

A Few Good Choices

Haoyu Song*

Thành Nguyen †

Young-San Lin‡

Abstract. *Condorcet winning sets* address the Condorcet paradox by selecting a small set of candidates—rather than a single winner—such that a majority prefers no unselected alternative over all members of the set. This notion extends to α -undominated sets, which require the same property to hold for any α -fraction of voters. Such sets are guaranteed to exist with constant size for any fixed α . However, the requirement that an outsider be preferred to *every* member of the set can be overly restrictive and difficult to justify in many applications. Motivated by this, we introduce a more flexible notion: (t, α) -undominated sets. Here, each voter compares an outsider to their t -th most preferred member of the set, and the set is undominated if no outsider is preferred by more than an α -fraction of voters. This framework subsumes prior definitions, recovering Condorcet winning sets when $(t = 1, \alpha = 1/2)$ and α -undominated sets when $t = 1$, and introduces a new, tunable notion of collective acceptability for $t > 1$.

We establish three main results. First, for all values of t and α , there exists a (t, α) -undominated set of size $O(t/\alpha)$. Second, as t grows large, the minimum size of such sets approaches t/α , establishing the asymptotic optimality of our bound. Finally, in the special case $t = 1$, we improve the upper bound on the size of α -undominated sets given by Charikar, Lassota, Pamakrishnan, Vetta, and Wang (STOC 2025); in particular, we show that a Condorcet winning set of five candidates always exists, improving their bound of six.

1 Introduction Selecting a fair outcome from diverse preferences is a central problem in social choice, with applications in voting, public goods provision, fairness in machine learning, and AI alignment. Yet, the generality of the social choice framework often leads to strong impossibility results, motivating the search for compromise or approximate solutions. A common response is to select multiple candidates instead of a single winner to accommodate conflicting preferences and ensure broader acceptability. For example, in contexts like prize distribution, awarding a group of candidates may better reflect the preferences of the judging panel when there is no clear front-runner. Similarly, in committee selection, slightly increasing the committee size can lead to more effective and representative outcomes. Even when only one candidate must ultimately be chosen, presenting a shortlist for further evaluation is common and can improve the decision-making process.

The idea of selecting more than one candidate has deep roots in voting theory, motivated in large part by impossibility results that emerge when restricting attention to a single winner. A prominent example is the Condorcet paradox (Condorcet, 1785), which illustrates that majority preferences can be cyclical, preventing any candidate from defeating all others in pairwise comparisons. More strikingly, generalized versions reveal that in some voting instances, no single candidate can even secure support from a positive fraction of voters (Charikar et al., 2025a). Despite these challenges, producing a voting outcome that garners support from a sufficiently large fraction of the electorate remains a central objective. One approach to achieving this is to relax the single-winner constraint and instead allow for the selection of multiple candidates.

One such relaxation extends the notion of a Condorcet winner to a Condorcet winning set (committee), as introduced by Elkind et al. (2015). A Condorcet winning committee is a set of candidates such that no candidate outside the set is preferred to *all* members of the set by a majority of voters. Identifying the minimal number k for which a Condorcet winning committee of size k always exists in all elections remains an open problem. The current best known upper bound is 6 established by Charikar et al. (2025a), while the lower bound is 3.

The Condorcet winning set, as described above, is not the only way to define outcomes with majority support. It boils down to how the comparison between an alternative and a set of alternatives are defined for a group of

*Department of Computer Science, Purdue University, song522@purdue.edu

†Daniels School of Business, Purdue University, nguye161@purdue.edu

‡Melbourne Business School, y.lin@mbs.edu

voters with heterogeneous preferences. For example, [McGarvey \(1953\)](#) studies an alternative definition, requiring that for every candidate outside the selected set, there exists at least one candidate within the set that gains majority support over the external candidate. This is a stronger requirement, as it demands that all voters compare the outside candidate with a single member of the committee. However, [McGarvey \(1953\)](#) shows that such a committee of constant size does not necessarily exist. [Charikar et al. \(2025b\)](#); [Bourneuf et al. \(2025\)](#) show that if the majority condition is relaxed to $\frac{1}{2} - \epsilon$ for any positive ϵ , a constant-size committee will always exist. However, a drawback of this stronger requirement is that it typically leads to very large bounds on the committee size for small ϵ , limiting practical applicability.¹

A more serious limitation of both concepts discussed above is their reliance on the assumption that each voter compares the outside alternative with only a single member of the selected committee. This assumption becomes problematic in the context of large committees. For example, if an alternative is excluded from a committee of 20 members, it is not well justified to attribute this outcome to a preference for just one selected candidate. More plausibly, the majority supports multiple committee members over the outside alternative. Ignoring this aggregated preference risks misrepresenting the collective rationale behind the committee's composition. A particularly relevant application is the selection of papers for a conference or research proposals for funding, where decisions are based on diverse and sometimes conflicting expert evaluations. In these settings, the selected set of candidates corresponds to the accepted papers or funded projects. To justify the rejection of a given submission, it is not sufficient to argue that it is weaker than the single best entry. Rather, the decision should be evaluated in the context of the entire selected set. For example, one may want to assess whether the rejected paper had less support than the median accepted submission, or little support compared with those in the top quartile.

Our Contributions The goal of this paper is to introduce a general framework for analyzing Condorcet-style winning sets—one that allows voters to compare a single alternative against a set of selected candidates. This framework involves two key generalizations. First, it extends the notion of how a voter may prefer a committee over an alternative outside the committee. Second, it generalizes the concept of a *majority* to any fixed fraction of the electorate. Specifically, we define a (t, α) -undominated set (committee) as a subset C of candidates with $|C| \geq t$, such that for every candidate $a \notin C$, at least a $1 - \alpha$ fraction of voters prefer at least t members of C over a . The notion of (t, α) -undominance captures both key generalizations. The central question we study is:

What is the minimal size of a (t, α) -undominated set that exists in all voting instances?

A special case arises when $t = 1$, in which case a $(1, \alpha)$ -undominated set is simply referred to as an α -undominated set, introduced in [Elkind et al. \(2015, 2011\)](#). When $t = 1$ and $\alpha = 1/2$, the concept coincides with the classical Condorcet winning set. To the best of our knowledge, the case of $t > 1$ has not been previously explored.

In this paper, we establish three main results. We begin with the case $t = 1$ and prove the following theorem in Section 3.

THEOREM 1.1. *Given $k \in \mathbb{N}$ and $\beta \in (0, 1)$, a $(\beta + (1 - \beta)^k)$ -undominated set of size k exists.*

This theorem improves the undominance ratio in [Charikar et al. \(2025a\)](#) for all k , establishes the existence of a Condorcet winning set of size 5, and tightens the bound of 6 in that paper.

We study the case $t \geq 2$ in Section 4 and prove the following result.

THEOREM 1.2. *Given $\alpha \in (0, 1]$, there exists a set of size $\lceil \delta(t) \cdot \frac{t}{\alpha} \rceil$ which is (t, α) -undominated. Here, $1 < \delta(t) \leq 4.75$ for all integers $t \geq 2$ and $\delta(t) \rightarrow 1$ as $t \rightarrow \infty$.*

Since the exact formulation of $\delta(t)$ is complicated, we choose to instead compute the values numerically. For instance, we find that $\delta(2) = 4.75$, $\delta(3) = 4.11$, and $\delta(4) = 3.69$. Overall, the result implies that selecting $O(t/\alpha)$ candidates suffices to guarantee a (t, α) -undominated committee.

We also establish the following lower bound. Combined with [Theorem 1.2](#), this yields an asymptotically tight bound of $\frac{t}{\alpha}$ for the minimum committee size.

THEOREM 1.3. *If $k \geq \frac{t+1}{\alpha} - 1$, then there exist elections that admit no (t, α) -undominated set of size k .*

¹For instance, with $\epsilon = 0.01$, [Bourneuf et al. \(2025\)](#) gives a bound of 5,227,032, while [Charikar et al. \(2025b\)](#) yields a bound of 3,927.

