Probabilistic Assignment: A Two-fold Axiomatic Approach

Wonki Jo Cho*

July 2, 2012

Job Market Paper

Abstract

We study the problem of allocating objects by means of probabilistic mechanisms. Each agent has strict preferences over objects and ex post receives exactly one object. A standard approach in the literature is to extend agents' preferences over objects to preferences over lotteries defined on those objects, using the first-order stochastic dominance criterion, or the *sd*-extension. In a departure from this practice, we work with a general mapping, called an extension, from preferences over objects to preferences over lotteries. Our objective is to connect the extension operator theory in Cho (2012) with probabilistic assignment while maintaining an axiomatic perspective on each. Our main results concern axioms associated with the *sd*-, *dl*-, and *ul*-extensions. First, *sd-efficiency*, *dl-efficiency*, and *ul-efficiency* are equivalent (axioms are prefixed by the extensions they are associated with). Also, for each $e \in \{sd, dl, ul\}$, *e-adjacent strategy-proofness* is equivalent to *e-strategy-proofness*. For each $e \in \{sd, dl, ul\}$, a rule is *e-strategy-proof* iff the welfare of each agent weakly decreases as he reports preference relations that are further and further away from the truth. Moreover, *sd-strategy-proofness* is equivalent to *dl-strategy-proofness* and *ul-strategy-proofness*. Finally, we propose a family of rules that generalizes the serial rule.

Journal of Economic Literature Classification Numbers: C70, C78, D61, D63

Key Words: probabilistic assignment; generalized serial rules; random priority rule; e-efficiency; e-no-envy; e-strategy-proofness

*Department of Economics, University of Rochester, Rochester, NY 14627, USA; wonki.cho@rochester.edu.
I am deeply grateful to William Thomson for his guidance and encouragement. I also wish to thank P. Barelli, S. Govindan, B.-G. Ju, A. Kochov, R. Pancs, J. Sethuraman, and seminar participants at Korea University, Seoul National University, the Society for Economic Design Conference 2011, and the 23rd International Conference on Game Theory for their comments.

1 Introduction

We study the problem of allocating indivisible commodities called *objects* among agents by means of probabilistic mechanisms. Each agent has strict preferences over objects and ex post receives exactly one object. Going beyond deterministic assignments to probabilistic ones affords the advantage of converting objects into de facto perfectly divisible commodities. This raises the hope for fair allocation that might have been otherwise impossible were only deterministic assignments considered. Due in large part to the latter fact, lotteries are frequently used in real life. Examples include on-campus housing allocation in colleges and student placement in public schools.

This problem is known as "probabilistic assignment", and we focus on the ordinal approach to the problem (Bogomolnaia and Moulin, 2001).[1] According to the ordinal approach, agents submit their preferences over objects, and an assignment is selected based on this information. But what agents receive cannot be evaluated according to the elicited preferences, and therefore, we cannot properly speak of properties of assignment rules. Bogomolnaia and Moulin (2001) circumvent this problem by extending agents' preferences over objects to preferences over lotteries using the first-order stochastic dominance criterion.[2] We call this procedure the *sd*-extension.

Once we use the *sd*-extension, it is automatically embedded in axioms on assignment rules—e.g., efficiency, no-envy, and strategy-proofness—and affects their content. While many authors follow this practice and use the *sd*-extension, they do not provide much justification for it, except for the well-known fact that relates the *sd*-extension to von Neumann-Morgenstern (vNM) preferences (Remark 3.1). Moreover, they do not explore other possibilities of extending preferences over objects to preferences over lotteries.

Our objective is to address these issues. Instead of taking the *sd*-extension for granted, we consider a general mapping from preferences over objects to preferences over lotteries. Such a mapping is called an extension operator, or simply an extension, and the *sd*-extension is an example. Cho (2012) develops an axiomatic analysis of extensions, and here we apply it to probabilistic assignment. We examine standard axioms on assignment rules, but from the perspective of a general extension, and rules are evaluated on the grounds of these new criteria.

To facilitate the presentation of our approach, we briefly introduce the concepts and the results in Cho (2012). The paper studies axioms on extensions, proposes new extensions, investigates their relations, and identifies extensions satisfying certain combinations of axioms. Some of the axioms are quite familiar from the theory of binary relations; e.g., an extension is *complete* if it

[1] Another approach is the cardinal one, which allows agents to express their von Neumann-Morgenstern preferences over lotteries (Hylland and Zeckhauser, 1979).

[2] While Bogomolnaia and Moulin (2001) are the first to use the *sd*-extension in the context of probabilistic assignment, it was adopted much earlier in the context of probabilistic public choice. For instance, Gibbard (1977), in effect, applies the *sd*-extension in defining strategy-proofness.

associates with each preference relation over objects a complete preference relation over lotteries. Other axioms are new and specific to the model; e.g., an extension is *monotone* if it associates with each preference relation over objects a preference relation over lotteries such that given any lottery, a positive transfer of probability from a less preferred object to a preferred one leads to a preferred lottery.

Among the extensions proposed in Cho (2012), most notable are the downward and upward lexicographic dominance extensions, or the *dl*- and *ul*-extensions for short, respectively. According to the *dl*-extension, each agent compares lotteries in a lexicographic fashion, starting from the most preferred object. Given two lotteries, the probabilities for the most preferred object are compared first. If one lottery assigns a higher probability than the other, then the former is preferred. If the two probabilities are equal, then the probabilities for the second most preferred object are compared next, and the lottery with a higher probability is preferred. If the two probabilities are equal again, the probabilities for the third most preferred object are compared, and so on.

On the other hand, the *ul*-extension performs lexicographic comparison in the opposite way. It first looks at the probabilities for the least preferred object and the lottery with a lower probability is preferred. If the two probabilities are equal, then the probabilities for the second least preferred object are compared next, and the lottery with a lower probability is preferred, and so on.

The two extensions are both *monotone*, *complete*, and *transitiv*e (see Table 1 for more on this). As is apparent from the definition, the *dl*- and *ul*-extensions are intimately related, and their relation is formally captured by the notion of duality. To define duality, consider two extensions and a preference relation over objects. Take the inverse of the preference relation over objects and extend it according to the first extension.[3] If the inverse of the resulting preference relation over lotteries is the same as the preference relation over lotteries that the second extension associates with the original preference relation over objects, then the two extensions are dual. An extension is *self-dual* if it is the dual of itself. The *dl*- and *ul*- extensions are dual and the *sd*-extension is *self-dual*.

To turn to probabilistic assignment, our axioms on assignment rules are formulated on the basis of a general extension e, and therefore, they are all prefixed by the notation e; e.g., *e-efficiency*, *e-no-envy*, *e-strategy-proofness*, and so on. We show that if an extension e satisfies some axioms, then *e-efficiency* is equivalent to *sd-efficiency* (Theorem 4.1 and Corollaries 4.1-4.2). All extensions proposed in Cho (2012), including the *dl*- and *ul*-extensions, satisfy those axioms. Moreover, for any extension e "contained" in the *dl*- or *ul*-extensions, regardless of whether it meets the axioms or not, *e-efficiency* coincides with *sd-efficiency* (Corollary 4.3).[4] Therefore, existing characterizations of

[3]Given an arbitrary set X and a binary relation B over X, the inverse of B, denoted B^{-1}, is defined as follows: for each pair $x, y \in X$, $x\, B^{-1}\, y$ iff $y\, B\, x$.

[4]For an arbitrary set X, a binary relation over X can be seen as a subset of X^2. Given extensions e and $\hat{e}$, e

sd-efficiency (Abdulkadiroğlu and Sönmez, 2003; and McLennan, 2002) also apply to *e-efficiency*.

With regard to strategy-proofness, we consider a strategic requirement called *e-adjacent strategy-proofness*. It says that no agent is better off reporting a preference relation obtained by switching only two objects that are adjacent in his true preference rankings. While the axiom is much weaker than *e-strategy-proofness* in general, for each $e \in \{sd, dl, ul\}$, the two are equivalent (Theorem 4.2). This result becomes useful when we check *e-strategy-proofness*. Clearly, it also applies to deterministic rules and to the probabilistic public choice model (Gibbard, 1977). Similar results of Sato (2010) and Carroll (2012) obtain as a special case of ours.

The equivalence of strategy-proofness and adjacent strategy-proofness allows us to better understand how strategy-proof rules behave. We show that for each $e \in \{sd, dl, ul\}$, a rule is *e-strategy-proof* iff it is "lie monotone" when viewed as a function of one agent's report; i.e., when the reports of all but one agent are fixed, the welfare of the remaining agent decreases in the weak sense as he reports preference relations that are further and further away from the truth (Theorem 4.3). Thus, if our objective is to design a strategy-proof rule, we *should* distinguish "small" lies from "big" ones by punishing the latter more severely. To the best of our knowledge, this monotonicity property of strategy-proof rules has not been explored previously in the context of probabilistic assignment or probabilistic public choice.

The equivalence of strategy-proofness and adjacent strategy-proofness also yields an interesting connection among strategy-proofness notions associated with the *sd*-, *dl*-, and *ul*-extensions. It is clear that *sd-strategy-proofness* implies *dl-strategy-proofness* and *ul-strategy-proofness*. Unless there are only three objects, the converse cannot be established by appealing to the definition. However, using the characterization of the three strategy-proofness notions (Remark 4.4), we show that the converse also holds. Thus, *sd-strategy-proofness* is equivalent to the combination of *dl-strategy-proofness* and *ul-strategy-proofness* (Theorem 4.4). This decomposition result is quite surprising in light of the fact that each of *dl-strategy-proofness* and *ul-strategy-proofness* is a weakest strategy-proofness notion among all we consider.

To assess the (probabilistic) serial and random priority rules, we consider axioms associated with various extensions and look for the strongest requirements they satisfy (or equivalently, the weakest they violate). Concerning efficiency, the serial rule satisfies *sd-efficiency*, and hence all efficiency concepts derived from the extensions proposed in Cho (2012); by contrast, the random priority rule fails all these efficiency concepts (Theorem 4.5). With regard to no-envy, the serial rule satisfies *sd-no-envy*, the strongest no-envy notion, whereas the random priority rule satisfies *dl-no-envy*, a weakest notion (Theorem 4.6). The latter rule, however, outperforms the former rule

is contained in $\hat{e}$ if for each preference relation over objects, the preference relation over lotteries that e associates with the preference relation over objects is contained in the preference relation over lotteries that $\hat{e}$ associates with the preference relation over objects.

as far as strategy-proofness is concerned: the random priority rule satisfies *sd-strategy-proofness*, the strongest strategy-proofness concept, and the serial rule only satisfies *dl-strategy-proofness*, a weakest concept (Theorem 4.7). These observations yield the by-product that Bogomolnaia and Moulin's (2001) impossibility result, namely that no rule satisfies *sd-efficiency*, *equal treatment of equals*, and *sd-strategy-proofness*, hinges critically on the extension chosen.[5]

Finally, we propose a family of rules that generalizes the idea underlying the serial rule. Each rule in the family is defined by an algorithm similar to the one that defines the serial rule. The key difference is that probability shares of objects are distributed over time at a speed that may vary across objects and time (but not across agents). Each generalized serial rule is associated with an allocation speed function defined over the set of objects and the time horizon (see Section 4.3 for a formal definition). We show that for each speed function, the generalized serial rule associated with it satisfies *dl-efficiency* and *dl-no-envy* (Theorems 4.5-4.6); however, *dl-strategy-proofness* may be violated if the speed varies too greatly across objects (Example 4.1).

As is true of any economic model, many results in probabilistic assignment are driven by the assumptions the modeller makes, and the connection between the two reveals itself most clearly when different assumptions are imposed. In light of this principle, the contribution of the current work lies in studying probabilistic assignment under varying assumptions (extensions) and formalizing their effect on analysis. This "two-fold axiomatic approach" allows us to provide a new perspective and isolate driving factors in existing results.

The rest of the paper proceeds as follows. Section 2 discusses related literature. Section 3 summarizes some of the extension operator theory in Cho (2012), and Section 4 applies it to probabilistic assignment. Finally, Section 5 concludes.

2 Related Literature

The first model of probabilistic assignment is due to Hylland and Zeckhauser (1979). They assume that agents have vNM preferences and propose a rule that first entitles each agent to a budget and then lets them trade probability shares of the objects in a pseudo-market mechanism. This rule satisfies (ex ante) efficiency and no-envy, but not strategy-proofness. In fact, no rule meets the three requirements (Zhou, 1990).

In contrast with this cardinal approach, recent works adopt the ordinal framework that only consider agents' preferences over objects but not lotteries. To enable the agents to evaluate lotteries, Bogomolnaia and Moulin (2001) extend preferences over objects to preferences over lotteries

[5] *Equal treatment of equals* requires that agents with the same preference relation receive the same lottery up to indifference.

by the *sd*-extension. They then define *sd-efficiency*, a concept that is intermediate in strength between ex post and ex ante efficiency.[6] They also propose the serial rule and study its properties, together with the random priority rule. Subsequent works generalize their model in several directions to allow for the following possibilities: (i) there are multiple copies of each object; (ii) agents may choose not to receive any object, that is, to receive a "null" object; (iii) agents may receive more than one object; (iv) agents may have indifferences among some objects; and (v) agents privately own fractions of objects.[7] Yet in all the papers, the use of the *sd*-extension remains an assumption that goes unchallenged, and in this respect, our theory based on a general extension offers a new approach.

The notion of *sd-efficiency* evolved to become a topic of independent interest. Abdulkadiroğlu and Sönmez (2003) characterize *sd-efficiency* by means of a dominance concept defined over sets of assignments. McLennan (2002) proves a welfare theorem involving *sd-efficiency*, using a separating hyperplane theorem for polyhedra, and Manea (2008) provides an alternative constructive proof. Katta and Sethuraman (2006) study *sd-efficiency* in a more general setup that permits indifferences among objects. Liu and Pycia (2012) show that if a sequence of rules consists of uniform randomizations over efficient deterministic rules, then the sequence is asymptotically *sd-efficient*. Because for some extensions, *e-efficiency* is equivalent to *sd-efficiency* (Theorem 4.1 and Corollaries 4.1-4.3), the aforementioned papers can also be viewed as providing additional properties of *e-efficiency* as well.

