

Basteck, Christian

Working Paper

An Axiomatization of the Random Priority Rule

Discussion Paper, No. 502

Provided in Cooperation with:

University of Munich (LMU) and Humboldt University Berlin, Collaborative Research Center
Transregio 190: Rationality and Competition

Suggested Citation: Basteck, Christian (2025) : An Axiomatization of the Random Priority Rule, Discussion Paper, No. 502, Ludwig-Maximilians-Universität München und Humboldt-Universität zu Berlin, Collaborative Research Center Transregio 190 - Rationality and Competition, München und Berlin

This Version is available at:

<https://hdl.handle.net/10419/336552>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

An Axiomatization of the Random Priority Rule

Christian Basteck (WZB Berlin)

Discussion Paper No. 502

September 01, 2025

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 at most one. To ensure fairness in the absence of monetary transfers, we consider random assignments. Random Priority, also known as Random Serial Dictatorship, is characterized by equal-treatment-of-equals, ex-post efficiency and probabilistic monotonicity – whenever preferences change so that a given deterministic assignment is ranked higher by all agents, the probability of that assignment arising should not decrease. Probabilistic monotonicity implies strategy-proofness for random assignment problems and is equivalent to it on the universal domain of strict preferences; 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 leads 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.

[†] WZB Berlin Social Science Center, Berlin, Germany. christian.basteck@wzb.eu. I thank Nick Arnosti, Anna Bogomolnaia, Sean Horan, Bettina Klaus, Hervé Moulin, Szilvia Pápai, Juan Sebastián Pereyra, Ludvig Sinander, Alex Teytelboym, William Thomson, Alexander Westkamp and, in particular, Lars Ehlers for inspiring discussions and helpful comments. Financial support by the Deutsche Forschungsgemeinschaft through CRC TRR 190 (project number 280092119) is gratefully acknowledged.

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 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.

This paper identifies another normatively appealing, property: the Random Priority rule 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 equal-treatment-of-equals, (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 in a cardinal framework where preferences are described by von Neumann–Morgenstern utility functions.

Note that this 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 the universal domain of strict preferences,¹ probabilistic monotonicity and strategy-proofness are equivalent – hence our characterisation of random priority (also known as random serial dictatorship) 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 \[1977\]](#).

Since probabilistic monotonicity is stronger than strategy-proofness on the domain of assignment problems, one may ask whether it can be replaced by strategy-proofness in conjunction with an additional property. This is indeed the case: replacing probabilistic monotonicity with strategy-proofness and a weak, object-wise non-bossiness notion yields an alternative characterization of RP (Theorem 1).

¹In contrast, our setting assumes that agents are indifferent between different assignments as long as they themselves receive the same object.

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.² Following up on this surprising result, several authors have shown how randomization over certain deterministic rules, so as to symmetrize the treatment of agents, yields the Random Priority rule [[Pathak and Sethuraman, 2011](#), [Lee and Sethuraman, 2011](#), [Carroll, 2014](#), [Bade, 2020](#)]:³ in the most general formulation, due to [Bade \[2020\]](#), symmetrizing any strategy-proof, non-bossy, and efficient deterministic rule yields RP .

Interestingly, for *deterministic* rules, strategy-proofness and non-bossiness are equivalent to group-strategy-proofness [[Pycia and Ünver, 2023](#)], to Maskin monotonicity [[Takamiya, 2007](#)],⁴ and thus also to probabilistic monotonicity.

As randomization over efficient, strategy-proof, and non-bossy (and hence Maskin monotonic) deterministic rules preserves both (ex post) efficiency, and probabilistic monotonicity, our characterization implies the previous equivalence results from [Abdulkadiroğlu and Sönmez \[1998\]](#) to [Bade \[2020\]](#) (Corollary 1). Moreover, it generalizes these earlier results in that our framework allows for multiple copies of objects, say multiple seats at the same school, as well as the presence (or absence) of outside options.

[Erdil \[2014\]](#) was the first to show that RP is not characterized by equal-treatment-of-equals, ex-post efficiency and strategy-proofness in the presence of outside options – a long-standing conjecture since [Bogomolnaia and Moulin \[2001\]](#) proved this to be the case for 3 agents and 3 acceptable objects. [Basteck and Ehlers \[2025\]](#) show RP not to be characterized by these properties even in the absence of outside options. Hence, to characterize RP in conjunction with equal-treatment-of-equals and ex post efficiency, we are forced to strengthen strategy-proofness, for example to probabilistic monotonicity.

[Pycia and Troyan \[2024\]](#) 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

²For this, they assume an equal number of agents and distinct objects.

³All these assume a unit supply of objects, i.e., do not consider the possibility of copies of objects between which agents are indifferent.

⁴[Pycia 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 \[1974\]](#)).

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 [2024] 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 Further, while Pycia and Troyan [2024] consider a setting with an equal number of agents and objects and assume that all objects have unit supply, the characterization presented in this paper allows also for multiple copies of objects, the presence (or absence) of outside options, as well as an unequal number of objects and agents.

Finally, (weak) Maskin monotonicity allows to characterize two of the most prominent families of *deterministic* assignment rules. Kojima and Manea [2010] characterize deferred acceptance rules (DA) 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 rules (TTC) by efficiency, weak Maskin monotonicity, independence of irrelevant rankings, and mutual best. Since TTC-rules satisfy Maskin monotonicity [Pápai, 2000], one can replace weak Maskin monotonicity by monotonicity in the characterizations. Note that while DA and TTC satisfy Maskin monotonicity – and hence probabilistic monotonicity – neither are OSP-implementable in general.

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.

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). For each object $x \in O$, let q_x denote the number of copies. If one wants to allow for the possibility that some agents are not assigned any proper object, we may do so by introducing a null-object, denoted $\emptyset \in O$,⁸ and set $q_\emptyset = |N|$. Thus, it is without loss of generality to assume $\sum_{x \in O} q_x \geq |N|$.

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

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., DA with Ergin-acyclic priorities [Ergin, 2002] or that satisfies bounded invariance Basteck and Ehlers [2024].

⁸One may also interpret \emptyset as representing agent-specific outside options.

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 $R = (R_1, \dots, R_n)$. Refer to \mathcal{R} as the unrestricted domain of strict preferences. In contrast, if O includes the null object \emptyset and if, *in addition*, we restrict agents preferences by requiring that \emptyset is ranked last by each agent, we denote this restricted domain as $\underline{\mathcal{R}}$. For a given object x and given R_i we write $U(x, R_i) = \{y \in O | y P_i x\}$ for i 's (strict) upper contour set at x and denote the (strict) lower contour set as $L(x, R_i) = \{y \in O | x P_i y\}$.