This asymptotic bound provides an interpretation of fairness in group selection and highlights a sharp trade-off. Specifically, a candidate can be legitimately excluded only if fewer than an α -fraction of voters rank them above the committee's most-preferred α -fraction. In other words, exclusion is justified when the candidate fails to outperform the committee's top tier in the eyes of a sufficiently large portion of the electorate.

A particularly transparent case arises when $\alpha = 1/2$: a candidate is excluded only if a majority of voters prefer the median member of the committee over that candidate. More generally, the parameter α serves a dual role—governing both the benchmark for comparison and the threshold of support—highlighting a delicate balance between depth of preference and breadth of agreement. Crucially, the $\frac{t}{\alpha}$ bound is not only intuitive, but also asymptotically optimal; no smaller committee can satisfy this level of undominance.

Technical Overview The technical contribution of our paper lies in adapting the Lindahl Equilibrium with Ordinal preferences (LEO), recently introduced by [Nguyen and Song \(2024\)](#), to our setting. In this formulation, LEO consists of a continuous distribution of income (or tokens) assigned to each voter, along with personalized prices for each voter-candidate pair. Voters consume probabilistically—that is, they receive a lottery over candidates determined by their income distribution and individual prices. A centralized agent then selects a lottery over candidates that maximizes total revenue. A condition analogous to market clearing must be satisfied: the consumption of each voter must closely match the selected lottery. Discrepancies are allowed only in cases where the price of a candidate is zero.

This framework is inspired by the classical Lindahl equilibrium, where voters have convex preferences over lotteries, receive a fixed endowment of tokens (normalized to one), and choose lotteries based on personalized prices. The key difference is to replace deterministic incomes with randomized ones, which allows us to accommodate purely ordinal preferences without relying on a cardinal utility representation or convex relaxation.

To construct a Condorcet winning set, we introduce a *threshold distribution*, which is a continuous approximation of a discrete distribution placing probability β on 0 and $1 - \beta$ on 1. This construction enables us to show that if a candidate falls within a voter's top $1 - \beta$ percentile of the lottery constructed from LEO, the price they face for that candidate is close to 1. We leverage this fact and the property of a LEO to bound the number of voters who place an “outside” candidate within their top $1 - \beta$ percentile. To construct a Condorcet winning set, we carefully choose the parameter β and then randomly select k candidates from the support of the LEO outcome.

Compared to [Charikar et al. \(2025a\)](#), our approach not only yields a stronger result but is also simpler and more amenable to generalization—particularly in the construction of (t, α) -undominated sets. In the generalized setting—where comparisons between the selected set and an outside candidate depend on the top t members of the set—we modify LEO into a scaled version, which we call SLEO (Scaled Lindahl Equilibrium with Ordinal preferences). The key motivation for this modification is to ensure that each voter's consumption is spread across $O(t)$ candidates, rather than concentrating on their single most preferred option. This diffusion allows us to bound the price of any outside candidate that is ranked above one of the top t candidates in the selected set.

However, this modification introduces an extra layer of complexity: the SLEO outcome is no longer a lottery over committees but a fractional allocation across individual candidates, with total weight at most equal to the committee size. To handle this, we take two additional technical steps: (a) we convert the fractional solution into a randomized committee of the desired size using dependent rounding with negative correlation, and (b) we implement a more refined, adaptive iterative procedure to construct the final committee while preserving the desired properties.

Related Literature This paper contributes to the literature on (computational) social choice by addressing several fundamental questions related to committee selection and preference aggregation.

Condorcet winning set One of the most well-known properties in this literature is the Condorcet property, which requires that the chosen outcome be preferred by a majority over any other alternative. However, such an outcome is often unattainable ([Condorcet, 1785](#)). To address this, the literature restricts preference domains or adopts weaker axioms, such as Condorcet consistency, which selects a Condorcet winner when one exists². We follow [Elkind et al. \(2015\)](#) to extend the set of possible outcomes to sets of candidates. [Elkind et al. \(2015\)](#) show that a Condorcet winning set of size at most the logarithmic of the number of candidates always exists. A constant-size Condorcet winning set (of size 32) was implicitly established in [Jiang et al. \(2020\)](#), recently improved

²Randomization is also widely studied to resolve the Condorcet paradox; see [Brandt \(2017\)](#) for a survey.

to 6 by Charikar et al. (2025a). Our results further reduce this bound to 5, although a gap remains between this and the lower bound of 3.

Undominated and dominating set The notion of majority support naturally extends to any α fraction, giving rise to the concept of an α -undominated set, as studied by Charikar et al. (2025a). In our framework, this corresponds to the special case $t = 1$. Even in this case, our approach yields tighter bounds for all α compared with prior work.

As discussed in the introduction, a stronger notion is that of an α -dominating set: a set C is α -dominating if, for every $a \notin C$, there exists $b \in C$ such that at least an α fraction of voters prefer b to a . For $\alpha \geq \frac{1}{2}$, the smallest such set can be arbitrarily large (McGarvey, 1953). In contrast, recent work (Bourneuf et al., 2025; Charikar et al., 2025b) shows that for any $\alpha < \frac{1}{2}$, constant-size α -dominating sets exist, with size depending on $\epsilon = \frac{1}{2} - \alpha$.

Our notion of a (t, α) -undominated set is incomparable to α -domination. It is stronger in requiring each outside candidate to be compared against the top t candidates in the set, but weaker in allowing each voter to compare with a different subset of t candidates. A key advantage of our notion is that a constant-size (t, α) -undominated set exists for any fixed t and $\alpha \in (0, 1)$.

Other concepts Several alternatives to Condorcet winners have been proposed, but many suffer from similar limitations: non-existence or the requirement to select an unbounded number of candidates, as in α -dominating sets for $\alpha \geq \frac{1}{2}$. For example, Fishburn (1981) introduced the idea of a *majority committee*—a set preferred by a majority over any other set of the same size—though this requires preferences over committees and may not always exist. The Smith set (Smith, 1973), a minimal set beating all outsiders in pairwise comparisons, offers a stronger form of the Condorcet principle. The *bipartisan set* (Laffond et al., 1993), defined via maximal lotteries, is a subset of the Smith set and captures randomized generalizations. We also refer the reader to Brandt et al. (2016) for a comprehensive survey of related concepts based on tournament graph structures.

Committee Selection with Ranking Preference The problem of selecting multiple candidates is closely related to committee selection and participatory budgeting, particularly through the concept of the *core*. A key distinction is that participatory budgeting typically imposes stronger core constraints, allowing deviations toward *subsets* of candidates. This requires extending preferences to sets of candidates, often leading to weaker approximation guarantees than those based on Condorcet-style conditions. For instance, Jiang et al. (2020) prove a 16-approximation for core outcomes under ranking-based preferences, later improved to 9.8217 by Charikar et al. (2025a) and to 5.1 by Nguyen and Song (2024). However, these only imply a bound of 11 on the size of a Condorcet-winning set—substantially weaker than our bound of 5. More generally, our approach yields tighter bounds for α -undominated sets for any $\alpha \geq 0.093$.

Our techniques build on Nguyen and Song (2024), who introduce Lindahl Equilibrium with Ordinal preferences (LEO) for general monotonic participatory budgeting. In contrast, our single-winner setting permits a sharper analysis and stronger bounds. Moreover, their approach does not directly yield a (t, α) -undominated set in our setting without modifying the LEO framework.

Intransitive dice The connection between voting theory and the Condorcet paradox has a surprising parallel in the phenomenon of *intransitive dice*, as observed in Charikar et al. (2025a,b). In such a setup, each die has a higher chance of beating the next in a cycle—yet paradoxically, the last die outperforms the first. This curious behavior, popularized by Gardner (1970), has inspired a substantial body of mathematical research. For a deeper exploration of this link, see Charikar et al. (2025b). Uncovering the full extent of its relevance to the generalization we study in this paper presents an intriguing avenue for future investigation.

