STABLE SOCIAL CHOICE

THÀNH NGUYEN

Abstract. Stability—resistance to majority deviations—is central to fair and robust social choice, yet unattainable under many deterministic and probabilistic rules. We establish a possibility result by allowing randomization while restricting deviations to deterministic outcomes. We introduce two stability concepts for probabilistic social choice—ex-ante and ex-post stability—capturing majority protection before and after the lottery's realization. We show that standard probabilistic rules, including random dictatorship, Borda rules, and maximal lotteries, fail to satisfy these criteria. To address this gap, we propose a new mechanism that achieves both forms of stability while preserving efficiency. Our framework further accommodates individualized deviation thresholds, extending its applicability to resource allocation and stable matching problems, including those with externalities.

Date: August 2025.

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

1. Introduction

Social choice functions, which aggregate individual preferences into collective decisions, are a cornerstone of economic theory with wide-ranging applications. They influence everything from elections and public project selection, resource allocation to emerging challenges in AI-human alignment. No mechanism can fully satisfy all individuals due to the heterogeneity of preferences. However, a well-designed social choice mechanism should aggregate broad-based preferences, ensuring that the outcome aligns with the interests of a sufficiently large subset of participants.

This requirement can be formalized through the concept of stability in social choice, which prevents any large enough group from blocking the chosen outcome because they favor an alternative. Equivalently, compared to any alternative option, the final social choice must secure support from a critical mass of participants. A classical and well-studied notion of stability is that the outcome should have majority support, as captured by the concept of a Condorcet winner. However, the famous Condorcet paradox illustrates that such an outcome need not always exist. More strikingly, extensions of the Condorcet example show that for every potential choice, there exists an alternative that is preferred by all but one agent. As the number of voters increases, this implies that no outcome can consistently secure support from any fixed fraction of the electorate.

This instability underscores the challenge that cyclic social preferences pose in designing social choice mechanisms. While often considered an inherent feature of collective decision-making, there is little consensus on how best to address it. For instance, the Condorcet consistency condition guarantees the selection of a Condorcet winner, but only when one exists. Imposing structured preference restrictions, such as single-peaked preferences in single-dimensional domains (Black et al. 1958), mitigates the issue but does not apply to complex applications like machine learning, where preferences rarely conform to such neat structures. Attempts to address these challenges through randomization have also proven insufficient. As shown by Aziz (2015), Zeckhauser (1969), there may not always be a Condorcet winner lottery. Even relaxing the majority condition to hold only in expectation (Brandl et al. 2016, Fishburn 1984, Kreweras 1965) remains inadequate, as such approaches can still produce outcomes susceptible to majority deviations.

This paper proposes an alternative approach to overcoming the impossibility of a stable social choice. We allow randomized social choice, but restrict deviations to be deterministic. Randomization is essential not only for enhancing fairness and stability ex ante but also for supporting and justifying deterministic outcomes ex post. The goal of this paper is to define and analyze the stability of probabilistic choice in both stages: the commitment stage, when the planner announces a lottery, and the implementation stage, when the outcome is realized. In each stage, a majority could block the outcome in favor of an alternative. Our main result constructs a lottery that satisfies stability both ex-ante and ex-post. Moreover, we extend the analysis beyond the majority condition and apply it to several market design problems.

When a social planner commits to a lottery outcome, the majority may prefer a different alternative. This is a standard Condorcet phenomenon that persists even when lottery games are considered (Aziz 2015). The literature adopts a weaker notion of stability: the outcome lottery is preferred by the expected majority (Fishburn 1984). We choose a different approach. Our notion of ex-ante stability addresses this by restricting the set of possible deviating alternatives to

deterministic outcomes. Behaviorally, people typically find it easier to compare a lottery with a deterministic outcome than to compare two different lotteries (Brandl et al. 2023). This makes it easier for voters to reach a consensus on a deterministic outcome they prefer over a lottery. Consequently, ex-ante stability, by limiting deviations to deterministic alternatives, serves as a minimal, yet crucial, requirement for robust voting rules. Surprisingly, we are not aware of any existing probabilistic rule that satisfies our concept of ex ante stability. It is not immediately clear whether such a lottery exists.

Once the lottery is implemented, the realization of the probabilistic outcome may also become unstable, an issue that has been largely overlooked in the literature. Our approach measures ex-post instability by comparing the realized outcome to the committed lottery. Specifically, the lottery is viewed as the expected outcome for voters. If the realization fails to meet the expectations of the majority and there exists an alternative that does, which we prove always occurs, then the realization is considered unstable. Ex-post stability ensures that each realization remains stable.

Ensuring both stability properties strengthens probabilistic social choice and provides a justification for deterministic outcomes. Unlike the classical literature on social choice, which analyzes axioms for deterministic outcomes in isolation, our approach considers an outcome 'justifiable if it emerges as an *ex-post stable* realization of a *ex-ante stable* lottery agreed upon by participants.

Specifically, we define a lottery *ex-ante stable* if no majority strongly prefers any deterministic outcome over it, and *ex-post stable* if the majority prefers weakly every possible realization of the lottery. A voter is said to strongly (weakly) prefer a deterministic outcome over a lottery if the probability of strongly (weakly) favoring the deterministic outcome over the lottery's realized outcomes is at least $1/2$, a threshold that can be generalized to individualized values for each voter.

We first prove that prevailing probabilistic social-choice rules—random dictatorship, Borda, and maximal lotteries—fail to satisfy either ex-ante or ex-post stability. To overcome this, we construct a lottery that meets both stability criteria for every preference profile, while remaining stochastically efficient and Condorcet-consistent. Furthermore, we show that the $\frac{1}{2}$ threshold in our stability condition is tight: lowering it would eliminate the guarantee that a stable lottery against majority deviations exists.

Second, we extend the classic majority condition to a richer notion of fairness in probabilistic social choice, defined through a stability criterion: an outcome is fair if no sufficiently large group of voters can point to an alternative that offers them a sufficiently better outcome. This formulation captures a natural trade-off: smaller coalitions must realize substantial gains to justify deviation, while larger coalitions may be swayed by more modest improvements. Specifically, we generalize the majority condition by introducing an individual parameterized condition based on $\alpha_i \in (0,1)$ for each voter $i \in 1,\dots,n$. A voter i *α-(weakly) prefers* a deterministic alternative over a lottery if the probability that they (weakly) prefer it to the lottery's outcome is at least α_i. We define the weight of the voter i as $\frac{1}{1-\alpha_i}$ and replace the majority condition with any group of voters whose weighted sum exceeds n. The majority condition corresponds to the special case where $\alpha_i = \frac{1}{2}$ for all i. In the general case, voters who require a higher threshold to deviate from the randomized outcome are assigned a higher weight. This result captures a trade-off between a voter's willingness to accept the lottery and the extent to which their preferences influence the outcome. For example,

an $\alpha = \frac{3}{4}$ means that a group of voters will only deviate if for each voter, the alternative ranks above the third quartile relative to the probabilistic choice. This corresponds to a weight of $\frac{1}{1-\alpha} = 4$, implying that just one-fourth of the voters are sufficient to justify deviation. In contrast, $\alpha = \frac{1}{4}$ sets a lower threshold for improvement—the alternative only needs to exceed the first quartile—but requires at least three-fourths of the voters to agree on the deviation.

Application to Market Design. Social choice provides a general framework capable of capturing a wide range of applications, including those in market design. In particular, our use of general thresholds allows the framework to accommodate flexible coalition structures. Social outcomes are treated as abstract objects that can represent general preferences, including those with externalities and complementarities—settings that are often challenging for the existing market design literature.

We illustrate three applications of our solution concept. First, in resource allocation problems where agents' preferences depend not only on their own consumption but also on the consumption of others. Second, in the stable roommate problem, where it is well known that a stable deterministic matching may not exist. Third, in school choice settings with peer effects among students. In each case, our notion of randomized stable social choice provides a natural interpretation as a randomized, stable, and fair solution.

The construction of our probabilistic choice function shares conceptual similarities with the Borda count: each voter assigns points to alternatives, with more preferred alternatives receiving higher points. The lottery is then selected from the set of alternatives with the highest total points. However, unlike traditional Borda count, the points assigned are endogenous, depend on the chosen lottery, and are determined by a scoring function. This creates an endogenous cycle: the lottery is selected based on the aggregate scores, while the scores themselves are assigned based on the chosen lottery. The existence of such a lottery is established using a fixed-point argument.

The flexibility of our approach lies in the design of the scoring function. In particular, we show that the maximal lottery (Fishburn 1984) arises as a special case when the scoring function is linear. By carefully designing the scoring function, we can identify a stable lottery that satisfies our desired stability properties and generalizes to individual probability thresholds.

Related literature. Social choice theory is a fundamental topic in political science, economics, and, more recently, in computer science, with an extensive literature. It is well-established that randomization plays a crucial role in satisfying fairness and addressing strategic issues (see Brandt (2017) for a recent survey). A common approach in the literature is to define the axiomatic properties and characterize the choices that satisfy them or to show their incompatibilities.[1]

A central desideratum in social choice is majority support, or stability. Yet a stable outcome may fail to exist when there is no Condorcet winner. The literature addresses this by either restricting preferences or relaxing majority. The former imposes single-peakedness (Black et al. 1958), but this assumption is often unrealistic in multidimensional settings. The latter includes the Maximal Lottery rule (Brandl et al. 2016, Fishburn 1984, Kreweras 1965), which ensures weak majority support in expectation but fails to guarantee stability when counting actual voters (Brandt 2017). Relatedly,

[1]Strategy-proofness, a stringent condition, can only be satisfied by the random dictatorship mechanism (Gibbard 1977, Satterthwaite 1975). As a result, we do not focus on this issue in this paper.

Zeckhauser (1969) shows that extending preferences to lotteries via von Neumann–Morgenstern utility eliminates the Condorcet winner, even under single-peakedness, while Aziz (2015) shows that stochastic dominance preserves Condorcet winners under single-peaked preferences.