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 $R'_i | O \setminus \{x\} = R_i | O \setminus \{x\}$, $R'_i | O \setminus \{y\} = R_i | O \setminus \{y\}$ 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 $\mu : N \rightarrow O$ such that each object $x \in O$ is assigned to at most q_x agents, i.e., $|\{i \in N | \mu_i = x\}| \leq q_x$.¹² 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 $j \in N$. We write $\mathcal{PO}(R)$ for 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 *equal-treatment-of-equals* 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, different deterministic assignments that merely permute the assignments of agents with identical preferences arise with the same probability.¹³

A random assignment rule f maps preference profiles to random assignment, i.e., $f : \mathcal{R} \rightarrow \Delta(\mathcal{M})$. For a given $R \in \mathcal{R}$ and $\mu \in \mathcal{M}$, let $f_\mu(R)$ denote the probability assigned to μ under $f(R)$. We say that f is *ex-post efficient*, if $f(R)$ is ex-post efficient for every $R \in \mathcal{R}$. It satisfies *equal-treatment-of-equals*, if $f(R)$ satisfies

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 $i \in N$.

¹³Alternatively, if one does not wish to distinguish between random assignments that are welfare equivalent, one may define *equal-treatment-of-equals* with respect to agents' individual object assignment probabilities rather than the probabilities of choosing complete assignments – see Appendix A.2.

equal-treatment-of-equals for every $R \in \mathcal{R}$. It is deterministic, if $|supp(f(R))| = 1$ for every $R \in \mathcal{R}$.

Besides these intra-profile axioms we will consider four inter-profile axioms that describe how random assignments should respond to changes in agents' preferences. The first such axioms demands that an assignment should be chosen with higher probability as it becomes more preferred in the eyes of all agents. For that 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 $f_\mu(R') \geq f_\mu(R)$.

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

Probabilistic monotonicity strengthens (stochastic dominance) strategy-proofness (see Proposition 2). To define the latter formally, let p^{ia} denote the probability with which agent i is assigned object a under random assignment p , and let $p^i = (p^{ia})_{a \in O}$ denote i 's individual random assignment. 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 $i \in N$. 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 $f^i(R'_i, R_{-i})$.¹⁵

For deterministic assignment rules, non-bossiness demands that if a change in an agents' preferences does not change their own assigned object, it should not change the overall assignment of objects to (other) agents. Our next axiom generalizes this requirement to random assignment rules: A random assignment rule f is *object-wise non-bossy* if for all $z \in O$, $R, R'_i \in \mathcal{R}$ we have

$$f^{iz}(R) = f^{iz}(R'_i, R_{-i}) \implies f_\mu(R) = f_\mu(R'_i, R_{-i}), \forall \mu \in \{\mu' \in \mathcal{M} | \mu'_i = z\}.$$

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 $f(R') = \mu$.

¹⁵Sd-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.

In words, if the change in one agent's preferences does not affect the probability with which they are assigned a particular object z , then, conditional on her being assigned that object, it does not affect the probability distribution over (all agents') assignments. Note that for deterministic assignment rules, object-wise non-bossiness is equivalent to non-bossiness, since it has implications only if there is some z such that $f^{iz}(R) = f^{iz}(R'_i, R_{-i}) = 1$, in which case it demands that the unique assignment μ for which $f_\mu(R) = 1$, arises with probability 1 also under $f(R'_i, R_{-i})$.

Weakening the axiom in that it only needs to hold when the change in an agents' preferences is due to a swap of two adjacently ranked objects, we say that a random assignment rule f is *weakly object-wise non-bossy*¹⁶ if for all $z \in O$, $R \in \mathcal{R}$, and $R'_i \in N_{R_i}$ we have

$$f^{iz}(R) = f^{iz}(R'_i, R_{-i}) \implies f_\mu(R) = f_\mu(R'_i, R_{-i}), \forall \mu \in \{\mu' \in \mathcal{M} \mid \mu'_i = z\}.$$

The last intra-profile axiom was introduced by Gibbard [1977] in a setting where agents have strict preferences over all possible outcomes.¹⁷ A random assignment rule f is *pairwise responsive* if for all $x, y \in O$, $R \in \mathcal{R}$, $R'_i \in N_{R_i}$ such that xP_iy and yP'_ix , and $\mu \in \mathcal{M}$ we have

$$\mu_i \notin \{x, y\} \implies f_\mu(R) = f_\mu(R'_i, R_{-i}).$$

In words, swapping two adjacently ranked objects in the preference order of some agent may affect the probability of assignments where that agent is assigned one of the two objects but should not change the probability of any assignment where the agent is assigned some other object.

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.¹⁸ Then the Random Priority rule is defined by $RP(R) = \frac{1}{n!} \sum_{\triangleright \in \Pi} f^\triangleright(R)$ for all $R \in \mathcal{R}$. Note that any f^\triangleright satisfies Maskin monotonicity [Kojima and Manea, 2010], and hence, RP satisfies probabilistic monotonicity. Analogously, one finds that as any f^\triangleright is strategy-proofness, so is RP .

¹⁶For deterministic, *strategy-proof* assignment rules, weak object-wise non-bossiness can again be shown to be equivalent to non-bossiness.

¹⁷In contrast, in the assignment problem agents have strict preferences over their own assignment.

¹⁸For 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 $f_{i_l}^\triangleright(R)$).

3. RESULTS

3.1. Axiomatization. We are now set to state our main result: ex-post efficiency, equal-treatment-of-equals and some (combination of) inter-profile axiom(s) characterize the Random Priority rule. Recall that \mathcal{R} denotes the unrestricted domain of strict preferences while, in the presence of a null-object, $\underline{\mathcal{R}} \subsetneq \mathcal{R}$ is restricted by the assumption that all agents rank the null-object last.¹⁹

Theorem 1. *Consider a random assignment rule f on \mathcal{R} (or $\underline{\mathcal{R}}$), satisfying ex post efficiency and equal-treatment-of-equals. The following statements are equivalent:*

The random assignment rule f

- (1) satisfies probabilistic monotonicity.*
- (2) satisfies both strategy-proofness and weak object-wise non-bossiness.*
- (3) satisfies pairwise responsiveness.*
- (4) is the Random Priority rule (on \mathcal{R} or $\underline{\mathcal{R}}$, respectively).*