Computation This paper is primarily concerned with existence results. The analysis of computational aspects is deferred to future work. In particular, for undominated sets, Cheng et al. (2019) present an algorithm to compute a 16-approximation for core outcomes under ranking-based preferences, which can be used to compute an α -undominated set of size $\frac{16}{\alpha}$. In our ongoing work, we improve upon this result by showing that LEO with uniform distribution can be used to compute an α -undominance set of size $\frac{9.82}{\alpha}$. For generalized (t, α) -undominated sets, the LP approach is not sufficient. However, it is possible to use a convex program to produce a (t, α) -undominated set of size $O(\frac{t}{\alpha})$, albeit with a worse constant factor.

Organization In Section 2, we introduce necessary notations and our main technical tool, Lindahl Equilibrium with Ordinal Preferences (LEO). In Section 3, we prove that five candidates suffice for a Condorcet winning set. In Section 4, we prove that $O(t/\alpha)$ candidates are sufficient and required to ensure a (t, α) -undominated set. Section 5 concludes.

2 Notation and Preliminaries Throughout the paper, we use bold lowercase letters such as \mathbf{x}, \mathbf{y} to denote a vector and regular lowercase letters such as $x_{i,j}, y_k$ to denote a coordinate of a vector.

Condorcet Winning Sets and Undominated Sets Let $V = \{1, 2, \dots, n\}$ be the set of voters (agents), and let A be the set of m alternatives (candidates). Each $v \in V$ has a strict preference relation \succ_v over the alternatives in A . $a \succeq_v b$ means that either $a = b$ or agent v prefers a to b . An election instance is denoted by $\mathcal{E} = (V, A, \succ)$.

Given an election $\mathcal{E} = (V, A, \succ)$, a voter $v \in V$, a set of candidates $C \subseteq A$, and an outside candidate $a \in A \setminus C$, we write $a \succ_v C$ to indicate that v strictly prefers a to all alternatives in C , i.e., $a \succ_v c$ for all $c \in C$. A Condorcet winning set is a set of candidates $C \subseteq A$ such that for every candidate $a \in A \setminus C$, the minority of the voters prefer a to any candidates in C , i.e., $|\{v \in V \mid a \succ_v C\}| \leq \lfloor \frac{n}{2} \rfloor$ for all $a \in A$. A generalization of the Condorcet winning set is the notion of α -undominated set.

DEFINITION 2.1. A subset $C \subseteq A$ is an α -undominated set if for every candidate $a \in A \setminus C$,

$$|\{v \in V : a \succ_v C\}| \leq \alpha n.$$

We note that a $\frac{1}{2}$ -undominated set is equivalent to a Condorcet winning set.

Lindahl Equilibrium with Ordinal Preferences We adapt the concept of Lindahl Equilibrium under Ordinal Preferences (LEO), introduced in [Nguyen and Song \(2024\)](#), to our setting. As discussed in the introduction, LEO itself is an ordinal adaptation of the classical Lindahl equilibrium ([Foley, 1970](#)).

We assume that each voter selects only individual candidates or an outside option \emptyset . Recall that each $v \in V$ is associated with a strict preference \succ_v over A . We extend the use of \succ_v over $A \cup \{\emptyset\}$ by having \emptyset as the least preferred, while the ranking over A remains the same.

Each voter v is endowed with a random income \mathcal{I}_v supported on $[0, 1]$ and has a *personalized* price vector $\mathbf{p}_v \in \mathbb{R}_+^{|A \cup \{\emptyset\}|}$. We use $p_{v,a}$ to denote the personalized price of $a \in A \cup \{\emptyset\}$ for voter v with $p_{v,\emptyset} = 0$. We assume that voters share a common income distribution \mathcal{I} .

Under a fixed income $b \in \mathbb{R}_+$, the voter v 's demand is $\max_{\succ_v} \{a : a \in A \cup \{\emptyset\} \text{ and } p_{v,a} \leq b\}$, which is her most preferred and affordable alternative within her income. Notice that when the prices of all candidates are greater than b , the demand will be the outside option, \emptyset . Given the random income distribution \mathcal{I} , voter v 's *random demand* is defined as

$$\mathcal{D}_v(\mathbf{p}_v, \mathcal{I}) := \left\{ \max_{\succ_v} \{a : a \in A \cup \{\emptyset\} \text{ and } p_{v,a} \leq b\} \mid b \sim \mathcal{I} \right\}.$$

Under the assumptions that the random income distribution is non-negative with probability 1 and $p_{v,\emptyset} = 0$, the resulting random demand is non-empty and constitutes a lottery on $A \cup \{\emptyset\}$.

Lindahl equilibrium also involves a centralized agent called the *producer*, who decides on the quantity of each candidate in A to produce. From the producer's perspective, there is a unit cost associated with each candidate $a \in A$. The producer has a total budget of $B > 0$ and can only produce the quantity $\mathbf{z} \in \mathbb{R}_+^{|A|}$ satisfying $\mathbf{1}^T \mathbf{z} = B$. Thus, the budget B determines the size of the committee. The producer aims to maximize total revenue. Therefore, the set of optimal strategies for the producer is

$$\mathcal{P}(B, \mathbf{p}) := \arg \max_{\mathbf{z} \in \mathbb{R}_+^{|A|}} \sum_{a \in A} \left(\sum_{v \in V} p_{v,a} \right) \cdot z_a \text{ subject to } \mathbf{1}^T \mathbf{z} = B.$$

We have the following definition of the Lindahl equilibrium with ordinal preference (LEO) for our setting adapted from [Nguyen and Song \(2024\)](#).

DEFINITION 2.2. Given a random income \mathcal{I} supported on $[0, 1]$ and a total budget $B > 0$, the Lindahl equilibrium with ordinal preference (LEO) $(\mathbf{x}, \mathbf{y}, \mathbf{p})$ consists of personalized prices $\mathbf{p}_v \in [0, 1]^{|A \cup \{\emptyset\}|}$ with $p_{v,\emptyset} = 0$, individual consumptions $\mathbf{x}_v \in [0, 1]^{|A \cup \{\emptyset\}|}$ for each voter $v \in V$, and a common allocation $\mathbf{y} \in [0, B]^{|A|}$ such that the following holds:

- (1) for each $a \in A \cup \{\emptyset\}$, $x_{v,a} = \Pr[\mathcal{D}_v(\mathbf{p}_v, \mathcal{I}) = a]$;
- (2) for each $a \in A$, $x_{v,a} \leq y_a$; if the strict inequality holds then $p_{v,a} = 0$;

(3) $\mathbf{y} \in \mathcal{P}(B, p)$.

Compared to the traditional Lindahl equilibrium with convex and continuous preferences, there are two key differences in our setting. In the classical Lindahl equilibrium, each voter has a fixed, constant income or tokens, normalized to 1, and there is a common allocation \mathbf{y} along with a set of individual prices such that: (i) the common allocation is the optimal allocation for each voter given their individual prices, and (ii) \mathbf{y} is also the revenue-maximizing allocation for the producer.

In our formulation, the first key departure from the classical setting is that each voter's income is modeled not as a fixed constant but as a random variable with a continuous distribution. This modification is essential because we work with ordinal preferences over discrete choices; introducing income uncertainty ensures that each voter's resulting random demand varies continuously with prices, without requiring an extension of preferences from discrete choices to lotteries. Second, we relax the requirement that each individual's demand must match the common allocation. Instead, we adopt a condition analogous to market clearing, formalized as condition (2) in Definition 2.2.

Together, these two modifications are essential for applying Kakutani's fixed-point theorem and establishing the existence of a LEO.

THEOREM 2.3. (*Nguyen and Song (2024)*) *For any $b > 0$ and distribution D with support on $[0, 1]$ and a continuous cumulative function, there exists a LEO with $B = b$ and a random income $\mathcal{I} = D$.*

3 Five Candidates Suffice to Win a Voter Majority Our main result in this section is the following theorem.

THEOREM 1.1. *Given $k \in \mathbb{N}$ and $\beta \in (0, 1)$, a $(\beta + (1 - \beta)^k)$ -undominated set of size k exists.*

Given $k \in \mathbb{N}$, we define $\alpha(k) := \inf_{\beta \in (0, 1)} \{\beta + (1 - \beta)^k\}$. From Theorem 1.1, we know that a $\alpha(k)$ -undominated set of size k exists. We compare our $\alpha(k)$ with the one from Theorem 5 of Charikar et al. (2025a) in Table 3.1.

k	2	3	4	5	6	7	8
$\alpha(k)$ from Theorem 1.1	0.750	0.615	0.528	0.465	0.418	0.380	0.350
ratios from Charikar et al. (2025a)	0.798	0.674	0.589	0.526	0.477	0.438	0.406

Table 3.1: Bounds on the undominance ratio for size k .

