEDGREEDYCAPTURE over V with panel size k and each individual initially assigned to their own group. If $P \subseteq V$ is a panel such that, for every $G \in \mathcal{G}$, there exists $p \in P$ with $d(p, c_G) \leq \alpha \cdot r_G$, then P satisfies $(\alpha + 3)$ -PFC.*

Proof. Let $S \subseteq V$ be an arbitrary set of individuals with $|S| \geq n/k$. Let $y^* \in \operatorname{argmin}_{y \in V} \max_{u \in S} d(u, y)$ be the point that minimizes the maximum distance to any point in S , and define $r = \max_{u \in S} d(u, y^*)$ as this optimal radius. Let $v^* \in S$ be the farthest point from y^* , so $d(y^*, v^*) = r$.

We first claim there exists a group $G \in \mathcal{G}$ with $r_G \leq r$ and some $v \in S$ satisfying $d(v, c_G) \leq r_G$. Indeed, if no such group existed, then when MODIFIEDGREEDYCAPTURE reaches radius r , all points in S would not be disregarded yet. Since $|S| \geq n/k$, the greedy algorithm would form a group containing some point from S at a radius of at most r , leading to a contradiction.

For this group G with witness $v \in S$, we have

$$d(y^*, c_G) \leq d(y^*, v) + d(v, c_G) \leq r + r_G \leq 2r,$$

since $d(y^*, v) \leq r$ by the definition of y^* . Thus

$$d(v^*, c_G) \leq d(v^*, y^*) + d(y^*, c_G) \leq r + 2 \cdot r = 3 \cdot r.$$

By assumption, there exists $p \in P$ with $d(c_G, p) \leq \alpha \cdot r_G \leq \alpha \cdot r$. Therefore,

$$d(v^*, p) \leq d(v^*, c_G) + d(c_G, p) \leq 3 \cdot r + \alpha \cdot r = (\alpha + 3) \cdot r.$$

Since $r = \min_{y \in V} \max_{u \in S} d(u, y)$, we have shown that for $v^* \in S$ and $p \in P$, we have $d(v^*, p) \leq (\alpha + 3) \cdot r$. As S was arbitrary, P satisfies $(\alpha + 3)$ -PFC. $\square$

Proof of Theorem 5.1. We establish both properties of the theorem. For the first property, we see that each prefix panel $P_{\leq t}$ satisfies 19-PFC by combining Lemma 5.3 and Lemma 5.4.

For the second property, we analyze the selection probability for each individual. If an individual appears in a group $G \in \mathcal{G}^\ell$, then it appears in some panel with probability $\sum_{t=1}^\ell k_t/n$ since each such group $\mathcal{G}^\ell$ contains exactly $n/\sum_{t=1}^\ell k_t$ individuals, and one of them is chosen uniformly at random. For individuals not covered by groups in $\mathcal{G}^\ell$, the number of such individuals equals $n/\sum_{t=1}^\ell k_t$ times the number of empty panel seats, and they are sampled with equal probability to fill the remaining positions. Hence, each individual appears in some panel with probability $\sum_{t=1}^\ell k_t/n$.

The polynomial time complexity follows from the fact that the MODIFIEDGREEDYCAPTURE algorithm runs in polynomial time [Chen et al., 2019], and both the CONSTRUCTCHAIN and CHAINBASEDREPRESENTATION algorithms have predetermined loop sizes that are polynomial in the input size. $\square$

It remains an open question whether one can design a distribution over a sequence of panels such that every prefix satisfies a constant-factor approximation of the proportionality notion PRF, which, in contrast to PFC, requires that each subset of individuals receive not just one representative, but a number of representatives proportional to its size.

Open Question 3: Does there exist a distribution over a sequence of panels, $P_1, \dots, P_\ell$ such that each prefix panel $P_{\leq t}$ satisfies $O(1)$ -PRF?

6 Conclusion

Permanent citizens' assemblies represent a promising democratic innovation by ensuring that minority voices can persist over time, even when too small to warrant representation in any single panel. We formalize this challenge as a problem of *temporal sortition* in metric spaces and make progress on it by proposing three algorithms, each providing different guarantees for panel-level and prefix-level representation.