Our approach allows for randomized outcomes, but restricts the alternatives that the majority can propose to deterministic ones. Behaviorally, individuals tend to find it easier to reach a consensus on a deterministic alternative. Thus, ex-ante stability, which limits deviations to deterministic outcomes, serves as a minimal yet essential requirement for robust voting rules. Surprisingly, this notion has not been thoroughly studied, and existing rules such as Borda count, random dictator, and maximal lottery fail to satisfy it. We are the first to show that such an ex-ante stable lottery always exists.

Conceptually, our paper takes a distinct approach to probabilistic social choice by incorporating both ex-ante and ex-post stability conditions. This not only advances the study of probabilistic social choice from an ex-ante perspective but also provides a novel framework for justifying deterministic social outcomes. Specifically, in our framework, a deterministic outcome is considered "justifiable" if it corresponds to a *stable* realization of an *ex-ante stable* probabilistic choice that was previously agreed upon by voters. Our approach provides a dynamic way to reason about deterministic social choice, moving beyond traditional static axiomatic properties.

The paper is organized as follows. Section 2 introduces the notation and illustrates the instability of existing voting rules and the tightness of our stability definition. Section 3 presents our new social choice function and proves its stability and efficiency. Section 4 generalizes to individual probability thresholds and explores applications. Section 5 concludes. The appendix contains omitted proofs and provides a computational algorithm for constructing a stable lottery.

2. Model and Examples

Let $V = \{1, 2, \dots, n\}$ denote the set of voters (agents) and A the set of m alternatives. Each voter $i \in V$ has a preference relation $\succeq_i$ over A, which may include ties. We write $a \succ_i b$ to indicate that voter i strictly prefers a to b, and $a \succeq_i b$ to indicate that i weakly prefers a to b. A *preference profile* is the collection of all voters' preferences, denoted by $\{\succeq_i\}_{i=1}^n$.

An outcome $a \in A$ is called a *Condorcet winner* if, for every other outcome $b \in A \setminus \{a\}$, a strict majority of voters weakly prefer a to b. Equivalently, a beats every other outcome in pairwise majority comparisons.

A lottery over A is a probability distribution $\vec{p} = (p_1, p_2, \dots, p_m)$ such that $\sum_{j=1}^m p_j = 1$ and $p_j \geq 0$ for all $j = 1, 2, \dots, m$. The set of all lotteries is denoted by $\Delta(A)$. For ease of presentation, we also write a lottery as a convex combination of alternatives. For example $\frac{1}{2}a + \frac{1}{2}b$ denotes the lottery with half probability for each outcome a and b.

A probabilistic choice function is a mapping from a profile of preferences to a lottery in $\Delta(A)$.

Definition 2.1. *Given a preference ranking $\succeq$, an alternative $a \in A$ and a lottery $\vec{p} \in \Delta(A)$, we say a is* strictly preferred *to $\vec{p}$, denoted by $a \succ^{(\frac{1}{2})} \vec{p}$, if*

$$\sum_{j \in A : a \succ j} p_j \geq \frac{1}{2}.$$

We say a is weakly preferred *to $\vec{p}$, denoted by $a \succeq^{(\frac{1}{2})} \vec{p}$, if*

$$\sum_{j \in A: a \succeq j} p_j \geq \frac{1}{2}.$$

Thus, a deterministic alternative is *preferred* to the lottery when the probability that it *beats* the lottery's outcome is at least $\frac{1}{2}$; it is *weakly preferred* when the probability that it either beats or ties the lottery's outcome is at least $\frac{1}{2}$. The threshold of $1/2$ is chosen to mirror the standard majority-stability condition. In a later section, we show that this threshold is tight under the majority condition, and we extend the definition to allow arbitrary (possibly voter-specific) probability thresholds for other requirements.

Definition 2.2. *Given a preference profile $\{\succeq_i\}_{i=1}^n$, a lottery $\vec{p} \in \Delta(A)$ is* ex-ante stable *if*

$$|\{i \in V \text{ s.t: } a \succ_i^{(\frac{1}{2})} \vec{p} \ \}| \leq \lfloor \frac{n}{2} \rfloor \text{ for all } a \in A.$$

Ex-ante stability prevents a majority from favoring a specific deterministic alternative over the lottery. In essence, ex-ante stability ensures the lottery is not biased toward any particular alternative, maintaining a balanced outcome that avoids strong opposition from any group of voters prior to realization.

Definition 2.3. *Given a preference profile $\{\succeq_i\}_{i=1}^n$, a lottery $\vec{p} \in \Delta(A)$ is* ex-post stable *if*

$$|\{i \in V \text{ s.t: } a \succeq_i^{(\frac{1}{2})} \vec{p} \ \}| \geq \lceil \frac{n}{2} \rceil \text{ for every } a \in A \text{ with } p_a > 0.$$

The ex-post stability condition ensures that no realization of the lottery will disappoint the majority of voters when compared to the probabilistic outcome. This is justified by the fact, that for any lottery, there exists a deterministic outcome that is weakly preferred by the majority.[2] Thus, ex-post stability is desirable because it prevents strong dissatisfaction among voters once the lottery is realized, minimizing the desire to change the outcome ex-post and enhancing the credibility of the lottery implementation.

Both ex-ante and ex-post stability are desirable properties in social choice. We define a lottery as *stable* if it satisfies both conditions.

Stable lotteries generalize the notion of a Condorcet winner to probabilistic outcomes. A stable lottery ensures that no alternative is strictly preferred by a majority of voters. In particular, if a stable lottery is deterministic—assigning a single outcome $a \in A$ with probability one—then ex-ante stability implies that for any other outcome $b \in A$, at most $\lfloor n/2 \rfloor$ voters strictly prefer b. Hence, a is a Condorcet winner. Formally:

Proposition 2.1. *A deterministic outcome is a stable lottery if and only if it is a Condorcet winner.*

Proof. As argued above, if a deterministic outcome is a stable lottery, then it is a Condorcet winner. The reverse also holds: it is straightforward to verify that a Condorcet winner is stable both ex-ante and ex-post. □

[2]One can show this by linearity of expectation when compare a lottery with itself.

As in the market design literature, stability is defined with respect to the admissible set of blocking coalitions. In our framework, a blocking coalition compares a lottery with deterministic alternatives using the $\succ^{(\frac{1}{2})}$ relation. We first show that existing mechanisms fail to satisfy either ex ante or ex post stability, motivating the introduction of our new social choice function in the next section. We then prove that even slight relaxations of the blocking definition eliminate stable outcomes. More generally, a broad class of "complete" extensions of preferences to lotteries also fails to guarantee stability.

Instability of Existing Mechanisms

Via counter examples, we show that Random Dictatorship, Borda's Rule, and the Maximal Lottery—fail to satisfy either ex-ante or ex-post stability.

Random Dictatorship (RD). The probabilistic social choice function most studied is Random dictatorship (RD), where a voter is selected uniformly at random and their preferred alternative is chosen. The following example show that RD is not stable.

Example 2.1. *There are three voters. The first two voters have $a \succ_{1,2} b$, while the third voter has $b \succ_3 a$. Random Dictatorship (RD) results in lottery $\frac{2}{3}a + \frac{1}{3}b$.*

This lottery is stable ex ante but not stable ex post because only the third voter weakly prefers b to $\frac{2}{3}a + \frac{1}{3}b$.

In this example, the only stable lottery is the one that assigns probability 1 to a.

The following example demonstrates that RD is not exante stable.

Example 2.2. *Consider three voters, $1, 2, 3$, and four alternatives: a, b, c, o. The preferences are as follows:*

$$a \succ_1 o \succ_1 b \succ_1 c, \quad b \succ_2 o \succ_2 c \succ_2 a, \quad c \succ_3 o \succ_3 a \succ_3 b.$$

The Random Dictatorship (RD) mechanism produces the lottery $\tilde{X} = \frac{1}{3}a + \frac{1}{3}b + \frac{1}{3}c$. However, it is not ex-ante stable. All voters strictly prefer the alternative o to this lottery, as the probability

$$\Pr(o \succ_i \tilde{X}) = \frac{2}{3}, \quad \forall i \in \{1, 2, 3\}.$$

In this example, the lottery that assigns probability 1 to o is stable.

Borda's Rule. The traditional Borda rule (or Borda count) is a ranked voting method named after Jean-Charles de Borda, designed to aggregate individual preferences into a collective decision. The process operates as follows: Each voter ranks the alternatives in preference order. Points are assigned based on these rankings. For m alternatives: the top-ranked alternative receives $m - 1$ points, the second-ranked alternative receives $m - 2$ points, and so on, until the last-ranked alternative, which receives 0 points. The points for each alternative are then summed across all voters. The alternative with the highest total score is the winner. Borda rule can be adapted into a probabilistic social choice mechanism. The first one, $Borda_{\max}$ yields all lotteries that randomize over alternatives with maximal Borda score. The second one, $Borda_{\text{pro}}$, involves much more randomness and assigns probabilities to the alternatives that are proportional to their Borda

scores. The following example, taken from Brandt (2017), demonstrates that neither $Borda_{\max}$ nor $Borda_{\text{pro}}$ is exante stable.

EXAMPLE 2.3. *There are 5 voters: 3 of them have the preference $a \succ b \succ c$, and 2 of them have $b \succ c \succ a$. The Borda scores for a, b, and c are 6, 7, and 2, respectively. The $Borda_{max}$ rule returns b, which is not stable because the majority prefers the alternative a. On the other hand, $Borda_{pro}$ returns the lottery $\left(\frac{6}{15}a + \frac{7}{15}b + \frac{2}{15}c\right)$. This is also not stable because for the first 3 voters, the probability that they prefer a is $\frac{9}{15} > \frac{1}{2}$.*

In this example, c is dominated by b, so an efficient lottery will not include c. The unique stable lottery is a with probability 1.

Similarly to example 2.1, it is not hard to construct an example where $Borda_{\text{pro}}$ are not ex post stable.

Maximal Lottery. The Maximal Lottery (ML) is based on pairwise comparisons of alternatives according to voters' preferences. First, construct a matrix where each entry represents the majority margin between two alternatives. For alternatives x and y, the entry $M(x, y)$ is the number of voters who prefer x to y, minus the number of voters who prefer y to x. A maximal lottery is a probability distribution $\vec{p}$ over the alternatives such that $(\vec{p})^T M \geq \vec{0}$.