Moreover, we obtain the following result as a direct consequence of Theorem 1.1.

COROLLARY 3.1. *There exists a 5-candidate Condorcet winning set for any election.*

To prove Theorem 1.1, we use LEO with a total budget $B = 1$ and a specially designed class of income distributions, which we call *threshold distributions*. The class of threshold distributions $\mathcal{I}^{\beta, \varepsilon}$ has two parameters, β and ε , and serves as a smoothed approximation of a Bernoulli distribution that assigns income 1 with probability β and 0 with probability $1 - \beta$. The small parameter ε ensures that the distribution is continuous.

DEFINITION 3.2. *Given $\beta, \varepsilon \in (0, 1)$, a distribution \mathcal{I} supported on $[0, 1]$ belongs to the class of threshold distributions $\mathcal{I}^{\beta, \varepsilon}$ if the following conditions are satisfied:*

1. *The random variable $X \sim \mathcal{I}$ has a continuous CDF $F(x) = \Pr[X \leq x]$.*
2. *$F(1 - \varepsilon) = 1 - \beta$, i.e., $\Pr[X \geq 1 - \varepsilon] = \beta$.*
3. *The expected value of X satisfies $\mathbb{E}[X] \leq \beta$, that is, $\int_0^1 F(x) dx \leq \beta$.*

It is straightforward to construct a distribution belonging to $\mathcal{I}^{\beta, \varepsilon}$ for any $\beta, \varepsilon \in (0, 1)$. Consider the distribution $\mathcal{I}^{\beta, 0}$, which assigns probability β to 1 and $1 - \beta$ to 0. While this distribution satisfies the last two properties, it is not continuous. However, a slight perturbation by ε ensures continuity.

CLAIM 3.3. *For any $\beta, \varepsilon \in (0, 1)$, there exists a distribution supported on $[0, 1]$ in the class $\mathcal{I}^{\beta, \varepsilon}$.*

Proof. Consider the following piecewise linear CDF F .

$$F(x) = \begin{cases} 0 & \text{if } x \in [0, 1 - \beta - \delta], \\ \frac{1-\varepsilon}{\delta} [x - (1 - \beta - \delta)] & \text{if } x \in [1 - \beta - \delta, 1 - \beta], \\ \frac{\varepsilon}{\mu} \left[x - \left(1 - \beta - \frac{\mu(1-\varepsilon)}{\varepsilon} \right) \right] & \text{if } x \in [1 - \beta, 1 - \beta + \mu], \\ 1 & \text{if } x \in [1 - \beta + \mu, 1], \end{cases}$$

where $\delta \in (0, 1 - \beta]$, $\mu \in (0, \beta]$, and $\delta(1 - \varepsilon) \leq \mu\varepsilon$. Clearly, the first two conditions in Definition 3.2 are satisfied. The third condition is satisfied because $\delta(1 - \varepsilon) \leq \mu\varepsilon$. \square

Now, we first observe that when $B = 1$, the common allocation \mathbf{y} of a LEO coincides with all individual consumptions and forms a lottery over the alternatives.

CLAIM 3.4. *Let $\beta, \varepsilon \in (0, 1)$ be given and $(\mathbf{x}, \mathbf{y}, \mathbf{p})$ be a LEO with $B = 1$ and a random income $\mathcal{I} \in \mathcal{I}^{\beta, \varepsilon}$. Then, $\sum_{a \in A} y_a = 1$ and $x_{v, a'} = y'_a$ for all $a' \in A, v \in V$.*

Proof. First, by definition, $\sum_{a \in A} y_a = B = 1$. Second, suppose $x_{v, a'} < y'_a$ for some $a' \in A$. Then, by condition (2) of the LEO, it must be that $p_{v, a'} = 0$. Moreover, $\sum_{a \in A} x_{v, a} < 1$, which implies that voter v consumes \emptyset with positive probability. This contradicts the definition of random demand for v , because if $p_{v, a'} = 0$, then v should not have chosen \emptyset , since a' was free and strictly preferred. \square

To prove Theorem 1.1, we consider a random committee formed by independently sampling k alternatives from the lottery induced by the LEO mechanism with total budget $B = 1$ and a random income $\mathcal{I} \in \mathcal{I}^{\beta, \varepsilon}$. We then show, via an expectation argument, that there exists at least one realization of this random process that satisfies the conditions of Theorem 1.1.

The argument bounding the number of voters who may prefer an outside candidate proceeds in two parts. First, we define a notion of *coverage*: intuitively, a voter is said to be covered by the committee if at least one of its members lies within the top β percentile of the voter's preferences under the LEO lottery. The probability that a voter is covered can be bounded based on β , and by applying a standard averaging argument, we obtain a deterministic committee that fails to cover only a small fraction of voters. This yields the first bound.

Second, for voters who are covered, we show that only a limited number can strictly prefer some outside candidate over the entire committee. If too many do, the total price they assign to that candidate would exceed the total available income. However, a key property of the LEO is that the total price any candidate can receive across all voters is bounded above by the total income, which is at most βn . This provides the second bound. The sum of the two bounds then gives the desired result.

We start by introducing the notion of *boundary candidates*.

DEFINITION 3.5. *Let $\delta \in (0, 1)$, a voter $v \in V$ and her personalized price vector $\mathbf{p}_v \in [0, 1]^{|A \cup \{\emptyset\}|}$ with $p_{v, \emptyset} = 0$ be given. The δ -boundary candidate $a_{v, \delta}$ of v is her most preferred option in $A \cup \{\emptyset\}$ with price at most $1 - \delta$, that is, $a_v := \max_{\succsim_v} \{a \in A \cup \{\emptyset\} \mid p_{v, a} \leq 1 - \delta\}$.*

Now we are ready to define *covered* and *uncovered* voters.

DEFINITION 3.6. *Let $\beta, \varepsilon \in (0, 1)$ be given and $(\mathbf{x}, \mathbf{y}, \mathbf{p})$ be a LEO with a random income $\mathcal{I} \in \mathcal{I}^{\beta, \varepsilon}$. A voter $v \in V$ is said to be covered by a set of candidates $C \subseteq A$ if there is a candidate $c \in C$ such that $c \succeq_v a_{v, \varepsilon}$, and uncovered by C if otherwise.*

As outlined previously in this section, we first establish the existence of a set of candidates that covers a significant number of voters.

CLAIM 3.7. *Let $\beta, \varepsilon \in (0, 1)$ be given and $(\mathbf{x}, \mathbf{y}, \mathbf{p})$ be a LEO with $B = 1$ and a random income $\mathcal{I} \in \mathcal{I}^{\beta, \varepsilon}$. For any $k \geq 1$, there exists a candidate set C of size at most k s.t. the number of voters not covered by C is at most $(1 - \beta)^k$.*

Proof. From Claim 3.4, we regard \mathbf{y} as a lottery for selecting a single candidate. We sample k times independently according to \mathbf{y} . Let the chosen random candidate in the i -th iteration be \tilde{a}_i . Let $\tilde{C} = \{\tilde{a}_i \mid i \in [k]\}$ be the random candidate set (not a multi-set). We now bound the probability that a realization of \tilde{C} does not cover an arbitrary voter $v \in V$.

