
An Axiomatization of the Random Priority Rule

Christian Basteck (WZB Berlin)

Discussion Paper No. 502

April 09, 2024

AN AXIOMATIZATION OF THE RANDOM PRIORITY RULE

CHRISTIAN BASTECK[†]

ABSTRACT. We study the problem of assigning indivisible objects to agents where each is to receive one object. To ensure fairness in the absence of monetary compensation, we consider random assignments. Random Priority, also known as Random Serial Dictatorship, is characterized by symmetry, ex-post efficiency and probabilistic (Maskin) monotonicity – whenever preferences change so that a given deterministic assignment is ranked weakly higher by all agents, the probability of that assignment being chosen should be weakly larger. Probabilistic monotonicity implies strategy-proofness for random assignment problems and is equivalent on a general social choice domain; for deterministic rules it coincides with Maskin monotonicity.

Keywords: Random Assignment; Random Priority; Random Serial Dictatorship; Ex-Post Efficiency; Probabilistic Monotonicity; Maskin Monotonicity

JEL codes: C70, C78, D63

1. INTRODUCTION

Many allocation problems require us to assign indivisible objects to agents such that each agent receives at most one object – public housing associations assign apartments to tenants, education administrations match teachers to schools, and municipalities assign daycare spots to children. In the absence of compensating transfers, and in light of the indivisible nature of objects, a desire to treat agents equally forces us to consider randomized assignments.

A random assignment rule prescribes a lottery over deterministic assignments for any possible profile of agents' (strict) preferences over objects. Arguably one of the simplest such rules is known as the Random Priority rule (*RP*, also known as random serial dictatorship) implemented for example by the following extensive form mechanism: order agents uniformly at random and let each agent, one after another according to the realized ordering, choose their most-preferred among all objects still available.

It is readily seen to be fair, efficient, and incentive compatible: it treats agents with identical preferences equally from an ex ante point of view, all assignments

[†]WZB Berlin, Berlin, Germany. christian.basteck@wzb.eu

I thank Lars Ehlers for many inspiring discussions and helpful comments. Financial support by the Deutsche Forschungsgemeinschaft through CRC TRR 190 (project number 280092119) is gratefully acknowledged.

arising with positive probability are Pareto-efficient, and choosing the most-preferred available object according to one's true preferences is a dominant strategy – thus, if the rule is implemented as a direct mechanism, where agents are asked to report their preferences before the mechanism chooses optimally on their behalf, it is strategy-proof.

Moreover, it satisfies another normatively appealing, property: it is responsive to agents preferences, in that an assignment is more likely to be chosen once agents consider it more preferred. More precisely, our new axiom, probabilistic (Maskin) monotonicity, considers different preference profiles R, R' and an assignment μ where for each agent an assignments preferred to μ under R' is also preferred to μ under R – hence, compared to R , the assignment μ has moved up in each agents' ranking of assignments. In any such situation, probabilistic monotonicity requires that the assignment μ arises with weakly larger probability under R' than under R .

As our main result, we find that the random priority rule is the unique random assignment rule satisfying symmetry, (ex post) efficiency, and probabilistic monotonicity. Moreover, this characterization holds both in an ordinal framework, where agents preferences are represented by rank order list of objects, as well as a cardinal framework where preferences are described by von Neumann–Morgenstern utility functions.

Note that the characterisation does not invoke strategy-proofness. This is because probabilistic monotonicity can be seen as a natural strengthening of strategy-proofness, as we argue in our second proposition. In fact, on a general social choice domain with strict preferences, probabilistic monotonicity and strategy-proofness are equivalent – hence our characterisation of random priority may be seen as the counterpart to the classic characterisation of random dictatorship as the only random social choice rule on such domains satisfying (ex post) efficiency and strategy-proofness due to [Gibbard \[1977\]](#).

Many authors have recognized the exceptional position that RP occupies among random assignment rules. [Abdulkadiroğlu and Sönmez \[1998\]](#) and [Knuth \[1996\]](#) independently showed that it is not only implemented by a random serial dictatorship as described above but equivalently as the core from random endowments.¹ Following up on this surprising result, several authors have shown how randomization over certain deterministic rules, so as to symmetrize the treatment agents, yields the random priority rule [[Pathak and Sethuraman, 2011](#), [Lee and Sethuraman, 2011](#), [Carroll, 2014](#), [Bade, 2020](#)]:² in the most general formulation, due to [Bade \[2020\]](#), symmetrizing any strategy-proof, non-bossy, and ex post efficient rule yields RP .

¹For this, they assume an equal number of agents and distinct objects.

²All these assume a unit supply of objects, i.e., do not consider the possibility of copies of objects between which agents are indifferent.

Interestingly, strategy-proofness and non-bossiness are equivalent to group-strategy-proofness [Pycia and Ünver, 2023] and thus equivalent to Maskin monotonicity [Takamiya, 2007].³

Thus, as randomization over Maskin monotonic and ex post efficient rules yields a rule that still satisfies probabilistic monotonicity (and preserves ex post efficiency), our characterization implies and unifies the previous equivalence results from Abdulkadiroğlu and Sönmez [1998] to Bade [2020].

Further, Bade [2016] shows any ex post efficient and symmetric random assignment rule to necessarily violate group-strategy-proofness⁴ so that randomizing over group strategy proof (and ex post efficient) rules to arrive at a symmetrized random assignment rule entails a loss of group strategy-proofness. In contrast, Maskin monotonicity, equivalent to group strategy proofness, naturally generalizes to probabilistic monotonicity and is in that form preserved under randomization.

Erdil [2014] was the first to show that RP is not characterized by symmetry, ex-post efficiency and strategy-proofness in the presence of outside options – a longstanding conjecture since Bogomolnaia and Moulin [2001] proved this to be the case for 3 agents and 3 acceptable objects. Basteck and Ehlers [2024] show RP not to be characterized by these properties even in the absence of outside options. Hence, to characterize RP in conjunction with symmetry and ex post efficiency, we are forced to strengthen strategy-proofness, for example to probabilistic monotonicity.

Pycia and Troyan [2021] consider another strengthening of strategy-proofness – they show that an obviously strategy-proof mechanism is symmetric and (ex post) efficient if and only if it implements RP . In other word, their characterization strengthens strategy-proofness to obvious strategy proofness (OSP)[Li, 2017]. OSP and probabilistic monotonicity are logically independent. Note that the characterization or RP provided by Pycia and Troyan [2021] refers to the *possibility* of implementing it in an obviously strategy-proof way – if instead agents are asked to submit their preferences in advance before the outcome is then computed by means of the random priority rule, obvious strategy proofness will no longer be satisfied.⁵ Nonetheless, the latter procedure still ensures probabilistic monotonicity.

Finally, (weak) Maskin monotonicity allows to characterize two of the most prominent families of deterministic assignment rules. Kojima and Manea [2010] characterize

³Pycia and Ünver [2023] and Takamiya [2007] consider domains that are general enough to allow for the possibility of multiple object copies; the same is the case for our characterization. For a setting with unit supplies, the first equivalence was shown by Pápai [2000] (allowing for an unequal number of objects and agents) and the second by Takamiya [2001] (assuming an equal number of objects and agents, i.e., a housing market as in Shapley and Scarf [1974]).