The existence of maximal lotteries is ensured by von Neumann's Minimax Theorem (Fishburn 1984). The following example demonstrates that ML is not ex-ante nor expost stable.

EXAMPLE 2.4. *There are three alternatives a, b, and c. 20% of the population have preferences $a \succ_1 b \succ_1 c$, 35% have preferences $b \succ_2 c \succ_2 a$, and 45% have preferences $c \succ_3 a \succ_3 b$. Normalize the population to 100, the M matrix is:*

	a	b	c
a	0	30	-60
b	-30	0	10
c	60	-10	0

The resulting maximal lottery is $\tilde{X} = 0.1a + 0.6b + 0.3c$. This lottery is not ex-ante stable because $\Pr(a \succ_1 \tilde{X}) = 0.9$ and $\Pr(a \succ_3 \tilde{X}) = 0.6$; consequently, 65% of the population strictly prefer a over the lottery.

This lottery is also not ex-post stable because c is weakly preferred to $\tilde{X}$ only when the preference is $c \succ a \succ b$, but that is true for only 45% of the population.

In this example, our construction gives the following stable lottery $\frac{1}{6}a + \frac{4}{9}b + \frac{7}{18}c = 0.167a + 0.444b + 0.389c$.

EXAMPLE 2.5. *Continuation of Example 2.4. Here we assume that α, β, and γ represent the fractions of the population with preferences 1, 2, and 3, respectively, and that no single preference holds a majority. The maximal lottery is unique and given by $(1 - 2\gamma)a + (1 - 2\alpha)b + (1 - 2\beta)c$. Our construction gives the lottery $xa + yb + zc$, where x, y, z is the solution of the following linear*

program. (See Section C for its derivation.)

$$(\frac{1}{2}-\alpha)x+\gamma z=\frac{\gamma}{2}$$
$$(\frac{1}{2}-\beta)y+\alpha x=\frac{\alpha}{2}$$
$$(\frac{1}{2}-\gamma)z+\beta y=\frac{\beta}{2}.$$

The table below compares the values of the maximal lottery and our proposed construction.

α,β,γ	***Maximal Lottery***	***Our Stable Solution***
0.25; 0.375; 0.375	*0.25a+ 0.5b+ 0.25c*	*0.3a+ 0.4b+ 0.3c*
0.2; 0.4; 0.4	*0.2a+ 0.6b+ 0.2c*	*0.286a+ 0.428b+ 0.286c*
0.1; 0.45; 0.45	*0.1a+ 0.8b +0.1c*	*0.265a +0.47b+ 0.265c*
...	...	...
0.01; 0.495; 0.495	*0.01a+ 0.98b+0.01c*	*0.251a +0.498b+ 0.251c*

TABLE 1. Comparison of Maximal Lottery and Our Solution for different values of α,β, and γ. None of the maximal lotteries is exante stable.

1/2 THRESHOLD IS TIGHT

We now show that it is not possible to improve on the $\frac{1}{2}$-threshold in the preference relation $\succ^{(\frac{1}{2})}$ while maintaining the existence of ex-ante stable lotteries.

Given the set of alternatives A and $p=(p_j)_{j\in A}$ a lottery. Fix $\varepsilon>0$ and define

$$a \succ_i^* p \quad\Longleftrightarrow\quad \sum_{j\in A:\, a\succ_i j} p_j \;\geq\; \frac{1}{2}-\varepsilon. \tag{1}$$

We construct an example in which **no** lottery exists that is not majority-dominated—under $\succ^*$—by some deterministic alternative.

Consider $2n+1$ voters and $2n+1$ alternatives arranged in the standard Condorcet cycle, so that each voter's ranking is a cyclic shift of the next. Label alternatives $1,\dots,2n+1$, where voter i's most favorite alternative is i. Assume a lottery p satisfies ex-ante stability under $\succ^*$(1).

For alternative 1, the voters numbered $1, 2n{+}1, 2n, \dots, n$—a strict majority—rank 1 ahead of alternatives $2,\dots,n-1$. Preventing this majority from preferring 1 to the lottery requires

$$p_2+p_3+\cdots+p_{n-1} \;<\; \frac{1}{2}-\varepsilon. \tag{2}$$

Applying the same reasoning to alternative 2 yields

$$p_3+p_4+\cdots+p_n \;<\; \frac{1}{2}-\varepsilon, \tag{3}$$

and continuing in this manner produces $2n+1$ distinct inequalities analogous to (2) and (3).

Summing these $n-2$ inequalities gives

$$(n-2)\sum_{j=1}^{2n+1} p_j \;<\; (2n+1)\left(\frac{1}{2}-\varepsilon\right).$$

Because the probabilities sum to 1, this simplifies to

$$\frac{n-2}{2n+1} < \frac{1}{2} - \varepsilon. \tag{4}$$

The left-hand side of (4) converges to $\frac{1}{2}$ as $n \to \infty$. Hence, for any fixed $\varepsilon > 0$, inequality (4) fails once n is large enough. Therefore no lottery can satisfy (1) with a threshold strictly below $\frac{1}{2}$ in sufficiently large electorates.

This implies that the one-half benchmark is tight: any attempt to strengthen majority protection by replacing $\frac{1}{2}$ with $\frac{1}{2} - \varepsilon$ (however small) is infeasible in large Condorcet-cycle electorates.

Impossibility Result for complete relations

Using the preference relation $\succ^{(\frac{1}{2})}$ is a natural choice for comparing lotteries. However, one may ask why we do not instead rely on other standard preference models over lotteries. Two well studied preferences are *pairwise comparison (PC)* or *expected utility (EU)*.

- **Pairwise Comparison (PC):** A lottery x is preferred to a lottery y if the probability that a random draw from x defeats a random draw from y (according to the majority relation over deterministic alternatives) exceeds the probability that a draw from y defeats one from x.

  When this pairwise comparison is restricted to comparing a lottery and a deterministic alternative, it yields the relation $\succeq^{PC1}$, introduced by Brandl et al. (2023), defined as follows: a deterministic alternative a is preferred to a lottery $\mathbf{q}$ under PC1 if

  $$a \succeq^{PC1} \mathbf{q} \quad \text{iff} \quad \sum_{j \in A: a \succ j} q_j \geq \sum_{j \in A: j \succ a} q_j.$$

  Equivalently, the PC1 comparison between a deterministic alternative a and a lottery $\mathbf{p}$ can be rewritten as:

  $$a \succ^{PC1} \mathbf{p} \quad \text{iff} \quad \sum_{j \in A: a \succ j} p_j + \frac{1}{2} \sum_{k \in A: a \sim k} p_k \geq \frac{1}{2}.$$

  It is straightforward to verify that PC1 implies weak preference ($\succeq^{(\frac{1}{2})}$), while strong preference ($\succ^{(\frac{1}{2})}$) implies PC1.
- **Expected Utility (EU):** Each voter has a von Neumann–Morgenstern utility function over alternatives. A lottery x is preferred to a lottery y if the expected utility of x exceeds that of y.

Both PC and EU preferences induce complete preference relations and satisfy a regularity condition (defined below). However, we show that these relations are too strong to guarantee the existence of a Condorcet-winning lottery—even when comparisons are restricted to deterministic alternatives. In other word, no stable lottery may exist under these preferences.

Definition 2.4. *A transitive relation $\trianglelefteq$ over lotteries $\Delta(A)$ is called* regular *if it satisfies the following properties for all $a, b \in A$ such that b is strictly prefered to a, that is, $a \triangleleft b$:*

- ***Completeness:*** *For every pair of lotteries $x, y \in \Delta(A)$, either $x \trianglelefteq y$, $y \trianglelefteq x$, or $x = y$.*

- ***Monotonicity under replacement:*** *If a lottery $x \in \Delta(A)$ assigns positive probability to a, then replacing a with b (while preserving the rest of the lottery) yields a strictly preferred lottery.*
- ***Continuity of improvement:*** *If a lottery x assigns probability at least $\delta > 0$ to a, then there exists $\epsilon_\delta > 0$ such that, after replacing a with b in x to obtain a new lottery y, any lottery y' that is within ϵ_δ distance of y (in the L_2 norm) is also strictly preferred to x.*

Definition 2.5. *A lottery $x \in \Delta(A)$ is a* Condorcet winner *under $\trianglelefteq$ if no other lottery is preferred to it by a majority of voters. A lottery $x \in \Delta(A)$ is ex-ante stable under $\trianglelefteq$ if no deterministic alternative (i.e., pure outcome in A) is preferred to it by a majority of voters.*

With this, we obtain the following result; the proof is provided in the appendix.

Proposition 2.2. *If $\trianglelefteq$ is regular, then there exists a voting instance where no stable lottery exists. Consequently, no Condorcet winner exists either.*

Both pairwise comparison relations over lotteries and expected utility preferences are regular. Consequently, a stable lottery—such as a Condorcet winner—does not necessarily exist under these preference models.

3. Main Result

This section constructs a stable lottery. We first introduce a general framework for defining the probabilistic choice function and show that, under a specific scoring function, the induced lottery is stable. We then establish existence via a fixed-point theorem and conclude by proving the efficiency of stable lotteries.

General Construction. Our construction is similar to the $Borda_{max}$ rule: Each voter assigns points to each alternative, where lower ranked alternatives receive fewer points. The points are then summed across all voters and a lottery is chosen over the alternatives with the maximal total score. However, there is one key difference: the way voters distribute points is endogenous. It depends on the outcome lottery itself and an exogenous scoring function that we will specify later.

- **Scoring Function**: Let $f : [0,1] \to [0,1]$ be a monotone, integrable function, which we refer to as the *scoring function*. We do not require f to be strictly monotonic or continuous, but we assume that the integral $\int_I f(t)\,dt$ exists for any subinterval $I \subset [0,1]$.
- **Individual Scores**: For a lottery $\vec{p}$ and agent i, we define $\Pi_i^j(f, \vec{p})$ as the *individual score* for agent i over each alternative $j \in A$ as follows.