The present paper is also related to the literature that examines the equivalence of strategy-proofness and seemingly weaker versions of it, such as adjacent strategy-proofness. Sato (2010) studies deterministic mechanisms in the context of the classical voting model, and identifies a condition on the preference domain for the equivalence to hold. His results, in effect, strengthen the Gibbard-Satterthwaite theorem (Gibbard, 1973; Satterthwaite, 1975). Carroll (2012) pursues a similar question but in the realm that allows probabilistic mechanisms. His results pertain to many interesting preference domains, such as expected utility, single-peaked, and single-crossing preferences. Our result that *dl*- and *ul-adjacent strategy-proofness* are equivalent to *dl*- and *ul-strategy-proofness*, respectively (Theorem 4.2), complements these papers.

[6] Bogomolnaia and Moulin (2001) call it *ordinal efficiency*.

[7] The following papers explore these variations in different combinations: Bogomolnaia and Heo (2011), Che and Kojima (2010), Hashimoto et al. (2012), Katta and Sethuraman (2006), Kesten (2009), Kojima (2009), Liu and Pycia (2012), and Yilmaz (2009, 2010).

3 Extension Operator Theory

Before studying probabilistic assignment, we first address the issue of extending preferences over objects to preferences over lotteries. This section summarizes the results in Cho (2012) that are necessary for our analysis in Section 4. We refer the reader to the original paper for a more complete discussion and omitted proofs.

3.1 The Model

Throughout Section 3, we imagine a single agent facing the problem of extending his preferences over deterministic outcomes to preferences over lotteries on them. Let $\boldsymbol{A} \equiv \{1, \cdots, m\}$ be the set of indivisible commodities called *objects*. Assume that $m \geq 3$.[8] Objects are denoted by k, ℓ, k', ℓ', and so on. The agent is equipped with a complete, transitive, and anti-symmetric preference relation $\boldsymbol{R}$ over A. Let $\boldsymbol{P}$ and $\boldsymbol{I}$ be the strict preference and indifference relations, respectively, associated with R. Let $\boldsymbol{\mathcal{R}(A)}$ be the class of all complete, transitive, and anti-symmetric preference relations over A. Also, for each $R \in \mathcal{R}(A)$ and each $k \in \{1, \cdots, , m\}$, let $\boldsymbol{k(R)}$ be the object ranked kth according to R.

Let $\boldsymbol{\Delta A}$ be the set of all lotteries over A. Lotteries are denoted by π, π', and so on. For each $\pi \in \Delta A$ and each $k \in A$, let π_k be the probability that lottery π assigns to object k. Let $\boldsymbol{\mathcal{R}(\Delta A)}$ be the set of all reflexive relations over ΔA. Members of $\mathcal{R}(\Delta A)$ are called preferences over lotteries. All restrictions on $\mathcal{R}(\Delta A)$ other than reflexivity—e.g., completeness and transitivity—are formulated as axioms in the next subsection.

Our objective is to provide a systematic procedure of evaluating lotteries, taking preferences over objects as a primitive. Since A can be embedded in ΔA, this entails "extending" preferences over A to ΔA. Formally, an **extension operator**, or simply an **extension**, is a mapping from $\mathcal{R}(A)$ to $\mathcal{R}(\Delta A)$. Given an extension $e : \mathcal{R}(A) \to \mathcal{R}(\Delta A)$, for each $R \in \mathcal{R}(A)$, let $\boldsymbol{R^e} \equiv e(R)$ be the preference relation over ΔA that e assigns to R. In line with the previous notation, the strict preference and indifference relations associated with R^e are denoted by $\boldsymbol{P^e}$ and $\boldsymbol{I^e}$, respectively. Also, given a property p that applies to binary relations over $\mathcal{R}(\Delta A)$, we say that **$\boldsymbol{e}$ satisfies property $\boldsymbol{p}$** if for each $R \in \mathcal{R}(A)$, R^e satisfies property p; e.g., e is transitive if for each $R \in \mathcal{R}(A)$, R^e is transitive. Since each preference relation in $\mathcal{R}(\Delta A)$ is reflexive, it follows by definition that each extension is reflexive.

Later in the section, we define a notion that relates binary relations over a set, and to that end, further notation is needed. Let X be an arbitrary set and B a binary relation over X. Let

[8] When $m = 2$, the problem is not so interesting because there is only one extension satisfying *monotonicity*, the axiom that we take most basic.

$\boldsymbol{B^{-1}}$ be the inverse of B; i.e., for each pair $x, y \in X$, $x\, B^{-1}\, y$ iff $y\, B\, x$. Also, noting that binary relations over X can equivalently be treated as subsets of X^2, given binary relations B and B' over X, let $\boldsymbol{B \cap B'}$ be the binary relation over X such that for each pair $x, y \in X$, $x\, B \cap B'\, y$ iff $x\, B\, y$ and $x\, B'\, y$. Similarly, let $\boldsymbol{B \cup B'}$ be the binary relation over X such that for each pair $x, y \in X$, $x\, B \cup B'\, y$ iff $x\, B\, y$ or $x\, B'\, y$. Finally, $\boldsymbol{B}$ **is contained in** $\boldsymbol{B'}$, denoted $\boldsymbol{B \subseteq B'}$, iff for each pair $x, y \in X$, $x\, B\, y$ implies $x\, B'\, y$. The relations $B \subsetneq B'$ and $B = B'$ are defined in the standard way.

3.2 Axioms on Extension Operators

Now we introduce axioms on extensions. Let e be an extension. Consider a lottery, and suppose that some probability is transferred from a less preferred object to a preferred object, with the probabilities for all other objects fixed. It appears intuitive that the lottery so obtained should be considered more desirable than the initial lottery. Our first axiom formalizes this idea. For each $R \in \mathcal{R}(A)$ and each pair $\pi, \pi' \in \Delta A$, $\boldsymbol{\pi}$ **is a monotone improvement over** $\boldsymbol{\pi'}$ **for** $\boldsymbol{R}$, written as $\boldsymbol{\pi\, M_R\, \pi'}$, if there are $k, \ell \in A$ such that (i) $k\, P\, \ell$; (ii) for each $h \in A \backslash \{k, \ell\}$, $\pi_h = \pi'_h$; and (iii) $\pi_k > \pi'_k$ and $\pi_\ell < \pi'_\ell$. The following axiom requires that for each $R \in \mathcal{R}(A)$ and each pair $\pi, \pi' \in \Delta A$, if π is a monotone improvement over π' for R, then π should be preferred to π' according to R^e.

Monotonicity: For each $R \in \mathcal{R}(A)$ and each pair $\pi, \pi' \in \Delta A$, if $\pi\, M_R\, \pi'$, then $\pi\, P^e\, \pi'$.

A consequence of *monotonicity* is that degenerate lotteries are ranked in the same way as objects are according to the underlying preferences over objects.

Our next axiom requires that an extension associate with each preference relation over objects a complete preference relation over lotteries, so that all lotteries are comparable.

Completeness: For each $R \in \mathcal{R}(A)$ and each pair $\pi, \pi' \in \Delta A$, either $\pi\, R^e\, \pi'$ or $\pi'\, R^e\, \pi$.

The following axiom says that an extension should associate with each preference relation over objects a transitive preference relation over lotteries.

Transitivity: For each $R \in \mathcal{R}(A)$ and each triple $\pi, \pi', \pi'' \in \Delta A$, if $\pi\, R^e\, \pi'$ and $\pi'\, R^e\, \pi''$, then $\pi\, R^e\, \pi''$.

Next is an axiom that deals with indifferences over lotteries. It requires that two lotteries be indifferent if and only if they are, in fact, the same lottery.

Anti-symmetry: For each $R \in \mathcal{R}(A)$ and each pair $\pi, \pi' \in \Delta A$, if $\pi\, I^e\, \pi'$, then $\pi = \pi'$.

Now we present two axioms concerning the issue of how probabilities for preferred and less preferred objects are related. First, let $R \in \mathcal{R}(A)$ and $\pi, \pi' \in \Delta A$. Suppose that π is preferred to π' according to R^e and that π' assigns a higher probability to some object $k \in A$ than π does. The next axiom says that in such a case, π should assign a higher probability than π' does to some object that is preferred to object k according to R.

Probability Compensation for Preferred Objects: For each $R \in \mathcal{R}(A)$ and each pair $\pi, \pi' \in \Delta A$, if $\pi \, P^e \, \pi'$ and there is $k \in A$ such that $\pi_k < \pi'_k$, then there is $\ell \in A$ such that $\ell \, P \, k$ and $\pi_\ell > \pi'_\ell$.

The next axiom is similar to *probability compensation for preferred objects*. Suppose that π is preferred to π' according to R^e and that π' assigns a *lower* probability to some object $k \in A$ than π does. Then we require that π assign a *lower* probability than π' does to some object that is less preferred than object k according to R.

Probability Compensation for Less Preferred Objects: For each $R \in \mathcal{R}(A)$ and each pair $\pi, \pi' \in \Delta A$, if $\pi \, P^e \, \pi'$ and there is $k \in A$ such that $\pi_k > \pi'_k$, then there is $\ell \in A$ such that $k \, P \, \ell$ and $\pi_\ell < \pi'_\ell$.

All axioms introduced so far are punctual; they have no bite on how the images under an extension of two preference relations over objects are related. The next axiom pertains to an issue of this nature, and thus, it is a "relational" axiom. To introduce it, we define the notion of duality. Given an extension e, **the dual of $\boldsymbol{e}$**, denoted $\boldsymbol{e^d}$, is the extension such that for each $R \in \mathcal{R}(A)$, $R^{e^d} = ((R^{-1})^e)^{-1}$ (or equivalently, $(R^{e^d})^{-1} = (R^{-1})^e$). It is easy to see that $(e^d)^d = e$ and that the dual of e is unique; therefore, we can unambiguously say that two extensions are dual. The next axiom requires that an extension be the dual of itself; i.e., $e = e^d$. In words, this means the following. Consider a preference relation $R \in \mathcal{R}(A)$ and its inverse, R^{-1}. Since R and R^{-1} rank objects in opposite ways, it is reasonable to expect a similar connection to hold between the preference relations obtained by extending R and R^{-1} by e; i.e., R^e and $(R^{-1})^e$ rank lotteries in opposite ways.

Self-duality: For each $R \in \mathcal{R}(A)$, $(R^{-1})^e = (R^e)^{-1}$.

Our last axiom is about the "difference" between the *sd*-extension and the extension under consideration. As Proposition 3.2 below shows, among the class of *monotone* and *transitive* extensions, the *sd*-extension is "minimal". Therefore, given an extension, it is interesting to ask how much it differs from the *sd*-extension. The next axiom says that if lottery π is at least as desirable as lottery π' according to R^e and the opposite is true according to $(R^{-1})^e$, then they are at least comparable according to R^{sd} (or equivalently, according to $(R^{-1})^{sd}$).

Condition C: For each $R \in \mathcal{R}(A)$ and each pair $\pi, \pi' \in \Delta A$, if $\pi \, R^e \, \pi'$ and $\pi' \, (R^{-1})^e \, \pi$, then $\pi \, R^{sd} \, \pi'$ or $\pi' \, R^{sd} \, \pi$.

Axioms	sd	dl	ul	sd^k	sd_k
Monotonicity	+	+	+	+	+
Completeness	−	+	+	−	−
Transitivity	+	+	+	+	+
Anti-symmetry	+	+	+	+	+
PCPO	+	+	−	+	−
PCLPO	+	−	+	−	+
Self-duality	+	−	−	−	−
Condition C	+	−	−	−	−

Table 1: **Axioms and extension operators.** In the table, "*PCPO*" stands for "*probability compensation for preferred objects*" and "*PCLPO*" for "*probability compensation for less preferred objects*".

Simply put, if an extension satisfies condition C, then it does not differ from the sd-extension too much.

3.3 An Inventory of Extension Operators

Below are some of the extensions proposed in Cho (2012). Table 1 summarizes their properties.

3.3.1 sd-extension

We first present an extension that occupies the dominant position in the literature on probabilistic assignment: the (first-order) stochastic dominance extension, or the ***sd*-extension** for short. Most papers on probabilistic assignment are based on this notion.[9] For each $R \in \mathcal{R}(A)$ and each pair $\pi, \pi' \in \Delta A$, π is at least as desirable as π' according to R^{sd} iff for each $k \in \{1, \cdots, m\}$, the sum of probabilities that π assigns to the k most preferred objects is at least as large as the corresponding sum of probabilities that π' assigns.

***sd*-extension:** For each $R \in \mathcal{R}(A)$ and each pair $\pi, \pi' \in \Delta A$, $\pi \, R^{sd} \, \pi'$ if for each $k \in \{1, \cdots, m\}$, $\sum_{h=1}^{k} \pi_{h(R)} \geq \sum_{h=1}^{k} \pi_{h(R)}$.

The sd-extension is *transitive* and *anti-symmetric*, but it is not *complete*. Therefore, for each $R \in \mathcal{R}(A)$, R^{sd} is a partial order. The prevalent use of the sd-extension is due to its connection with expected utility preferences. However, as Proposition 3.2 below shows, the appeal of the sd-extension lies at a more fundamental level that has nothing to do with the expected utility property. Remark 3.1 elaborates on this point.

[9]See, for instance, Bogomolnaia and Heo (2011), Bogomolnaia and Moulin (2001, 2002), Che and Kojima (2010), Hashimoto et al. (2012), Heo (2010), Kasajima (2011), Katta and Sethuraman (2006), Kesten (2009), Kojima (2009), Kojima and Manea (2010), Liu and Pycia (2012), and Yilmaz (2009, 2010).

3.3.2 Lexicographic Family

Now we introduce two extensions that are new to the literature. First, the downward lexicographic dominance extension, or simply the ***dl*-extension**, compares lotteries in a lexicographic fashion, as follows. Given two lotteries, the lottery that assigns a higher probability to the most preferred object is preferred; if the two lotteries assign equal probability, then the lottery that assigns a higher probability to the second most preferred object is preferred; if the two lotteries assign equal probability again, then the probabilities for the third most preferred object are compared, and so on.

***dl*-extension:** For each $R \in \mathcal{R}(A)$ and each pair $\pi, \pi' \in \Delta A$, $\pi \, R^{dl} \, \pi'$ if either (i) there is $k \in \{1, \cdots, m\}$ such that for each $h \leq k-1$, $\pi_{h(R)} = \pi'_{h(R)}$ and $\pi_{k(R)} > \pi'_{k(R)}$; or (ii) $\pi = \pi'$.