⁴Zhang [2019] strengthens this implication by showing that such rules are strongly group manipulable under some mild additional fairness axioms.

⁵For example, this procedure is used to assign students to secondary schools in the city of Amsterdam: <https://schoolkeuze020.nl/plaatsing/>

deferred acceptance (with respect to some priority) by non-wastefulness, weak Maskin monotonicity,⁶ and population monotonicity. Moreover, if weak Maskin monotonicity is strengthened to Maskin monotonicity, the only assignment rules satisfying the three properties are efficient deferred acceptance rules.⁷ Morrill [2013] characterizes top-trading-cycles (with respect to any priority) by efficiency, weak Maskin monotonicity, independence of irrelevant rankings, and mutual best. Note that, a fortiori, both deferred acceptance and top trading cycle rules satisfy Maskin monotonicity and hence probabilistic monotonicity. However, neither satisfies OSP.

The paper is organized as follows: Section 2 introduces the model; Section 3 presents the results and discusses probabilistic monotonicity in more detail; Section 4 concludes with some remarks on non-symmetric assignment rules.

2. MODEL

2.1. Agents, objects, preferences. Let N be a set of agents and O be a set of distinct objects (or object types), possibly including a null-object, denoted by \emptyset .⁸ For each object $x \in O$, let q_x denote the number of copies; if $\emptyset \in O$, we set $q_\emptyset = |N|$. Taking into account a possible null-object, it is without loss of generality to assume $\sum_{x \in O} q_x \geq |N|$.

Each agent $i \in N$ holds strict preferences over the set of objects O , represented by a linear order⁹ R_i whose asymmetric part we denote by P_i .¹⁰ Let \mathcal{R}^i denote the set of all strict preferences of agent i over O and let $\mathcal{R} = \times_{i \in N} \mathcal{R}^i$ denote the set of all preference profiles $R = (R_1, \dots, R_n)$. For a given object x and given R_i we write $U(x, R_i) = \{y \in O \mid y P_i x\}$ for i 's (strict) upper contour set at x and denote the (strict) lower contour set as $L(x, R_i) = \{y \in O \mid x P_i y\}$.

For any $R_i \in \mathcal{R}^i$ and any $x, y \in O$ adjacent in R_i ¹¹ we say that R'_i differs from R_i by a swap of x and y if $R'_i \setminus O \setminus \{x\} = R_i \setminus O \setminus \{x\}$, $R'_i \setminus O \setminus \{y\} = R_i \setminus O \setminus \{y\}$ and $y R'_i x \Leftrightarrow x R_i y$. Let $N_{R_i} \subset \mathcal{R}^i$ denote the set of preferences in the neighbourhood of R_i , i.e., the set of all preferences that differ from R_i by a swap of two objects adjacent in R_i .

2.2. Random Assignments. A (deterministic) assignment is a mapping $\mu : N \rightarrow O^{12}$ such that each object $x \in O$ is assigned to at most q_x agents, i.e., $|\{i \in N \mid \mu_i = x\}| \leq q_x$.

⁶An assignment rule satisfies weak Maskin monotonicity if an assignment that is chosen at some preference profile R and ranked weakly higher by all agents at another profile R' will either still be chosen or replaced by another assignment that Pareto-dominates it, i.e., ranked even higher by all agents.

⁷I.e., deferred acceptance rules where priorities are Ergin-acyclic [Ergin, 2002] or that satisfy bounded invariance Basteck and Ehlers [2024].

⁸Thus, our setup allows agents to not receive any (proper) object, be it because there are not enough objects or because agents can reject an assigned object that they deem unacceptable in comparison to some outside option.

⁹I.e., complete, transitive, and antisymmetric ($x R_i y$ and $y R_i x$ together imply $x = y$).

¹⁰That is, $x P_i y$ whenever $x R_i y$ and $x \neq y$.

¹¹Two objects $x, y \in O$ are adjacent in R_i , if $U(x, R_i) \setminus \{x, y\} = U(y, R_i) \setminus \{x, y\}$.

¹²We will write μ_i instead of $\mu(i)$ for any $i \in N$.

Let \mathcal{M} denote the set of all assignments. An assignment $\mu \in \mathcal{M}$ is efficient under R if there exists no $\mu' \in \mathcal{M}$ such that $\mu'_i R_i \mu_i$ for all $i \in N$ and $\mu'_j P_j \mu_j$ for some $j \in N$. Let $\mathcal{PO}(R)$ denote the set of all assignments efficient under R .

Let $\Delta(\mathcal{M})$ denote the set of all probability distributions over \mathcal{M} to which we refer to as random assignments. For a given $p \in \Delta(\mathcal{M})$ and $\mu \in \mathcal{M}$, let p_μ denote the probability assigned to μ under p and refer to $\text{supp}(p) = \{\mu \in \mathcal{M} | p_\mu > 0\}$ as the support of p . We say that a random assignment $p \in \Delta(\mathcal{M})$ is *ex-post efficient* under R if all assignments in its support are efficient, i.e., $\text{supp}(p) \subseteq \mathcal{PO}(R)$.

A random assignment p satisfies (*assignment-)*symmetry if for any two agents $i, j \in N$ such that $R_i = R_j$ and any two deterministic assignments μ, μ' such that $\mu_i = \mu'_j$, $\mu_j = \mu'_i$ and $\mu_k = \mu'_k$ for all $k \in N \setminus \{i, j\}$ we have $p_\mu = p_{\mu'}$. In other words, deterministic assignments that merely permute the roles of agents with identical preferences arise with the same probability.¹³

A random assignment rule f maps preference profiles to random assignment, i.e., $f : \mathcal{R} \rightarrow \Delta(\mathcal{M})$. We say that f is ex-post efficient, if $f(R)$ is ex-post efficient for every $R \in \mathcal{R}$. It is symmetric, if $f(R)$ is symmetric for every $R \in \mathcal{R}$. It is deterministic, if $|\text{supp}(f(R))| = 1$ for every $R \in \mathcal{R}$.

For our last axiom, consider two preference profiles $R, R' \in \mathcal{R}$ and a deterministic assignment μ . We say that R' is a μ -monotonic transformation of R if for agents' upper contour sets we have $U(\mu_i, R'_i) \subseteq U(\mu_i, R_i)$ for all i , i.e., if all agents rank their assignment under μ weakly higher under R'_i than under R_i . We say that a random assignment rule f satisfies *probabilistic (Maskin) monotonicity* if for any $R, R' \in \mathcal{R}$ and $\mu \in \mathcal{M}$ such that R' is a μ -monotonic transformation of R we have $f(R')_\mu \geq f(R)_\mu$.

Note that if f is deterministic, probabilistic monotonicity reduces to Maskin monotonicity for deterministic, single valued rules.¹⁴

One of the most prominent random assignment rules is the Random Priority rule (*RP*), also known as random serial dictatorship. For that let \triangleright denote a strict priority order over N and let Π denote the set of all strict priority orders. Given $\triangleright \in \Pi$, let f^\triangleright denote the deterministic priority (or serial dictatorship) rule where agents are assigned their most-preferred among all available objects in order of their priority.¹⁵ Then the Random Priority rule is defined by $RP(R) = \frac{1}{n!} \sum_{\triangleright \in \Pi} f^\triangleright(R)$ for all $R \in \mathcal{R}$.