The characterization based on probabilistic monotonicity might be seen as the most compelling on normative grounds, as the axiom requires a rule to be responsive to changes in agents' preferences: an assignment should arise with weakly higher probability as it becomes more preferred in the eyes of agents. The second characterization may be seen as useful in that it pins down the difference between the Random Priority rule and other random assignment rules that satisfy equal-treatment-of-equals, ex post efficiency and strategy-proofness. For example Erdil [2014] constructs a rule on the domain of random assignment problems with outside options that (weakly) stochastically R -dominates the random priority rule at every profile R in the domain. Similarly, Basteck and Ehlers [2025] construct a rule for assignment problems without outside options that dominates the Random Priority rule for some profiles and yields assignments where agents receive a higher ranked object on average – if one wants to further explore possibilities of welfare improvements among random assignment rules that satisfy equal-treatment-of-equals, ex post efficiency, and strategy-proofness, Theorem 1 delineates the set of rules to consider: those that violate weak object-wise non-bossiness. Last, the characterization based on pairwise responsiveness pinpoints the precise property – implied by both (1) probabilistic monotonicity and (2) strategy-proofness in conjunction with weak object-wise non-bossiness – that is used in the main step of the proof.

The proof proceeds as follows. We first verify that (4) \implies (1), (1) \implies (2), and (2) \implies (3) – these steps are straightforward and relegated to appendix. The crucial step is to verify that (3) \implies (4) – which we sketch below for the case of

¹⁹While we defined a random assignment rule, as well our axioms, with respect to \mathcal{R} , all definitions carry over to the domain $\underline{\mathcal{R}}$ by simply replacing \mathcal{R} by $\underline{\mathcal{R}}$.

an equal number of agents and objects and with unit supply of each object, i.e., $|N| = |O| = n$ and $q_x = 1$ for all $x \in O$.

First, towards a contradiction, suppose $f \neq RP$ and introduce some additional notation for the set of assignments that arise with different probability under f and RP – for any $R \in \mathcal{R}$,

$$\mathcal{M}^\#(R) := \{\mu \in \mathcal{M} \mid f_\mu(R) \neq RP_\mu(R)\}.$$

Second, observe that ex-post efficiency implies $f_\mu(R) = 0 = RP_\mu(R)$ whenever μ is inefficient (given R). Hence, whenever $\mu \in \mathcal{M}^\#(R)$ for some μ and R , we know that μ is efficient at R . Third, recall that any assignment μ , efficient at R , is the outcome of some (not necessarily unique) deterministic priority rule at R where agents choose their most preferred object among those still available according to some deterministic priority order.

As $f \neq RP$, there exists some R and $\mu \in \mathcal{M}^\#(R)$. By the above, we know that μ must be efficient and hence the outcome under some deterministic priority rule. Let i be the agent who chooses last under the associated priority order and let $\mu_i = z \in O$. Since i chose last, we know that $\mu_j P_j z$ for all $j \neq i$.

Step A: Now consider R' where all $j \neq i$ move z to the bottom of their preference order, in a series of pairwise swaps of adjacent objects. Since these do not involve μ_j , pairwise responsiveness implies that $f_\mu(R') = f_\mu(R) \neq RP_\mu(R) = RP_\mu(R')$. Hence, f and RP differ not only for some profile $R \in \mathcal{R}$ but also for some profile in a restricted domain where most agents agree on the ranking of some object z , namely

$$\tilde{\mathcal{R}} = \{R \in \mathcal{R} \mid \forall j \neq i, x \in O : x R_j z\}.$$

Step B: Next, we will show that we may further narrow down this subdomain while preserving $f \neq RP$ (for some profile of the subdomain) – either there exist profiles where i also ranks z last and for which f and RP differ, or we can fix agent i 's preferences and know that under any assignment arising with different probability under f and RP (at some profile of the remaining subdomain), i will be assigned z .

Case (i): Suppose there exists $R \in \tilde{\mathcal{R}}$ and $\mu' \in \mathcal{M}^\#(R)$ such that $\mu'_i P_i z$. Then we may move z to the bottom of i 's preference order, in a series of pairwise swaps of adjacent objects. Call the new profile R^1 . Since the swaps did not involve μ'_i , pairwise responsiveness implies that $f_{\mu'}(R^1) = f_{\mu'}(R) \neq RP_{\mu'}(R) = RP_{\mu'}(R^1)$. Hence, f and RP differ for some profile in the restricted domain where all agents, including i , rank object z last, namely

$$\mathcal{R}^1 = \{R \in \tilde{\mathcal{R}} \mid \forall x \in O : x R_i z\}.$$

Case (ii): Suppose there exists *no* $R \in \tilde{\mathcal{R}}$ and $\mu' \in \mathcal{M}^\#(R)$ such that $\mu'_i P_i z$. If instead there exist $R \in \tilde{\mathcal{R}}$ and $\mu' \in \mathcal{M}^\#(R)$ such that $z P_i \mu'_i$, we may move z to the top

of i 's preference order, in a series of pairwise swaps of adjacent objects. Call the new profile R' . Since these swaps did not involve μ'_i , pairwise responsiveness implies that $f_{\mu'}(R') \neq RP_{\mu'}(R')$. Hence, by ex-post efficiency of f and RP , μ' is efficient at R' . Yet, there is a Pareto-improving trade where i exchanges μ'_i for z with j for whom $\mu'_j = z$ and who, as any $j \neq i$, ranks z last – a contradiction. Thus, for any $R \in \tilde{\mathcal{R}}$ and $\mu' \in \mathcal{M}^\#(R)$ we have $\mu'_i = z$. To narrow down further a subdomain for which $f \neq RP$, pick any $R^1 \in \tilde{\mathcal{R}}$ for which $\mathcal{M}^\#(R^1)$ is non-empty and define

$$\mathcal{R}^1 = \{R \in \tilde{\mathcal{R}} \mid R_i = R_i^1\}.$$

Regardless of whether we arrived at \mathcal{R}^1 via case (i) or (ii), we can then consider $R \in \mathcal{R}^1$ and $\mu \in \mathcal{M}^\#(R)$, note that μ is efficient and hence the outcome of some priority rule at R , and identify the last agent according to the priority order, who chose an object $y \neq z$ – call that agent j . Analogous to Step *A* above, we may then move y down to a position immediately above z for most agents – if Case (i) applied in the above, we may do so for all agents other than j , if Case (ii) applied, we may do so for all agents other than i and j .