The next extension also performs lexicographic comparison, though in the opposite direction. The upward lexicographic dominance extension, or the ***ul*-extension**, first looks at the probabilities for the least preferred object, then those for the second least preferred object, and so on. But since the probabilities for less preferred objects are compared first, the lottery that assigns a lower probability to them are preferred.

***ul*-extension:** For each $R \in \mathcal{R}(A)$ and each pair $\pi, \pi' \in \Delta A$, $\pi \, R^{dl} \, \pi'$ if either (i) there is $k \in \{1, \cdots, m\}$ such that for each $h \geq k+1$, $\pi_{h(R)} = \pi'_{h(R)}$ and $\pi_{k(R)} < \pi'_{k(R)}$; or (ii) $\pi = \pi'$.

As is transparent from the definitions, the *dl*- and *ul*-extensions are similar, but at the same time, they are diametrically opposite to each other. They are similar in that lexicographic comparison is used; they are opposite in that one maximizes the probabilities for preferred objects whereas the other minimizes the probabilities for less preferred objects. This observation can be formally stated using the notion of duality: the *dl*- and *ul*-extensions are dual. Other properties of the *dl*- and *ul*-extensions are as follows. They are both *transitive* and *anti-symmetric*, and in contrast with the *sd*-extension, they are *complete* as well. Therefore, for each $R \in \mathcal{R}(A)$, R^{dl} and R^{ul} are linear (complete, transitive, and anti-symmetric) orders on ΔA.

An essential feature of the *dl*- and *ul*-extensions is that probabilities are compared in a lexicographic manner. In fact, this idea can be generalized by specifying (i) the order in which probabilities enter the lexicographic comparison (i.e., the probabilities for which objects are used as the first criterion, the second criterion, and so on); and (ii) whether a lottery with a higher or lower probability is preferred at each level of the lexicographic comparison. A **lexicographic extension** is an extension defined this way, and we call them collectively the **lexicographic family**. However, there are lexicographic extensions that are not monotone, and we can identify a sufficient and necessary condition for a lexicographic extension to be monotone. Finally, each lexicographic extension gives preferences over lotteries that can be represented by a lexicographic expected utility function (Hausner, 1954; Chipman, 1960).

3.3.3 sd^k- and sd_k-extensions

Next, we define extensions having the flavor of bounded rationality. To motivate the concept underlying them, we draw a parallel with the sd-extension. Let $R \in \mathcal{R}(A)$ and $\pi, \pi' \in \Delta A$. Recall that to declare $\pi \, R^{sd} \, \pi'$, the following $m-1$ inequalities should hold:

$$\begin{aligned} \pi_{1(R)} &\geq \pi'_{1(R)}; \\ \pi_{1(R)} + \pi_{2(R)} &\geq \pi'_{1(R)} + \pi'_{2(R)}; \\ &\vdots \\ \pi_{1(R)} + \cdots + \pi_{(m-1)(R)} &\geq \pi'_{1(R)} + \cdots + \pi'_{(m-1)(R)}. \end{aligned} \tag{1}$$

For each $k \in \{1, \cdots, m-1\}$, the stochastic dominance top-k extension, or the ***sd*^{k}-extension**, first checks whether the top k of these inequalities hold, with at least one of them being strict. If that is the case, then π is preferred to π' according to R^{sd^k}. If there are two strict inequalities pointing in opposite directions, then the two lotteries are uncomparable according to R^{sd^k}. Only if the top k inequalities all hold with equality, the $k+1$st inequality is considered, and if it holds strictly, then π is preferred to π' according to R^{sd^k}. If it holds with equality, the $k+2$nd inequality is considered, and so on. The process continues in this lexicographic manner until either one inequality is found to hold strictly or the two lotteries are shown equal.

***sd*^{k}-extension:** For each $R \in \mathcal{R}(A)$ and each pair $\pi, \pi' \in \Delta A$, $\pi \, R^{sd^k} \, \pi'$ if (i) for each $\ell \leq k$, $\sum_{h=1}^{\ell} \pi_{h(R)} \geq \sum_{h=1}^{\ell} \pi'_{h(R)}$, with at least one strict inequality; or (ii) there is $\hat{k} \geq k+1$ such that for each $\ell \leq \hat{k}-1$, $\pi_{\ell(R)} = \pi'_{\ell(R)}$ and $\pi_{\hat{k}(R)} > \pi'_{\hat{k}(R)}$; or (iii) $\pi = \pi'$.

To justify the sd^k-extension, imagine an agent who places a high value on the k most preferred objects and little on the remaining objects. In that case, what matters most is the probabilities assigned to the k most preferred objects. Therefore, when comparing lotteries, the agent may find it enough to look at only the top k inequalities and to consider the probability for the $k+1$st most preferred object only if his preferences are not determined by doing so. In the literature, Sen (2011) considers an extension similar to the sd^k-extension. In proving a stronger version of the random dictatorship theorem (Gibbard, 1977), he assumes that agents only check the top two or three inequalities and do not take into account the probabilities for the remaining objects.

To introduce another variant of the sd-extension, rewrite the inequalities in (1) as follows:

$$\begin{aligned}
\pi_{m(R)} + \cdots + \pi_{3(R)} + \pi_{2(R)} &\leq \pi'_{m(R)} + \cdots + \pi'_{3(R)} + \pi'_{2(R)}; \\
\pi_{m(R)} + \cdots + \pi_{3(R)} &\leq \pi'_{m(R)} + \cdots + \pi'_{3(R)} \qquad ; \\
&\vdots \\
\pi_{m(R)} &\leq \pi'_{m(R)}.
\end{aligned} \tag{2}$$

Therefore, the sd-extension can equivalently be viewed as minimizing the sum of probabilities for less preferred objects, and we define an extension exploiting this idea. For each $k \in \{1, \cdots, m-1\}$, the stochastic dominance bottom-k extension, or the $\boldsymbol{sd_k}$**-extension**, is similar to the sd^k-extension but begins with the bottom k of the inequalities in (2). It checks the $k+1$st inequality from the bottom—and others—only if all of them hold with equality.

$\boldsymbol{sd_k}$**-extension**: For each $R \in \mathcal{R}(A)$ and each pair $\pi, \pi' \in \Delta A$, $\pi \, R^{sd_k} \, \pi'$ if (i) for each $\ell \leq k$, $\sum_{h=1}^{\ell} \pi_{(m+1-h)(R)} \leq \sum_{h=1}^{\ell} \pi'_{(m+1-h)(R)}$, with at least one strict inequality; or (ii) there is $\hat{k} \geq k+1$ such that for each $\ell \leq \hat{k}-1$, $\pi_{(m+1-\ell)(R)} = \pi'_{(m+1-\ell)(R)}$ and $\pi_{(m+1-\hat{k})(R)} < \pi'_{(m+1-\hat{k})(R)}$; or (iii) $\pi = \pi'$.

As is the case with the sd-extension, for $k \geq 2$, the sd^k- and sd_k-extensions are *transitive* and *anti-symmetric*, but they are not *complete* (though "less incomplete" than the sd-extension). If $k = 1$, then they are the same as the dl- and ul-extensions, respectively, and thus, they are *complete*. Further, the connection between the sd^k- and sd_k-extensions is similar to that between the dl- and ul-extensions: the sd^k- and sd_k-extensions are dual.

3.4 Results

First, we explore "inclusion" relations among extensions. To define the latter notion, let e and $\hat{e}$ be extensions. Then $\boldsymbol{e}$ **is contained in** $\boldsymbol{\hat{e}}$, denoted $\boldsymbol{e \subseteq \hat{e}}$, if for each $R \in \mathcal{R}(A)$, $R^e \subseteq R^{\hat{e}}$. The relations $e \subsetneq \hat{e}$ and $e = \hat{e}$ are defined in the obvious manner. One can interpret inclusion among extensions using the "finer-than" relation. If $e \subsetneq \hat{e}$, then $\hat{e}$ is more capable of comparing lotteries than e is; i.e., there are $R \in \mathcal{R}(A)$ and $\pi, \pi' \in \Delta A$ such that π and π' are uncomparable according to R^e but comparable according to $R^{\hat{e}}$. Now Proposition 3.1 shows the inclusion relation among the sd-, dl-, ul-, sd^k- and sd_k-extensions.

Proposition 3.1. *The following inclusion relations hold:*

(1) $sd \subsetneq sd^{m-2} \subsetneq sd^{m-3} \subsetneq \cdots \subsetneq sd^2 \subsetneq sd^1 = dl$.

(2) $sd \subsetneq sd_{m-2} \subsetneq sd_{m-3} \subsetneq \cdots \subsetneq sd_2 \subsetneq sd_1 = ul$.

The above proposition shows, in particular, that the *sd*-extension is contained in the *dl*-, *ul*-, sd^k-, and sd_k-extensions. In fact, a much stronger statement is true:

Proposition 3.2. *Let e be a* monotone *and* transitive *extension. Then $sd \subseteq e$.*

Remark 3.1. Let $\boldsymbol{\mathcal{U}}$ be the class of all expected utility preferences over ΔA. For each $R \in \mathcal{R}(A)$ and each $u \in \mathcal{U}$, u is **consistent with $\boldsymbol{R}$** if the restriction of u to A is the same as R. Many authors (e.g., Bogomolnaia and Moulin, 2001; Che and Kojima, 2010; Kojima, 2009) justify the use of the *sd*-extension by invoking the following theorem: for each $R \in \mathcal{R}(A)$ and each pair $\pi, \pi' \in \Delta A$, $\pi \, R^{sd} \, \pi'$ iff for each $u \in \mathcal{U}$ such that u is consistent with R, $u(\pi) \geq u(\pi')$. However, the reference to the theorem implicitly assumes that agents behave within the expected utility framework, satisfying the independence and continuity axioms. By contrast, Proposition 3.2 reveals the appeal of the *sd*-extension relying on two, more elementary properties: as long as a *monotone* and *transitive* extension e is used, if $\pi \, R^{sd} \, \pi'$, then $\pi \, R^{e} \, \pi'$. Therefore, the *sd*-extension is *minimal* among all extensions satisfying the axioms. $\triangle$

Among lexicographic extensions, of particular interest are the *dl*- and *ul*-extensions. They order lotteries linearly. Noting this linearity, one may ask whether they are "maximal" extensions with respect to some axioms. The next proposition answers this question in the affirmative.

Proposition 3.3. *(1) Let e be an extension satisfying* anti-symmetry *and* probability compensation for preferred objects. *Then $e \subseteq dl$.*

(2) Let e be an extension satisfying anti-symmetry *and* probability compensation for less preferred objects. *Then $e \subseteq ul$.*

Now we apply the notion of duality to the extensions introduced in Section 3.3.

Proposition 3.4. *(1) The sd-extension is* self-dual.
(2) The dl- and ul-extensions are dual.
(3) For each $k \in \{1, \cdots, m-1\}$, the sd^k- and sd_k-extensions are dual.

Given extensions e and $\hat{e}$, let $\boldsymbol{e \cap \hat{e}}$ denote the extension defined by, for each $R \in \mathcal{R}(A)$, $R^{e \cap \hat{e}} \equiv R^{e} \cap R^{\hat{e}}$. The operation $\boldsymbol{e \cup \hat{e}}$ is defined similarly, with $\cup$ in place of $\cap$. It is easy to check that for each extension e, $e \cap e^d$ is *self-dual*. Since the *sd*-extension is the minimal extension satisfying *monotonicity* and *transitivity* (Proposition 3.2) and since it is *self-dual*, we ask when the $e \cap e^d$-extension, in fact, coincides with the *sd*-extension. The following proposition provides a sufficient condition.

Proposition 3.5. *Let e be an extension satisfying* monotonicity, transitivity, *and* condition C. *Then $e \cap e^d = sd$.*

4 Probabilistic Assignment

4.1 The Model

Having addressed the issue of extending preferences over objects to preferences over lotteries, we now introduce more agents into the model and study probabilistic assignment. This section uses many concepts and notations from Section 3. Let $\boldsymbol{A} \equiv \{1, \cdots, n\}$ be the set of objects and $\boldsymbol{N} \equiv \{1, \cdots, n\}$ the set of agents. Assume that $n \geq 2$, and note that there are the same number of agents and objects. We denote objects by k, ℓ, k', ℓ', and so on, and agents by i, j, i', j', and so on. For each $i \in N$, agent i has a complete, transitive, and anti-symmetric preference relation $\boldsymbol{R_i}$ over A.[10] Let $\boldsymbol{P_i}$ and $\boldsymbol{I_i}$ be the strict preference and indifference relations, respectively, associated with R_i. Let $\boldsymbol{\mathcal{R}(A)}$ be the set of all such preferences. Since the sets of agents and objects are fixed throughout, an **economy** is a profile of preferences $\boldsymbol{R} \equiv (R_i)_{i \in N}$. Let $\boldsymbol{\mathcal{R}(A)^N}$ denote the set of all economies.

Given an economy $R \in \mathcal{R}(A)^N$, a (feasible) probabilistic assignment, or simply an **assignment**, for R is a profile of lotteries $\boldsymbol{\pi} \equiv (\pi_i)_{i \in N}$ such that (i) for each $i \in N$, $\pi_i \in \Delta A$; and (ii) for each $k \in A$, $\sum_{i \in N} \pi_{ik} = 1$. We call π_i agent i's lottery. Let $\boldsymbol{\Pi}$ be the set of all assignments. If, for each $i \in N$, π_i is a degenerate lottery, then π is a deterministic assignment. By the Birkhoff-von Neumann theorem (Birkhoff, 1946; von Neumann, 1953), each probabilistic assignment can be written as a convex combination of deterministic assignments.[11] An (assignment) **rule** is a mapping from $\mathcal{R}(A)^N$ to Π that associates with each economy an assignment. Our generic notation for a rule is the letter φ.

Two rules have been studied extensively in the literature: the serial and random priority rules, denoted $\boldsymbol{S}$ and $\boldsymbol{RP}$, respectively. We refer the reader to Bogomolnaia and Moulin (2001) for the definition of these rules. In Section 4.3, we propose a family of rules that generalizes the idea underlying the serial rule.