¹³One may also define symmetry, often referred to as equal-treatment-of-equals, with respect to agents' individual object assignment probabilities rather than the probabilities of choosing complete assignment – see Section 2.3.

¹⁴A deterministic assignment rule f is said to satisfy Maskin monotonicity if for any two profiles $R, R' \in \mathcal{R}$ and assignment $\mu \in \mathcal{M}$ such that R' is a μ -monotonic transformation of R , we have that $f(R) = \mu$ implies $f(R') = \mu$.

¹⁵For any $R \in \mathcal{R}^N$ and $i_1 \triangleright i_2 \triangleright \dots \triangleright i_n$, i_1 receives her R_{i_1} -most-preferred object in O (denoted by $f_{i_1}^\triangleright(R)$), and for $l = 2, \dots, n$, i_l receives her R_{i_l} -most-preferred object in $O \setminus \{f_{i_1}^\triangleright(R), \dots, f_{i_{l-1}}^\triangleright(R)\}$ (denoted by $f_{i_l}^\triangleright(R)$).

Note that any f^\triangleright satisfies Maskin monotonicity [Kojima and Manea, 2010], and hence, RP satisfies probabilistic monotonicity.

2.3. Individual random assignments and welfare equivalence. Since agents' preferences are only over their own assigned object, the literature on random assignments has often focussed on agents' individual object assignment probabilities, rather than on the probabilities with which complete deterministic assignments are chosen.

Formally, for a given random assignment p , let p^{ia} denote the probability with which agent i is assigned object a and let $p^i = (p^{ia})_{a \in O}$ denote i 's individual random assignment. Two random assignments p, q are said to be (ex-ante) welfare-equivalent if $p^i = q^i$ for all $i \in N$, i.e., if they agree on agents' individual object assignment probabilities.¹⁶ Similarly, two random assignment rules, f and g , are welfare equivalent, if $f(R)$ and $g(R)$ are welfare equivalent for every R .

A random assignment p is said to satisfy individual-assignment-symmetry if for all $i, j \in N$, $R_i = R_j$ implies $p^i = p^j$. It is clear that any random assignment that satisfies assignment-symmetry (with respect to underlying deterministic assignments) also satisfies individual-assignment-symmetry (with respect to agents' individual object assignment probabilities). Moreover, for any random assignment that satisfies individual-assignment-symmetry, there exists a welfare equivalent assignment p' that satisfies assignment-symmetry. Hence, if one is only interested in pinning down random assignments up to welfare equivalence, rather than the exact convex combination of deterministic assignments, assignment-symmetry is no more restrictive than individual-assignment-symmetry.

Moreover, given any two random assignments $p, q \in \Delta(\mathcal{M})$ as well as some agent i 's preference R_i , we say that p^i stochastically R_i -dominates q^i if for all $x \in O$,

$$\sum_{y \in U(x, R_i)} p^{iy} \geq \sum_{y \in U(x, R_i)} q^{iy}.$$

A random assignment p stochastically R -dominates another random assignment q if p^i R_i -dominates q^i for all $i \in N$. It is stochastic dominance (sd)-efficient, given R , if there is no random assignment $q \neq p$ that stochastically R -dominates it.¹⁷

¹⁶Some papers directly define a random assignment as a matrix $(p^{ia})_{i \in N, a \in O}$ rather than as a convex combination of deterministic assignments. To the extent that one is interested in analysing random assignments only up to welfare equivalence, this is well defined as any such matrix $(p^{ia})_{i \in N, a \in O}$ can be represented as a convex combination of deterministic assignments by (a slight generalization of) the Birkhoff-von Neumann Theorem [Birkhoff, 1946].

¹⁷Bogomolnaia and Moulin [2001] refer to this as "ordinal efficiency". It is equivalent to (ex ante) Pareto-efficiency with respect to expected utilities for some von Neumann–Morgenstern utility function representation of agents' ordinal preferences over objects [McLennan, 2002].

A random assignment rule f is said to be *(sd)-strategy-proof* if for all $R \in \mathcal{R}$, all $i \in N$ and all $R'_i \in \mathcal{R}^i$, $f^i(R)$ stochastically R_i -dominates $f^i(R'_i, R_{-i})$.¹⁸ Since the definition relies only on agents' individual random assignments and preferences, every random assignment rule g that is welfare equivalent to a strategy-proof rule f will itself be strategy-proof. Moreover, strategy-proofness can be decomposed into the following three axioms [Mennle and Seuken, 2021]: A random assignment rule f satisfies

- *swap monotonicity* iff for all $i \in N$, $R \in \mathcal{R}$ and $R'_i \in N_{R_i}$ such that xP_iy and yP'_ix , we have (i) $f^{iy}(R'_i, R_{-i}) \geq f^{iy}(R)$ and (ii) $f^{iy}(R'_i, R_{-i}) = f^{iy}(R) \Rightarrow f^i(R'_i, R_{-i}) = f^i(R)$.
- *upper invariance* iff for all $i \in N$, $R \in \mathcal{R}$ and $R'_i \in N_{R_i}$ such that xP_iy and yP'_ix , we have $f^{iz}(R'_i, R_{-i}) = f^{iz}(R)$ for all $z \in U(x, R_i)$.
- *lower invariance* iff for all $i \in N$, $R \in \mathcal{R}$ and $R'_i \in N_{R_i}$ such that xP_iy and yP'_ix , we have $f^{iz}(R'_i, R_{-i}) = f^{iz}(R)$ for all $z \in L(y, R_i)$.
- *strategy proofness* iff it satisfies swap monotonicity, upper-, and lower invariance [Mennle and Seuken, 2021, Theorem 1].

3. RESULTS

3.1. Axiomatisation. We are now set to state our main result: symmetry, ex-post efficiency and probabilistic monotonicity characterize the Random Priority Rule.

Theorem 1. *Let f be a random assignment rule. Then f satisfies symmetry, ex post efficiency, and probabilistic monotonicity if and only if f is the Random Priority Rule.*

To prove the Theorem, let us first consider a preference profile R^* where all agents' preferences coincide. In that case symmetry and ex post efficiency require that $f(R^*)$ assigns equal probability to all Pareto efficient deterministic assignments – and hence agrees with $RP(R^*)$. We then proceed by induction: measuring the distance of a preference profile R to R^* by the number of necessary pairwise swaps of objects in agents' preference rankings to move from R to R^* – i.e., the Kemeny distance of R and R^* – and assuming that for all profiles with a smaller distance to R^* we know that f and RP coincide (induction hypothesis), we show that this implies $f(R) = RP(R)$ (induction step).