By Definition 3.5, $\sum_{a \succ_v a_{v,\varepsilon}} x_{v,a} \leq \Pr_{b \sim \mathcal{I}}(b < 1 - \varepsilon) = 1 - \beta$. Then, since the common allocation \mathbf{y} coincides with v 's individual consumption \mathbf{x}_v by Claim 3.4, we have $\sum_{a \succ_v a_{v,\varepsilon}} y_a \leq 1 - \beta$. Therefore, at each iteration i , the probability that v prefers a to the selected candidate $a_i \sim \tilde{a}_i$ is at most $1 - \varepsilon$. Because \tilde{C} selects candidates independently, the probability of v preferring $a_{v,\varepsilon}$ to all selected candidates is at most $(1 - \beta)^k$.

Thus, in expectation over \tilde{C} , there are at most $(1 - \beta)^k n$ uncovered voters. Hence, there must be a subset of candidates $C \subseteq A$ of size k such that at most $(1 - \beta)^k n$ voters are uncovered and this completes the proof. \square

We now bound the number of covered voters who prefer an outside candidate to all selected candidates in C .

CLAIM 3.8. *Let $\beta, \varepsilon \in (0, 1)$ be given and $(\mathbf{x}, \mathbf{y}, \mathbf{p})$ be a LEO with $B = 1$ and a random income $\mathcal{I} \in \mathcal{I}^{\beta, \varepsilon}$. Let C be a candidate set of size at most k satisfying the requirement of Claim 3.7. Then, for any $a \notin C$, the number of voters covered by C who prefer a to all candidates in C is at most $\frac{\beta n}{1 - \varepsilon}$.*

Proof. Under the LEO $(\mathbf{x}, \mathbf{y}, \mathbf{p})$, the expected revenue obtained by the producer is equal to the sum of the expected prices which voters pay for their individual consumptions. This is because by Claim 3.4, we have

$$\sum_{a \in A} \left(\sum_{v \in V} p_{v,a} \right) y_a = \sum_{a \in A} \sum_{v \in V} p_{v,a} x_{v,a} = \sum_{v \in V} \sum_{a \in A} p_{v,a} x_{v,a}.$$

Note that for each voter v , $\sum_{a \in A} p_{v,a} x_{v,a}$ equals the expected cost of her random demand under the LEO lottery. This expected cost must not exceed v 's expected income $\mathbb{E}_{b \sim \mathcal{I}}[b]$, which is at most β . Thus $\sum_{a \in A} (\sum_{v \in V} p_{v,a}) y_a \leq n\beta$.

We now show that the total price assigned to candidate a , i.e., $\sum_{v \in V} p_{v,a}$, is at most $n\beta$. Suppose, for the sake of contradiction, that $\sum_{v \in V} p_{v,a} > n\beta$. Then, the producer could choose the deterministic allocation with $y_a = 1$ and $y_{a'} = 0$ for all $a' \neq a$. The expected revenue of this allocation would exceed $n\beta$, which contradicts the fact that \mathbf{y} is the revenue-maximizing lottery. Therefore, we must have $\sum_{v \in V} p_{v,a} \leq n\beta$.

Notice that if a voter v is covered by C and $a \succ_v C$, then $a \succ_v a_{v,\varepsilon}$ by the transitivity of preference ordering. Then, by the definition of boundary candidates, $p_{v,a} > 1 - \varepsilon$. And we have already shown that the total price for each candidate $a \in A$ is at most $n\beta$. Therefore, among the covered voters, the number of people preferring a over C is at most $\frac{\beta n}{1 - \varepsilon}$. \square

Combining Claim 3.7 and Claim 3.8 yields an upper bound on the total number of voters preferring an outside candidate, where we consider the worst-case scenario in which all uncovered voters will deviate to the outside candidate.

COROLLARY 3.9. *Let $\beta, \varepsilon \in (0, 1)$ be given. For any $k \geq 1$, there exists a set of candidates C s.t. for any $a \notin C$, the number of voters preferring a to C is at most $(\frac{\beta}{1 - \varepsilon} + (1 - \beta)^k) \cdot n$.*

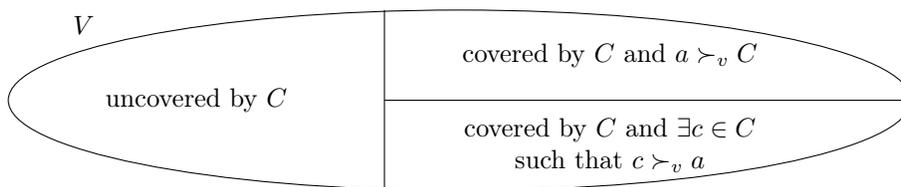


Figure 3.1: Voters $v \in V$ in 3 categories based on a fixed $a \in A \setminus C$: (1) uncovered by C , (2) covered by C and $a \succ_v C$, or (3) covered by C and $\exists c \in C$ such that $c \succ_v a$. Note that there are at most $\frac{\beta n}{1 - \varepsilon}$ voters in category (2) by Claim 3.8.

Our main theorem now follows easily from Corollary 3.9 by taking $\varepsilon \rightarrow 0$.

Proof of Theorem 1.1. For any $\varepsilon \in (0, 1)$, there exists a set of candidates C_ε with size k s.t. for any $a \notin C_\varepsilon$, the number of voters preferring a to C_ε is at most $(\frac{\beta}{1 - \varepsilon} + (1 - \beta)^k) \cdot n$ by Corollary 3.9. Because $\frac{\beta}{1 - \varepsilon} \rightarrow \beta$ as $\varepsilon \rightarrow 0$, there is a small $\varepsilon^* \in (0, 1)$ s.t. $\lfloor (\frac{\beta}{1 - \varepsilon^*} + (1 - \beta)^k) \cdot n \rfloor = \lfloor (\beta + (1 - \beta)^k) \cdot n \rfloor$. Since the number of deviating voters must be integral, it follows that the number of voters preferring a to C_{ε^*} is at most $(\beta + (1 - \beta)^k) \cdot n$, and therefore C_{ε^*} is the desired set.

4 Generalizing Condorcet Winning Sets This section studies a generalization of the Condorcet winning set. Specifically, we define a new method of comparing between a set of candidates and a single candidate outside the set, based not on the best candidate in the set but on the t -th best. We formalize this idea in the following definitions.

DEFINITION 4.1. *Given an integer $t \geq 1$, a voter $v \in V$, a subset of candidates $C \subseteq A$ with $|C| \geq t$, and a candidate $a \in A \setminus C$, we say that voter v t -prefers C to a and write $C \succ_v^t a$, if there are at least t candidates in C that v strictly prefers to a ; that is,*

$$C \succ_v^t a \quad \text{if and only if} \quad |\{c \in C : c \succ_v a\}| \geq t.$$

Otherwise, we say that v t -prefers a to C , and write $a \succ_v^t C$.

DEFINITION 4.2. *A subset $C \subseteq A$ with $|C| \geq t$ is a (t, α) -undominated set if for every candidate $a \in A \setminus C$, $|\{v \in V : a \succ_v^t C\}| \leq \alpha n$.*

A special case is when $t = 1$: in this case, $C \succ_v^1 a$ means that there is at least one candidate in C that is preferred to a , while $a \succ_v^1 C$ means that a is preferred to all candidates in C . Thus, when $t = 1$, the definition of a (t, α) -undominated set coincides with the notion of α -undominance introduced in Section 3. Our main result of this section is the following.

THEOREM 1.2. *Given $\alpha \in (0, 1]$, there exists a set of size $\lfloor \delta(t) \cdot \frac{t}{\alpha} \rfloor$ which is (t, α) -undominated. Here, $1 < \delta(t) \leq 4.75$ for all integers $t \geq 2$ and $\delta(t) \rightarrow 1$ as $t \rightarrow \infty$.*

We can numerically compute the values $\delta(2) = 4.75$, $\delta(3) = 4.11$, $\delta(4) = 3.69$, ..., $\delta(8) = 2.89$. The function δ converges to 1 as t becomes large. Moreover, we provide a lower bound for the size of (t, α) -undominated sets, derived by adapting the example provided in Charikar et al. (2025a).