  Each alternative j is represented by an interval of length p_j. In the case of ties among alternatives, they are merged into a single interval with a total length equal to the sum of their probabilities. These intervals are arranged consecutively within the interval $[0,1]$, sorted in increasing order of i's preferences. Let I_j denote the interval corresponding to alternative $j \in A$. We define $\Pi_i^j(f, \vec{p})$ as the average value of $f(.)$ in the interval I_j. Namely, if $p_j > 0$, then

$$\Pi_i^j(f, \vec{p}) := \frac{1}{p_j} \int_{I_j} f(t)dt.$$

When $p_j = 0$, I_j corresponds to a single point within $[0, 1]$, in which case $\Pi_i^j(f, \vec{p})$ is defined as the value of $f(.)$ at that point. (See Figure 1 for an illustration.)

- **Aggregate Scores** each alternative j's aggregate score is defined as

$$\Pi^j(f, \vec{p}) := \sum_{i \in V} \Pi_i^j(f, \vec{p})$$

Definition 3.1. *Given a scoring function $f(.)$, a lottery $\vec{p}$ is said to be **equilibrium lottery** associated with f if the total score of each alternative j in the support of $\vec{p}$ achieves the maximum. Mathematically, if $p_j > 0$, then*

$$\Pi^j(f, \vec{p}) \geq \Pi^{j'}(f, \vec{p}) \ \ \forall j' \in A.$$

We next show that an equilibrium lottery satisfies the Condorcet compatibility criterion.

Proposition 3.1. *Assume k is the Condorcet winner, then the lottery that assigns probability 1 to k is an equilibrium lottery for all scoring functions $f(.)$.*

Proof. Let k be the Condorcet winner, which means that k is preferred over all other alternatives by a majority of voters. The lottery that assigns probability 1 to k is an equilibrium lottery for any scoring function f. To prove this, we only need to show that, under this lottery, the total score of every other alternative is at most $n/2$.

If k is strictly preferred to some $a \in A$ by voter i, then the score for a assigned by voter i will be 0 due to the scoring rule. Since k is preferred over all other alternatives by a majority of voters, the total score for any other alternative is at most $n/2$.

□

Equilibrium lotteries of specific scoring functions. We show the property of equilibrium lottery associated with two specific scoring functions. First, if the scoring rule is a step function, the equilibrium lottery is stable. Second, if the scoring rule is linear, it coincides with the maximal lottery.

To prove these results, we use the following properties of an equilibrium lottery.

Lemma 3.1. *In an equilibrium lottery, the aggregate score of each alternative is at most $n \int_0^1 f(t)dt$, and each alternative in the support of the lottery attains exactly this score.*

Proof. In an equilibrium lottery, the expected score assigned by each voter i is

$$\sum_{a \in A} \Pi_i^a(f, \vec{p}) \cdot p_a = \int_0^1 f(t)dt$$

Therefore, the total aggregate avarage score across all voters is $n \int_0^1 f(t)dt$. Since the lottery selects alternatives that maximize the aggregate score, each alternative in its support must have an aggregate score equal to $n \int_0^1 f(t)dt$, while all other alternatives have aggregate scores at most $n \int_0^1 f(t)dt$. □

THEOREM 1. *Any equilibrium lottery associated with the step function $h(.)$ is ex ante and ex post stable. Here, $h(.)$ is defined by*

$$h(x) = \begin{cases} 0 & \text{if } 0 \leq x < 1/2, \\ 1 & \text{if } x \geq 1/2. \end{cases} \tag{5}$$

Proof. Consider an alternative a, and let $\vec{p}$ be an equilibrium lottery. Suppose that for some voter i, we have $a \succ_i^{1/2} \vec{p}$. Then, the interval corresponding to a lies entirely to the right of the midpoint $1/2$, which implies that the individual score assigned to a by voter i is 1. Consequently, the aggregate score of a is at most $n/2$, so there can be at most $n/2$ voters who rank a above $\vec{p}$ in the $\succ_i^{1/2}$ relation.

To see that $\vec{p}$ is ex-post stable, consider an alternative a in the support of $\vec{p}$. Suppose that voter i does not weakly prefer a to $\vec{p}$, i.e., $a \not\succeq_i^{1/2} \vec{p}$. Then, the interval corresponding to a lies entirely within the subinterval $[0, 1/2]$, and because $p_a > 0$, the individual score assigned to a is 0. Hence, the aggregate score of a is at most $n/2$, and so at most $n/2$ voters do not weakly prefer a to $\vec{p}$.

□

THEOREM 2. *An equilibrium lottery associated with $f(x) = x$ is a Maximal lottery.*

Proof. Given the scoring function $f(x) = x$, for a lottery p, the scoring rule for alternative $a \in A$ for voter i is exactly

$$\gamma_i^a = \sum_{j \in A: a \succ_i j} p_j + \frac{1}{2} p_a.$$

Since the equilibrium lottery maximizes the expected total score, the total score for any alternative is at most $n/2$. This shows that for all $a \in A$,

$$\sum_{i \in V} \Big(\sum_{j \in A: a \succ_i j} p_j + \frac{1}{2} \sum_{b: b \trianglerighteq a} p_b \Big) \leq n/2.$$

This is equivalent to

$$\sum_{i \in V} \Big(\sum_{j \in A: a \succ_i j} p_j - \sum_{k \in A: k \succ_i a} p_k \Big) \leq 0.$$

This is exactly the property of a maximal lottery.

□

EXISTENCE

We next establish the existence of an equilibrium lottery for any continuous scoring function, as well as for the step scoring function. To do so, we define an appropriate mapping and employ a fixed-point theorem to rigorously demonstrate the existence of such an equilibrium.

Fixed-Point Correspondence:
We define a correspondence $\Phi : \Delta(A) \times [0,1]^{nm} \rightrightarrows \Delta(A) \times [0,1]^{nm}$ that transforms a tuple $(\vec{p}, \vec{\gamma}_1, \ldots, \vec{\gamma}_n)$, where $\vec{p}$ is a lottery and $\vec{\gamma}_i$ are the individual score vectors for each agent i, into a

new tuple $(P, \vec{\theta}_1, \ldots, \vec{\theta}_n)$. The transformation is defined as follows:

$$\theta_i^j := \Pi_i^j(f, \vec{p}),$$

$$P := \arg\max_{\vec{q} \in \Delta(A)} (\sum_{i=1}^{n} \vec{\gamma}_i)^T \cdot \vec{q}.$$

We begin with the following result.

LEMMA 3.2. *For any continuous scoring function $f(.)$, Φ has a fixed point. Moreover a fixed point corresponds to an equilibrium lottery.*

Proof. First, the scoring rule is a continuous function because the scoring function f is continuous. Second, the selection of the lottery is based on the argument of the maximum of a linear function, so it is also a continuous mapping and the set of optimal lotteries is a convex set. Therefore, we can apply Kakutani's fixed-point theorem, which guarantees the existence of a fixed point.

Let $(\vec{p}, \vec{\gamma}_1, \ldots, \vec{\gamma}_n)$ be a fixed point of the mapping. By construction, each $\vec{\gamma}_i$ is the scoring vector of voter i, and $\vec{p}$ is the lottery that maximizes the aggregate expected score. □

Next, we show the existence of an equilibrium lottery for the step function. To do so, we construct a family of functions, denoted by g^ϵ, which depend on a sufficiently small parameter ϵ (here, we assume $\epsilon < 0.1$). The functions g^ϵ are required to satisfy the following properties. (Figure 1 illustrates $g^\epsilon(.)$)

- continuous, strictly increasing in [0,1]
- $g^\epsilon(1/2) = 1 - \epsilon$
- $\int_0^1 g^\epsilon(t)dt = 1/2$

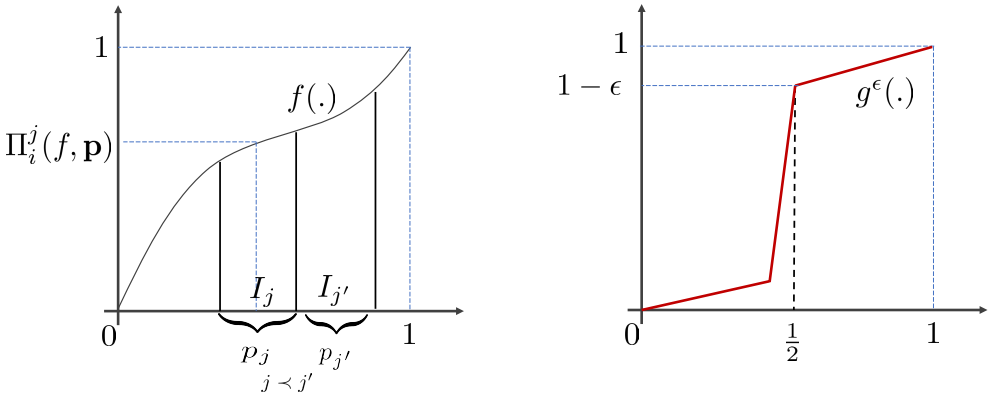


FIGURE 1. Scoring functions

Taking the limit as $\epsilon \to 0$, boundedness and compactness ensure that any sequence of lotteries contains a convergent subsequence. Hence a limit lottery always exists.

DEFINITION 3.2. *$\vec{p}$ is called a limit equilibrium lottery associated with $\{g^\epsilon\}_{\epsilon \to 0}$ if $\vec{p}$ is a converging lottery of a subsequence of equilibrium lotteries for $\{g^\epsilon\}$ as ϵ approaches 0.*

We show next that the limit equilibrium lottery is itself an equilibrium lottery for the step-scoring rule.

THEOREM 3. *The limit equilibrium associated with $\{g^\epsilon\}_{\epsilon\to 0}$ corresponds to equilibrium lottery associated with step function $h(.)$ defined in (5).*

Proof. Let $\vec{p}$ be a limit equilibrium lottery, and let γ_i^a be voter i's limit score for alternative $a \in A$.

As $\epsilon \to 0$, the scoring rules g^ϵ converge to the step rule, so the limit scores γ_i^a coincide with the scores based on the step function. For every $\epsilon > 0$, an equilibrium lottery for g^ϵ assigns each alternative an aggregate score no greater than $n/2$, with equality for every alternative in its support. This property persists in the limit, so $\vec{p}$ remains supported only on alternatives that attain the maximum aggregate score. Consequently, $\vec{p}$ is an equilibrium lottery for the step function scoring rule.