To illustrate the induction step, suppose, towards a contradiction, that we have a profile R where $f(R) \neq RP(R)$. Since the probabilities over all Pareto efficient assignments sum to one, there must be a Pareto efficient assignment μ that is chosen

¹⁸Sd-strategy-proofness is equivalent to the requirement that for any von Neumann–Morgenstern utility function compatible with a given ordinal ranking of objects, submitting the true ordinal ranking maximizes an agent's expected utility.

with higher probability under f than under RP , i.e., for which $f_\mu(R) > RP_\mu(R)$, as well as another Pareto efficient assignment μ' for which $f_{\mu'}(R) < RP_{\mu'}(R)$. Since both assignments are Pareto efficient, they must differ in the assignment of at least two agents. Hence, to illustrate the proof by means a minimal example, assume that for two agents $i, j \in N$ we have $\mu_i = \mu'_j = x$, $\mu_j = \mu'_i = y$ while for all $k \in N \setminus \{i, j\}$ we have $\mu_k = \mu'_k$ and for all $\hat{\mu} \in \mathcal{M} \setminus \{\mu, \mu'\}$ we have $f_{\hat{\mu}}(R) = RP_{\hat{\mu}}(R)$.

Then we must have that i 's upper contour set $U(x, R_i)$ is ranked as under R_i^* , i.e., $R_i|U(x, R_i) = R_i^*|U(x, R_i)$ – otherwise we could move to $R' = (R'_i, R_{-i})$ where compared to R_i , R'_i reorders objects in $U(x, R_i)$ in the same order as they appear in R_i^* . Since R' is closer to R^* than R while probabilistic monotonicity ensures that $f_\mu(R') = f_\mu(R) \neq RP_\mu(R) = RP_\mu(R')$ this contradicts the induction hypothesis. In the same way, via μ' , we find that $R_i|U(y, R_i) = R_i^*|U(y, R_i)$. Symmetrically, for the lower contour sets of x at R_i , we find that $R_i|L(x, R_i) = R_i^*|L(x, R_i)$ and $R_i|L(y, R_i) = R_i^*|L(y, R_i)$. Thus, either $R_i = R_i^*$ or R_i differs from R_i^* by a pairwise swap of x and y .

The same is true for j – either $R_j = R_j^* = R_i^*$ or R_j differs from R_j^* by a pairwise swap of x and y . Now, since in μ we have $\mu_i = x$ and $\mu_j = y$ while the assignment of x and y is reversed under μ' – and since both assignments are Pareto efficient – we must have that i and j rank x and y in the same order. Thus $R_i = R_j$. But then symmetry or RP demands that $f_\mu(R) > RP_\mu(R) = RP_{\mu'}(R) > f_{\mu'}(R)$, contradicting symmetry of f .

To complete the proof we establish the induction step for general μ and μ' , i.e., assignments that arise with different probabilities under f and RP and that may differ by more than the assignment of two agents and objects. Moreover, we verify the converse direction, i.e., the fact that $RP(R)$ satisfies not only symmetry and ex post efficiency but also probabilistic monotonicity. Both can be found in the Appendix.

3.2. Expected utility and preference intensities. Instead of assuming that agents' preferences are given by linear orders over the set of objects O , we may also assume that they are described by von Neumann–Morgenstern utility functions, i.e., that there exist $u_i : O \rightarrow \mathbb{R}$ for each $i \in N$. Depending on the interpretation attached, we may thus capture preference intensities or complete agents' preferences over lotteries of objects by assuming that they compare different lotteries by their corresponding expected utility. Restricting attention to strict preferences over objects, we require that $u_i(x) \neq u_i(y)$ for all $i \in N$, $x, y \in O$, and $x \neq y$. Let \mathcal{U}^i denote the set of all such utility functions of agent i and let $\mathcal{U} = \times_{i \in N} \mathcal{U}^i$ denote the set of all utility profiles $u = (u_1, \dots, u_n)$. The associated strict preference relations and preference profiles are denoted R_i^u and $R^u = (R_1^u, \dots, R_n^u)$.¹⁹

¹⁹I.e. for all $x, y \in O$ we have $xR_i^u y \Leftrightarrow u_i(x) \geq u_i(y)$.

A random assignment rule on the domain \mathcal{U} then maps utility profiles to random assignments, i.e., $f : \mathcal{U} \rightarrow \Delta(\mathcal{M})$. It satisfies probabilistic monotonicity, if probabilistic monotonicity is satisfied for the associated strict preferences $R^{u^{20}}$ and ex post efficient if $f(u)$ is ex post efficient for all $u \in \mathcal{U}$.

A random assignment p is symmetric with respect to $u \in \mathcal{U}$ if for any two agents $i, j \in N$ such that $u_i = u_j$ and any two deterministic assignments μ, μ' such that $\mu_i = \mu'_j$, $\mu_j = \mu'_i$ and $\mu_k = \mu'_k$ for all $k \in N \setminus \{i, j\}$ we have $p_\mu = p_{\mu'}$. In other words, deterministic assignments that merely permute the roles of agents with identical utility functions arise with the same probability. Note that this allows for unequal treatment of agents with identical ordinal preferences over objects, $R_i^u = R_j^u$, and thus weakens symmetry as defined with respect to R . A random assignment rule is symmetric on \mathcal{U} if $f(u)$ is symmetric for any $u \in \mathcal{U}$.

Despite the fact that symmetry with respect to u is weaker than symmetry with respect to R^u , it gives rise to a characterization of the Random Priority rule RP on the domain of (profiles of) von Neumann–Morgenstern utility functions in analogy to Theorem 1.

Proposition 1. *Let f be a random assignment rule on \mathcal{U} . Then f satisfies symmetry on \mathcal{U} , ex post efficiency, and probabilistic monotonicity if and only if f is the Random Priority Rule.*

Proof. Crucially, observe that probabilistic monotonicity implies ordinality of f , i.e., that for any $\mu \in \mathcal{M}$ and any two $u, u' \in \mathcal{U}$ such that $R^u = R^{u'}$, we have $f_\mu(u) = f_\mu(u')$: since R^u is a μ -monotonic transformation of $R^{u'}$, probabilistic monotonicity demands $f_\mu(u) \geq f_\mu(u')$ and, by a symmetric argument, $f_\mu(u) \leq f_\mu(u')$.

Hence f takes into account only agents' ordinal preferences R^u rather than the richer information encoded in their von Neumann–Morgenstern utility functions u . Thus, to be symmetric with respect to any $u \in \mathcal{U}$, f has to be symmetric with respect to the associated preference profiles R^u . The claim then follows from Theorem 1. \square

3.3. Probabilistic monotonicity. Given the crucial role of probabilistic monotonicity in the characterizations above, we complement Theorem 1 and Proposition 1 with some observations on the nature of the axiom. In particular we find that, analogous to strategy-proofness, probabilistic monotonicity can be decomposed into swap monotonicity, upper-, and lower invariance, where each of the three axioms is strengthened in that we require it to apply to the probability weights of separate deterministic assignment, rather than to sums of probability weights over several deterministic assignments (that correspond to agents' individual object assignment

²⁰I.e., f satisfies probabilistic monotonicity if for any $u, u' \in \mathcal{U}$ and $\mu \in \mathcal{M}$ such that $R^{u'}$ is a μ -monotonic transformation of R^u , we have $f_\mu(u') \geq f_\mu(u)$.

probabilities). Hence, for random assignment problems, probabilistic monotonicity implies strategy-proofness.