THEOREM 1.3. *If $k \geq \frac{t+1}{\alpha} - 1$, then there exist elections that admit no (t, α) -undominated set of size k .*

Given Theorem 1.3, we can see that Theorem 1.2 is asymptotically tight for every α as t becomes large; that is, the size of a (t, α) -undominated set is on the order of t/α . For a small value of t , the expression for $\delta(t)$, which is derived from the Chernoff bound, is somewhat complex. The exact formulation is provided in the proof at the end of this section.

The detailed proofs of Theorem 1.2 and Theorem 1.3 can be found in the full version of the paper, and we describe the main idea here. We start by highlighting the challenge in generalizing the result in Section 3. A natural but flawed idea is to apply the $t = 1$ case iteratively: construct an α -undominated subcommittee, remove its members, and repeat t times. The union has size $O(t)$, and one might expect that each outside candidate is dominated by t candidates—one from each subcommittee. However, this fails: although each subcommittee may contain a preferred candidate, the supporting voters may differ, and their intersection can be small or empty. Thus, we cannot ensure that enough voters prefer t members of the union over any outsider.

To prove Theorem 1.2, we adopt a strategy similar to that used in Section 3. We begin with a LEO-style fractional solution and iteratively construct an integral solution from it. However, a direct application of the LEO approach from the previous section does not suffice in our setting. We will need to make three key modifications.

First, in order to ensure a bound on the number of voters who have fewer than t good candidates in the selected set, we must modify the LEO so that voters “consume” multiple candidates in the equilibrium. In standard LEO, each voter consumes at most one unit of candidates in total. To address this, we introduce a scaled version of LEO, which we call s -SLEO. This variant replaces the condition $\mathbf{x}_v \leq \mathbf{y}$ with $s\mathbf{x}_v \leq \mathbf{y}$ for some scaling factor $s = O(t)$. Together with an additional constraint that \mathbf{y} is at most $\mathbf{1}$, this modification ensures that each voter effectively allocates their spending across their top t candidates.

Second, the fractional solution obtained is no longer a lottery over individual candidates, since we are not selecting a single candidate but a set of candidates. Therefore, we need to convert the fractional solution into a lottery over sets of candidates. We achieve this using dependent rounding with negative correlation.

Third, as we will show later, one-shot dependent rounding is not sufficient to achieve the desired result for certain ranges of α and t . A more complex, iterative procedure is required to construct the final outcome.

Scaled LEO and Dependent Rounding We modify LEO to a version called scaled LEO (SLEO) as follows. First, we modify the optimal strategy for the producer:

$$\mathcal{P}^S(B, \mathbf{p}) := \arg \max_{\mathbf{z} \in \mathbb{R}_+^{|A|}} \left(\sum_{a \in A} \sum_{v \in V} p_{v,a} \right) z_a \text{ subject to } \mathbf{1}^T \mathbf{z} = B \text{ and } \mathbf{z} \leq \mathbf{1}.$$

Note that in this setting, unlike LEO, we explicitly enforce that the producer allocates *at most one unit* of each candidate. We introduce the following notion of SLEO.

DEFINITION 4.3. *Given a random income \mathcal{I} supported on $[0, 1]$, a total budget $B > 0$, and $s > 0$, an s -SLEO $(\mathbf{x}, \mathbf{y}, \mathbf{p})$ consists of personalized prices $\mathbf{p}_v \in [0, 1]^{|A \cup \{\emptyset\}|}$ with $p_{v, \emptyset} = 0$, individual consumptions $\mathbf{x}_v \in [0, 1]^{|A \cup \{\emptyset\}|}$ for each $v \in V$, and a common allocation $\mathbf{y} \in [0, 1]^{|A|}$ such that*

- (1) $x_{v,a} = \Pr[\mathcal{D}_v(\mathbf{p}_v, \mathcal{I}) = a]$ for each $a \in A \cup \{\emptyset\}$, $v \in V$;
- (2) For each $a \in A$, $s \cdot x_{v,a} \leq y_a$; if the strict inequality holds then $p_{v,a} = 0$;
- (3) $\mathbf{y} \in \mathcal{P}^S(B, \mathbf{p})$.

The existence of the SLEO is a slight adjustment to the proof of Theorem 2.3 and can be found in the full version of this paper.

THEOREM 4.4. *For any $b \leq |A|$, $s > 0$ and distribution D with support on $[0, 1]$ and a continuous cumulative function, then there exists an s -SLEO with $B = b$ and a random income $\mathcal{I} = D$.*

It is important to note that the common allocation \mathbf{y} of a SLEO may not be a lottery over single candidates. Instead, we will convert \mathbf{y} into a lottery over committees of size either $\lceil B \rceil$ or $\lfloor B \rfloor$ using the following dependent rounding process described in Gandhi et al. (2006).

PROPOSITION 4.5. (Gandhi et al., 2006) *Given $\mathbf{y} \in [0, 1]^m$ s.t. $\mathbf{1}^T \mathbf{y} = B$, there exists a distribution \tilde{Y} over $\{0, 1\}^m$ satisfying the following properties:*

- (1) *Preservation of marginals:* $\mathbb{E}_{\mathbf{Y} \sim \tilde{Y}}[Y_k] = y_k$ for all $k \in [m]$.
- (2) *Preservation of weights:* $\mathbf{1}^T \mathbf{Y} \in \{\lceil B \rceil, \lfloor B \rfloor\}$ with probability 1.
- (3) *Negative correlation between entries:* for any $W \subseteq [m]$, $\Pr_{\mathbf{Y} \sim \tilde{Y}}[\bigwedge_{k \in W} Y_k = 0] \leq \prod_{k \in W} (1 - y_k)$.

Our construction is based on a randomized committee, which we obtain by applying the dependent rounding procedure in Gandhi et al. (2006) to the fractional allocation resulting from an SLEO.

DEFINITION 4.6 (Randomized (B, s, ε) -Lindahl Committee). *Let $s, B > 0$ and $\varepsilon \in (0, 1)$ be given. A randomized (B, s, ε) -Lindahl committee is a random subset $C \subseteq A$ of size at most $\lceil B \rceil$ constructed as follows:*

1. *Compute an s -SLEO with budget B , where each voter's income distribution is the uniform distribution over the interval $[1 - \varepsilon, 1]$, denoted as $\mathbf{U}[1 - \varepsilon, 1]$. Let $\mathbf{y} \in [0, 1]^{|A|}$ be the resulting common fractional allocation over the set of alternatives A .*
2. *Apply the dependent rounding procedure as described in Gandhi et al. (2006) to \mathbf{y} to obtain a committee C of size at most $\lceil B \rceil$ by Proposition 4.5.*

A key property of the randomized Lindahl committee is that in order for each voter to be assigned a set of t good candidates in the selected committee, it suffices to set the budget B and the scaling factor s to $O(t)$. Formally, a set of t candidates in C is good for a voter v if any candidate $a \notin C$ that v prefers over all members of the set has an individual Lindahl price satisfying $p_{v,a} \geq 1 - \varepsilon$. This property is useful because it limits the number of voters who can strongly prefer excluded candidates. Specifically, if a candidate a is excluded from the randomized Lindahl committee with positive probability, then its total price across all voters is at most $s \cdot \frac{t}{B}$, the average total price per candidate. This implies a bound on the number of voters who t -prefer a candidate outside the committee.

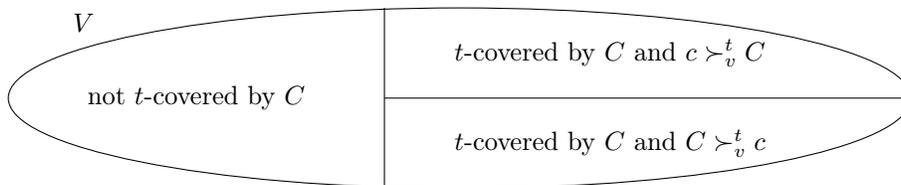


Figure 4.1: Voters $v \in V$ in 3 categories based on a fixed $c \in A \setminus C$: (1) not t -covered by C , (2) t -covered by C and $c \succ_v^t C$, or (3) t -covered by C and $C \succ_v^t c$. Note that there are at most $\omega(\gamma, t) \cdot |V|$ voters in category (1) by Claim 4.9 and at most $\frac{\gamma^t}{B(1-\varepsilon)} \cdot |V|$ voters in category (2) by Claim 4.10.