Since agents have preferences over objects and receive lotteries over them, we need to make an assumption as to their preferences over lotteries.[12] A standard approach in the literature is to assume that the latter preferences are obtained by extending preferences over objects using the *sd*-extension, but in this paper, we work with a general extension. This alternative has two advantages. First, because all axioms on rules and results are stated for a general extension, we

[10] Note the change in notation: preference relations and lotteries now have subscripts to indicate the identity of agents.

[11] Budish et al. (2011, Theorem 1) and Kojima and Manea (2010, Proposition 1) provide a generalization of this result.

[12] It may well be the case that each agent has a preference relation defined on the set of lotteries over objects, but the modeller, for whatever reason, only has access to the restriction of those preference relations to degenerate lotteries (objects). In either case, how agents evaluate lotteries remains an assumption on the part of the modeller.

can always substitute in an extension of our interest and reinterpret their implications. Second, abstracting away from the specifics of extensions, we can better reveal the factors that drive the results.

Now let $\boldsymbol{\mathcal{R}(\Delta A)}$ be the set of all reflexive relations over ΔA. Members of $\mathcal{R}(\Delta A)$ are called preferences over lotteries. An extension is a mapping from $\mathcal{R}(A)$ to $\mathcal{R}(\Delta A)$. Let e be an extension. We assume that for each $i \in N$, agent i with preference relation $R_i \in \mathcal{R}(A)$ evaluates lotteries according to $R_i^e \in \mathcal{R}(\Delta A)$.

To introduce further notation, for each $R \in \mathcal{R}(A)^N$, let $\boldsymbol{R^{-1}} \equiv (R_i^{-1})_{i \in N}$. Denote by $(P_i^{-1})^e$ and $(I_i^{-1})^e$ the strict preference and indifference relations, respectively, associated with $(R_i^{-1})^e$.

4.2 Axioms on Assignment Rules

Our first axiom is efficiency, and to state it, we first define the notion of Pareto dominance. Given an economy $R \in \mathcal{R}(A)^N$ and assignments $\pi, \pi' \in \Pi$, $\boldsymbol{\pi}$ ***e*-Pareto dominates** $\boldsymbol{\pi'}$ **for** $\boldsymbol{R}$ if (i) for each $i \in N$, $\pi_i \, R_i^e \, \pi_i'$; and (ii) for some $i \in N$, $\pi_i \, P_i^e \, \pi_i'$. An assignment is ***e*-efficient for** $\boldsymbol{R}$ if no other assignment e-Pareto dominates it for R. The following axiom requires that for each economy, a rule select an *e-efficient* assignment.

***e*-Efficiency:** For each $R \in \mathcal{R}(A)^N$, $\varphi(R)$ is *e-efficient* for R.

Next is a fairness axiom that originates in Tinbergen (1953) and Foley (1967). It says that no agent should prefer someone else's lottery to his own.

***e*-No-envy:** For each $R \in \mathcal{R}(A)^N$ and each pair $i, j \in N$, $\varphi_i(R) \, R_i^e \, \varphi_j(R)$.

When extension e is not *complete*, *e-no-envy* may be violated for two reasons: (i) there is an agent who prefers someone else's lottery to his; or (ii) there is an agent who finds his lottery uncomparable to someone else's. Therefore, we can formulate a weaker version of *e-no-envy* that only excludes (i).

***e*-Weak No-envy:** For each $R \in \mathcal{R}(A)^N$ and each pair $i, j \in N$, if $\varphi_j(R) \, R_i^e \, \varphi_i(R)$, then $\varphi_i(R) \, I_i^e \, \varphi_j(R)$.

The next requirement concerns the strategic behavior of agents. When agents' preferences are private information, they may find it in their interest to misrepresent their preferences and manipulate the rule. The following axiom requires immunity to such behavior; i.e., whatever other agents' announcements are, no agent ever profits from lying about his preferences.

***e*-Strategy-proofness:** For each $R \in \mathcal{R}(A)^N$, each $i \in N$, and each $R_i' \in \mathcal{R}(A)$, $\varphi_i(R) \, R_i^e \, \varphi_i(R_i', R_{-i})$.

Next, we weaken *e-strategy-proofness* in two directions. First, as is the case with *e-no-envy*, if e is not *complete*, a violation of *e-strategy-proofness* can arise due to uncomparability of lotteries.

The next axiom relaxes the requirement of *e-strategy-proofness* by allowing for such cases: the lottery an agent receives by telling any lie is either at most as desirable as or uncomparable to the lottery he receives by telling the truth.

***e*-Weak Strategy-proofness:** For each $R \in \mathcal{R}(A)^N$, each $i \in N$, and each $R'_i \in \mathcal{R}(A)$, if $\varphi_i(R'_i, R_{-i})\, R^e_i\, \varphi_i(R)$, then $\varphi_i(R)\, I^e_i\, \varphi_i(R'_i, R_{-i})$.

A number of authors study this weaker version of *e-strategy-proofness*, using the *sd*-extension (e.g., Bogomolnaia and Moulin, 2001; Kojima, 2009). As established in Proposition 4.7, the requirement of *e-weak strategy-proofness* is indeed very weak, especially when the *sd*-extension is used.

Another weakening of *e-strategy-proofness* constrains the set of possible lies that an agent can report. As is clear from the definition, *e-strategy-proofness* imagines a situation in which an agent can announce any preference relation in $\mathcal{R}(A)$. An alternative to this scenario is to assume that each agent chooses a preference relation from a smaller subset of $\mathcal{R}(A)$, and require a restricted form of immunity to misrepresentation. A natural candidate for the set of possible lies consists of those preferences that are "close" to the true preferences. To formalize this idea of closeness, we metrize the space of preferences as follows. First, for each pair $R_0, R'_0 \in \mathcal{R}(A)$, $\boldsymbol{R'_0}$ **is an adjacent-pair-switch transformation of** $\boldsymbol{R_0}$ if there is $\hat{k} \in \{1, \cdots, n\}$ such that (i) $\hat{k}(R_0) = (\hat{k}+1)(R'_0)$; (ii) $(\hat{k}+1)(R_0) = \hat{k}(R'_0)$; and (iii) for $k \in \{1, \cdots, n\} \setminus \{\hat{k}, \hat{k}+1\}$, $k(R_0) = k(R'_0)$.[13] That is, R'_0 is obtained from R_0 by switching only the rankings of the $\hat{k}$th and $\hat{k}+1$st most preferred objects. Now define a metric $d(\cdot, \cdot)$ on $\mathcal{R}(A)$: for each pair $R_0, R'_0 \in \mathcal{R}(A)$, let $\boldsymbol{d(R_0, R'_0)}$ be the smallest number of adjacent-pair-switch transformations that are necessary to change R_0 to R'_0.[14] In view of the metric $d(\cdot, \cdot)$, if $d(R_0, R'_0) = 1$, we also say that $\boldsymbol{R'_0}$ **is adjacent to** $\boldsymbol{R_0}$.

Figure 1 shows $\mathcal{R}(A)$ as metrized by $d(\cdot, \cdot)$ when $A \equiv \{1, 2, 3, 4\}$. In the figure, "1432", for instance, stands for the preference relation R_0 such that $1\, P_0\, 4\, P_0\, 3\, P_0\, 2$. Since there are four objects, each preference relation has three adjacent ones, and they are each connected by a single arc (straight line). The distance between preference relations 1432 and 2314, say, is the smallest number of arcs that we go through when traveling from 1432 to 2314, using only the arcs indicated in the figure. There are multiple paths achieving that smallest number, and the path $1432 - 1423 - 1243 - 2143 - 2134 - 2314$ is one of them. Therefore, $d(1432, 2314) = 5$. Note incidentally that the representation in Figure 1 partially orders $\mathcal{R}(A)$, with 1234 and 4321 being the smallest and greatest elements, respectively.

With this notion of distance in mind, *e-strategy-proofness* can be viewed as permitting the possibility that each agent can submit any preference relation, regardless of how far it is from his

[13]Preference relations and lotteries that are not associated with any particular agent have the subscript "0".

[14]The metric $d(\cdot, \cdot)$ is known as the Kemeny metric (Kemeny, 1959; Kemeny and Snell, 1962).

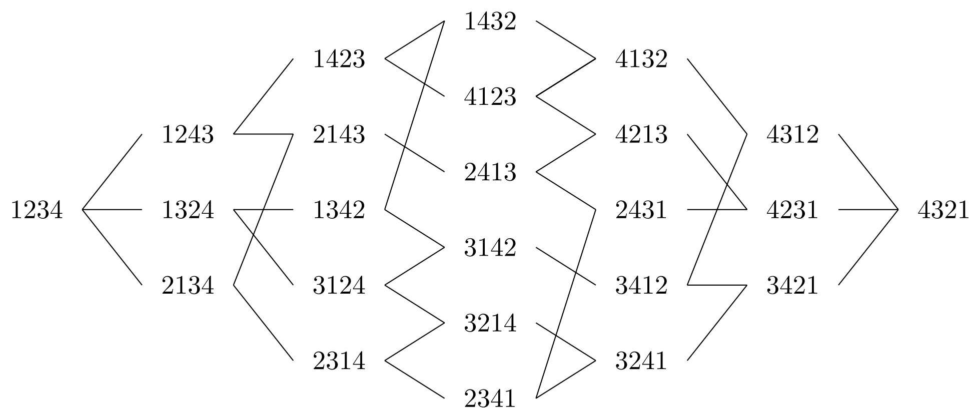

Figure 1: **Metrizing the space of preference relations.** Let $A \equiv \{1, 2, 3, 4\}$. In the figure, "1432", for instance, stands for the preference relation $R_0 \in \mathcal{R}(A)$ such that $1\, P_0\, 4\, P_0\, 3\, P_0\, 2$. Adjacent preference relations are connected by a single arc (straight line). The distance between two preference relations is the smallest number of arcs that we go through when traveling from one preference relation to the other, using only the arcs indicated in the figure. For example, $d(1432, 2314) = 5$.

true preference relation according to $d(\cdot, \cdot)$.[15] However, there are situations in which agents are constrained to announce a preference relation that is more or less close to the truth. This may be because preposterous lies are not so credible or because agents can tell a lie only about the part of private information that has not been disclosed yet.[16] Therefore, it is interesting to consider a strategic requirement that specifies the set of preference relations from which agents can choose. An extreme case is to assume that each agent announces a preference relation that is closest to his true preference relation according to metric $d(\cdot, \cdot)$. The following axiom requires that no agent be better off announcing a preference relation that is adjacent to the truth.

e-Adjacent Strategy-proofness: For each $R \in \mathcal{R}(A)^N$, each $i \in N$, and each $R'_i \in \mathcal{R}(A)$ such that R'_i is adjacent to R_i, $\varphi_i(R)\, R^e_i\, \varphi_i(R'_i, R_{-i})$.

While one might be able to provide plausible examples in which *e-adjacent strategy-proofness* is compelling, our interest primarily concerns the ease that it affords when we check the stronger axiom, *e-strategy-proofness*. As Theorem 4.2 shows below, for some extensions, *e-adjacent strategy-proofness*, in fact, implies *e-strategy-proofness*, so that the toil of verifying the latter property can be reduced greatly.

Next are two new axioms. While appearing to hold little normative appeal, as shown in Section 4.4, they are closely related to *e-no-envy* and *e-strategy-proofness* and reveal the hidden

[15] While we mechanically defined the metric $d(\cdot, \cdot)$, it can be derived as a consequence of a list of axioms on metrics over $\mathcal{R}(A)$; see Kemeny (1959) for details. Thus, the further away a preference relation is from the true preference relation according to $d(\cdot, \cdot)$, the bigger a lie it is.

[16] See Thomson (2011) for a detailed discussion.

content of the latter two axioms. In the first axiom, which is a variant of *e-no-envy*, we consider a situation in which an agent reports the inverse of his true preference relation, and examine the assignment that a rule recommends from the perspective of envy. The axiom requires that the agent find his lottery least desirable, according to his true preference relation, among all the lotteries in the assignment.

***e*-Complete Envy for Inverse Reports:** For each $R \in \mathcal{R}(A)^N$ and each pair $i, j \in N$, $\varphi_j(R_i^{-1}, R_{-i}) \, R_i^e \, \varphi_i(R_i^{-1}, R_{-i})$.

The second axiom looks at the same situation from the viewpoint of incentives, and therefore, it is related to *e-strategy-proofness*. It says that the lottery that an agent receives by reporting the inverse of his true preference relation is least desirable, according to the true preference relation, among all the lotteries that he can receive by reporting any preference relation.

***e*-Lowest Welfare for Inverse Reports:** For each $R \in \mathcal{R}(A)^N$, each $i \in N$, and each $R_i' \in \mathcal{R}(A)$, $\varphi_i(R_i', R_{-i}) \, R_i^e \, \varphi_i(R_i^{-1}, R_{-i})$.

It is not so clear why we should impose the above two axioms on rules. However, it turns out that imposing *sd-no-envy* and *sd-strategy-proofness* amounts to imposing *sd-complete envy for inverse reports* and *sd-lowest welfare for inverse reports*, respectively. In fact, the former two axioms have even stronger implications (Propositions 4.5 and 4.9).

4.3 Generalized Serial Rules

To introduce the family of generalized serial rules, we first define "allocation speed functions". An **(allocation) speed function** is a mapping $\sigma : A \times \mathbb{R}_+ \to \mathbb{R}_{++}$ such that for each $k \in A$, (i) $\sigma(k, \cdot)$ is measurable; and (ii) there is $\bar{t} < \infty$ satisfying $\int_0^{\bar{t}} \sigma(k,t)dt = 1$. Let $\boldsymbol{\Sigma}$ be the family of all such functions.