Next, analogous to Step *B* above, we may move down y to just above z also for j , as in Case (i), or fix the preferences of j , knowing that at all remaining preference profiles where f and RP differ, j will be assigned y , as in Case (ii) – in both cases we have arrived at a further restricted preference domain \mathcal{R}^2 . Define $O_2 := \{z, y\}$ as the set of objects we pushed down in the two iterations, N_2^{ii} the set of agents for whom Case (ii) applied (so $N_2^{\text{ii}} \subseteq \{i, j\}$) and define $N_2^{\text{i}} := N \setminus N_2^{\text{ii}}$ as the complement. Importantly, for all $k \in N_2^{\text{i}}$, and $R \in \mathcal{R}^2$ we know that $oR_k y R_k z$ for all $o \in O \setminus O_2$.

Now, consider \mathcal{R}^{n-1} arrived at by iterating the steps above, together with the associated sets of objects O_{n-1} and agents $N_{n-1}^{\text{ii}}, N_{n-1}^{\text{i}}$. Take any $R \in \mathcal{R}^{n-1}$ for which $\mathcal{M}^\#(R)$ is non-empty. Since the total probabilities over all assignments sum to one, there exist $\mu, \mu' \in \mathcal{M}^\#(R)$ for which $f_\mu(R) > RP_\mu(R)$ and $f_{\mu'}(R) < RP_{\mu'}(R)$. Now, by construction, $\mu_k = \mu'_k$ for all $k \in N_{n-1}^{\text{i}}$. Hence, μ and μ' differ only in the assignment of agents $k \in N_{n-1}^{\text{i}}$. However, all these agents rank the same set of $n-1$ objects at the bottom, and do so in the same order, so that $R_k = R_l$ for all $k, l \in N_{n-1}^{\text{i}}$. Thus, equal-treatment-of-equals for RP implies $f_\mu(R) > RP_\mu(R) = RP_{\mu'}(R) > f_{\mu'}(R)$, which violates equal-treatment-of-equals for f – a contradiction that concludes the sketch of proof for the case of $q_x = 1$ for all $x \in O$ (no copies of objects) and $\sum_{x \in O} q_x = |O| = |N|$. For the detailed proof of the general case, allowing for an unequal number of objects and agents as well as multiple copies per object, see the appendix. There we also explain how the proof may be further generalized to a domain where all agents are assumed to rank the same (null) object at the bottom

3.2. Unifying and generalizing previous equivalence results. [Abdulkadiroğlu and Sönmez \[1998\]](#) proved that RP is equivalent to the core from random endowments.²⁰ Thus, in a setting with an equal number of agents and objects, it arise not only from a (uniform) randomization over $|N|!$ many deterministic priority rules (where agents chose the most preferred among remaining objects according to a deterministic priority order), but equivalently from a randomization over $|N|!$ many TTC-rules that permute agents' roles by permuting their initial endowments. Note that any deterministic priority rule [[Kojima and Manea, 2010](#)] as well as any TTC-rule [[Pápai, 2000](#)], satisfies Maskin monotonicity and hence probabilistic monotonicity. Randomizing over $|N|!$ many permutations of each rule, yields a random assignment rule that preserves probabilistic monotonicity as well as ex post efficiency. Moreover, by construction, it also satisfies equal-treatment-of-equals. But then, by Theorem 1, it must coincide with RP .

In the same way, [Pathak and Sethuraman \[2011\]](#), [Lee and Sethuraman \[2011\]](#), [Carroll \[2014\]](#), [Bade \[2020\]](#) consider progressively larger classes of deterministic rules and show that for each such rule, permuting agents' roles and randomizing over all permutations of the rule yields RP . Since in all these results, the initial deterministic rules under consideration are Maskin monotonic, the results follow from Theorem 1 in the same way as described above.

To state this formally, let f be a deterministic assignment rule on \mathcal{R} or $\underline{\mathcal{R}}$ and define φ on the same domain by permuting agents roles in f and randomizing over all these permutations uniformly, i.e., for all preference profiles R and assignments μ

$$\varphi_{\mu}(R) = \frac{1}{|N|!} \sum_{\pi \in \Pi} f_{\mu \cdot \pi}(R_{\pi(1)}, R_{\pi(2)}, \dots, R_{\pi(|N|)}),$$

where Π denotes the set of all permutations of agents in N ²¹ and $\mu \cdot \pi$ denotes the accordingly permuted assignment, i.e., $(\mu \cdot \pi)_i = \mu_{\pi(i)}$ for all $i \in N$.

Corollary 1. *Consider a deterministic assignment rule f on \mathcal{R} or $\underline{\mathcal{R}}$ and the associated random assignment rule φ . If f is Maskin-monotonic (or equivalently strategy-proof and non-bossy) and ex post efficient, then φ is the Random Priority rule RP .*

Note that Theorem 1, and hence Corollary 1, allow for multiple copies of objects, thus generalizing the existing equivalence results.²² Moreover, Theorem 1 improves on these earlier results also in that it dispenses with the restriction to assignment

²⁰[Knuth \[1996\]](#) independently proved an analogous result.

²¹I.e., all bijections $\pi : N \rightarrow N$.

²²[Abdulkadiroğlu and Sönmez \[1998\]](#), [Pathak and Sethuraman \[2011\]](#), [Lee and Sethuraman \[2011\]](#), [Carroll \[2014\]](#) consider an equal number of agents and objects and no copies. [Bade \[2020\]](#) allows for an unequal number of agents and objects as well as a null-object in abundant supply with the restriction that the null-object is always ranked last – this corresponds to the domain $\underline{\mathcal{R}}$ under the restriction that \emptyset is the only object with multiple copies.

rules that arise from randomization over deterministic rules and instead considers random assignment rules directly. For example, Pycia and Ünver [2015] show that the set of strategy-proof random assignment rules is strictly larger than the set of randomizations over deterministic strategy-proof assignment rules, so that, in general, restricting to randomizations over deterministic rules is not innocuous when trying to identify random assignments satisfying appealing properties.

3.3. 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 $R^u = (R_1^u, \dots, R_n^u)$.²³

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 is ex post efficient if $f(u)$ is ex post efficient for all $u \in \mathcal{U}$ and satisfies probabilistic monotonicity, if probabilistic monotonicity is satisfied for the associated strict preferences R^u .²⁴

A random assignment p satisfies equal-treatment-of-equals 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 equal-treatment-of-equals as defined with respect to R . A random assignment rule satisfies equal-treatment-of-equals on \mathcal{U} if $f(u)$ satisfies equal-treatment-of-equals for any $u \in \mathcal{U}$.

Despite the fact that equal-treatment-of-equals with respect to u is weaker than equal-treatment-of-equals 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.

²³I.e. for all $x, y \in O$ we have $xR_i^u y \Leftrightarrow u_i(x) \geq u_i(y)$.

²⁴I.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 $f_\mu(u') \geq f_\mu(u)$.