□

EFFICIENCY

We begin with two notions of efficiency. The first is ex-post efficiency.

DEFINITION 3.3 (Ex post Efficiency). *Given a preference profile $\{\succeq_i\}_{i=1}^n$, a lottery $\vec{p} \in \Delta(A)$ is ex post efficient if whenever $a \succeq_i b$ for all $i \in V$ and $a \succ_j b$ for at least one voter $j \in V$, then $p_b = 0$.*

The second notion of efficiency is stronger. It is based on partial orders of stochastic dominance over lotteries. In particular, given a preference $\succeq$, and two lotteries $\vec{p}, \vec{q} \in \Delta(A)$, we say $\vec{p}$ stochastically dominates $\vec{q}$, and denoted by $\vec{p} \succeq^{sd} \vec{q}$, if for all $a \in A$

$$\sum_{k \in A: k \succeq a} p_k \geq \sum_{k \in A: k \succeq a} q_k.$$

If at least one of these inequalities is strict, we say that $\vec{p}$ strictly stochastically dominates $\vec{q}$.

DEFINITION 3.4 (Stochastic Dominance (SD) Efficiency). *Given a preference profile $\{\succeq_i\}_{i=1}^n$, a lottery $\vec{p} \in \Delta(A)$ is SD-efficient if, for every other lottery $\vec{p} \neq \vec{q} \in \Delta(A)$, there exists a voter $i \in V$ such that, according to their preference, $\vec{p}$ strictly dominates $\vec{q}$.*

It is well known that SD efficiency implies ex-post efficiency (see for example Brandt (2017)). Our next theorem shows that any equilibrium lottery satisfies the stronger notion of efficiency.

THEOREM 4. *For any strictly increasing scoring function f, an equilibrium lottery associated with f is SD efficient. Moreover, given an equilibrium lottery associated with the step scoring function, we can construct an equilibrium lottery that has identical scoring for each voter and each candidate and is SD efficient.*

The full proof is provided in the Appendix. To see the idea, first, suppose $\vec{p}$ is an equilibrium lottery under a strictly increasing scoring function f. If there were a lottery $\vec{q}$ that SD-dominates $\vec{p}$, then $\vec{q}$ would yield a strictly higher total expected score, contradicting the fact that $\vec{p}$ maximizes the average score. Hence, every such equilibrium lottery is SD-efficient.

Now consider the case of a step function h. Since h is not strictly increasing, some equilibrium lotteries may fail to be SD-efficient. To recover efficiency, we proceed in two steps.

(1) If $\vec{q}$ SD-dominates an equilibrium lottery $\vec{p}$, then $\vec{q}$ is also an equilibrium lottery and yields the same individual scores as $\vec{p}$.
(2) Every equilibrium lottery with these scores can be described by a linear program.

Thus, obtaining an SD-efficient equilibrium lottery amounts to solving this linear program while maximizing any aggregate utility consistent with the voters' ordinal preferences.

Moreover, in appendix C, we show how an equilibrium lottery with the step scoring function can be obtained via a set of linear programs. Although the number of programs may be exponential in the worst case, the approach is algorithmically practical when the number of alternatives is a small constant.

4. General Threshold and Applications

In this section, we extend blocking conditions beyond majority rule and apply them to market design applications.

We introduce an individual parameter for each voter that serves as a threshold for comparing a deterministic outcome with a lottery. Specifically, each voter i is associated with a parameter $0 < \alpha_i < 1$, and we define the weight of voter i as $w_i := \frac{1}{1-\alpha_i}$. For any set of voters $S \subseteq V$, the total weight is given by $w_S := \sum_{i \in S} w_i$.

Definition 4.1. *We say that $a \in A$ is strictly α_i-preferred to $\vec{p}$ and write $a \succ_i^{\alpha_i} \vec{p}$ if*

$$\sum_{j \in A: a \succ_i j} p_j \geq \alpha_i.$$

We say that a is weakly α_i-preferred to $\vec{p}$, denoted as $a \succeq_i^{\alpha_i} \vec{p}$, if

$$\sum_{j \in A: a \succeq_i j} p_j \geq \alpha_i.$$

Definition 4.2. *Given a preference profile $\{\succeq_i\}_{i=1}^n$ and individual threshold vector $\vec{\alpha}$, $\vec{p} \in \Delta(A)$ is ex-ante stable with respect to $\vec{\alpha}$ if for all $a \in A$*

$$\sum_{i: a \succ_i^{\alpha_i} \vec{p}} \frac{1}{1-\alpha_i} \leq n.$$

$\vec{p}$ is ex-post stable with respect to $\vec{\alpha}$ if

$$\sum_{i: a \succeq_i^{\alpha_i} \vec{p}} \frac{1}{1-\alpha_i} \geq n \text{ for all } a : p_a > 0.$$

Notice that if $\vec{p}$ is stable ex ante with respect to α, then it remains stable ex ante for any larger α. Similarly, if $\vec{p}$ is ex post stable with respect to α, then it remains ex post stable for any smaller α.

Our main theorem is the following.

Theorem 5. *For any $\vec{\alpha} \in (0,1)^n$, there exists a lottery that is both ex ante and ex post stable with respect to $\vec{\alpha}$. Moreover, the stable lottery is SD efficient.*

Theorem 5 illustrates the trade-off between the degree of "pickiness" in the acceptance by voters of a social outcome and the extent to which their preferences are considered. A lower α_i means

that a voter is willing to deviate even when the probability of an alternative being better is low. This results in a lower weight of $\frac{1}{1-\alpha_i}$.

For example, when $\alpha = 0.1$, a group can block a lottery only if at least 90% of the population ranks a deterministic outcome above their 10th percentile. In contrast, with $\alpha = 0.9$, it is sufficient for just 10% of the population to block the lottery.

The proof of Theorem 5 is analogous to Theorem 3, and the formal proof is provided in the appendix. Here, we briefly describe the construction of the scoring function used in the proof. Specifically, the scoring functions need to be modified to capture individual thresholds. Let $h_i^\epsilon : [0,1] \to [0,1]$ satisfy the following conditions:

- h_i^ϵ is continuous, strictly increasing
- $h_i^\epsilon(\alpha_i) = 1 - \epsilon$
- $\int_0^t h_i^\epsilon(t)d(t) = 1 - \alpha_i$.

We construct an equilibrium lottery where the scoring rule for each voter i is based on the scoring function $\frac{1}{1-\alpha_i} h_i^\epsilon(\cdot)$. It is straightforward to see that the fixed-point argument still holds, and therefore we obtain the existence of an equilibrium lottery. With this, we maintain the following property: at the equilibrium lottery, the average score for each voter is 1, and the total average score will be n.

We define the limit lottery in the same way, by taking a subsequence as $\epsilon \to 0$ such that the corresponding equilibrium lotteries converge. In the limit, the lottery corresponds to an equilibrium of a step function

$$g_{\alpha_i}(x) = \begin{cases} 0 & \text{for } 0 \le x < \alpha_i, \\ \frac{1}{1-\alpha_i} & \text{for } x \ge 1/2. \end{cases} \tag{6}$$

Thus, an alternative that exceeds the α_i threshold will have a score of $\frac{1}{1-\alpha_i}$, while any alternative that falls below the threshold will have a score of 0. The rest of the proofs follow the same logic as in Theorem 3 (See the Appendix).

Applications

Next, we discuss the applications of Theorem 5 to market design problems.

4.1. **Fair Allocations.** Consider an application in fair resource allocation. Suppose there are m goods and n agents, each with a preference ranking over the bundles they may receive. This setting is a special case of our framework, where each outcome corresponds to a partition of the goods into n disjoint sets $(S_1, \ldots, S_n)$, with S_i representing the bundle assigned to agent i. In the standard literature, each agent has weak preferences over joint allocations and is concerned only with their own assigned bundle.

Our approach is more general: it allows us to capture externalities, where an agent's utility can depend on other agents' allocations, since we assume that preferences can be defined over the entire allocation $(S_1, \ldots, S_n)$.

Even without externalities, our mechanism provides new insights. In particular, when $\alpha_i = \alpha > 1 - \frac{1}{n}$, it follows that for all i, $\frac{1}{1-\alpha_i} > n$, which implies that no blocking coalition of any size can form. Consequently, no agent strictly prefers any deterministic outcome over the lottery. This ensures that each agent receives their top-ranked outcome with probability at least $\frac{1}{n}$.

Notice that the random serial dictatorship (RSD) mechanism guarantees this property. However, it is not efficient in the stochastic dominance order. Below is the example from Bogomolnaia and Moulin (2001) to show it.

EXAMPLE 4.1. *Consider the following example in a unit-demand setting from Bogomolnaia and Moulin (2001). There are four agents,* $\{1, 2, 3, 4\}$*, and four goods,* $\{a, b, c, d\}$*, with each agent demanding exactly one good.*

Agents 1 *and* 2 *have preferences*

$$\{a\} \succ_{1,2} \{b\} \succ_{1,2} \{c\} \succ_{1,2} \{d\},$$

while agents 3 *and* 4 *have preferences*

$$\{b\} \succ_{3,4} \{a\} \succ_{3,4} \{d\} \succ_{3,4} \{c\}.$$

Under RSD, agents 1 *and* 2 *each receive (ex ante) a* $\frac{5}{12}$ *share of goods* a *and* c*, and a* $\frac{1}{12}$ *share of goods* b *and* d*. Symmetrically, agents* 3 *and* 4 *each receive a* $\frac{5}{12}$ *share of goods* b *and* d*, and a* $\frac{1}{12}$ *share of goods* a *and* c*. This allocation is not efficient under first-order stochastic dominance, since, for example, agent 1 would prefer to exchange their probability share of good* b *with agent 3 for a higher probability of receiving good* a*.*

By contrast, our mechanism guarantees efficiency. There are multiple efficient equilibrium lotteries, one of which randomizes between the two assignments

$$(1{-}a,\ 2{-}c,\ 3{-}b,\ 4{-}d) \quad \textit{and} \quad (1{-}c,\ 2{-}a,\ 3{-}d,\ 4{-}b),$$

each with probability $\frac{1}{2}$*. Thus, in this example, agents* 1 *and* 2 *each receive (ex ante) an equal* $\frac{1}{2}$ *share of goods* a *and* c*, while agents* 3 *and* 4 *each receive a* $\frac{1}{2}$ *share of goods* b *and* d*.*

4.2. **Stable Roommates.** The stable roommate problem, introduced in Gale and Shapley (1962) and is well studied. It generalizes the stable marriage problem and is well known for the fact that a stable solution need not exist. Irving (1985) provides an algorithm to determine whether such a solution exists. However, when no stable solution exists, the natural question is what notion of stability should apply. We show that stable lotteries offer a randomized solution concept for this setting.