Beyond the questions highlighted above, several intriguing directions remain open. For instance, can we ensure representation not just for prefix panels, but for *any* consecutive subsequence of panels? Additionally, is it possible to achieve meaningful guarantees without knowing in advance the total number or sizes of panels?

References

- Alberto Alemanno. Towards a permanent citizens’ participatory mechanism in the eu. Technical Report PE 735.927, European Parliament, Policy Department for Citizens’ Rights and Constitutional Affairs, 2022. URL [https://www.europarl.europa.eu/RegData/etudes/STUD/2022/735927/IPOL_STU\(2022\)735927_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2022/735927/IPOL_STU(2022)735927_EN.pdf).
- Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2):461–485, 2017.
- Haris Aziz, Barton E. Lee, Sean Morota Chu, and Jeremy Vollen. Proportionally representative clustering. In *Proceedings of the 20th Conference on Web and Internet Economics (WINE)*, 2024.
- Carmel Baharav and Bailey Flanigan. Fair, manipulation-robust, and transparent sortition. In *Proceedings of the 25th ACM Conference on Economics and Computation*, pages 756–775, 2024.
- Evripidis Bampis, Bruno Escoffier, and Sasa Mladenovic. Fair resource allocation over time. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, page 766–773, Richland, SC, 2018.
- Laurent Bulteau, Noam Hazon, Rutvik Page, Ariel Rosenfeld, and Nimrod Talmon. Justified representation for perpetual voting. *IEEE Access*, 9:96598–96612, 2021. doi: 10.1109/ACCESS.2021.3095087.
- Ioannis Caragiannis and Shivika Narang. Repeatedly matching items to agents fairly and efficiently. In *Proceedings of the 16th International Symposium on Algorithmic Game Theory (SAGT)*, page 347–364, Berlin, Heidelberg, 2023. doi: 10.1007/978-3-031-43254-5_20. URL https://doi.org/10.1007/978-3-031-43254-5_20.
- Ioannis Caragiannis, Evi Micha, and Jannik Peters. Can a few decide for many? the metric distortion of sortition. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, volume 235, pages 5660–5679, 2024.
- Xingyu Chen, Brandon Fain, Liang Lyu, and Kamesh Munagala. Proportionally fair clustering. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pages 1032–1041, 2019.
- Tamal K. Dey, Alfred Rossi, and Anastasios Sidiropoulos. Temporal Clustering. In *Proceedings of the 25th Annual European Symposium on Algorithms (ESA)*, volume 87, pages 34:1–34:14, Dagstuhl, Germany, 2017a. doi: 10.4230/LIPIcs.ESA.2017.34. URL <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ESA.2017.34>.
- Tamal K. Dey, Alfred Rossi, and Anastasios Sidiropoulos. Temporal Hierarchical Clustering. In *Proceedings of the 25th Annual European Symposium on Algorithms (ESA)*, volume 92, pages 28:1–28:12, Dagstuhl, Germany, 2017b. doi: 10.4230/LIPIcs.ISAAC.2017.28. URL <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ISAAC.2017.28>.
- Virginie Do, Matthieu Hervouin, Jérôme Lang, and Piotr Skowron. Online approval committee elections. In *Proceedings of the 31th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 251–257, 7 2022. doi: 10.24963/ijcai.2022/36. URL <https://doi.org/10.24963/ijcai.2022/36>. Main Track.
- S. Ebadian and E. Micha. Boosting sortition via proportional representation. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 667–675, 2025.
- S. Ebadian, G. Kehne, E. Micha, A. D. Procaccia, and N. Shah. Is sortition both representative and fair? In *Proceedings of the 36th Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 25720–25731, 2022.