A random assignment rule f satisfies

- *assignment swap monotonicity* iff for all $i \in N$, $R \in \mathcal{R}$ and $R'_i \in N_{R_i}$ such that $xP_i y$ and $yP'_i x$, we have (i) $f_\mu(R'_i, R_{-i}) \geq f_\mu(R)$ for all μ with $\mu_i = y$, and (ii) $f_\mu(R'_i, R_{-i}) = f_\mu(R)$ for all μ with $\mu_i = y$ implies $f(R'_i, R_{-i}) = f(R)$.
- *upper assignment invariance* iff for all $i \in N$, $R \in \mathcal{R}$ and $R'_i \in N_{R_i}$ such that $xP_i y$ and $yP'_i x$, we have $f_\mu(R'_i, R_{-i}) = f_\mu(R)$ for all μ where $\mu_i \in U(x, R_i)$.
- *lower assignment invariance* iff for all $i \in N$, $R \in \mathcal{R}$ and $R'_i \in N_{R_i}$ such that $xP_i y$ and $yP'_i x$, we have $f_\mu(R'_i, R_{-i}) = f_\mu(R)$ for all μ where $\mu_i \in L(y, R_i)$.

Proposition 2. *A random assignment satisfies probabilistic monotonicity if and only if it satisfies assignment swap monotonicity, upper assignment invariance, and lower assignment invariance.*

Proof. First observe that any violation of upper- or lower assignment invariance also constitutes a violation of probabilistic monotonicity. The same holds if assignment swap monotonicity is violated in that for some $i \in N$, $R \in \mathcal{R}$, $R'_i \in N_{R_i}$ such that $xP_i y$ and $yP'_i x$, and μ with $\mu_i = y$ we have $f_\mu(R'_i, R_{-i}) < f_\mu(R)$. If instead $f_\mu(R'_i, R_{-i}) = f_\mu(R)$ for all μ with $\mu_i = y$, yet $f(R'_i, R_{-i}) \neq f(R)$, there is some μ' with $\mu'_i \neq y$ such that $f_{\mu'}(R'_i, R_{-i}) > f_{\mu'}(R)$ – again, this constitutes a violation of probabilistic monotonicity. Hence, to satisfy probabilistic monotonicity, a random assignment rule must satisfy all three properties.

For the other direction, consider a violation of probabilistic monotonicity, i.e., two preference profiles $R, R' \in \mathcal{R}$ and an assignment $\mu \in \mathcal{M}$ such that R' is a μ -monotonic transformation of R ²¹ for which we have $f_\mu(R') < f_\mu(R)$. Note that we can move from R to R' in a sequence of profiles $R = R^1, R^2, R^3, \dots, R^m = R'$ where at each step two consecutive profiles differ only by a pairwise swap of two adjacent objects in one agents' preferences and where the successor profile is a μ -monotonic transformation of its predecessor. Thus, at some point of the sequence, there are R^k and R^{k+1} such that $f_\mu(R^{k+1}) < f_\mu(R^k)$. Let i be the agent for whom R^k_i and R^{k+1}_i differ and denote the two objects adjacent in R^k_i that are swapped as we move to R^{k+1}_i as x and y ; without loss of generality assume that $xR^k_i y$ and $yR^{k+1}_i x$. Since R^{k+1} is a μ -monotonic transformation of R^k , we have $\mu_i \neq x$.

Now, if $\mu_i \in U(x, R^k_i)$, the fact that $f_\mu(R^{k+1}) < f_\mu(R^k)$ implies that f violates upper assignment invariance. Similarly, if $\mu_i \in L(y, R^k_i)$, we find that f violates lower assignment invariance while if $\mu_i = y$, f violates assignment swap monotonicity.

²¹I.e., $U(\mu_i, R'_i) \subseteq U(\mu_i, R_i)$ for all $i \in N$.

Hence, any rule that violates probabilistic monotonicity must violate one of the three properties. \square

Corollary 1. *Probabilistic monotonicity implies strategy-proofness. The converse does not hold: there exist strategy-proof random assignment rules for which the rule itself, as well as any other random assignment rule welfare equivalent to it, fails to satisfy probabilistic monotonicity*

The first part is an immediate consequence of Proposition 2 above and Theorem 1 in [Mennle and Seuken, 2021], that decomposes strategy-proofness into (i) swap monotonicity, (ii) upper invariance, and (iii) lower invariance, each of which is weaker than their counterpart referring to assignments. The fact that probabilistic monotonicity is in fact stronger than strategy-proofness (rather than equivalent to) follows from Theorem 1 above as well as Proposition 3 in [Erdil, 2014] that shows *RP* not to be characterized by symmetry, ex post efficiency, and strategy-proofness. While the construction in Erdil [2014] relies on the possibility that agents may remain unassigned (or, equivalently, assumes that there is a (null-)object with at least $|N|$ copies), Basteck and Ehlers [2024] show that *RP* is not characterized by symmetry,²² ex post efficiency, and strategy-proofness even in the absence of outside options.

If we do not confine ourselves to the domain of random assignment problems – where agents have strict preferences over their own assigned object but are otherwise indifferent regarding the assignment – but instead consider general social choice problems with strict preferences over all possible outcomes, probabilistic monotonicity and strategy proofness are equivalent. Formally, let N be the finite set of agents, O be the finite set of outcomes, let R_i denote agent i 's strict preferences over all outcomes, and denote a (strict) preference profile as $R = (R_i)_{i \in N}$. The set of all possible profiles is denoted \mathcal{R} and a (random) social choice rule is a mapping $f : \mathcal{R} \rightarrow \Delta(O)$ where $\Delta(O)$ denotes the set of probability distributions over O .²³ Refer to this as a general social choice domain (given N and O). A random social choice rule f is strategy-proof if for all R_i, R'_i and R_{-i} we find that $f(R_i, R_{-i})$ stochastically R_i -dominates $f(R'_i, R_{-i})$.²⁴ Equivalently, Gibbard [1977] defines strategy-proofness such that, for any underlying von Neumann–Morgenstern utility function compatible with R_i , reporting R_i truthfully maximizes an agent's expected utility. A random social choice rule satisfies probabilistic monotonicity if for any two R, R' and $o \in O$, such that R'

²²Both Erdil [2014] and Basteck and Ehlers [2024] consider individual-assignment-symmetry, but the constructed rules also satisfy symmetry with respect to complete assignments.

²³Gibbard [1977] refers to these as *decision schemes*.

²⁴I.e., if $\sum_{y \in U(x, R_i)} p_{iy} \geq \sum_{y \in U(x, R'_i)} q_{iy}$ where $U(x, R_i) = \{y \in O \setminus \{x\} | y R_i x\}$ denotes the strict upper contour set at x .

is an o -monotonic transformation of R we find that $f(R')$ chooses o with weakly higher probability than $f(R)$.

Fact 1. *A random social choice rule f , defined on a general social choice domain, is strategy-proof if and only if it satisfies probabilistic monotonicity.*

This follows, upon closer inspection, from Lemma 2 in Gibbard [1977] (where ‘non-perverseness’ corresponds to swap monotonicity and ‘localizedness’ can be shown to be equivalent to lower- and upper- invariance).²⁵ Hence, in the classic characterization of Random Dictatorship as the only social choice rule on a general social choice domain satisfying ex post efficiency and strategy-proofness,²⁶ we may replace the last axiom by the equivalent requirement of probabilistic monotonicity. Interestingly, as we leave the domain of general social choice problems (with strict preferences) and move to the domain of random assignment problems, probabilistic monotonicity and strategy-proofness diverge – and it turns out that only probabilistic monotonicity, rather than strategy-proofness, yields a characterization of the Random Priority Rule in analogy to Gibbard’s and Sonnenschein’s characterization of Random Dictatorship.