Proposition 1. *Let f be a random assignment rule on \mathcal{U} . Then f satisfies equal-treatment-of-equals 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 satisfy equal-treatment-of-equals with respect to any $u \in \mathcal{U}$, f has to satisfy equal-treatment-of-equals with respect to the associated preference profiles R^u . The claim then follows from Theorem 1. \square

3.4. Probabilistic monotonicity vs. strategy-proofness. 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 and its relation to strategy-proofness. For that, observe that 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].

Swap monotonicity demands that if an agent swaps to adjacent objects in their preference order, then they should receive the object that is moved up in their ranking with weakly higher probability. In other words, the *sum* over the probabilities of all assignments where they receive the object should be weakly larger. More demanding, probabilistic monotonicity demands that the probability of *each* assignment where they receive the object that has moved up should increase – instead of comparing the sum of probabilities associated with certain assignments, we now compare each summand. Strengthening the other two axioms analogously, we arrive at the following: 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 $\mu_i \in L(y, R_i)$.

Note that pairwise responsiveness is equivalent to the conjunction of upper and lower assignment invariance. Keeping the two separate, however, allows us to formulate the next proposition in close analogy to [Mennle and Seuken, 2021, Theorem 1].

Proposition 2. *A random assignment rule 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 ²⁵ 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.

²⁵I.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 2. *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 that yields the same individual object assignment probabilities, 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] and Theorem 1 in [Basteck and Ehlers, 2025] which show RP not to be characterized by equal-treatment-of-equals, ex post efficiency, and strategy-proofness. While the construction by 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 [2025] show that RP is not characterized by equal-treatment-of-equals ex post efficiency, and strategy-proofness even in the absence of such outside options.

3.5. Probabilistic monotonicity on the universal domain of strict preferences.

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 on the universal domain of strict preferences, probabilistic monotonicity and strategy-proofness are equivalent. Formally, let N be a finite set of agents, O be a finite set of outcomes, let R_i denote agent i 's strict preferences over all outcomes, and denote a (strict) preference profile as $R = (R_i)_{i \in N}$. Let \mathcal{R} denote the domain of all such preference profiles. A random social choice rule²⁶ is a mapping $f : \mathcal{R} \rightarrow \Delta(O)$ where $\Delta(O)$ denotes the set of probability distributions over O ; it is ex post efficient if it assigns positive probability only to Pareto efficient outcomes; it 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 $f(R'_i, R_{-i})$.²⁷ 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'

²⁶Gibbard [1977] refers to these as *decision schemes*.

²⁷I.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 $f(R)$. A random social choice rule is a random dictatorship if it picks an individual at random, independently of the reported preference profile, to then select that individual's most preferred outcome.

A random social choice rule f , defined on the universal domain of strict preferences, is strategy-proof if and only if it satisfies probabilistic monotonicity – this follows, upon closer inspection, from Lemma 2 in [Gibbard \[1977\]](#).²⁸ Hence, in the classic characterization of Random Dictatorship as the only random social choice rule on the domain of strict preferences satisfying ex post efficiency and strategy-proofness,²⁹ we may replace the last axiom by the equivalent requirement of probabilistic monotonicity.

Fact 1 ([Gibbard \[1977\]](#)). *Let f be an ex post efficient random social choice rule on the universal domain of strict preferences, with $|O| \geq 3$. The following statements are equivalent:*

The random social choice rule f

- (1) satisfies probabilistic monotonicity.*
- (2) satisfies strategy-proofness.*
- (3) is a random dictatorship.*

Interestingly, as we leave the universal domain of strict preferences and move to the domain of random assignment problems where agents are indifferent between assignments as long as they award them the same object, probabilistic monotonicity and strategy-proofness are no longer equivalent – and it is probabilistic monotonicity, rather than strategy-proofness, that yields a characterization of the Random Priority rule (also known as Random *Serial* Dictatorship) in analogy to [Gibbard's](#) and [Sonnenschein's](#) characterization of Random Dictatorship.

3.6. Probabilistic monotonicity vs. obvious strategy-proofness. 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.³⁰ We find that both are logically independent: there exist OSP-implementable random assignment rules that violate probabilistic monotonicity and vice versa.

²⁸For 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.

²⁹[Gibbard \[1977\]](#) 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.

³⁰More 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.

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 objects³¹ 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 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 $aR'_1 cR'_1 b$, $bR'_2 aR'_2 c$, and $R'_3 = R_3$, 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]. Thus, while OSP-implementability and probabilistic monotonicity both strengthen strategy-proofness, probabilistic monotonicity is still satisfied by some of the most prominent matching mechanisms while OSP-implementability is too demanding.

4. CONCLUDING REMARKS

Besides being easily implementable and commonly used, the Random Priority rule (RP) is the only random assignment rule satisfying probabilistic monotonicity, equal-treatment-of-equals, 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 on the domain of assignment problems (Proposition 2).

The fact that our characterisation of RP relies on strengthening strategy-proofness may be considered particularly compelling given that strategy-proofness, together with equal-treatment-of-equals and ex post efficiency, is insufficient to characterize RP [Erdil, 2014, Basteck and Ehlers, 2025]³² – there are other random assignment

³¹The 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 .

³²The construction by Erdil [2014] relies on the existence of outside options; Basteck and Ehlers [2025] provide a construction without outside options.

rules on the domain of ordinal preferences that satisfy strategy-proofness, equal-treatment-of-equals and ex post efficiency but are not welfare equivalent to RP . By Theorem 1 these rules must then violate weak object-wise non-bossiness.

The characterization of the Random Priority rule given in Theorem 1 is the first known characterization for settings with multiple copies of objects, such as for example multiple seats at different schools.³³ Moreover, as spelled out in Corollary 1, it implies and generalizes several equivalence results according to which randomizing over certain appealing deterministic assignment rules yields RP [Abdulkadiroğlu and Sönmez, 1998, Pathak and Sethuraman, 2011, Lee and Sethuraman, 2011, Carroll, 2014, Bade, 2020]. Instead, Theorem 1 characterizes RP directly among all random assignment rules, rather than among those that are constructed as randomizations over deterministic rules – an important distinction since the set of random assignment rules satisfying a particular property may be strictly larger than the set of rules that arises from randomizations over deterministic rules, each of which satisfies the property in question.³⁴

Finally, it may seem surprising that the natural generalization of Maskin monotonicity to probabilistic settings offered by probabilistic monotonicity, has so far escaped the attention of researchers.³⁵ One would hope that the current paper is only the first among many to fruitfully apply this notion in the analysis of general social choice and mechanism design problems.

APPENDIX A.