Also note that, as compared to Section 3, we have switched from threshold distributions to $\mathbf{U}[1 - \varepsilon, 1]$ in Definition 4.6. To clarify, recall that in Section 3, the rationale of using threshold distributions is to ensure that the common allocation \mathbf{y} is a lottery over candidates so that independent sampling can be applied. In this section, however, we apply dependent rounding, which does not require \mathbf{y} to be a lottery over candidates, and it then becomes more convenient to just use $\mathbf{U}[1 - \varepsilon, 1]$.

We start with the following claim bounding the probability that a voter has less than t good candidates in the random Lindahl committee.

CLAIM 4.7. *Let an integer $t \geq 2$, $\gamma \geq 1$ and $\varepsilon \in (0, 1)$ be given. Let $(\mathbf{x}, \mathbf{y}, \mathbf{p})$ be a γt -SLEO with random income distribution $\mathcal{I} = \mathbf{U}[1 - \varepsilon, 1]$ and a budget $B \geq \gamma t$. Let \tilde{C} be the randomized $(B, \gamma t, \varepsilon)$ -Lindahl committee derived from Definition 4.6. Then, for every voter $v \in V$, if $x_{v, \emptyset} = 0$, the following holds:*

$$\Pr_{C \sim \tilde{C}} \left[\left| C \cap \{a \in A : x_{v,a} > 0\} \right| < t \right] \leq \omega(\gamma, t), \text{ where } \omega(\gamma, t) := \left(\frac{e^{-\frac{\gamma t - t + 1}{\gamma t}}}{\left(\frac{t-1}{\gamma t}\right)^{\frac{t-1}{\gamma t}}} \right)^{\gamma t}.$$

Next, we keep the definition of boundary candidate the same as in Section 3 and introduce the notion of t -covering.

DEFINITION 4.8. *Let $\varepsilon \in (0, 1)$ be given and $(\mathbf{x}, \mathbf{y}, \mathbf{p})$ be a SLEO with a random income $\mathcal{I} = \mathbf{U}[1 - \varepsilon, 1]$. We say that a voter $v \in V$ is t -covered by a candidate set $C \subseteq A$ if $C \succ_v^t a_{v, \varepsilon}$.*

Using Claim 4.7, we now show the existence of a committee which t -covers a significant number of voters and contains all candidates a with $y_a = 1$.

CLAIM 4.9. *Let an integer $t \geq 2$, $\gamma \geq 1$ and $\varepsilon \in (0, 1)$ be given and $(\mathbf{x}, \mathbf{y}, \mathbf{p})$ be a γt -SLEO with a random income $\mathcal{I} = \mathbf{U}[1 - \varepsilon, 1]$ and $B \geq \gamma t$. Then, there exists a committee C of size at most $\lceil B \rceil$ such that:*

1. *the number of voters that are not t -covered by C is at most $\omega(\gamma, t) \cdot |V|$,*
2. *C contains every alternative a with $y_a = 1$.*

The following claim states that any committee containing all candidates a with $y_a = 1$ must satisfy that only a small number of t -covered voters prefer an outside option to the committee.

CLAIM 4.10. *Let an integer $t \geq 2$, $\gamma \geq 1$ and $\varepsilon \in (0, 1)$ be given and $(\mathbf{x}, \mathbf{y}, \mathbf{p})$ be a γt -SLEO for a set of voters V with a random income $\mathcal{I} = \mathbf{U}[1 - \varepsilon, 1]$ and $B \geq \gamma t$. Let C be a committee that includes all candidates a s.t. $y_a = 1$. Then, the following statements are true.*

(i) $\sum_{v \in V} p_{v,c} \leq \gamma t \cdot |V| \cdot \frac{1-\varepsilon/2}{B}$.

(ii) *There are at most $\frac{\gamma^t}{B(1-\varepsilon)} \cdot |V|$ voters who are t -covered by C and t -prefer c to C .*

Claims 4.9 and 4.10 are illustrated in Figure 4.1. Together, these results establish the existence of a committee C of size at most $\lceil B \rceil$ that limits the number of voters who prefer $c \in A \setminus C$ to C . Specifically, Claim 4.9 provides an upper bound of $\omega(\gamma, t) \cdot |V|$ on the number of voters who are not t -covered by C . Claim 4.10, in turn, shows

that for any candidate $c \in A \setminus C$, among the voters who are t -covered by C , at most $\frac{\gamma t}{B(1-\varepsilon)} \cdot |V|$ of them t -prefer c over C .

As in the result from Section 3, combining these two bounds yields a bound on the total number of voters who might t -prefer c to C . To ensure that this number is less than $\alpha \cdot |V|$, it is necessary to choose γ and B as functions of α and t . Combining Claims 4.9 and 4.10 implies the following.

CLAIM 4.11. *Let $\alpha \in (0, 1)$, an integer $t \geq 2$, and $\varepsilon > 0$ be given. If there exists a $\gamma \geq 1$ and an integer $B \geq \gamma t$ s.t. $\alpha \geq \frac{\gamma t}{B} + \omega(\gamma, t)$, then a $(t, \frac{\alpha}{1-\varepsilon})$ -undominated set of size B exists.*

Claim 4.11 allows us to derive bounds for (t, α) -undominated sets for certain values of t and α . In general, when α is relatively large, the bound in Claim 4.11 can be used to obtain a committee of size at most $\delta(t) \cdot t$, as stated in Theorem 1.2. Furthermore, as we show in the proof of Theorem 1.2, the bound approaches the optimal value of $\frac{t}{\alpha}$ as $t \rightarrow \infty$.

However, this method has limitations. It does not guarantee the construction of a (t, α) -undominated set of the desired size order $O(\frac{t}{\alpha})$ when $\frac{1}{\alpha}$ is big compared to t . For example, one can show that in order to be a (t, α) -undominated set, B needs to be in the order $\frac{1}{\alpha} \cdot \log(\frac{1}{\alpha})$, which may well exceed $O(\frac{t}{\alpha})$ when $\log(\frac{1}{\alpha})$ is relatively large to t .

To obtain a constant bound for all t and α , we incorporate SLEO into the framework of Jiang et al. (2020) to derive a new iterative method for constructing the committee, which yields a tighter bound when α is small. By combining the bound from Claim 4.11 with the new bound described below, we show that the resulting committee satisfies the desired size requirement of $O(t/\alpha)$.

The detailed proofs of Claim 4.7, Claim 4.9, Claim 4.10, and Claim 4.11 can be found in the full version of this paper.

Iterative Algorithm To overcome the issue described above, we give an iterative process that repeatedly applies Claim 4.9 and Claim 4.10, each time restricting attention to the set of voters who are not t -covered by the current committee C . Note that the number of such voters decreases by a constant factor in each iteration (specifically, by $\omega(\gamma, t)$). The process will continue until “almost” all voters are covered, thus solving the issue of the one-step rounding.

To maintain control over the total committee size, we also reduce the budget B by a constant factor at each step, ensuring that the sequence $\{B_i\}_{i=0}^{\infty}$ forms a geometric progression. This prevents the total size of the output committee from growing unbounded across iterations.

The precise iterative algorithm is described in Algorithm 4.1. In this algorithm, the parameters $\gamma, \tau \geq 1$ are selected later to minimize the approximation factor $\delta(t)$ in Theorem 1.2.