Finally, let us compare probabilistic monotonicity to obvious-strategy-proofness (OSP) [Li, 2017]. Both strengthen strategy-proofness, with the caveat that OSP applies not to random assignment rules per se, but to possible extensive form mechanisms implementing them.²⁷ We find that both are logically independent: there exist OSP-implementable random assignment rules that violate probabilistic monotonicity and vice versa.

Example 1. An OSP-implementable (random) assignment rule violating probabilistic monotonicity. Consider $N = \{1, 2, 3\}$, $O = \{a, b, c\}$ and a rule f that (i) awards 1 their most preferred object among O according to R_1 , (ii) if 1’s second and third most preferred objects are ranked alphabetically, i.e., aR_1b , aR_1c , or bR_1c , then 2 receives their most preferred among the remaining objects²⁸ while otherwise (iii) 3 receives their most preferred of the remaining objects before (iv) the last remaining agent is assigned the last remaining object. This rule can be readily implemented via

²⁵For deterministic, singleton valued social choice rules on the domain of strict preferences, Muller and Satterthwaite [1977] show strategy-proofness to be equivalent to ‘strong positive association’, i.e., Maskin monotonicity.

²⁶Gibbard [1977] reports this result as Corollary 1 and credits Hugo Sonnenschein. It is straightforward to see that once we assume anonymity, i.e., require agents to be treated symmetrically, the randomization needs to be uniform.

²⁷More precisely, a random assignment rule is OSP-implementable, if there exists an extensive form mechanism that implements the rule in obviously dominant strategies. Since any obviously dominant strategy is, a fortiori, a dominant strategy, any OSP-implementable rule is strategy proof.

²⁸The example does not rely on randomization, but instead constructs a deterministic rule. It can be readily generalized to non-degenerate random assignment rules, e.g., by making the probability with which 2 gets to choose dependent on R_1 .

sequential barter [Bade and Gonczarowski, 2016] and is hence OSP-implementable (as well as efficient). To see that it violates probabilistic monotonicity, consider the preference profile R such that $aR_i bR_i c$ for all $i \in N$. Then f returns the assignment $\mu = (\mu_1, \mu_2, \mu_3) = (a, b, c)$ (with probability 1). If instead we have R' such that $aR'_1 cR'_1 b$, while as before $aR'_i bR'_i c$ for all $i \in \{2, 3\}$, then f returns $\mu' = (a, c, b)$ – despite the fact that R' constitutes a μ -monotonic transformation of R .

For the other direction, Li [2017] shows that the core in Shapley-Scarf housing markets [Shapley and Scarf, 1974], while implementable in dominant strategies by a top-trading-cycle mechanism, is not OSP-implementable. Nonetheless, it is (Maskin) monotonic [Sönmez, 1996, Takamiya, 2001]. Similarly, deferred acceptance satisfies monotonicity [Kojima and Manea, 2010] but fails to be OSP-implementable [Ashlagi and Gonczarowski, 2018].

4. CONCLUDING REMARKS

Besides being easily implementable and commonly used, the Random Priority Rule (RP) is the only random assignment rule satisfying probabilistic monotonicity, symmetry, and ex post efficiency (Theorem 1). Probabilistic monotonicity has not been considered in the literature so far, but constitutes a natural strengthening of strategy-proofness (Proposition 2) on the domain of assignment problems.

The fact that our characterisation of RP relies on strengthening strategy proofness may be considered particularly compelling given that strategy proofness, together with symmetry and ex post efficiency, is insufficient to characterise RP [Erdil, 2014, Basteck and Ehlers, 2024]²⁹ – there are other random assignment rules on the domain of ordinal preferences that satisfy strategy-proofness, symmetry and ex post efficiency but are not welfare equivalent to RP .

Our characterizations of Random Priority assumes that agents with identical preferences should be treated symmetrically, which may be viewed as a fairness requirement and indeed the main motivation to consider randomization. However, depending on the application, equal treatment may not be desired – instead we may want to prioritize some agents over others and restrict symmetry to apply only between agents of equal priority. For example, dormitories may be assigned based on seniority, or school seats by favouring applicants with better grades.

If such priorities are given by a weak priority order of agents \succeq and we consider fairness to require that (i) no agent of strictly higher priority should (ex post) envy an agent of lower priority (i.e., prefer the other agent's assignment over their own), and (ii) between agents of equal priority assignment symmetry should hold, then our

²⁹The construction by Erdil [2014] relies on the existence of outside options; Basteck and Ehlers [2024] eliminate this domain restriction.

characterizations are readily adapted in that a random assignment rule is ex post efficient, probabilistically monotonic and fair in the above sense if and only if it is a non-uniform random priority rule where we randomize uniformly over the subset of deterministic priority rules compatible with \succeq . Put differently, ties in \succeq need to be broken uniformly at random if one wants to satisfy all three requirements.

It seems plausible that probabilistic monotonicity can similarly inform market designers' choices when agents' priorities are weak and *object specific*, e.g., in cases where some agent should be prioritized over another at some school while the reverse may hold at another – a setting in which the question of how to break ties has long been debated [Erdil and Ergin, 2008, Abdulkadiroğlu et al., 2009, Ashlagi et al., 2019, De Haan et al., 2023].

APPENDIX

Proof of Theorem 1. We first establish that the random priority rule RP satisfies the three properties in question. For symmetry and ex post efficiency this is immediate; for probabilistic monotonicity, this follows from the fact that RP is a convex combination of the $n!$ serial dictatorships f^\succ associated with all possible strict priority orders over agents, $\succ \in \Pi$. Each such deterministic rule f^\succ constitutes a Pareto efficient deferred acceptance rule (with objects' priorities over agents uniform) and hence satisfies Maskin monotonicity [Kojima and Manea, 2010]. As probabilistic (Maskin) monotonicity is preserved for convex combinations, this completes the *if* part of the proof.

For the other direction consider a preference profile where all agents share the same preferences, i.e., R^* such that $R_i^* = R_j^*$ for all $i, j \in N$. Then symmetry and ex post efficiency of f and RP imply that $f(R^*) = RP(R^*)$ as each Pareto efficient assignment is chosen with equal probability under both rules. Now, for an arbitrary preference profile $R \in \mathcal{R}$ let d_R denote the Kemeny-distance between R and R^* , i.e., the number of pairwise swaps of adjacent objects that are necessary to move from R to R^* . Define $\mathcal{R}^\#$ as the set of those profiles for which $f(\cdot)$ and $RP(\cdot)$ differ and, towards a contradiction, assume that $\mathcal{R}^\# \neq \emptyset$. Let δ be the minimal distance to R^* among profiles in $\mathcal{R}^\#$, i.e.,

$$\delta = \min_{R \in \mathcal{R}^\#} d_R,$$

and define \mathcal{R}^δ as the set of profiles $R \in \mathcal{R}^\#$ with $d_R = \delta$.