A.1. Proof of Theorem 1.

Proof. (4) \implies (1): We first establish that the Random Priority rule RP satisfies probabilistic monotonicity: this follows from the fact that RP is a convex combination of the $n!$ serial dictatorships f^\triangleright associated with all possible strict priority orders over agents, $\triangleright \in \Pi$. Each such deterministic rule f^\triangleright constitutes a Pareto efficient deferred acceptance rule (with objects' priorities over agents given by \triangleright) and hence satisfies Maskin monotonicity [Kojima and Manea, 2010]. As probabilistic (Maskin) monotonicity is preserved for convex combinations, RP satisfies probabilistic monotonicity.

³³Pycia and Troyan [2024] assume an equal number of objects and agents with all objects in unit supply.

³⁴For example, Pycia and Ünver [2015] show that the set of strategy-proof random assignment rules is strictly larger than the set of randomizations over deterministic strategy-proof assignment rules.

³⁵As explained in section 3.5, Gibbard [1977] does implicitly analyse probabilistic monotonicity, but does so in a setting where it is equivalent to strategy-proofness (and hence refers to it only under that name).

(1) \implies (2): Similarly, strategy-proofness of RP follows from the fact that each f^\triangleright is strategy-proof. For weak object-wise non-bossiness, consider any swap of two objects in some agents preference order, i.e., $R \in \mathcal{R}$ (resp. $\underline{\mathcal{R}}$) and $R'_i \in N_{R_i}$ (resp. $N_{R_i} \cap \underline{\mathcal{R}}$) such that $xP_i y$ and $yP'_i x$. Now, suppose that $f^{iz}(R) = f^{iz}(R'_i, R_{-i})$ for some arbitrary $z \in O$. We have to show that $f_\mu(R) = f_\mu(R'_i, R_{-i})$ for all $\mu \in \{\mu' \in \mathcal{M} \mid \mu'_i = z\}$.

If $z = y$, then (R'_i, R_{-i}) is a μ -monotonic transformation of R , so we have $f_\mu(R) \leq f_\mu(R'_i, R_{-i})$ for each μ with $\mu_i = z$. Moreover, if for some μ' with $\mu'_i = z$ we had $f_{\mu'}(R) < f_{\mu'}(R'_i, R_{-i})$, then $\sum_{\mu: \mu_i = z} f_\mu(R) < \sum_{\mu: \mu_i = z} f_\mu(R'_i, R_{-i})$ contradicting $f^{iz}(R) = f^{iz}(R'_i, R_{-i})$. Analogously if $z = x$.

If instead $z \neq x, y$, then any μ with $\mu_i = z$, (R'_i, R_{-i}) is a μ -monotonic transformation of R and vice versa. Hence $f_\mu(R) \leq f_\mu(R'_i, R_{-i}) \leq f_\mu(R)$ by probabilistic monotonicity.

(2) \implies (3): For pairwise responsiveness, consider any swap of two objects in some agents preference order, i.e., $R \in \mathcal{R}$ (resp. $\underline{\mathcal{R}}$) and $R'_i \in N_{R_i}$ (resp. $N_{R_i} \cap \underline{\mathcal{R}}$) such that $xP_i y$ and $yP'_i x$. By strategy-proofness, this does not affect i 's assignment probability for any other object $z \neq x, y$. By weak object-wise non-bossiness we have $f_\mu(R) = f_\mu(R'_i, R_{-i})$ for all $\mu \in \mathcal{M}$ s.t. $\mu_i = z$. Thus, f satisfies pairwise responsiveness.

(3) \implies (4): Let f be a random assignment rule satisfying ex post efficiency, equal-treatment-of-equals, and pairwise responsiveness.

First, recall the notation $\mathcal{M}^\#(R) := \{\mu \in \mathcal{M} \mid f_\mu(R) \neq RP_\mu(R)\}$. Second, by ex post efficiency of f and RP , $\mu \in \mathcal{M}^\#(R)$ implies that μ is efficient at R . Third, recall that any assignment μ , efficient at R , is the outcome of some deterministic priority rule at R where agents choose their most preferred object among those still available according to some deterministic priority order.

Now, towards a contradiction, suppose $f \neq RP$, i.e., that there exists some R and $\mu \in \mathcal{M}^\#(R)$. We argue by induction. Suppose that for $k \geq 0$ we have a triple $\mathcal{R}^k \subseteq \mathcal{R}$, $O_k \subseteq O$, and $N_k \subseteq N$ such that,

- there exists some $R \in \mathcal{R}^k$ and $\mu \in \mathcal{M}$ such that $f_\mu(R) \neq RP_\mu(R)$,
- $|O_k| = k$ (so increasing in k),
- all agents in N_k rank all objects in O_k at the bottom and in the same order, i.e.,
 - $\forall j \in N_k, o \notin O_k, o' \in O_k : oP_j o'$,
 - $\forall j, l \in N_k, o, o' \in O_k : oR_j o' \Leftrightarrow oR_l o'$,
- for all other agents, $i \in N \setminus N_k$, preferences are fixed and they receive the same object from O_k under any assignment for which f and RP differ, i.e.,
 - $\forall R, R' \in \mathcal{R}^k : R_i = R'_i$,
 - $\forall R \in \mathcal{R}^k, \mu, \mu' \in \mathcal{M}^\#(R) : \mu_i = \mu'_i \in O_k$.

We want to show that there exists an analogous triple for $k' = k + 1$ with $O_{k+1} \supset O_k$ and $N_{k+1} \subseteq N_k$ until we find a triple that allows us to derive a contradiction: intuitively, for any $R \in \mathcal{R}^k$ where there is a difference between f and RP , that difference must be limited to the assignment of objects to agents in N_k – and since their preferences are increasingly aligned as O_k expands, we eventually arrive at a contradiction to equal-treatment-of-equals.

For the induction base, we consider $k = 0$, $\mathcal{R}^k = \mathcal{R}$, $O_k = \{\}$, and $N_k = N$ – and are relieved to find that there is nothing to show as the existence of some $R \in \mathcal{R}$ and $\mu \in \mathcal{M}$ such that $f_\mu(R) \neq RP_\mu(R)$ was the starting assumption of our proof by contradiction.