Let $\sigma \in \Sigma$. We define the **generalized serial rule associated with $\boldsymbol{\sigma}$**, denoted $\boldsymbol{S^\sigma}$, by the **generalized simultaneous consumption algorithm associated with $\boldsymbol{\sigma}$**. The latter algorithm is similar to the simultaneous consumption algorithm that Bogomolnaia and Moulin (2001) propose. To illustrate, let us treat each object as a continuum of measure 1, consisting of "probability shares" of the object. At time $t = 0$, each agent goes to the object that he most prefers. For each $(k, t) \in A \times \mathbb{R}_+$, at time t, the probability shares of object k are allocated at speed $\sigma(k, t)$ to those agents who have come to object k. When object k is exhausted and agent i, say, has received probability shares of measure less than 1, he moves on to the object, say ℓ, that he most prefers among the remaining ones, and receives the probability shares of object ℓ at speed $\sigma(\ell, t)$. The process continues in the same manner, and it terminates when probability shares of each object are fully allocated. Note that for each $(k,t) \in A \times \mathbb{R}_+$, the speed $\sigma(k,t)$ of object k at time t is independent of the identity of the agents.

To define the algorithm formally, for each $R \in \mathcal{R}(A)^N$, each $B \subseteq A$, and each $k \in B$, let $\boldsymbol{N^*(R,B,k)} \equiv \{i \in N : \text{ for each } \ell \in B,\ k\, R_i\, \ell\}$. Let $A^0 \equiv A$, $t^0 \equiv 0$, and $\pi^0 \equiv (0)_{i \in N, k \in A}$. For each $s \in \mathbb{N}$, given $(A^{s-1}, t^{s-1}, \pi^{s-1})$, define (A^s, t^s, π^s) recursively as follows. For each $k \in A^{s-1}$, let

$$t^s(k) \equiv \inf \left\{ t \in \mathbb{R}_+ : |N^*(R, A^{s-1}, k)| \cdot \int_{t^{s-1}}^{t} \sigma(k,\tau)\, d\tau + \sum_{i \in N} \pi_{ik}^{s-1} = 1 \right\}$$

if $N^*(R, A^{s-1}, k) \neq \emptyset$; and $t^s(k) \equiv \infty$ otherwise. Let $t^s \equiv \min_{k \in A^{s-1}} t^s(k)$; $A^s = A^{s-1} \setminus \{k \in A^{s-1} : t^s(k) = t^s\}$; and $\pi^s \equiv (\pi_{ik}^s)_{i \in N, k \in A}$, where for each $i \in N$ and each $k \in A$,

$$\pi_{ik}^s \equiv \begin{cases} \pi_{ik}^{s-1} + \int_{t^{s-1}}^{t^s} \sigma(k,t)\, dt & \text{if } i \in N^*(R, A^{s-1}, k); \\ \pi_{ik}^{s-1} & \text{otherwise.} \end{cases}$$

By condition (ii) in the definition of speed functions, it follows that for each $s \in \mathbb{N}$, $t^s < \infty$. Also, for each $s \in \mathbb{N}$, $A^{s-1} \subsetneq A^s$, so that for each $s \in \mathbb{N}$ with $s \geq n$, $A^s = \emptyset$ and $\pi^s = \pi^n$. Therefore, $\pi^n \in \Pi$, and π^n is the assignment that S^σ chooses for R; i.e., $S^\sigma(R) = \pi^n$.

If $\sigma \in \Sigma$ is such that for each pair $k, \ell \in A$ and each $t \in \mathbb{R}_+$, $\sigma(k,t) = \sigma(\ell,t)$, then S^σ coincides with the serial rule. As is apparent from its definition, a generalized serial rule differs from the serial rule in that objects may not be treated in the same manner, and this leads to an asymmetric treatment of agents in the following sense. Consider two objects, say objects 1 and 2. Let $\sigma \in \Sigma$ be such that for each $t \in \mathbb{R}_+$, $\sigma(1,t) > \sigma(2,t)$. Then in the generalized simultaneous consumption algorithm associated with σ, the probability shares of object 1 are distributed at a higher rate than the probability shares of object 2 are, and therefore, those agents who prefer object 1 to object 2 are treated more favorably than those who prefer object 2 to object 1. This, however, does not mean that S^σ favors any particular agent across all economies. The identity of agents whom the rule favors depends on the economy at hand.

4.4 Results

4.4.1 Efficiency

With a variety of extensions available, it is important to explore how efficiency concepts based on different extensions are related. In pursuing this question, we first recall the inclusion relation and the duality operation defined over extensions, and study their consequences on efficiency. The following proposition shows logical relations among various efficiency concepts.

Proposition 4.1. *Let e and $\hat{e}$ be extensions such that $e \subseteq \hat{e}$.*

(1) $\hat{e}$-efficiency *implies* e-efficiency.

(2) Let $\tilde{e}$ be an extension such that $e \subseteq \tilde{e} \subseteq \hat{e}$. If $\hat{e}$-efficiency *is equivalent to* e-efficiency, *then so is* $\tilde{e}$-efficiency.

(3) For each $R \in \mathcal{R}(A)^N$ and each $\pi \in \Pi$, π is e-efficient *for R iff there is no assignment $\pi' \in \Pi$ such that π e^d-Pareto dominates π' for R^{-1}.*

Proof. Part (1) follows by definition.

To prove part (2), suppose that $\hat{e}$*-efficiency* is equivalent to e*-efficiency*. First, by part (1), $\tilde{e}$*-efficiency* implies e*-efficiency*. Further, since e*-efficiency* is equivalent to $\hat{e}$*-efficiency*, which implies $\tilde{e}$*-efficiency*, e*-efficiency* also implies $\tilde{e}$*-efficiency*. Thus, $\tilde{e}$*-efficiency* is equivalent to e*-efficiency*.

To prove part (3), first assume that π is e*-efficient* for R. Suppose, on the contrary, that there is $\pi' \in \Pi$ such that π e^d-Pareto dominates π' for R^{-1}. Then for each $i \in N$, $\pi_i \, (R_i^{-1})^{e^d} \, \pi'_i$, and for some $j \in N$, $\pi_j \, (P_j^{-1})^{e^d} \, \pi'_j$. By duality, for each $i \in N$, $\pi_i \, (R_i^e)^{-1} \, \pi'_i$, and $\pi_j \, (P_j^e)^{-1} \, \pi'_j$, so that π' e-Pareto dominates π for R, a contradiction. The converse can be proved similarly. □

Next, we turn to a characterization of efficiency. Since we have an inventory of extensions, one might suspect that the choice of extension e varies the content of e*-efficiency*. It turns out, however, that the *sd*-, *dl*-, *ul*-, sd^k- and sd_k-extensions—and many others that we have not defined—all give rise to the same efficiency concept. To prove this result, for each $R \in \mathcal{R}(A)^N$ and each $\pi \in \Pi$, define a binary relation $\boldsymbol{\tau(R, \pi)}$ over A as follows: for each pair $k, \ell \in A$, let $k \, \tau(R, \pi) \, \ell$ if there is $i \in N$ such that $k \, P_i \, \ell$ and $\pi_{i\ell} > 0$. Bogomolnaia and Moulin (2001) show that *sd-efficiency* is equivalent to the acyclicity of $\tau(\cdot, \cdot)$, but the logical relation is much more general. To see this, we first identify necessary and sufficient conditions for e*-efficiency*.

Theorem 4.1. *Let e be an extension. Let $R \in \mathcal{R}(A)^N$ and $\pi \in \Pi$.*

(1) Assume that e satisfies monotonicity. *If π is* e-efficient *for R, then $\tau(R, \pi)$ is acyclic.*

(2) Assume that e satisfies anti-symmetry *and either* probability compensation for preferred objects *or* probability compensation for less preferred objects. *If $\tau(R, \pi)$ is acyclic, then π is* e-efficient *for R.*

Proof. The simple proof of part (1) is omitted. To prove part (2), let e be as in the theorem. Let $R \in \mathcal{R}(A)^N$ and $\pi \in \Pi$. Assume that $\tau(R, \pi)$ is acyclic. Suppose, on the contrary, that π is not e*-efficient* for R; i.e., there is $\pi' \in \Pi$ such that for each $i \in N$, $\pi'_i \, R_i^e \, \pi_i$, and for some $i_1 \in N$, $\pi'_{i_1} \, P_{i_1}^e \, \pi_{i_1}$. We distinguish two cases.

Case 1: e satisfies *probability compensation for preferred objects*.

By *probability compensation for preferred objects*, there are $k_1, k_2 \in A$ such that $k_2 \, P_{i_1} \, k_1$, $\pi_{i_1 k_1} > \pi'_{i_1 k_1}$, and $\pi_{i_1 k_2} < \pi'_{i_1 k_2}$. Thus, $\pi_{i_1 k_1} > 0$, so that $k_2 \, \tau(R, \pi) \, k_1$. Now because $\pi_{i_1 k_2} < \pi'_{i_1 k_2}$, by feasibility, there is $i_2 \in N$ such that $\pi_{i_2k_2} > \pi'_{i_2k_2}$. This implies, in particular, that $\pi_{i_2} \neq \pi'_{i_2}$, and by *anti-symmetry*, $\pi'_{i_2} \, P^e_{i_2} \, \pi_{i_2}$. By *probability compensation for preferred objects*, there is $k_3 \in A$ such that $k_3 \, P_{i_2} \, k_2$ and $\pi_{i_2k_3} < \pi'_{i_2k_3}$. Thus, because $\pi_{i_2k_2} > 0$, $k_3 \, \tau(R,\pi) \, k_2$. Continuing this process, by finiteness of A, we can construct a cycle of $\tau(R,\pi)$, a contradiction.

Case 2: e satisfies *probability compensation for less preferred objects*.

By *probability compensation for less preferred objects*, there are $k_1, k_2 \in A$ such that $k_1 \, P_{i_1} \, k_2$, $\pi_{i_1k_1} < \pi'_{i_1k_1}$, and $\pi_{i_1k_2} > \pi'_{i_1k_2}$. Thus, $\pi_{i_1k_2} > 0$, so that $k_1 \, \tau(R,\pi) \, k_2$. Now because $\pi_{i_1k_2} > \pi'_{i_1k_2}$, by feasibility, there is $i_2 \in N$ such that $\pi_{i_2k_2} < \pi'_{i_2k_2}$. This implies, in particular, that $\pi_{i_2} \neq \pi'_{i_2}$, and by *anti-symmetry*, $\pi'_{i_2} \, P^e_{i_2} \, \pi_{i_2}$. By *probability compensation for less preferred objects*, there is $k_3 \in A$ such that $k_2 \, P_{i_2} \, k_3$ and $\pi_{i_2k_3} > \pi'_{i_2k_3}$. Thus, $\pi_{i_2k_3} > 0$, so that $k_2 \, \tau(R,\pi) \, k_3$. Continuing this process, by finiteness of A, we can construct a cycle of $\tau(R,\pi)$, a contradiction. $\square$

Combining parts (1) and (2) of Theorem 4.1 gives the following corollary.

Corollary 4.1. *Let e be an extension satisfying* monotonicity, anti-symmetry*, and either* probability compensation for preferred objects *or* probability compensation for less preferred objects. *Then for each $R \in \mathcal{R}(A)^N$ and each $\pi \in \Pi$, π is* e-efficient *for R iff $\tau(R,\pi)$ is acyclic.*

Remark 4.1. Since *sd-efficiency* is equivalent to the acyclicity of $\tau(\cdot,\cdot)$, it follows that for extension e satisfying the axioms listed in Corollary 4.1, alternative characterizations of *sd-efficiency* in Abdulkadiroğlu and Sönmez (2003) and McLennan (2002) also characterize *e-efficiency.* $\triangle$

Since the *sd*-, *dl*-, *ul*-, sd^k- and sd_k-extensions all satisfy the axioms listed in Corollary 4.1, we obtain the following corollary.

Corollary 4.2. *The following efficiency concepts are equivalent:* *sd*-efficiency*,* *dl*-efficiency*,* *ul*-efficiency*,* sd^k-efficiency*, and* sd_k-efficiency*, where $k \in \{1, \cdots, n-1\}$.*

Now consider a *monotone* and *transitive* extension e. Assume, in addition, that either $e \subseteq dl$ or $e \subseteq ul$. By Proposition 3.2, it follows that $sd \subseteq e$, so that either $sd \subseteq e \subseteq dl$ or $sd \subseteq e \subseteq ul$. Then part (2) of Proposition 4.1 and Corollary 4.2 imply that *e-efficiency* is also equivalent to *sd-efficiency*. Thus, we have proved the following:

Corollary 4.3. *Let e be an extension satisfying* monotonicity, transitivity*, and either $e \subseteq dl$ or $e \subseteq ul$. Then* e-efficiency *is equivalent to* *sd*-efficiency.

4.4.2 No-envy

Next, we study logical relations among various no-envy concepts. Since no-envy requires that no agent prefer someone else's lottery to his own, as an extension becomes finer, the no-envy notion associated with it becomes weaker. The straightforward proof of this fact is omitted.

Proposition 4.2. *Let e and $\hat{e}$ be extensions such that $e \subseteq \hat{e}$. Then* e-no-envy *implies* $\hat{e}$-no-envy.

The extensions introduced in Section 3.3 are related by the inclusion relation $\subseteq$ (Proposition 3.1). Thus, we obtain a corollary to the above proposition that relates the no-envy concepts associated with them.

Corollary 4.4. *The following logical relations hold:*
(1) sd-no-envy $\Longrightarrow$ sd^{n-2}-no-envy $\Longrightarrow$ sd^{n-3}-no-envy $\Longrightarrow \cdots \Longrightarrow$ sd^1-no-envy $=$ dl-no-envy.
(2) sd-no-envy $\Longrightarrow$ sd_{n-2}-no-envy $\Longrightarrow$ sd_{n-3}-no-envy $\Longrightarrow \cdots \Longrightarrow$ sd_1-no-envy $=$ ul-no-envy.

Now we present a proposition that connects *sd-no-envy* and *sd-weak no-envy* with *e-no-envy* and *e-weak no-envy.* Among all no-envy and weak no-envy concepts associated with *monotone*, *transitive*, and *anti-symmetric* extensions, *sd-no-envy* is the strongest and *sd-weak no-envy* is the weakest. This result is due in large part to the fact that the sd-extension is minimal among all *monotone* and *transitive* extensions (Proposition 3.2).

Proposition 4.3. *Let e be a* monotone*,* transitive*, and* anti-symmetric *extension. Then the following logical relations hold:*
sd-no-envy $\Longrightarrow$ *e-no-envy* $\Longrightarrow$ *e-weak no-envy* $\Longrightarrow$ *sd-weak no-envy.*

Proof. Let e be as in the proposition. By Proposition 3.2, $sd \subseteq e$. Then the first implication follows from Proposition 4.2, and the second holds by definition.