Consider a set $N = \{1, 2, \ldots, n\}$ of agents. Each agent $i \in N$ has a preference relation $\succ_i$ over potential roommates $j \in N \setminus \{i\}$, together with the option of remaining unmatched. A (deterministic) roommate matching is a partition of N into disjoint pairs and, if n is odd, possibly one unmatched agent.

A pair of agents $\{i, j\}$ is said to *block* a deterministic matching M if both strictly prefer being matched with each other to their assigned outcome in M, i.e., $j \succ_i M(i) \quad \text{and} \quad i \succ_j M(j)$, where $M(i)$ denotes the roommate of i in matching M (or the unmatched option if i is unmatched). A deterministic matching M is *stable* if no pair of agents blocks it.

EXAMPLE 4.2. *To illustrate the nonexistence of a stable deterministic matching, consider three agents a, b, and c with cyclic preferences: agent a prefers b to c, agent b prefers c to a, and agent c prefers a to b. Every agent prefers a match over staying unmatched. In any deterministic matching, one agent must remain unmatched. For example, if a is matched with b, then c is left unmatched and would rather match with b, while b also prefers c to a; hence, the pair (b, c) would block this outcome. A similar blocking argument applies no matter how the matching is arranged. Even the outcome where all agents remain unmatched is unstable, since any two agents would prefer to match with each other. Thus, in this simple three-agent setting, no stable deterministic roommate matching exists.*

The existing literature (Biró and Fleiner 2016, Teo and Sethuraman 1998) studies fractional stable roommate matchings, where each agent is assigned fractional shares of matches with others, with total probability bounded by one. Such fractional solutions, however, may not correspond to a distribution over deterministic matchings. For example, in the three-agent instance above, assigning the matchings ab, bc, and ca each with probability $1/2$ is a valid fractional matching but cannot be realized as a distribution over deterministic matchings.

Rounding these fractional solutions to obtain near-feasible implementations, as discussed in Nguyen and Vohra (2018, 2024), may require matching more than two agents together in the roommate setting, which may not be feasible in practice.

To distinguish our work from this literature, we introduce the solution concept of a *randomized matching*. Specifically, a randomized matching is a probability distribution over deterministic, feasible roommate matchings. We denote a randomized matching by $\tilde{M}$.

A pair of agents $\{i, j\}$ is said to α-*block* a fractional matching $\tilde{M}$ at threshold $\alpha \in [0, 1]$ if

$$\Pr_{M\sim\tilde{M}}\left(j \succ_i M(i)\right) \geq \alpha \quad \text{and} \quad \Pr_{M\sim\tilde{M}}\left(i \succ_j M(j)\right) \geq \alpha.$$

With this, we have the following definition of a stable randomized roommate matching.

DEFINITION 4.3 (Stable Randomized Roommate Matching). *A probability distribution over deterministic roommate matchings is* α-stable *if there is no α-blocking pair of agents.*

Applying Theorem 5, where the set of alternatives is the set of deterministic matchings and with $\alpha = 1 - \frac{2}{n} + \epsilon$, we obtain:

COROLLARY 4.1. *Given any roommate instance, for every $\epsilon > 0$, there exists a $1 - \frac{2}{n} + \epsilon$-stable randomized roommate matching.*

In plain terms, this corollary means that in any roommate instance, no pair of agents can both prefer to be matched with each other over their assigned partners with probability greater than $1 - \frac{2}{n}$. Equivalently, each agent is matched to partners that they like almost as much as any other possible partner, up to a margin of roughly $2/n$.

We provide an example with $n = 3$ agents, showing that a $(\frac{1}{3} + \varepsilon)$-stable randomized roommate matching exists, and prove that this bound is tight.

EXAMPLE 4.3. ***Continuation of Example 4.2.*** *Consider agents a, b, and c with cyclic preferences. First, note that the randomized matching which pairs any two agents with probability $1/3$ is $(1/3+\varepsilon)$-stable.*

Now consider any other randomized matching. We will show that, in every case, there exists a blocking coalition of size at least $1/3$.

Let the probabilities of the matching be (p_{ab}, p_{bc}, p_{ca}), where $p_{ab}+p_{bc}+p_{ca}=1$. (The argument also holds if the sum is strictly less than 1.)

By symmetry, it suffices to consider two representative cases:

Case 1.

$$p_{ab}=\tfrac{1}{3}+\varepsilon_1+\varepsilon_2, \quad p_{bc}=\tfrac{1}{3}-\varepsilon_1, \quad p_{ca}=\tfrac{1}{3}-\varepsilon_2,$$

with $\varepsilon_1, \varepsilon_2 \geq 0$. Here, the pair (b,c) would strictly prefer to match with each other rather than follow the lottery, forming a blocking coalition at threshold $\alpha = 1/3$.

Case 2.

$$p_{ab}=\tfrac{1}{3}-\varepsilon_1-\varepsilon_2, \quad p_{bc}=\tfrac{1}{3}+\varepsilon_1, \quad p_{ca}=\tfrac{1}{3}+\varepsilon_2,$$

with $\varepsilon_1, \varepsilon_2 \geq 0$. In this scenario, the pair (a,b) strictly prefers to match together, forming a blocking coalition at threshold $\alpha = 1/3$.

4.3. **Stable Matching with Peer Effects.** We consider a setting with m schools and $n = km$ students, where each school is to be assigned exactly k students. This assumption is made for simplicity and without loss of generality, as dummy students can be added to satisfy the capacity constraints if needed.

Each school has preferences over the set of students assigned to it, while students have preferences both over which school they attend and over the peers they are matched with, capturing peer effects. School preferences are completely general, allowing for complementarities among students.

Under such general preferences, a stable matching may not exist. We therefore apply Theorem 5 to construct a randomized stable matching.

We use notation similar to above. For a matching M that assigns k students to each school, let $M(i)$ denote the set of students assigned to school i, and for each student j, let $M(j)$ denote the school j is assigned to together with the set of students assigned to the same school.

Let $\tilde{M}$ denote a randomized matching, that is, a probability distribution over deterministic matchings M.

DEFINITION 4.4. *A randomized matching $\tilde{M}$ is* (β,γ)-stable *if there does not exist any school i and group of k students J such that*

$$\Pr_{M\sim\tilde{M}}\big(J \succ_i M(i)\big) \geq \beta,$$
$$\Pr_{M\sim\tilde{M}}\big((i,J) \succ_j M(j)\big) \geq \gamma \quad \textit{for all } j \in J.$$

In other words, no school and group of students prefer to deviate from $\tilde{M}$ with probability at least β and γ, respectively.

In other words, the school i and a group of k students J is a block if school i prefers the group of k students J to its assignment in the random matching with probability at least β, and each

student $j \in J$ prefers being assigned to school i with the group J to their assignment in the random matching with probability at least γ.

Applying Theorem 5, we take the set of alternatives to be all possible matchings that assign k students to each school, with thresholds $\alpha_i = \beta$ for each school and $\alpha_j = \gamma$ for each student.

COROLLARY 4.2. *There exists a (β, γ)-stable randomized matching if*

$$\frac{1}{1-\beta} + \frac{k}{1-\gamma} > (k+1)m.$$

This result allows for a trade-off between β and γ, reflecting the incentives of schools and students in forming a blocking coalition. For example, to maintain stronger incentives for students, one can decrease γ and increase β, which permits blocking coalitions where students can deviate under a lower probability threshold.

In the special case $\beta = \gamma$, the condition reduces to

$$\beta = \gamma > 1 - \frac{1}{m}.$$

An important feature of this condition is that the parameters β and γ depend only on the number of schools, not on the number of students. This implies that stability is preserved even as the market grows large in the number of students.

For example, with 10 schools, there exists a randomized matching such that no school and group of students can jointly improve their assignment beyond the 90^{th} percentile of their respective randomized outcomes, even under very general preferences and an arbitrarily large number of students.

5. CONCLUSION

This paper introduces a new notion of stability in probabilistic social choice, addressing both ex-ante and ex-post perspectives. We show that none of the existing rules satisfy this stability criterion and propose a new rule that is both efficient and stable.

The flexibility of our framework allows modeling individual preferences over both probabilistic and deterministic outcomes, and it has several applications in market design, particularly in settings with complex preferences involving externalities and complementarities.

Several directions remain open for future research, including the axiomatic characterization of stable social choice, computational challenges, practical implementation, empirical evaluation, and comparison with alternative methods.

REFERENCES

Aziz, H. (2015). Condorcet's paradox and the median voter theorem for randomized social choice. *Economics Bulletin 35*(1), 745–749.

Biró, P. and T. Fleiner (2016). Fractional solutions for capacitated ntu-games, with applications to stable matchings. *Discrete Optimization 22*, 241–254.

Black, D. et al. (1958). The theory of committees and elections.

Bogomolnaia, A. and H. Moulin (2001). A new solution to the random assignment problem. *Journal of Economic Theory 100*(2), 295–328.

Brandl, F., F. Brandt, and H. G. Seedig (2016). Consistent probabilistic social choice. *Econometrica 84*(5), 1839–1880.

Brandl, F., P. Lederer, and S. Warut (2023). Incentives in social decision schemes with pairwise comparison preferences. *Games and Economic Behavior 142*, 266–291.

Brandt, F. (2017). Rolling the dice: Recent results in probabilistic social choice. *Trends in computational social choice*, 3–26.

Fishburn, P. C. (1984). Probabilistic social choice based on simple voting comparisons. *The Review of Economic Studies 51*(4), 683–692.