Algorithm 4.1 Iterative Rounding with SLEO

Input: A set V of n voters, $\alpha \leq 1$, an integer $t \geq 2$, $\gamma, \tau \geq 1$ with $\omega(\gamma, t)\tau < 1$, $\varepsilon \in (0, 1)$

Output: A $(t, \frac{\alpha}{1-\varepsilon})$ -undominated set C

$V_0 \leftarrow V$, $t \leftarrow 0$, $B_0 \leftarrow \gamma \cdot \frac{1}{1-\omega(\gamma, t)\tau} \cdot \frac{t}{\alpha}$

while $|B_i| \geq \gamma t$ **do**

$B_i \leftarrow \frac{B_0}{\tau^i}$

 Generate a γt -SLEO $(\mathbf{x}, \mathbf{y}, \mathbf{p})$ for V_i with a budget B_i and random income $\mathcal{I} = \mathbf{U}[1 - \varepsilon, 1]$

 Find a set C_i of size at most $\lceil B_i \rceil$ which:

- (i) contains all candidates a with $y_a = 1$ and
- (ii) does not cover at most $\omega(\gamma, t) \cdot |V_i|$ voters

$V_{i+1} \leftarrow V_i \setminus \{v \in V_i : C_i \text{ } t\text{-covers } v\}$

$C \leftarrow C \cup C_i$

$i \leftarrow i + 1$

return C

end while

The algorithm begins by considering the full set of voters and computing a γt -SLEO with a total budget B_0 . It then converts this outcome into a randomized Lindahl committee. By Claim 4.9, there exists a realization of this randomized committee that fails to cover a certain fraction of the voters. We then restrict our attention to

the uncovered voters and apply the same procedure, but with a reduced budget: $B_1 = B_0/\tau$. This step may again leave a portion of voters uncovered.

This process is repeated iteratively, each time applying the method to the currently uncovered voters and reducing the budget geometrically. The iteration continues until the current budget B_i becomes smaller than γt , at which point the process terminates.

To see the correctness of the algorithm, a crucial observation is that because $|V_i|$ decreases *simultaneously* with B_i for each candidate a , the number of t -covered voters preferring a to C_i also forms a geometric sequence.

CLAIM 4.12. *For each step $i \geq 0$ of Algorithm 4.1, let $S(C_i, V_i)$ denote the set of voters among V_i t -covered by C_i . Then for any $a \notin C_i$, it holds*

$$|v \in S(C_i, V_i) : a \succ_v^t C_i| \leq \frac{\gamma t}{1 - \varepsilon} \cdot [\omega(\gamma, t)\tau]^i \cdot \frac{|V|}{B_0}.$$

CLAIM 4.13. *Given a set V of n voters, $\alpha \leq 1$, integer $t \geq 2, \gamma, \tau \geq 1$ with $\omega(\gamma, t)\tau < 1$, and $\varepsilon \in (0, 1)$, the output of Algorithm 4.1 is a committee which is $(t, \frac{\alpha}{1-\varepsilon})$ -undominated and of size at most $\frac{\gamma\tau}{\tau-1} \cdot \frac{1}{1-\omega(\gamma, t)\tau} \cdot \frac{t}{\alpha} + \log_\tau(\frac{1}{1-\omega(\gamma, t)\tau} \cdot \frac{1}{\alpha})$.*

The detailed proofs of Claim 4.12 and Claim 4.13 can be found in the full version of this paper. Theorem 1.2 then follows from Claim 4.13 and Claim 4.11. In particular, the function $\delta(t)$ in Theorem 1.2 is derived by optimizing the bounds established in Claim 4.13 and Claim 4.11 and then taking their minimum. The undominance ratio α is obtained by letting $\varepsilon \rightarrow 0$.

5 Conclusion This paper systematically studies the problem of selecting undominated committees in social choice, aiming to overcome the Condorcet paradox by allowing for the selection of multiple candidates. Our framework not only generalizes the majority condition to accommodate any arbitrary fraction of voter support, but more importantly, introduces a richer comparison principle: evaluating an outside candidate against the t -th most-preferred member of the selected committee, rather than just the top-ranked one. This generalized notion has practical implications, providing a more nuanced basis for excluding candidates through comparisons with median or quantile-level members of the committee. It also reveals a striking structural property: in the limit, a tight lower bound emerges that reflects the dual role of the threshold parameter, simultaneously capturing the breadth of voter support and the depth of preference comparisons.

Future work includes deriving tighter bounds for finite t and different ways of comparing a set of candidates to a single one. More broadly, our results are based on a novel approach grounded in Lindahl equilibrium with ordinal preferences and its extensions. This technique is general and flexible, making it applicable beyond the specific context of committee selection. For instance, it has been applied to participatory budgeting (Nguyen and Song (2024)), where resources must be allocated fairly across competing projects without relying on cardinal utilities. It may also be applied to public goods provision, where the goal is to ensure fair and efficient outcomes based on individuals' ranked preferences over outcomes. Furthermore, our framework has potential implications in algorithmic fairness, particularly in settings where decisions must be made under ordinal input data, such as in hiring, admissions, or recommendation systems. By bridging concepts from market equilibrium and ordinal social choice, the approach offers a unifying lens for designing fair mechanisms across a wide range of collective decision-making problems.

References

- Romain Bourneuf, Pierre Charbit, and Stéphan Thomassé. 2025. A Dense Neighborhood Lemma: Applications of Partial Concept Classes to Domination and Chromatic Number. *arXiv preprint arXiv:2504.02992* (2025).
- Felix Brandt. 2017. Rolling the dice: Recent results in probabilistic social choice. *Trends in computational social choice* (2017), 3–26.
- Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. 2016. *Handbook of computational social choice*. Cambridge University Press.
- Moses Charikar, Alexandra Lassota, Prasanna Ramakrishnan, Adrian Vetta, and Kangning Wang. 2025a. Six Candidates Suffice to Win a Voter Majority. In *STOC 2025*. <https://arxiv.org/pdf/2411.03390>

- Moses Charikar, Prasanna Ramakrishnan, and Kangning Wang. 2025b. Approximately Dominating Sets in Elections. *arXiv preprint arXiv:2504.20372* (2025).
- Yu Cheng, Zhihao Jiang, Kamesh Munagala, and Kangning Wang. 2019. Group Fairness in Committee Selection. In *EC 19: Proceedings of the 2019 ACM Conference on Economics and Computation* (Phoneix, Arizona, USA). Association for Computing Machinery, 263–278. <https://dl.acm.org/doi/10.1145/3328526.3329577>
- Nicholas de Condorcet. 1785. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie royale.
- Edith Elkind, Jérôme Lang, and Abdallah Saffidine. 2011. Choosing collectively optimal sets of alternatives based on the condorcet criterion. In *Proceedings of the Twenty-Second international joint conference on Artificial Intelligence-Volume Volume One*. 186–191.
- Edith Elkind, Jérôme Lang, and Abdallah Saffidine. 2015. Condorcet winning sets. *Social Choice and Welfare* 44, 3 (2015), 493–517.
- Peter C Fishburn. 1981. Majority committees. *Journal of Economic Theory* 25, 2 (1981), 255–268.
- Duncan K Foley. 1970. Lindahl's Solution and the Core of an Economy with Public Goods. *Econometrica* 38, 1 (1970), 66–72.
- Rajiv Gandhi, Samir Khuller, Parthasarathy Srinivasan, and Aravind Srinivasan. 2006. Dependent rounding and its applications to approximation algorithms. *J. ACM* 53, 3 (2006), 324–360.
- Martin Gardner. 1970. Paradox of nontransitive dice and elusive principle of indifference. *Scientific American* 223, 6 (1970), 110.
- Zhihao Jiang, Kamesh Munagala, and Kangning Wang. 2020. Approximately Stable Committee Selection. In *STOC 2020: Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing* (Chicago, Illinois, USA). Association for Computing Machinery, New York, NY, USA, 463–472. <https://dl.acm.org/doi/abs/10.1145/3357713.3384238>
- Gilbert Laffond, Jean-Francois Laslier, and Michel Le Breton. 1993. The bipartisan set of a tournament game. *Games and Economic Behavior* 5, 1 (1993), 182–201.
- David C McGarvey. 1953. A theorem on the construction of voting paradoxes. *Econometrica: Journal of the Econometric Society* (1953), 608–610.
- Thanh Nguyen and Haoyu Song. 2024. Approximate core of Participatory Budgeting. *Manuscript* (2024).
- John H Smith. 1973. Aggregation of preferences with variable electorate. *Econometrica: Journal of the Econometric Society* (1973), 1027–1041.