To prove the third implication, let φ be a rule satisfying *e-weak no-envy.* Let $R \in \mathcal{R}(A)^N$ and $i, j \in N$. If $\varphi_j(R) \, R_i^{sd} \, \varphi_i(R)$, then since $sd \subseteq e$, $\varphi_j(R) \, R_i^e \, \varphi_i(R)$, so that by *e-weak no-envy,* $\varphi_i(R) \, I_i^e \, \varphi_j(R)$. Now by *anti-symmetry* of e, $\varphi_i(R) = \varphi_j(R)$. Thus, φ satisfies *sd-weak no-envy.* □

Next, we reinterpret *e-complete envy for inverse reports* in the context of *e-no-envy.* The following proposition shows that the two axioms are related by the duality of extensions.

Proposition 4.4. *Let e be an extension. Then* e-no-envy *is equivalent to* e^d-complete envy for inverse reports.

Proof. Let φ be a rule satisfying *e-no-envy.* Let $R \in \mathcal{R}(A)^N$ and $i, j \in N$. Consider an economy (R_i^{-1}, R_{-i}). By *e-no-envy,* $\varphi_i(R_i^{-1}, R_{-i}) \, (R_i^{-1})^e \, \varphi_j(R_i^{-1}, R_{-i})$. Then by duality, $\varphi_j(R_i^{-1}, R_{-i}) \, R_i^{e^d}$ $\varphi_i(R_i^{-1}, R_{-i})$, so that φ satisfies *e^d-complete envy for inverse reports.* The converse can be proved similarly. □

The above proposition has an immediate application for the sd-extension: because the sd-extension is *self-dual*, *sd-no-envy* is, in fact, equivalent to *sd-complete envy for inverse reports.* Moreover, this observation can be generalized to "decompose" *sd-no-envy* into two axioms based on a general extension.

Proposition 4.5. *Let e be a* monotone *and* transitive *extension.*

*(1) sd-*no-envy *implies* e-no-envy *and* e-complete envy for inverse reports.

(2) If, in addition, e satisfies condition C, *then the converse of part (1) also holds.*

Proof. To show part (1), let φ be a rule satisfying *sd-no-envy*. First, by Proposition 4.3, φ satisfies *e-no-envy*. To show that φ also satisfies *e-complete envy for inverse reports*, let $R \in \mathcal{R}(A)^N$ and $i, j \in N$. Consider an economy (R_i^{-1}, R_{-i}). By *sd-no-envy*, $\varphi_i(R_i^{-1}, R_{-i})\,(R_i^{-1})^{sd}\,\varphi_j(R_i^{-1}, R_{-i})$. Since e^d is also *monotone* and *transitive*, Proposition 3.2 implies that $sd \subseteq e^d$, so that $\varphi_i(R_i^{-1}, R_{-i})$ $(R_i^{-1})^{e^d}\,\varphi_j(R_i^{-1}, R_{-i})$. By duality, $\varphi_j(R_i^{-1}, R_{-i})\,R_i^e\,\varphi_i(R_i^{-1}, R_{-i})$. Thus, φ satisfies *e-complete envy for inverse reports*.

To show part (2), let φ be a rule satisfying *e-no-envy* and *e-complete envy for inverse reports*. Let $R \in \mathcal{R}(A)^N$ and $i, j \in N$. By *e-no-envy*, $\varphi_i(R)\,R_i^e\,\varphi_j(R)$. Also, by *e-complete envy for inverse reports*, $\varphi_j(R)\,(R_i^{-1})^e\,\varphi_i(R)$, so that by duality, $\varphi_i(R)\,R_i^{e^d}\,\varphi_j(R)$. Thus, $\varphi_i(R)\,R_i^{e \cap e^d}\,\varphi_j(R)$. Since, by Proposition 3.5, $sd = e \cap e^d$, it follows that $\varphi_i(R)\,R_i^{sd}\,\varphi_j(R)$. Therefore, φ satisfies *sd-no-envy*. □

Remark 4.2. By Proposition 4.4, the above proposition still holds if we replace *e-complete envy for inverse reports* by e^d*-no-envy*. △

4.4.3 Strategy-proofness

Now we present results on strategy-proofness, starting with the one that concerns inclusion of extensions. As is the case with *e-no-envy*, as extension e becomes finer, *e-strategy-proofness* becomes weaker.

Proposition 4.6. *Let e and $\hat{e}$ be extensions such that $e \subseteq \hat{e}$. Then* e-strategy-proofness *implies* $\hat{e}$-strategy-proofness.

Since the *sd*-, *dl*-, *ul*-, sd^k-, and sd_k-extensions are related by the inclusion relation $\subseteq$ (Proposition 3.1), Proposition 4.6 yields a corollary that relates strategy-proofness concepts based on them.

Corollary 4.5. *The following logical relations hold:*
(1) sd-strategy-proofness $\Longrightarrow$ sd^{n-2}-strategy-proofness $\Longrightarrow$ sd^{n-3}-strategy-proofness $\Longrightarrow \cdots \Longrightarrow$ sd^1-strategy-proofness $=$ dl-strategy-proofness.
(2) sd-strategy-proofness $\Longrightarrow$ sd_{n-2}-strategy-proofness $\Longrightarrow$ sd_{n-3}-strategy-proofness $\Longrightarrow \cdots \Longrightarrow$ sd_1-strategy-proofness $=$ ul-strategy-proofness.

Next, we discuss *sd-strategy-proofness* and *sd-weak strategy-proofness* in connection with the corresponding axioms associated with a general extension. It turns out that *sd-strategy-proofness* is the strongest and *sd-weak strategy-proofness* is the weakest among all strategic requirements based on *monotone*, *transitive*, and *anti-symmetric* extensions. The proof of the following result is similar to that of Proposition 4.3 and is omitted.

Proposition 4.7. *Let e be a* monotone*,* transitive*, and* anti-symmetric *extension. Then the following logical relations hold:*

sd-strategy-proofness $\Longrightarrow$ e-strategy-proofness $\Longrightarrow$ e-weak strategy-proofness $\Longrightarrow$ sd-weak strategy-proofness.

It is clear by definition that *e-adjacent strategy-proofness* weakens *e-strategy-proofness*, but there are extensions such that the two are, indeed, equivalent. The following theorem shows that this is true for the *sd*-, *dl*-, and *ul*-extensions. The proof is relegated to Appendix A.

Theorem 4.2. *For each $e \in \{sd, dl, ul\}$,* e-adjacent strategy-proofness *is equivalent to* e-strategy-proofness.

Remark 4.3. In Appendix A, we identify sufficient conditions on the preference domain under which *e-adjacent strategy-proofness* is equivalent to *e-strategy-proofness*. Depending on the extension chosen, different conditions are needed to guarantee the equivalence, but the present preference domain $\mathcal{R}(A)$ satisfies all of them, and hence, Theorem 4.2 follows.

In the mechanism design literature, requirements in the spirit of adjacent strategy-proofness are also known as local incentive compatibility, and several papers investigate its sufficiency for global incentive compatibility, or strategy-proofness. Carroll (2012) studies various kinds of preference domains, and provides conditions on them for the sufficiency to hold. Among others, he shows that *sd-adjacent strategy-proofness* is equivalent to *sd-strategy-proofness* on preference domains satisfying a certain regularity condition. His condition is more restrictive than the one we identify in Appendix A.

If the mechanism under consideration is deterministic, then for each $e \in \{sd, dl, ul\}$, e-(*adjacent*) *strategy-proofness* reduces to the same requirement, and Theorem 4.2 implies that adjacent strategy-proofness and strategy-proofness—as they are defined in the deterministic setup—are equivalent. This enables us to strengthen many impossibility results in social choice theory, including the Gibbard-Satterthwaite theorem (Gibbard, 1973; Satterthwaite, 1975). In fact, Sato (2010) proves the equivalence for the deterministic case and provides a sufficient condition on preference domains.

Finally, it remains an open question that for what extension e, *e-adjacent strategy-proofness* is equivalent to *e-strategy-proofness*. △

Remark 4.4. Theorem 4.2 yields a corollary on the behavior of *sd*-, *dl*-, and *ul-strategy-proof* rules. Let φ be a rule. In the statement below, we take arbitrary $R \in \mathcal{R}(A)^N$, $i \in N$, and $R'_i \in \mathcal{R}(A)$ such that R'_i is adjacent to R_i. To simplify notation, however, once such (R, i, R'_i) is chosen, (i) relabel objects so that $1\, P_i\, 2\, P_i\, \cdots\, P_i\, n$; (ii) let $k \in A$ be the object such that $(k+1)\, P'_i\, k$; and (iii) let $\pi \equiv \varphi(R)$ and $\pi' \equiv \varphi(R'_i, R_{-i})$.

(i) φ is *sd-strategy-proof*
iff for each $R \in \mathcal{R}(A)^N$, each $i \in N$, and each $R'_i \in \mathcal{R}(A)$ adjacent to R_i,
either (a) $\pi_i = \pi'_i$;
or (b) $\pi_{ik} > \pi'_{ik}$, $\pi_{i,k+1} < \pi'_{i,k+1}$, and for each $\ell \in A \backslash \{k, k+1\}$, $\pi_{i\ell} = \pi'_{i\ell}$.

(ii) φ is *dl-strategy-proof*
iff for each $R \in \mathcal{R}(A)^N$, each $i \in N$, and each $R'_i \in \mathcal{R}(A)$ adjacent to R_i,
either (a) $\pi_i = \pi'_i$;
or (b) $\pi_{ik} > \pi'_{ik}$, $\pi_{i,k+1} < \pi'_{i,k+1}$, and for each $\ell \in \{1, \cdots, k-1\}$, $\pi_{i\ell} = \pi'_{i\ell}$.

(iii) φ is *ul-strategy-proof*
iff for each $R \in \mathcal{R}(A)^N$, each $i \in N$, and each $R'_i \in \mathcal{R}(A)$ adjacent to R_i,
either (a) $\pi_i = \pi'_i$;
or (b) $\pi_{ik} > \pi'_{ik}$, $\pi_{i,k+1} < \pi'_{i,k+1}$, and for each $\ell \in \{k+2, \cdots, n\}$, $\pi_{i\ell} = \pi'_{i\ell}$. △

Theorem 4.2 provides an insight into how an *e-strategy-proof* rule, where $e \in \{sd, dl, ul\}$, behaves as we vary the announcement of one agent while fixing those of the other agents. To state this formally, we first define the notion of direct path. Let $R_0, R'_0 \in \mathcal{R}(A)$, and let $h \equiv d(R_0, R'_0)$. A **direct path from R_0 to R'_0** is the sequence of preference relations $\{R_0^0, R_0^1, \cdots, R_0^h\}$ in $\mathcal{R}(A)$ such that (i) $R_0^0 = R_0$ and $R_0^h = R'_0$; and (ii) for each $\tilde{h} \in \{0, \cdots, h-1\}$, $R_0^{\tilde{h}}$ and $R_0^{\tilde{h}+1}$ are adjacent. To illustrate, let $A \equiv \{1, 2, 3, 4\}$ and refer to Figure 1. Consider two preference relations 1234 and 1432. There are two direct paths from 1234 to 1432: $\{1234, 1243, 1423, 1432\}$ and $\{1234, 1324, 1342, 1432\}$. Now let $R_0 \in \mathcal{R}(A)$. Define an order $\geq_{R_0}$ over $\mathcal{R}(A)$ as follows: for each pair $R'_0, R''_0 \in \mathcal{R}(A)$, $R'_0 \geq_{R_0} R''_0$ if there is a direct path from R_0 to R''_0 containing R'_0. The asymmetric order $>_{R_0}$ associated with $\geq_{R_0}$ is defined in the obvious way. It is easy to check that $\geq_{R_0}$ is reflexive, anti-symmetric, and transitive, so that $(\mathcal{R}(A), \geq_{R_0})$ is a partially ordered set.

Now fix $R \in \mathcal{R}(A)^N$ and $i \in N$. Given a rule φ, we study the function $\varphi_i(\cdot, R_{-i}) : \mathcal{R}(A) \to \Delta A$, using the orders $\geq_{R_i}$ on $\mathcal{R}(A)$ and R_i^e on ΔA, in connection with strategy-proofness. Let $e \in \{sd, dl, ul\}$. If φ is *e-strategy-proof*, then $\varphi_i(\cdot, R_{-i})$ attains its maximum, according to R_i^e, at R_i, but *e-strategy-proofness* alone does not tell us how it behaves as we move away, according to $\geq_{R_i}$, from the true preference relation R_i. The next theorem shows that in fact, the function is monotone.

Theorem 4.3. *Let* $e \in \{sd, dl, ul\}$ *and let* φ *be a rule. Then* φ *is* e-strategy-proof *iff the following holds: for each* $R \in \mathcal{R}(A)^N$ *and each* $i \in N$*, the function* $\varphi_i(\cdot, R_{-i}) : (\mathcal{R}(A), \geq_{R_i}) \rightarrow (\Delta A, R_i^e)$ *is monotone; i.e., for each pair* $R_i', R_i'' \in \mathcal{R}(A)$ *such that* $R_i' \geq_{R_i} R_i''$*,* $\varphi_i(R_i', R_{-i})\, R_i^e\, \varphi_i(R_i'', R_{-i})$*.*

Proof. We only prove the "only if" part of the theorem for the sd-extension; a similar argument applies for the dl- and ul-extensions and the "if" part is obvious. Let $R \in \mathcal{R}(A)^N$ and $i \in N$. Let $R_i', R_i'' \in \mathcal{R}(A)$ be such that $R_i' \geq_{R_i} R_i''$. We may assume that $R_i' \neq R_i''$, so that $R_i' >_{R_i} R_i''$. Let $\{R_i^0, R_i^1, \cdots, R_i^m, \cdots, R_i^h\} \subseteq \mathcal{R}(A)$ be a direct path from R_i to R_i'' containing R_i' such that $R_i^m = R_i'$. For each $\tilde{h} \in \{0, 1, \cdots, h\}$, let $\pi_i^{\tilde{h}} \equiv \varphi_i(R_i^{\tilde{h}}, R_{-i})$.