Next, for any profile $R \in \mathcal{R}^\delta$ consider those assignments where f and PR differ, i.e., $\mathcal{M}^{R,+} := \{\mu | f_\mu(R) > RP_\mu(R)\}$ and $\mathcal{M}^{R,-} := \{\mu | f_\mu(R) < RP_\mu(R)\}$. Any two assignments $\mu \in \mathcal{M}^{R,+}$, $\mu' \in \mathcal{M}^{R,-}$ will differ in the assignment of some objects to some agents, i.e., there will be a set of agents $I_{\mu,\mu'} \subset N$ such that $\mu_i \neq \mu'_i$ if and only if $i \in I_{\mu,\mu'}$.

Finally, let us consider a specific $R \in \mathcal{R}^\delta$ with two assignments $\mu \in \mathcal{M}^{R,+}$, $\mu' \in \mathcal{M}^{R,-}$ such that $|I_{\mu,\mu'}|$ is minimal, i.e., such that for any $\tilde{R} \in \mathcal{R}^\delta$ and $\tilde{\mu} \in \mathcal{M}^{\tilde{R},+}$, $\tilde{\mu}' \in \mathcal{M}^{\tilde{R},-}$ we have $|I_{\mu,\mu'}| \leq |I_{\tilde{\mu},\tilde{\mu}'}|$. If \mathcal{R}^δ is indeed non-empty, such a profile R , together with two assignments, exists – conversely, towards the desired contradiction, we will show that such an R does not exist.

For that, let us first define the set of objects assigned to different agents in μ and μ' as $Z = \{z \in O \mid \exists i \in I_{\mu,\mu'} : z = \mu_i \text{ or } z = \mu'_i\}$. Moreover, let us label the objects in Z as $Z = \{z_1, z_2, \dots, z_m\}$ such that $z_p R_i^* z_q$ if and only if $1 \leq p \leq q \leq m$. For each agent $i \in I_{\mu,\mu'}$ denote their R_i -most preferred object among all objects in Z as \bar{z}^i , i.e., $\bar{z}^i R_i x$ for all $x \in Z$. Next, since f is ex post efficient and since $f_\mu(R) > RP_\mu(R) \geq 0$, μ is a Pareto-efficient assignment given R ; by a symmetric argument, so is μ' . Thus, there is some agent $i \in I_{\mu,\mu'}$ who receives their R_i -most preferred among all objects in Z under μ , i.e., $\mu_i = \bar{z}^i$: first, for each $k \in I_{\mu,\mu'}$ with $\mu_k \neq \bar{z}^k$ it must be that all copies of \bar{z}^k are assigned, i.e., $|\{l \in N \mid \mu_l = \bar{z}^k\}| = q_{\bar{z}^k}$ as otherwise there would be a possible Pareto improvement on μ ; second, this implies that there is some $l \in I_{\mu,\mu'} \setminus \{k\}$ for whom $\mu_l = \bar{z}^k$; third, if there was no agent $i \in I_{\mu,\mu'}$ receiving object \bar{z}^i under μ , there would then be a top-trading cycle of agents in $I_{\mu,\mu'}$ and objects Z that would allow for a Pareto-improvement – a contradiction. Analogously some agent $j \in I_{\mu,\mu'}$ receives their R_j -most preferred object among objects in Z in μ' , i.e., $\mu'_j = \bar{z}^j$. Note that $i \neq j$ as, by the definition of $I_{\mu,\mu'}$, $\mu_i \neq \mu'_i$. Thus $|I_{\mu,\mu'}| \geq 2$. We consider three cases.

Case 1: $\mu_i = \bar{z}^i = z_1$, $\mu'_j = \bar{z}^j = z_1$. Then $\mu'_i = z_p \neq z_1$ for some p , and $z_1 R_i^* z_p$. Since $R \in \mathcal{R}^\delta$, it must be that for the strict upper contour set $U(z_p, R_i)$ we have $R_i|U(z_p, R_i) = R_i^*|U(z_p, R_i)$ – otherwise we could move to R'_i by reordering objects in $U(z_p, R_i)$ as in R_i^* , reducing the Kemeny distance to $d_{R'} < \delta$ while keeping $f_{\mu'}(R'_i, R_{-i}) = f_{\mu'}(R) < RP_{\mu'}(R) = RP_{\mu'}(R'_i, R_{-i})$ (where the equalities follow from probabilistic monotonicity). Symmetrically, we find $R_i|L(z_1, R_i) = R_i^*|L(z_1, R_i)$ and hence $R_i = R_i^*$. In the same way, we find that $R_j = R_j^* = R_i^*$.

Now, if $|I_{\mu,\mu'}| = 2$ so that μ and μ' differ only in the assignment of two individuals, then $\mu'_i = \mu_j$ as Pareto efficiency demands that μ'_i is the R_i^* -most preferred object among all those still available given $\mu'_j = z_1$ and $\mu'_k = \mu'_k$, $k \neq i, j$ while the same holds for μ_j (given $\mu_i = z_1$ and $\mu_k = \mu'_k$ for all $k \neq i, j$). Hence this falls under the sub-case we considered in the main text.

If instead $|I_{\mu,\mu'}| \geq 3$ consider $\hat{\mu}$ that differs from μ' only in that the roles of i and j are reversed, i.e., $\hat{\mu}_i = \mu'_j$, $\hat{\mu}_j = \mu'_i$, and $\hat{\mu}_k = \mu'_k$ for $k \neq i, j$. By symmetry we still have $f_{\hat{\mu}}(R) = f_{\mu'}(R) < RP_{\mu'}(R) = RP_{\hat{\mu}}(R)$ so that $\hat{\mu} \in \mathcal{M}^{R,-}$, yet μ and $\hat{\mu}$ differ in the assignment of fewer agents than μ and μ' (as i is assigned the same object z_1 under both μ and $\hat{\mu}$) – hence μ and μ' were not chosen so as to minimize $|I_{\mu,\mu'}|$ after all, a contradiction.

Case 2: $\mu_i = \bar{z}^i = z_1$, $\mu'_j = \bar{z}^j = z_q \in Z$ with $q > 1$. As in case 1, this implies $R_i = R_i^*$ while for j we find that $R_j|L(z_q, R_i) = R_j^*|L(z_q, R_i)$. Hence z_1 is ranked adjacent and immediately below z_q in R_j . Moreover, this implies that in μ , j is assigned z_1 – otherwise $\mu_j = z_p$ ranked below z_1 in R_j and we could move to R'_j by reordering objects in $U(z_p, R_i)$ as in R_j^* , reducing the Kemeny distance to $d_{R'} < \delta$ while keeping $f_\mu(R'_j, R_{-j}) \neq RP_\mu(R'_j, R_{-i})$. Pareto efficiency of μ implies that all copies of $z_q = \bar{z}^j$ must be assigned to some agent in μ so that there is some $k \in I_{\mu, \mu'}$ for whom we have $\mu_k = z_q$. Further, once more by Pareto efficiency of μ , we have $z_q P_k z_1$, as otherwise a trade by j and k would constitute a Pareto improvement on μ . But then, as for j , we find that z_1 is ranked adjacent and immediately below z_q in R_k and that in μ' , k is assigned z_1 . As in Case 1, we may then consider $\hat{\mu}$ that differs from μ' only in that the roles of j and k are reversed, i.e., $\hat{\mu}_j = z_1 = \mu'_k$, $\hat{\mu}_k = z_q = \mu'_j$ and arrive at a contradiction in that $\hat{\mu} \in \mathcal{M}^{R,-}$, yet μ and $\hat{\mu}$ differ in the assignment of fewer agents than μ and μ' . In the same way, the case $\mu_i = z_p \in Z$ with $p > 1$, $\mu'_j = z_1$ can be shown to lead to a contradiction.