- Edith Elkind, Svetlana Obraztsova, and Nicholas Teh. Temporal fairness in multiwinner voting. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, 2024. doi: 10.1609/aaai.v38i20.30273. URL <https://doi.org/10.1609/aaai.v38i20.30273>.
- Edith Elkind, Svetlana Obraztsova, Jannik Peters, and Nicholas Teh. Verifying proportionality in temporal voting. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(13):13805–13813, Apr. 2025. doi: 10.1609/aaai.v39i13.33509. URL <https://ojs.aaai.org/index.php/AAAI/article/view/33509>.
- B. Flanigan, P. Gözl, A. Gupta, and A. D. Procaccia. Neutralizing self-selection bias in sampling for sortition. In *Proceedings of the 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 6528–6539, 2020.
- B. Flanigan, P. Gözl, A. Gupta, B. Hennig, and A. D. Procaccia. Fair algorithms for selecting citizens’ assemblies. *Nature*, 596:548–552, 2021.
- Bailey Flanigan, Jennifer Liang, Ariel D Procaccia, and Sven Wang. Manipulation-robust selection of citizens’ assemblies. In *Proceedings of the aai conference on artificial intelligence*, volume 38, pages 9696–9703, 2024.
- Sreenivas Gollapudi, Kostas Kollias, and Benjamin Plaut. Almost envy-free repeated matching in two-sided markets. In Xujin Chen, Nikolai Gravin, Martin Hoefer, and Ruta Mehta, editors, *Web and Internet Economics*, pages 3–16, Cham, 2020.
- Daniel Halpern, Ariel D Procaccia, Ehud Shapiro, and Nimrod Talmon. Federated assemblies. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI)*, volume 39, pages 13897–13904, 2025.
- Ayumi Igarashi, Martin Lackner, Oliviero Nardi, and Arianna Novaro. Repeated fair allocation of indivisible items. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, 2024. doi: 10.1609/aaai.v38i9.28837. URL <https://doi.org/10.1609/aaai.v38i9.28837>.
- Christopher Jung, Sampath Kannan, and Neil Lutz. Service in Your Neighborhood: Fairness in Center Location. In Aaron Roth, editor, *1st Symposium on Foundations of Responsible Computing (FORC 2020)*, volume 156, pages 5:1–5:15, Dagstuhl, Germany, 2020. doi: 10.4230/LIPIcs.FORC.2020.5. URL <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.FORC.2020.5>.
- Yusuf Kalayci, David Kempe, and Vikram Kher. Proportional representation in metric spaces and low-distortion committee selection. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI)*, volume 38, pages 9815–9823, 2024.
- Leon Kellerhals and Jannik Peters. Proportional fairness in clustering: a social choice perspective. In *Proceedings of the 38th Annual Conference on Neural Information Processing Systems (NeurIPS)*, NIPS ’24, Red Hook, NY, USA, 2025.
- Martin Lackner. Perpetual voting: Fairness in long-term decision making. *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 34(02):2103–2110, Apr. 2020. doi: 10.1609/aaai.v34i02.5584. URL <https://ojs.aaai.org/index.php/AAAI/article/view/5584>.
- Martin Lackner and Jan Maly. Proportional decisions in perpetual voting. *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 37(5):5722–5729, Jun. 2023. doi: 10.1609/aaai.v37i5.25710. URL <https://ojs.aaai.org/index.php/AAAI/article/view/25710>.
- Martin Lackner and Piotr Skowron. Approval-based committee voting. In *Multi-Winner Voting with Approval Preferences*, pages 1–7. Springer, 2022.
- E. Micha and N. Shah. Proportionally fair clustering revisited. In *Proceedings of the 47th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 85:1–85:16, 2020.
- Christoph Niessen and Min Reuchamps. Designing a permanent deliberative citizens’ assembly: The ostbelgien modell in belgium. 2019.
- Peter Stone. *The Luck of the Draw: The Role of Lotteries in Decision Making*. Oxford University Press, 2011.

Yohai Trabelsi, Abhijin Adiga, Sarit Kraus, S. S. Ravi, and Daniel J. Rosenkrantz. Resource sharing through multi-round matchings. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 2023. doi: 10.1609/aaai.v37i10.26380. URL <https://doi.org/10.1609/aaai.v37i10.26380>.

David Van Reybrouck. *Against Elections: The Case for Democracy*. The Bodley Head / Random House, 2016.