Consider π_i^m and π_i^{m+1}. Since R_i^m and R_i^{m+1} are adjacent, there are exactly two objects $k, k' \in A$ such that $k\, P_i^m\, k'$ and $k'\, P_i^{m+1}\, k$. Since $\{R_i^0, R_i^1, \cdots, R_i^m, \cdots, R_i^h\}$ is a direct path from R_i to R_i'', it follows that $k\, P_i\, k'$ and $k'\, P_i''\, k$. By the characterization of *sd-strategy-proofness* in Remark 4.4, either (i) $\pi_i^m = \pi_i^{m+1}$; or (ii) $\pi_{ik}^m > \pi_{ik}^{m+1}$, $\pi_{ik'}^m < \pi_{ik'}^{m+1}$, and for each $\ell \in A \backslash \{k, k'\}$, $\pi_{i\ell}^m = \pi_{i\ell}^{m+1}$. Thus, $\pi_i^m\, R_i^{sd}\, \pi_i^{m+1}$.

It is clear that for each $\tilde{h} \in \{m+1, \cdots, h\}$, the previous argument can be adapted to $\pi_i^{\tilde{h}}$ and $\pi_i^{\tilde{h}+1}$, showing that $\pi_i^{\tilde{h}}\, R_i^{sd}\, \pi_i^{\tilde{h}+1}$. Thus, $\pi_i^m\, R_i^{sd}\, \pi_i^h$. □

To illustrate Theorem 4.3, let $A \equiv \{1, 2, 3, 4\}$ and refer to Figure 1. Let $i \in N$ and $R_{-i} \in \mathcal{R}(A)^{N \backslash \{i\}}$. Suppose that agent i's true preference relation is $R_i \equiv 1234$. Consider a direct path from 1234 to 4321, $\{1234, 1243, 1423, 1432, 4132, 4312, 4321\}$. Note that $\geq_{R_i}$ completely orders this direct path. Let $e \in \{sd, dl, ul\}$. Suppose that agent i reports preference relations in the path, sequentially, starting from 1234. Theorem 4.3 says that agent i's welfare, as measured by R_i^e, decreases in the weak sense. In the case of the sd-extension, this, in particular, implies the comparability of all welfare levels attained along the path.

Another application of Theorem 4.2 is to use it to decompose *sd-strategy-proofness* into two strategic requirements. When there are just three objects, it follows by definition that *sd-strategy-proofness* is equivalent to the combination of *dl-strategy-proofness* and *ul-strategy-proofness*. On the other hand, if there are more than three objects, we cannot deduce the equivalence directly from the definition. However, the characterization of *sd-*, *dl-*, and *ul-strategy-proofness* in Remark 4.4 reveals that the equivalence still holds. An easy proof of the following theorem is omitted.

Theorem 4.4. sd-strategy-proofness *is equivalent to the combination of* dl-strategy-proofness *and* ul-strategy-proofness*.*

Now, to study *e-lowest welfare for inverse reports*, we first examine its connection with *e-strategy-proofness*. The next proposition shows that the two are related by the duality of extensions, as are *e-no-envy* and *e-complete envy for inverse reports*. The proof is similar to that of Proposition 4.4 and is omitted.

$$\text{(a)}\quad \textit{sd-SP} \iff \textit{sd-ASP} \iff \textit{dl-ASP} \text{ and } \textit{ul-ASP} \iff \textit{dl-SP} \text{ and } \textit{ul-SP}$$

$$\Updownarrow$$

$$\textit{sd-LWIR} \iff \textit{sd-ALWIR} \iff \textit{ul-ALWIR} \text{ and } \textit{dl-ALWIR} \iff \textit{ul-LWIR} \text{ and } \textit{dl-LWIR}$$

$$\text{(b)}\quad \textit{dl-SP} \iff \textit{dl-ASP}, \quad \textit{dl-SP} \Updownarrow \textit{ul-LWIR}, \quad \textit{ul-LWIR} \iff \textit{ul-ALWIR}$$

$$\text{(c)}\quad \textit{ul-SP} \iff \textit{ul-ASP}, \quad \textit{ul-SP} \Updownarrow \textit{dl-LWIR}, \quad \textit{dl-LWIR} \iff \textit{dl-ALWIR}$$

Figure 2: **Equivalence of various strategic requirements.** In the figure, "SP" stands for "strategy-proofness", "ASP" for "adjacent strategy-proofness", "LWIR" for "lowest welfare for inverse reports", and "ALWIR" for "adjacent lowest welfare for inverse reports" (defined in Appendix B). Panels (a)-(c) summarize some of the results in Theorems 4.2 and 4.4, and Propositions 4.8, B.1, and B.2.

Proposition 4.8. *Let* e *be an extension. Then* e-strategy-proofness *is equivalent to* e^d-lowest welfare for inverse reports.

While *e-lowest welfare for inverse reports* is primarily of technical interest, noting Propositions 4.2 and 4.8, one may ask whether an "adjacency version" of the axiom suffices for it. In fact, this is true, and a detailed discussion is in Appendix B.

Proposition 4.8 can be applied to decompose *sd-strategy-proofness* into two strategic requirements based on a general extension, namely, *e-strategy-proofness* and *e-lowest welfare for inverse reports.* The proof of the following proposition is similar to that of Proposition 4.5 and is omitted.

Proposition 4.9. *Let* e *be a* monotone *and* transitive *extension.*

(1) sd-strategy-proofness *implies* e-strategy-proofness *and* e-lowest welfare for inverse reports.

(2) If, in addition, e *satisfies* condition C*, then the converse of part (1) also holds.*

Remark 4.5. By Proposition 4.8, the above proposition still holds if we replace *e-lowest welfare for inverse reports* by e^d*-strategy-proofness.* $\triangle$

4.4.4 Assignment Rules

Now we assess the family of generalized serial rules and the random priority rule on the grounds of the axioms in Section 4.2. The first criterion is efficiency. While the random priority rule is not *sd-efficient*, for each $\sigma \in \Sigma$, the generalized serial rule associated with σ is *sd-efficient.* Further, applying the results on various efficiency concepts (Proposition 4.1 and Corollaries 4.1-4.3), we can state the (in)efficiency of these rules in more general terms.

Theorem 4.5. *(1) For each $\sigma \in \Sigma$ and each extension e satisfying* monotonicity*,* anti-symmetry*, and either* probability compensation for preferred objects *or* probability compensation for less preferred objects*, the generalized serial rule associated with σ is e-*efficient.

(2) For each $\sigma \in \Sigma$ and each extension e satisfying monotonicity*,* transitivity*, and either $e \subseteq dl$ or $e \subseteq ul$, the generalized serial rule associated with σ is* e-efficient.

(3) For each extension e satisfying monotonicity*, the random priority rule is not* e-efficient*.*

Proof. To prove part (1), let $\sigma \in \Sigma$ and let e be an extension satisfying the properties listed in part (1). Let $R \in \mathcal{R}(A)^N$ and $\pi \equiv S^\sigma(R)$. It suffices to show that $\tau(R, \pi)$ is acyclic, and this can be proved by an argument similar to that in Bogomolnaia and Moulin (2001, Theorem 1).

Part (2) follows from part (1) and Corollary 4.3. Finally, as for part (3), recall that the random priority rule, RP, is not *sd-efficient* (Bogomolnaia and Moulin, 2001). There is $R \in \mathcal{R}(A)^N$ such that $RP(R)$ is not *sd-efficient* for R. Let $\pi \equiv RP(R)$. By part (2) of Theorem 4.1, $\tau(R, \pi)$ is cyclic. Then by part (1) of Theorem 4.1, π is not *e-efficient* for R. Thus, RP is not *e-efficient.* □

Remark 4.6. While our model is a fixed population framework, we can state the inefficiency of the random priority rule in stronger terms by allowing n, the common number of agents and objects, to approach infinity. Combined with part (3) of Theorem 4.5, Manea (2009, Theorem 1) yields the following: for each extension e satisfying *monotonicity*, the fraction of economies for which the random priority rule selects an *e-efficient* assignment converges to zero as $n \to \infty$. △

Next is no-envy. The serial rule satisfies *sd-no-envy* (Bogomolnaia and Moulin, 2001). Since for each *monotone* and *transitive* extension e, *sd-no-envy* implies *e-no-envy*, the rule satisfies the strongest no-envy concept. On the other hand, the generalized serial rules violate *sd-no-envy* in general (see Remark 4.8 below), and the same is true for the random priority rule. Among the no-envy concepts in Corollary 4.4, the strongest that they satisfy is *dl-no-envy*.[17]

Theorem 4.6. *(1) For each* monotone *and* transitive *extension e, the serial rule satisfies* e-no-envy*.*

(2) For each $\sigma \in \Sigma$, the generalized serial rule associated with σ satisfies dl-no-envy*.*

(3) The random priority rule satisfies dl-no-envy*.*

Proof. Part (1) follows from Propositions 3.2 and 4.2 and the fact that the serial rule is *sd-no-envy*.

To show part (2), let $\sigma \in \Sigma$, $R \in \mathcal{R}(A)^N$, and $i, j \in N$ with $i \neq j$. Assume, without loss of generality, that $1 \, P_i \, 2 \, P_i \, \cdots \, P_i \, n$. Let $\pi \equiv S^\sigma(R)$. To show that $\pi_i \, R_i^{dl} \, \pi_j$, consider the generalized

[17] It is easy to construct an example with three objects in which the random priority rule violates *sd-no-envy*. Since (i) by definition, *sd-no-envy* is equivalent to the combination of *dl-* and *ul-no-envy* in the three-object case and (ii) the rule satisfies *dl-no-envy*, it follows that the violation of *sd-no-envy* is, in fact, due to the violation of *ul-no-envy*.

simultaneous consumption algorithm associated with σ, applied to R. Let s_1 be the step in which object 1 is exhausted; i.e., s_1 is such that $1 \in A^{s_1-1} \setminus A^{s_1}$.

Now we show that $\pi_{i1} \geq \pi_{j1}$. First, for each $s \leq s_1 - 1$, $i \in N^*(R, A^s, 1)$. Also, there is $\hat{t} \in [0, t^{s_1}]$ such that agent j consumes object 1 during the interval $[\hat{t}, t^{s_1})$. Thus,

$$\pi_{i1}^{s_1} = \int_0^{t^{s_1}} \sigma(1,t)\, dt \geq \int_{\hat{t}}^{t^{s_1}} \sigma(1,t)\, dt = \pi_{j1}^{s_1}. \tag{3}$$

Moreover, because object 1 is exhausted in Step s_1, $\pi_{i1} = \pi_{i1}^{s_1}$ and $\pi_{j1} = \pi_{j1}^{s_1}$, so that $\pi_{i1} \geq \pi_{j1}$.

If $\pi_{i1} > \pi_{j1}$, then $\pi_i \, P_i^{dl} \, \pi_j$. Assume, henceforth, that $\pi_{i1} = \pi_{j1}$. Since for each $t \in \mathbb{R}_+$, $\sigma(1,t) > 0$, Inequality (3) implies that in fact, $\hat{t} = 0$, so that for each $k \in A$, $1 \, R_j \, k$. Now let s_2 be the step in which each object in $\{1,2\}$ is exhausted; i.e., s_2 is such that $A^{s_2-1} \cap \{1,2\} \neq \emptyset$ and $A^{s_2} \cap \{1,2\} = \emptyset$.

To show that $\pi_{i2} \geq \pi_{j2}$, note that $s_1 \leq s_2$. If $s_1 = s_2$, then $\pi_{i2} = \pi_{j2} = 0$. If $s_1 < s_2$, then for each $s \in \{s_1, s_1+1, \cdots, s_2-1\}$, $i \in N^*(R, A^s, 2)$. Also, there is $t' \in [t^{s_1}, t^{s_2}]$ such that agent j consumes object 2 during the interval $[t', t^{s_2})$. Thus,

$$\pi_{i2}^{s_2} = \int_{t^{s_1}}^{t^{s_2}} \sigma(2,t)\, dt \geq \int_{t'}^{t^{s_2}} \sigma(2,t)\, dt = \pi_{j2}^{s_2}. \tag{4}$$

Moreover, because object 2 is exhausted in Step s_2, $\pi_{i2} = \pi_{i2}^{s_2}$ and $\pi_{j2} = \pi_{j2}^{s_2}$, so that $\pi_{i2} \geq \pi_{j2}$.

If $\pi_{i2} > \pi_{j2}$, then $\pi_i \, P_i^{dl} \, \pi_j$. Otherwise, we can repeat the above argument to eventually obtain that $\pi_i \, R_i^{dl} \, \pi_j$.

As for part (3), Bogomolnaia and Moulin (2001, Proposition 1) show that the random priority rule satisfies *sd-weak no-envy*. In fact, their proof can be adapted, with only minor changes, to prove a stronger statement that the random priority rule satisfies *dl-no-envy*. We omit the details. $\square$

Remark 4.7. In conjunction with Proposition 4.4, we can see how the generalized serial and random priority rules behave in light of *e-complete envy for inverse reports*: (1) for each *monotone* and *transitive* extension e, the serial rule satisfies *e-complete envy for inverse reports*; (2) for each $\sigma \in \Sigma$, the generalized serial rule associated with σ satisfies *ul-complete envy for inverse reports*; and (3) the random priority rule satisfies *ul-complete envy for inverse reports.* $\triangle$

Remark 4.8. If the speed function σ varies greatly across objects, then the generalized serial rule associated with σ may not satisfy *sd-no-envy*. To see this, let $N \equiv \{1,2,3\}$ and $A \equiv \{1,2,3\}$. Let $\varepsilon \in (0, \frac{1}{2})$ and let $\sigma \in \Sigma$ be such that for each $t \in \mathbb{R}_+$, $\sigma(1,t) = \varepsilon$ and $\sigma(2,t) = \sigma(3,t) = 1$. To show that S^σ is not *sd-no-envy*, consider $R \in \mathcal{R}(A)^N$ such that (i) $1 \, R_1 \, 2 \, R_1 \, 3$; and (ii) for each

R_1	R_2	R_3	R_4	R_5		R'_3
1	1	2	2	2		2
4	5	4	5	5		5
·	·	5	3	3		4
·	·	3	4	4		3
·	·	1	1	1		1

Figure 3: **A generalized serial rule may not be *dl-strategy-proof* (Example 4.1).** Let $N \equiv \{1, \cdots, 5\}$ and $A \equiv \{1, \cdots, 5\}$. Let $\sigma \in \Sigma$ be such that (i) for each $k \in \{1, 3, 4, 5\}$ and each $t \in \mathbb{R}_+$, $\sigma(k, t) = 1$; and (ii) for each $t \in \mathbb{R}_+$, $\sigma(2, t) = \frac{1}{3}$. Let $R \in \mathcal{R}(A)^N$ and $R'_3 \in \mathcal{R}(A)$ be as specified above. Let $\pi \equiv S^\sigma(R)$ and $\pi' \equiv S^\sigma(R'_3, R_{-3})$. Since $\pi'_3 \, P_3^{dl} \, \pi_3$, S^σ is not *dl-strategy-proof*.