Gale, D. and L. S. Shapley (1962). College admissions and the stability of marriage. *The American mathematical monthly 69*(1), 9–15.

Gibbard, A. (1977). Manipulation of schemes that mix voting with chance. *Econometrica: Journal of the Econometric Society*, 665–681.

Irving, R. W. (1985). An efficient algorithm for the "stable roommates" problem. *Journal of Algorithms 6*(4), 577–595.

Kreweras, G. (1965). Aggregation of preference orderings. In *Mathematics and Social Sciences I: Proceedings of the seminars of Menthon-Saint-Bernard, France (1–27 July 1960) and of Gösing, Austria (3–27 July 1962)*, pp. 73–79.

Nguyen, T. and R. Vohra (2018). Near-feasible stable matchings with couples. *American Economic Review 108*(11), 3154–3169.

Nguyen, T. and R. Vohra (2024). (near-) substitute preferences and equilibria with indivisibilities. *Journal of Political Economy 132*(12), 4122–4154.

Satterthwaite, M. A. (1975). Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of economic theory 10*(2), 187–217.

Teo, C.-P. and J. Sethuraman (1998). The geometry of fractional stable matchings and its applications. *Mathematics of Operations Research 23*(4), 874–891.

Zeckhauser, R. (1969). Majority rule with lotteries on alternatives. *The Quarterly Journal of Economics 83*(4), 696–703.

APPENDIX

Appendix A. Condorcet Example

The Condorcet paradox famously states that no deterministic outcome can be supported by a majority. The simplest example involves three candidates, a, b, and c, with three voters holding cyclic preferences:

$$a \succ_1 b \succ_1 c, \quad b \succ_2 c \succ_2 a, \quad c \succ_3 a \succ_3 b.$$

Clearly, no single candidate receives majority support.

This example extends to n voters $\{1, \ldots, n\}$ and n candidates $\{a_1, \ldots, a_n\}$, where preferences remain cyclic:

$$a_i \succ a_{i+1} \succ \cdots \succ a_n \succ a_1 \succ \cdots \succ a_{i-1}.$$

By symmetry, assume a_1 is chosen as the final outcome. Then, all voters except voter 1 prefer a_n over a_1. As n grows, no outcome secures the support of any constant fraction of voters, reinforcing the paradox.

Appendix B. Missing Proof

B.1. **Proof of Proposition 2.2.**

Proof. We first show that no Condorcet-winning lottery exists.

Consider the standard Condorcet cycle with three voters and three alternatives a, b, and c, where each pairwise majority preference forms a cycle (e.g., two voters prefer $a \lhd b$, $b \lhd c$, and $c \lhd a$). Suppose for contradiction that there exists a Condorcet-winning lottery $x = \alpha a + \beta b + \gamma c$. Then one of the coefficients must be at least $1/3$; without loss of generality, assume $\alpha \geq 1/3$.

Because two voters prefer b to a, we can replace the a-mass in x with b to obtain a new lottery:

$$y = (\alpha + \beta)b + \gamma c.$$

By the monotonicity condition and the preference of $a \lhd b$ for two voters, we have $x \lhd y$ for these voters.

Next, we show that there are voting instances in which no ex-ante stable lottery exists. We construct such an instance starting from the example above. By continuity, there exists $\varepsilon > 0$ such that any lottery within ε (in L_2 distance) of y is also preferred to x by these two voters. Now consider an ε-grid $\Phi \subset \Delta(A)$, meaning every lottery in $\Delta(A)$ has some lottery in Φ within ε L_2-distance. Consider a voting instance in which each element of Φ is treated as a deterministic alternative, and preferences remain unchanged. Note that constructing this voting instance relies on the assumption that the preference relation is complete.

Assume for contradiction that there exists an ex-ante stable lottery $x \in \Delta(\Phi)$ under these preferences. Then, by the argument above and the continuity assumption, there exists a lottery y such that every lottery within ε of y is strictly preferred to x by two voters. In particular, there exists such a lottery in Φ—contradicting the assumption that x is exante stable.

□

B.2. **Proof of Theorem 4.**

Proof. First, to show that an equilibrium lottery associated with f is SD efficient for any strictly increasing scoring function f, we assume the contrary. Suppose that there exists a lottery $\vec{q}$ that strictly dominates $\vec{p}$. Then, the total expected score of $\vec{q}$ would be strictly greater than that of $\vec{p}$, because the scoring function is strictly increasing. This leads to a contradiction, as $\vec{p}$ was selected to maximize the total average score.

Let $\vec{p}$ be an equilibrium lottery for a step function h. Because the score function is not strictly increasing, some equilibrium lotteries are not SD-efficient, but we can still select an efficient one. We do so in two stages. First, we establish that if a lottery $\vec{q}$ SD-dominates an equilibrium lottery $\vec{p}$, then $\vec{q}$ itself is an equilibrium lottery and yields exactly the same individual scores as $\vec{p}$. Second, we show that every equilibrium lottery with these fixed scores can be characterized by a linear program. Consequently, finding an SD-efficient equilibrium lottery reduces to solving this linear program while maximizing any aggregate utility functions that are consistent to the voters' ordinal preferences.

For a lottery $\vec{p}$ and a voter i, consider the interval representation described in Section 3. Let M_i be the half-point interval that is a positive-length interval containing the midpoint $1/2$, and a point strictly less than $1/2$. (If there are two such intervals containing $1/2$, choose the one on the left.) Denote the beginning and end of this interval as the position of the midpoint interval. This interval corresponds to an alternative or a tier of alternatives (in case of ties). We refer to these as the *mid-tier alternatives*.

Mathematically, for a voter i, M_i be the set of all alternatives in the same tier with an alternative a such that

$$\sum_{b \prec_i a} p_b \;<\; \tfrac{1}{2} \;\leq\; \sum_{b \preceq_i a} p_b.$$

Claim B.1. *Given an equilibrium lottery $\vec{p}$ according the step function $h(.)$, if $\vec{q}$ SD-dominates $\vec{p}$, then $\vec{q}$ is also an equilibrium lottery, and the scoring rules based on $\vec{p}$ and $\vec{q}$ are identical.*

Proof. First, since the scoring rule is monotonic, alternatives ranked lower receive weakly lower scores. Therefore, the score of $\vec{q}$ under $\Pi(f, \vec{p})$ is at least that of $\vec{p}$. However, $\vec{p}$ is chosen to maximize aggregate score, so equality must hold for every voter.

We now examine when equality can occur. Fix a voter i. Let $M_i \subseteq A$ be the set of mid-tier alternatives in $\vec{p}$, and let k be their rank. Let $\gamma := \Pi_i^a(h, \vec{p})$ be the score assigned to a mid-tier alternative $a \in M_i$. Note that $0 \leq \gamma < 1$.

Consider first the case where $\gamma > 0$. Voter i assigns one of three possible scores to the alternatives: 0, γ, or 1. Specifically, alternatives ranked below position k receive a score of 0, those ranked above k receive a score of 1, and the alternative at the mid-tier rank k receives a score of γ.

Since $\vec{q}$ SD-dominates $\vec{p}$, and the average score under $\vec{q}$ equals that under $\vec{p}$, the total probability assigned to each score—0, γ, and 1—must be identical under both $\vec{p}$ and $\vec{q}$. This implies that the score profile induced by $\vec{q}$, using the same step-function scoring rule, must match that induced by $\vec{p}$.

The same logic applies in the case where $\gamma = 0$. Here, the scoring function is binary: alternatives ranked at or below the mid-tier position receive 0, and those ranked strictly above receive 1. For

the average scores under $\vec{p}$ and $\vec{q}$ to match, and given that $\vec{q}$ SD-dominates $\vec{p}$, it must be that the total probability assigned to alternatives ranked at or below the mid-tier rank is exactly $1/2$ under both distributions. This again implies that the score profile under $\vec{q}$, according to the step-function scoring rule, is identical to that under $\vec{p}$.

Therefore, $\vec{q}$ has the same scoring rule and aggregate score as $\vec{p}$, so $\vec{q}$ is also an equilibrium lottery. □

Given a lottery $\vec{p}$ and voter i, we describe an LP that characterizes all equilibrium lottery associated with the step scoring function.

Consider the following linear program, denoted SCORELP$(\vec{p}, i)$, which characterizes all lotteries $\vec{q}$ for which $\Pi_i^a(h, \vec{p}) = \Pi_i^a(h, \vec{q})$ for all alternatives $a \in A$.

The first constraint requires $\vec{q}$ to be a valid lottery:

$$\sum_{a \in A} q_a = 1, \quad q_a \geq 0.$$

Given $M_i \subset A$, the set of mid-tier alternatives under $\vec{p}$ for voter i, the second constraint ensures the upper bound of the midpoint interval remains the same:

$$\sum_{a \succ_i M_i} q_a = \sum_{a \succ_i M_i} p_a.$$

The third constraint ensures the lower bound of the midpoint interval remains the same when the upper bound is not exactly $1/2$. That is, if $\sum_{a \succ_i M_i} p_a > 1/2$, then:

$$\sum_{a \prec_i M_i} q_a = \sum_{a \prec_i M_i} p_a.$$

Any $\vec{q}$ satisfying these constraints will have a scoring rule such that:

- Alternatives ranked above M_i receive score 1.
- If the midpoint interval touches $1/2$ exactly, then alternatives in M_i and below receive score 0.
- If the midpoint interval contains $1/2$ strictly, then alternatives in M_i receive score $\sum_{a \preceq_i M_i} p_a - 1/2$, and lower-ranked alternatives receive score 0.

Define

$$\text{SCORELP}(\vec{p}) := \bigcap_{i \in V} \text{SCORELP}(\vec{p}, i),$$

which captures all lotteries that yield the same scoring as $\vec{p}$ for all voters.

Next, for each $\vec{q} \in \text{SCORELP}(\vec{p})$, the scoring is fixed, so we can determine which alternatives have the highest aggregate score. Denote this set by $A^*(\vec{p})$. Add the constraint:

$$q_a = 0 \quad \text{for all } a \notin A^*(\vec{p}),$$

which restricts the support of $\vec{q}$ to the top-scoring alternatives.