For the induction step, suppose we have \mathcal{R}^k , O_k , and N_k as described above with $|O_k| < |O| - 1$ (so that for agents in N_k , preferences are not yet completely determined) and $|N_k| \geq 1$ (so that there is some agent whose preferences are not yet completely determined). Consider any $R \in \mathcal{R}^k$ and $\mu \in \mathcal{M}^\#(R)$. Since $\mu \in \mathcal{M}^\#(R)$, agents in $N \setminus N_k$ are assigned objects in O_k . Moreover μ is efficient and hence the outcome under some deterministic priority rule at R . Since agents in N_k rank objects in O_k last, they will choose objects in O_k only after all objects in $O \setminus O_k$ are assigned. Let $i \in N_k$ be the last agent to choose an object in $O \setminus O_k$ under the associated priority order and denote the object as $y = \mu_i$. Then, for all $j \in N_k$ who choose before i , we have $\mu_j R_j y P_j o$ for all $o \in O_k$ while for all $j \in N_k$ who choose after i we have $y P_j \mu_j$, $\mu_j \in O_k$. Set $N_k^y = \{j' \in N_k \mid \mu_{j'} = y\}$.

Now let all $j \in N_k \setminus N_k^y$ move y down in their preference order to a position immediately above objects in O_k , and do so in a series of pairwise swaps of adjacent objects. Call this new profile R' . Since swaps did not involve μ_j , pairwise responsiveness implies $f_\mu(R') = f_\mu(R) \neq RP_\mu(R) = RP_\mu(R')$. Hence, f and RP differ for some profile in a further restricted domain where weakly more than $|N_k| - |q_y|$ agents in N_k agree on the ranking of objects $O_k \cup \{y\}$ at the bottom, namely

$$\begin{aligned} \tilde{\mathcal{R}}^k &= \{R \in \mathcal{R}^k \mid \forall j \in N_k \setminus N_k^y, o \notin O_k \cup \{y\}, o' \in O_k \cup \{y\} : o P_j o' \\ &\quad \text{and } \forall j, l \in N_k \setminus N_k^y, o, o' \in O_k \cup \{y\} : o R_j o' \Leftrightarrow o R_l o'\}. \end{aligned}$$

To further narrow down the subdomain, denote agents in N_k^y as $N_k^y = \{j_1, \dots, j_{|N_k^y|}\}$ and consider j_1 . There are two cases:

Case (i): Suppose there exists $R \in \tilde{\mathcal{R}}^k$ and $\mu' \in \mathcal{M}^\#(R)$ such that $\mu'_{j_1} P_{j_1} y$ or $\mu'_{j_1} \in O_k$. Then we may move y down in j_1 's preference order, to a position immediately above objects in O_k , in a series of pairwise swaps of adjacent objects. Call the new profile $R^{k,1}$. Since the swaps did not involve μ'_{j_1} , pairwise responsiveness implies that

$f'_{\mu}(R^{k.1}) \neq RP_{\mu'}(R^{k.1})$. Hence, f and RP differ for some profile on

$$\begin{aligned} \tilde{\mathcal{R}}^{k.1} &= \{R \in \tilde{\mathcal{R}}^k \subseteq \mathcal{R}^k \mid \forall o \notin O_k \cup \{y\}, o' \in O_k \cup \{y\} : oP_{j_1} o' \\ &\text{and } \forall j \in N_k \setminus N_k^y, o, o' \in O_k \cup \{y\} : oR_j o' \Leftrightarrow oR_{j_1} o'\}. \end{aligned}$$

Case (ii): Suppose instead that for any $R \in \tilde{\mathcal{R}}^k$, $\mu' \in \mathcal{M}^\#(R)$ implies $yR_{j_1}\mu'_{j_1}$ and $\mu'_{j_1} \notin O_k$. If $yP_{j_1}\mu'_{j_1}$, we may move y to the top of j_1 's preference order, in a series of pairwise swaps. Call the new profile R' . Since the swaps did not involve μ'_{j_1} , we have $f_{\mu'}(R') \neq RP_{\mu'}(R')$. Hence, by ex-post efficiency of f and RP , μ' is efficient at R' . Hence, under μ' , all copies of y must be assigned to agents other than j_1 . Since $y \notin O_k$, it must be assigned to agents in N_k and since $|N_k^y \setminus \{j_1\}| < q_y$ there must be some agent $j \in N_k \setminus N_k^y$ who is assigned y (and ranks y immediately above all objects in O_k , i.e., below all other objects in $O \setminus O_k$) – but then there is a Pareto-improving trade between j and j_1 , a contradiction. Thus, for any $R \in \tilde{\mathcal{R}}^k$ and $\mu' \in \mathcal{M}^\#(R)$ we have $\mu'_{j_1} = y$. Now, pick any $R^{k.1} \in \tilde{\mathcal{R}}^k$ for which $\mathcal{M}^\#(R^{k.1}) \neq \emptyset$ and define

$$\tilde{\mathcal{R}}^{k.1} = \{R \in \tilde{\mathcal{R}}^k \subseteq \mathcal{R}^k \mid R_{j_1} = R_{j_1}^{k.1}\}.$$

Next, take j_2 . Again one of two cases applies, analogous to the two cases above where the initial subdomain is now $\tilde{\mathcal{R}}^{k.1}$ instead of $\tilde{\mathcal{R}}^k$ and the new subdomain is $\tilde{\mathcal{R}}^{k.2}$. Continue in this way to arrive at $\mathcal{R}^{k+1} := \tilde{\mathcal{R}}^{k.|\mathcal{N}_k^y|}$. To complete the induction step, set $O_{k+1} := O_k \cup \{y\}$ and remove all agents for which Case (ii) applied from N_k to arrive at N_{k+1} .

Finally, repeat the induction step until we arrive at $\mathcal{R}^{\bar{k}}$, $O_{\bar{k}}$, and $N_{\bar{k}}$ for which $\bar{k} = |O| - 1$ or for which $|N_{\bar{k}}| = 0$.

If $\bar{k} = |O| - 1$, take any $R \in \mathcal{R}^{\bar{k}}$ for which $\mathcal{M}^\#(R)$ is non-empty and consider $\mu, \mu' \in \mathcal{M}^\#(R)$ for which $f_\mu(R) > RP_\mu(R)$ and $f_{\mu'}(R) < RP_{\mu'}(R)$. By construction, $\mu_j = \mu'_j$ for all $j \notin N_{\bar{k}}$. Hence, μ and μ' differ only in the assignment of agents in $N_{\bar{k}}$. However, all these agents rank the same set of $|O| - 1$ objects at the bottom, and do so in the same order, so that $R_i = R_j$ for all $i, j \in N_{\bar{k}}$. Thus, equal-treatment-of-equals for RP implies $f_\mu(R) > RP_\mu(R) = RP_{\mu'}(R) > f_{\mu'}(R)$, which violates equal-treatment-of-equals for f – a contradiction.