Case 3: $\mu_i = \bar{z}^i = z_p \in Z$, $\mu'_j = \bar{z}^j = z_q \in Z$ with $p, q > 1$. As in case 2, we find that $R_j|L(z_q, R_i) = R_j^*|L(z_q, R_i)$, that z_1 is ranked adjacent and immediately below z_q in R_j , as well as that in μ , j is assigned z_1 . In the same way, we find that $R_i|L(z_p, R_i) = R_i^*|L(z_p, R_i)$, that z_1 is ranked adjacent and immediately below z_p in R_i , as well as that in μ' , i is assigned z_1 .

If $z_p = z_q$ then i and j have identical preferences. Thus, as in case 1, if $|I_{\mu, \mu'}| = 2$, this falls under the sub-case considered in the main text. If instead $|I_{\mu, \mu'}| \geq 3$ consider $\hat{\mu}$ that differs from μ' only in that the roles of i and j are reversed, i.e., $\hat{\mu}_i = \mu'_j$, $\hat{\mu}_j = \mu'_i$, and $\hat{\mu}_k = \mu'_k$ for $k \neq i, j$. By symmetry we still have $f_{\hat{\mu}}(R) = f_{\mu'} < RP_{\mu'} = RP_{\hat{\mu}}(R)$, yet μ and $\hat{\mu}$ differ in the assignment of fewer agents than μ and μ' – a contradiction.

If instead $z_p \neq z_q$, then as in Case 2 there exist another agent k in $I_{\mu, \mu'}$ with $\mu_k = z_q$. Since μ is Pareto efficient, and $z_q P_j z_1$, we have $z_q P_k z_1$. In the same way as before, this implies that z_1 is ranked adjacent and immediately below z_q in R_k and that $\mu'_k = z_1$. As in Case 2, we may then consider $\hat{\mu}$ that differs from μ' only in that the roles of j and k are reversed, i.e., $\hat{\mu}_j = z_1 = \mu'_k$, $\hat{\mu}_k = z_q = \mu'_j$ and arrive at a contradiction in that μ and $\hat{\mu}$ differ in the assignment of fewer agents than μ and μ' . \square

REFERENCES

Atila Abdulkadiroğlu and Tayfun Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3): 689–701, 1998.

- Atila Abdulkadiroğlu, Parag A Pathak, and Alvin E Roth. Strategy-proofness versus efficiency in matching with indifference: Redesigning the NYC high school match. *American Economic Review*, 99(5):1954–1978, 2009.
- Itai Ashlagi and Yannai A Gonczarowski. Stable matching mechanisms are not obviously strategy-proof. *Journal of Economic Theory*, 177:405–425, 2018.
- Itai Ashlagi, Afshin Nikzad, and Assaf Romm. Assigning more students to their top choices: A comparison of tie-breaking rules. *Games and Economic Behavior*, 115:167–187, 2019.
- Sophie Bade. Fairness and group-strategyproofness clash in assignment problems. *Journal of Economic Theory*, 165:257–262, 2016.
- Sophie Bade. Random serial dictatorship: the one and only. *Mathematics of Operations Research*, 45(1):353–368, 2020.
- Sophie Bade and Yannai A Gonczarowski. Gibbard-satterthwaite success stories and obvious strategyproofness. *arXiv preprint arXiv:1610.04873*, 2016.
- Christian Basteck and Lars Ehlers. On (constrained) efficiency of strategy-proof random assignment. *Available at SSRN*, 2024.
- Garrett Birkhoff. Three observations on linear algebra. *Univ. Nac. Tacuman, Rev. Ser. A*, 5:147–151, 1946.
- Anna Bogomolnaia and Hervé Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100:295–328, 2001.
- Gabriel Carroll. A general equivalence theorem for allocation of indivisible objects. *Journal of Mathematical Economics*, 51:163–177, 2014.
- Monique De Haan, Pieter A Gautier, Hessel Oosterbeek, and Bas Van der Klaauw. The performance of school assignment mechanisms in practice. *Journal of Political Economy*, 131(2):388–455, 2023.
- Aytek Erdil. Strategy-proof stochastic assignment. *Journal of Economic Theory*, 151:146–162, 2014.
- Aytek Erdil and Haluk Ergin. What’s the matter with tie-breaking? Improving efficiency in school choice. *The American Economic Review*, 98(3):669–689, 2008.
- Haluk I Ergin. Efficient resource allocation on the basis of priorities. *Econometrica*, 70(6):2489–2497, 2002.
- Allan Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3):665–681, 1977.
- Donald E Knuth. An exact analysis of stable allocation. *Journal of Algorithms*, 20(2):431–442, 1996.
- Fuhito Kojima and Mihai Manea. Axioms for deferred acceptance. *Econometrica*, 78(2):633–653, 2010.

- Thiam Lee and Jay Sethuraman. Equivalence results in the allocation of indivisible objects: A unified view. *Columbia University, Unpublished mimeo*, 2011.
- Shengwu Li. Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–3287, 2017.
- Andrew McLennan. Ordinal efficiency and the polyhedral separating hyperplane theorem. *Journal of Economic Theory*, 105:435–449, 2002.
- Timo Mennle and Sven Seuken. Partial strategyproofness: Relaxing strategyproofness for the random assignment problem. *Journal of Economic Theory*, 191:105–144, 2021.
- Thayer Morrill. An alternative characterization of top trading cycles. *Economic Theory*, 54:181–197, 2013.
- Eitan Muller and Mark A Satterthwaite. The equivalence of strong positive association and strategy-proofness. *Journal of Economic Theory*, 14(2):412–418, 1977.
- Szilvia Pápai. Strategyproof assignment by hierarchical exchange. *Econometrica*, 68(6):1403–1433, 2000.
- Parag A Pathak and Jay Sethuraman. Lotteries in student assignment: An equivalence result. *Theoretical Economics*, 6(1):1–17, 2011.
- Marek Pycia and Peter Troyan. A theory of simplicity in games and mechanism design. *Working paper series/Department of Economics*, (393), 2021.
- Marek Pycia and M Utku Ünver. Ordinal simplicity and auditability in discrete mechanism design. *Available at SSRN*, 2023.
- Lloyd Shapley and Herbert Scarf. On cores and indivisibility. *Journal of mathematical economics*, 1(1):23–37, 1974.
- Tayfun Sönmez. Implementation in generalized matching problems. *Journal of Mathematical Economics*, 26(4):429–439, 1996.
- Koji Takamiya. Coalition strategy-proofness and monotonicity in shapley–scarf housing markets. *Mathematical Social Sciences*, 41(2):201–213, 2001.
- Koji Takamiya. Domains of social choice functions on which coalition strategy-proofness and maskin monotonicity are equivalent. *Economics Letters*, 95(3):348–354, 2007.
- Jun Zhang. Efficient and fair assignment mechanisms are strongly group manipulable. *Journal of Economic Theory*, 180:167–177, 2019.