Define:

$$\text{EQUILIBRIUM}(\vec{p}) := \{\vec{q} \in \text{SCORELP}(\vec{p}) \mid q_a = 0 \text{ for all } a \notin A^*(\vec{p})\}.$$

Claim B.2. *Given an equilibrium lottery $\vec{p}$ associated with the step function $h(.)$, EQUILIBRIUM($\vec{p}$) characterizes all equilibrium lotteries associated with $h(.)$ and has the same scoring as $\vec{p}$.*

Let $u_i(\cdot)$ be a cardinal utility consistent with $\succeq_i$. Let $\vec{q} \in$ EQUILIBRIUM($\vec{p}$) be a lottery that maximizes $\sum_{i,a} u_i(a) \cdot q_a$. Then $\vec{q}$ is an SD-efficient equilibrium lottery.

Proof. It is clear that EQUILIBRIUM($\vec{p}$) has all the necessary and sufficient constraints of equilibrium lotteries associated with h and has the same scoring has $\vec{p}$.

To prove the second part of the claim, observe that since $\vec{q} \in$ EQUILIBRIUM($\vec{p}$), it is an equilibrium lottery. Suppose there exists a lottery $\vec{q}'$ that SD-dominates $\vec{q}$. Because $\vec{q}'$ is weakly better for all voters and strictly better for at least 1 voter, we have:

$$\sum_{i,a} u_i(a) \cdot q_a < \sum_{i,a} u_i(a) \cdot q'_a.$$

Moreover, by Claim 3.2, $\vec{q}'$ is also an equilibrium lottery with the same scoring rule, so $\vec{q}' \in$ EQUILIBRIUM($\vec{p}$), contradicting the assumption that $\vec{q}$ maximizes utility.

Therefore, $\vec{q}$ is an SD-efficient equilibrium lottery. □

□

B.3. **Proof of Theorem 5.** Let $\vec{p}$ be a limit equilibrium lottery for the subsequence of equilibrium lotteries with the scoring function h_i^ϵ, with $\epsilon \to 0$. Let γ_i^a be the corresponding score for alternative a given by voter i in the limit. Similarly to the proof of Theorem 3, we find that for every alternative, the total score is at most n. Moreover, if for voter i,

$$\sum_{j \in A: a \succ j} p_j \geq \alpha_i,$$

then $\gamma_i^a = \frac{1}{1-\alpha_i}$.

Let $S \subset V$ be the set of voters who strictly prefer a to $\vec{p}$. The total score for a is then at least $w_S := \sum_{i \in S} \frac{1}{1-\alpha_i}$. This cannot exceed n, so we establish $w_S \leq n$, thereby showing that the lottery is ex ante stable.

Next, for $a \in A$ such that $p_a > 0$, the total score of a is exactly n. If a voter i does not weakly prefer a to $\vec{p}$, it means

$$\sum_{j \in A: a \succeq j} p_j < \alpha_i.$$

This implies that in the limit, the score of a for voter i will be 0. On the other hand, if a voter i weakly prefers a to $\vec{p}$, the score assigned to a will be at most $\frac{1}{1-\alpha_i}$. Since the total score for a is n, we find that if S is the set of voters who weakly prefer a to $\vec{p}$, then

$$w_S := \sum_{i \in S} \frac{1}{1-\alpha_i} \geq n.$$

This shows that the limit lottery is ex-post stable.

The existence of an SD efficient lottery that has the same scoring as the limit lottery follows exactly the same reasoning as in Theorem 4.

Appendix C. Computation of Stable Lotteries

In this section, we show that, when preferences are strict, an equilibrium lottery for the step rule can be obtained through a set of linear programs. Although the number of programs may be exponential in the worst case, the approach is algorithmically practical when the number of alternatives is a small constant, yielding an effective procedure for computing a stable lottery.

As in the proof of the SD efficiency in the previous section, the argument relies crucially on the notion of a *mid-point alternative*. For a given lottery and voter i, an alternative a is the mid-point alternative if its interval in the lottery's interval representation straddles the probability $1/2$; that is,

$$\sum_{b \prec_i a} p_b \;<\; \tfrac{1}{2} \;\leq\; \sum_{b \preceq_i a} p_b.$$

Illustrative computation of an equilibrium lottery. We continue Example 2.4. Let the fractions of the population whose preference types $1, 2, 3$ be α, β, γ respectively, with no single group holding a majority: $\max\{\alpha, \beta, \gamma\} < \frac{1}{2}$.

Assume the equilibrium lottery places positive probability on every alternative,

$$(p_a, p_b, p_c) \quad \text{with} \quad p_a, p_b, p_c > 0, \;\; p_a + p_b + p_c = 1,$$

where p_a, p_b, p_c are the probabilities assigned to a, b, c. Under the step-score rule the midpoint alternative for types $1, 2, 3$ is b, c, a, respectively.

This assumption corresponds to the constraint that all $p_a, p_b, p_c < 1/2$.
Scores for alternative a:

- For voters of type 1 (α-fraction), a lies above their midpoint b, so it receives score 1.
- For voters of type 2 (β-fraction), a lies below their midpoint c, hence score 0.
- For voters of type 3 (γ-fraction), a is exactly their midpoint. Their step score equals the share of probability mass on alternatives that beat the midpoint minus the share below it, divided by the probability of the midpoint itself:

$$\frac{(p_a + p_b) - \frac{1}{2}}{p_a} = \frac{\frac{1}{2} - p_c}{p_a}.$$

Aggregating scores across all voters and imposing the equilibrium condition that every alternative attains the common score $\frac{1}{2}$ yields

$$\alpha \cdot 1 + \beta \cdot 0 + \gamma \cdot \frac{\frac{1}{2} - p_c}{p_a} \;=\; \frac{1}{2}.$$

Proceeding analogously for b and c gives

$$\alpha \cdot \frac{\frac{1}{2} - p_a}{p_b} + \beta \cdot 1 + \gamma \cdot 0 \;=\; \frac{1}{2}, \qquad \alpha \cdot 0 + \beta \cdot \frac{\frac{1}{2} - p_b}{p_c} + \gamma \cdot 1 \;=\; \frac{1}{2}.$$

Collecting all constraints, the equilibrium lottery is the solution to

$$\begin{aligned}\text{find } p_a, p_b, p_c \quad \text{s.t.} \qquad & p_a + p_b + p_c = 1, \qquad 0 < p_a, p_b, p_c < 1/2,\\ & \big(\tfrac{1}{2} - \alpha\big)p_a + \gamma p_c = \tfrac{\gamma}{2},\\ & \big(\tfrac{1}{2} - \beta\big)p_b + \alpha p_a = \tfrac{\alpha}{2},\\ & \big(\tfrac{1}{2} - \gamma\big)p_c + \beta p_b = \tfrac{\beta}{2}.\end{aligned}$$

Solving this linear program yields the unique equilibrium probabilities (p_a^*, p_b^*, p_c^*) whenever $\alpha, \beta, \gamma \in (0, \frac{1}{2})$ and $\alpha + \beta + \gamma = 1$.

General equilibrium-lottery program. Next, we show how to compute the equilibrium lottery for the step–score rule in a general setting with strict preferences.

Let the support of the equilibrium lottery be $A^+ \subseteq A$, and suppose we already know each voter's midpoint alternative under that lottery. For voter i let this midpoint be $m_i \in A^+$.
Individual scores. The midpoint fully determines voter i's scores on all other alternatives:

- Every alternative ranked *below* m_i receives score 0.
- Every alternative ranked *above* m_i receives score 1.
- The only unknown score is the one assigned to the midpoint itself. Using the step–score definition, it equals

$$\gamma_i^{m_i} \;=\; \frac{\frac{1}{2} \;-\; \sum_{a \succ_i m_i} p_a}{p_{m_i}}, \text{ equivalently:} \quad p_{m_i} \cdot \gamma_i^{m_i} = \frac{1}{2} \;-\; \sum_{a \succ_i m_i} p_a.$$

Consider the following linear program, whose variables are

$$\{p_a \mid a \in A^+\} \quad \text{and} \quad \{y_{i,a} \mid i \in V,\ a \in A^+\},$$

where $y_{i,a} = p_a\, \gamma_i^a$. There are $(n+1)|A^+|$ variables in total.

We have the following constraints.
Midpoint constraints: For every voter i let m_i denote her midpoint alternative. The support probabilities must satisfy

$$\sum_{a \in A^+ \,:\, a \prec_i m_i} p_a \;<\; \tfrac{1}{2} \;\le \sum_{a \in A^+ \,:\, a \preceq_i m_i} p_a, \qquad \forall i \in V.$$

Score-encoding constraints: Rather than constraining γ_i^a directly, we impose conditions on the products $y_{i,a} = p_a\, \gamma_i^a$:

$$y_{i,a} \;=\; \begin{cases} \frac{1}{2} \;-\; \sum_{b \succ_i m_i} p_b, & \text{if } a = m_i,\\ 0, & \text{if } a \prec_i m_i \text{ (score } \gamma_i^a = 0),\\ p_a, & \text{if } a \succ_i m_i \text{ (score } \gamma_i^a = 1). \end{cases} \qquad \forall i \in V,\ a \in A^+.$$

Aggregate-score constraint. For every $a \in A^+$, the alternative's aggregate score must equal $n/2$. In terms of the variables $y_{i,a} = p_a\,\gamma_i^a$, this is captured by

$$\sum_{i \in V} y_{i,a} \;=\; \frac{n}{2}\,p_a, \qquad \forall a \in A^+.$$

Finally, for any alternative $a \notin A^+$ the aggregate score must be at most $n/2$. Because every voter assigns either 0 or 1 to such an alternative, we do not need extra LP variables for these scores; it suffices to verify this consistency condition using the support constraint and the midpoint constraints already imposed.

Together with the usual probability constraints $\sum_{a \in A^+} p_a = 1$ and $p_a \geq 0$, these inequalities characterize all lotteries whose support lies in A^+ and whose individual scores match the step rule.

If this linear program is feasible, every feasible solution gives an equilibrium lottery that is supported on A^+. Conversely, any equilibrium lottery must satisfy the same constraints, so it solves the program.

Thus, once the support A^+ and each voter's midpoint m_i are known, verifying or computing an equilibrium reduces to solving a linear program.