So instead we must have $|N_{\bar{k}}| = 0$. In this case preferences of all agents are completely determined, i.e., $\mathcal{R}^{\bar{k}} = \{R\}$ is a singleton. Moreover since all agents are in $N \setminus N_{\bar{k}}$, they receive the same object across all $\mu \in \mathcal{M}^\#(R)$ – so that $\mathcal{M}^\#(R) = \{\mu\}$ is a singleton as well. But if $f_\mu(R) \neq RP_\mu(R)$ (while for all other assignments μ' we have $f_{\mu'}(R) = RP_{\mu'}(R)$), this contradicts the fact that under both $f(R)$ and $RP(R)$, the total probability of all assignments must sum to 1.

This completes the proof for the case of the unrestricted domain of strict preferences \mathcal{R} . If we consider random assignment rules on $\underline{\mathcal{R}}$ where all agents rank the same (null) object \emptyset last, the proof proceeds analogous, except that we may choose the induction base as $k = 0$, $\mathcal{R}^k = \mathcal{R}$, $O_k = \{\emptyset\}$, and $N_k = N$, i.e., we start ‘one step ahead’ in our attempt to narrow down the domain of preferences by making preferences as similar as possible for agents in N_k (while preserving $f \neq RP$). \square

A.2. Welfare equivalent random assignments. Since agents’ have preferences over their own assigned object – but are indifferent between assignments apart from that – 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.

Recall that p^{ia} denotes the probability with which agent i is assigned object a under a given random assignment p and that $p^i = (p^{ia})_{a \in O}$ denotes 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. Similarly, two random assignment rules, f and g , are welfare equivalent, if $f(R)$ and $g(R)$ are welfare equivalent for every R .³⁶

Welfare-equivalence gives rise to equivalence classes of random assignments that we denote as $[p] := \{q \in \Delta(\mathcal{M}) \mid \forall i \in N : q_i = p_i\}$ and that can be identified with the collection of individual object assignment probabilities $(p^{ia})_{i \in N, a \in O}$. A collection of individual object assignment probabilities $[p]$ is then typically said to be ex post efficient given R , if there *exists* some random assignment $q \in [p]$ that is ex-post-efficient with respect to R , i.e., if there exist a convex combination of deterministic, Pareto-efficient assignments giving rise to the individual object assignment probabilities $(p^{ia})_{i \in N, a \in O}$.

Moreover, for given R , $[p]$ is typically said to satisfy equal-treatment-of-equals if $R_i = R_j$ implies $p^{ia} = p^{ja}$ for all $i, j \in N$, $a \in O$. Equivalently, we can define equal-treatments-of-equals for a collection of individual object assignment probabilities in analogy to ex post efficiency, i.e., say that the property holds for $[p]$ if it holds for some $q^* \in [p]$:

Claim 1. *Consider a preference profile $R \in \mathcal{R}$ and a collection of individual assignment probabilities $[p]$. The following statements are equivalent:*

$$(i) \quad \forall i, j \in N, a \in O : \quad R_i = R_j \implies p^{ia} = p^{ja}.$$

³⁶Some 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].

(ii) there exists a random assignment $q^* \in [p]$ satisfying equal-treatment-of-equals.

Proof of Claim 1. (ii) \implies (i): Take $q \in [p]$ such that q satisfies equal-treatment-of-equals and suppose there are $i, j \in N$ for which $R_i = R_j$. Now for any a ,

$$q^{ia} = \sum_{\mu: \mu_i = a} q_\mu = \sum_{b \in O} \left(\sum_{\mu: \mu_i = a, \mu_j = b} q_\mu \right) = \sum_{b \in O} \left(\sum_{\mu: \mu_i = b, \mu_j = a} q_\mu \right) = \sum_{\mu: \mu_j = a} q_\mu = q^{ja}$$

where the central equality follows from the definition of equal-treatment-of-equals of random assignments.

(i) \implies (ii): Let R be a preference profile and p be a random assignment such that for all $i, j \in N$ and $R_i = R_j$ we have $p^i = p^j$. Moreover, let $[p]$ be the equivalence class of welfare equivalent random assignments. Further, let $M \subseteq N$ be maximal subset of agents with identical preferences, i.e., $R_i = R_j$ for all $i, j \in M$ and $R_k \neq R_i$ for all $k \notin M$. and let Π denote the set of all permutations of agents in M .³⁷ Now, for each assignment μ in the support of q , and each $\pi \in \Pi$, define μ^π as the assignment where $\mu_i^\pi = \mu_{\pi(i)}$. Last, construct the random assignment q_M from p by replacing each assignment μ with weight p_μ by the $|M|!$ many assignments μ^π , $\pi \in \Pi$, each with weight $\frac{p_\mu}{|M|!}$. Then, by construction, q_M satisfies equal-treatment-of-equals among all agents in M . Moreover $q_M^i = p^i$ for all $i \in M$ as well as all $i \notin M$ (since for them the permutation did not change their assignment, i.e., $\mu_i = \mu_i^\pi$ for all μ and $\pi \in \Pi$). Thus $q_M \in [p]$. Repeating the construction for all maximal sets of agents with identical preferences $M' \subseteq N$, $M' \neq M$, we arrive at a random assignment $q^* \in [p]$ that satisfies equal-treatment-of-equals for all $i, j \in N$. □

Hence, if one wants to distinguish between random assignments only up to welfare-equivalence, defining equal-treatment-of-equals with respect to the probability distribution over deterministic assignments, as we do in Section 2.2, is no more restrictive than defining equal-treatments-of-equals for a collection of individual assignment probabilities – whenever the latter satisfies equal-treatment-of-equals with respect to individual object assignment probabilities, we could just as well consider some decomposition into deterministic assignments that satisfies the equal-treatments-of-equals as defined in Section 2.2.

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.

³⁷A permutation π on M is a bijection $\pi: M \rightarrow M$.

- Itai Ashlagi and Yannai A Gonczarowski. Stable matching mechanisms are not obviously strategy-proof. *Journal of Economic Theory*, 177:405–425, 2018.
- 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.
- Christian Basteck and Lars Ehlers. On (constrained) efficiency of strategy-proof random assignment. *Econometrica*, 93(2):569–595, 2025.
- 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.
- Aytek Erdil. Strategy-proof stochastic assignment. *Journal of Economic Theory*, 151:146–162, 2014.
- 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.
- 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. The random priority mechanism is uniquely simple, efficient, and fair. *CEPR Discussion Papers*, (19689), 2024.
- Marek Pycia and M Utku Ünver. Decomposing random mechanisms. *Journal of Mathematical Economics*, 61:21–33, 2015.
- 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.