$i \in \{2, 3\}$, $2 \, R_i \, 1 \, R_1 \, 3$. Let $\pi \equiv S^\sigma(R)$. Then $\pi_1 = \left(\frac{1+2\varepsilon}{3}, 0, \frac{2-2\varepsilon}{3}\right)$ and $\pi_2 = \pi_3 = \left(\frac{1-\varepsilon}{3}, \frac{1}{2}, \frac{1+2\varepsilon}{6}\right)$. Since $\pi_{11} + \pi_{12} < \pi_{21} + \pi_{22}$, it is not the case that $\pi_1 \, R_1^{sd} \, \pi_2$, in violation of *sd-no-envy*. △

In terms of strategy-proofness, the random priority rule outperforms each generalized serial rule. The former is *sd-strategy-proof* (Bogomolnaia and Moulin, 2001). Combined with Propositions 3.2 and 4.6, this means that for each *monotone* and *transitive* extension e, the random priority rule is *e-strategy-proof*. By constrast, the serial rule is only *dl-strategy-proof* and in general, the generalized serial rules are not *dl-strategy-proof* (see Example 4.1 below).[18,19]

Theorem 4.7. *(1) The serial rule is* dl-strategy-proof.

(2) For each monotone *and* transitive *extension e, the random priority rule is* e-strategy-proof.

Remark 4.9. Applying Proposition 4.8, we obtain a result on *e-lowest welfare for inverse reports*: (1) the serial rule satisfies *ul-lowest welfare for inverse reports*; and (2) for each *monotone* and *transitive* extension e, the random priority rule satisfies *e-lowest welfare for inverse reports.* △

While the proof of part (1) in Theorem 4.7 is relegated to Appendix C, we convey the main intuition informally. By virtue of Theorem 4.2, it is enough to verify that the serial rule is *dl-adjacent strategy-proof.* Suppose that agent i with true preference relation R_i, say, reports a preference relation R'_i adjacent to R_i while all others announce R_{-i}. Let $k, \ell \in A$ be such that $k \, P_i \, \ell$ and $\ell \, P'_i \, k$. Now consider the simultaneous consumption algorithm that defines the serial rule, applied to R and (R'_i, R_{-i}). When agent i changes his announcement from R_i to R'_i, the probability that he receives object k cannot go up. And if that probability is unaffected, so is the whole lottery he receives. Thus, the serial rule is *dl-adjacent strategy-proof.*

Example 4.1. *A generalized serial rule may not be* dl-strategy-proof. Let $N \equiv \{1, \cdots, 5\}$ and $A \equiv \{1, \cdots, 5\}$. Let $\sigma \in \Sigma$ be such that (i) for each $k \in \{1, 3, 4, 5\}$ and each $t \in \mathbb{R}_+$, $\sigma(k, t) = 1$;

[18] By Theorem 4.4, *sd-strategy-proofness* is equivalent to the combination of *dl-strategy-proofness* and *ul-strategy-proofness.* Thus, similarly to Footnote 17, the serial rule is not *sd-strategy-proof* because it is not *ul-strategy-proof.*

[19] I thank Jay Sethuraman for providing an example where a generalized serial rule is not *dl-strategy-proof.*

and (ii) for each $t \in \mathbb{R}_+$, $\sigma(2,t) = \frac{1}{3}$. Let $R \in \mathcal{R}(A)^N$ be the economy specified in Figure 3 (the unspecified part of R_1 and R_2 can be completed in an arbitrary way). Let $\pi \equiv S^\sigma(R)$. It is easy to compute that $\pi_3 = (0, \frac{1}{3}, \frac{1}{6}, \frac{1}{2}, 0)$. Now consider agent 3. Let $R'_3 \in \mathcal{R}(A)$ be the preference relation specified in Figure 3. Let $\pi' \equiv S^\sigma(R'_3, R_{-3})$. Simple calculation shows that $\pi'_3 = (0, \frac{1}{3}, 0, \frac{1}{2}, \frac{1}{6})$, so that $\pi'_3 \, P^{dl}_3 \, \pi_3$. Thus, S^σ is not *dl-strategy-proof*. $\triangle$

5 Concluding Remarks

The growing literature on probabilistic assignment focuses on the ordinal approach and adopts the *sd*-extension. While the use of the *sd*-extension is well justified, having it as an unchallenged assumption limits the way we analyze assignment problems. If there is a single extension available, its effect on analysis cannot be well distinguished from the effects that other assumptions have, and the underlying intuition is obscured.

These limitations are what we hoped to highlight and overcome in this paper. Our two-fold axiomatic approach departs from the standard practice and explores the effects of extensions, using axioms. The greatest advantage to this approach is that it allows us to reexamine existing results from varying perspectives and to uncover and address issues that have hitherto been overlooked.

A Appendix: Sufficiency of Adjacent Strategy-proofness

This appendix provides conditions on the preference domain that guarantee the equivalence of *e-adjacent strategy-proofness* and *e-strategy-proofness* (Theorem 4.2). To that end, we define basic concepts and notations, and modify several existing ones. A **preference domain** is a subset of $\mathcal{R}(A)$ from which agents' preferences are drawn. Let $\boldsymbol{\mathcal{D}} \subseteq \mathcal{R}(A)$ be a preference domain. An (assignment) **rule** is a mapping from $\mathcal{D}^N$ to Π. Denote by φ a generic rule. Now redefine *e-strategy-proofness* by taking into account the fact that agents can only announce preference relations in $\mathcal{D}$.

***e*-Strategy-proofness:** For each $R \in \mathcal{D}^N$, each $i \in N$, and each $R'_i \in \mathcal{D}$, $\varphi_i(R) \, R^e_i \, \varphi_i(R'_i, R_{-i})$.

Redefine *e-adjacent strategy-proofness* similarly. To avoid ambiguity, whenever a rule is, for instance, *e-strategy-proof*, we say that it is *e-strategy-proof* on $\mathcal{D}^N$, indicating the preference domain under consideration.

Next are conditions on the preference domain. For each triple $R_0, R'_0, \hat{R}_0 \in \mathcal{R}(A)$, $\boldsymbol{\hat{R}_0}$ **is between** $\boldsymbol{R_0}$ **and** $\boldsymbol{R'_0}$ if for each pair $k, \ell \in A$ such that $k \, R_0 \, \ell$ and $k \, R'_0 \, \ell$, $k \, \hat{R}_0 \, \ell$.[20] Then the

[20] This notion of betweenness, due to Kemeny (1959), can alternatively be stated using direct path, a concept introduced in Section 4.4.3: for each triple $R_0, R'_0, \hat{R}_0 \in \mathcal{R}(A)$, $\hat{R}_0$ is between R_0 and R'_0 if there is a direct path

preference domain $\mathcal{D}$ is **convex** if for each pair $R_0, R'_0 \in \mathcal{D}$ and each $\hat{R}_0 \in \mathcal{R}(A)$ such that $\hat{R}_0$ is between R_0 and R'_0, $\hat{R}_0 \in \mathcal{D}$. If $\mathcal{D}$ is convex, then for each $e \in \{sd, dl, ul\}$, a rule is *e-adjacent strategy-proof* on $\mathcal{D}^N$ iff it is *e-strategy-proof* on $\mathcal{D}^N$. However, the convexity condition can be weakened substantially.

To define weaker conditions, further notation is needed. Let $R_0, R'_0 \in \mathcal{R}(A)$. Let $\boldsymbol{k^*(R_0|R'_0)} \equiv (\min\{1 \leq k \leq n : k(R_0) \neq k(R'_0)\})\,(R_0)$ be the object that R_0 ranks highest among those whose rankings differ according to R_0 and R'_0. Also, let $\boldsymbol{k_*(R_0|R'_0)} \equiv (\max\{1 \leq k \leq n : k(R_0) \neq k(R'_0)\})\,(R_0)$ be the object that R_0 ranks lowest among those whose rankings differ according to R_0 and R'_0. To illustrate the definitions, let $A \equiv \{1, 2, 3, 4\}$. Throughout examples in this appendix, write $k_1k_2k_3k_4$ for the preference relation in $\mathcal{R}(A)$ such that object k_1 is most preferred, object k_2 is second most preferred , and so on. Then $k^*(1234|1432) = 2$, $k^*(1432|1234) = 4$, $k_*(1234|1432) = 4$, and $k_*(1432|1234) = 2$.

Let $R_0, R'_0 \in \mathcal{R}(A)$, and let $h \equiv d(R_0, R'_0)$. The **preferred-objects-first (POF) direct path from $\boldsymbol{R_0}$ to $\boldsymbol{R'_0}$** is the direct path $\{R_0^0, R_0^1, \cdots, R_0^h\}$ from R_0 to R'_0 such that for each $\tilde{h} \in \{0, \cdots, h-1\}$, the ranking of object $k^*(R'_0|R_0^{\tilde{h}})$ is higher according to $R_0^{\tilde{h}+1}$ than according to $R_0^{\tilde{h}}$. The **less-preferred-objects-first (LOF) direct path from $\boldsymbol{R_0}$ to $\boldsymbol{R'_0}$** is the direct path $\{R_0^0, R_0^1, \cdots, R_0^h\}$ from R_0 to R'_0 such that for each $\tilde{h} \in \{0, \cdots, h-1\}$, the ranking of object $k_*(R'_0|R_0^{\tilde{h}})$ is lower according to $R_0^{\tilde{h}+1}$ than according to $R_0^{\tilde{h}}$. As we go along the POF direct path from R_0 to R'_0, the preference relations become closer to R'_0 in such a way that first, object $1(R'_0)$ moves to the first place in a sequence of adjacent-pair-switch transformations, and then object $2(R'_0)$ moves to the second place in a sequence of adjacent-pair-switch transformations, and so on. On the other hand, as we go along the LOF direct path from R_0 to R'_0, the preference relations become closer to R'_0 in such a way that first, object $n(R'_0)$ moves to the nth place in a sequence of adjacent-pair-switch transformations, and then object $(n-1)(R'_0)$ moves to the $n-1$st place in a sequence of adjacent-pair-switch transformations, and so on. For example, if $A \equiv \{1, 2, 3, 4\}$, then the POF direct path from preference relation 1234 to preference relation 2143 is the sequence $\{1234, 2134, 2143\}$; the LOF direct path from 1234 to 2143 is the sequence $\{1234, 1243, 2143\}$.[21]

The preference domain $\mathcal{D}$ satisfies the **POF direct path property** if for each pair $R_0, R'_0 \in \mathcal{D}$, the POF direct path from R_0 to R'_0 is in $\mathcal{D}$ (i.e., each preference relation in the path is in $\mathcal{D}$). Similarly, $\mathcal{D}$ satisfies the **LOF direct path property** if for each pair $R_0, R'_0 \in \mathcal{D}$, the LOF direct path from R_0 to R'_0 is in $\mathcal{D}$. Also, $\mathcal{D}$ satisfies the **PLOF direct path property** if for each pair $R_0, R'_0 \in \mathcal{D}$, either the POF or LOF direct path from R_0 to R'_0 is in $\mathcal{D}$.

from R_0 to R'_0 containing $\hat{R}_0$.

[21] In the example, reversing the sequence of preference relations in the POF direct path from 1234 to 2143 gives the LOF direct path from 2143 to 1234; similarly, reversing the sequence of preference relations in the LOF direct path from 1234 to 2143 gives the POF direct path from 2143 to 1234. However, this is not the case in general.

As can be seen easily, among the various requirements on the preference domain defined so far, convexity is the strongest and the PLOF direct path property is the weakest: convexity implies the POF and LOF direct path properties, each of which, in turn, implies the PLOF direct path property. However, the converse of none of these statements is true. To see this, let $A \equiv \{1, 2, 3, 4\}$, and consider the following preference domains: $\mathcal{D}_1 \equiv \{1234, 1324, 1342, 1432, 3142\}$; $\mathcal{D}_2 \equiv \{1234, 1243, 1423, 1432\}$; $\mathcal{D}_3 \equiv \{1234, 1324, 1342, 1432\}$; and $\mathcal{D}_4 \equiv \mathcal{D}_2 \cup \mathcal{D}_3$. Then $\mathcal{D}_1$ satisfies the PLOF direct path property, but neither the POF nor LOF direct path property; $\mathcal{D}_2$ satisfies the POF direct path property but it is not convex; $\mathcal{D}_3$ satisfies the LOF direct path property but it is not convex; and $\mathcal{D}_4$ is convex, thus satisfying the POF and LOF direct path properties.

Now we state the equivalence of *e-adjacent strategy-proofness* and *e-strategy-proofness*. As the following proposition shows, depending on the extension chosen, different properties on the preference domain suffice.

Proposition A.1. *Let $\mathcal{D}$ be a preference domain.*

*(1) Assume that $\mathcal{D}$ satisfies the PLOF direct path property. A rule is sd-*adjacent strategy-proof *on $\mathcal{D}^N$ iff it is sd-*strategy-proof *on $\mathcal{D}^N$.*

*(2) Assume that $\mathcal{D}$ satisfies the POF direct path property. A rule is dl-*adjacent strategy-proof *on $\mathcal{D}^N$ iff it is dl-*strategy-proof *on $\mathcal{D}^N$.*

*(3) Assume that $\mathcal{D}$ satisfies the LOF direct path property. A rule is ul-*adjacent strategy-proof *on $\mathcal{D}^N$ iff it is ul-*strategy-proof *on $\mathcal{D}^N$.*

To prove Proposition A.1, we introduce an auxiliary axiom called *e-within-m strategy-proofness*, where $m \in \mathbb{N}$. It requires that no agent benefit from reporting a preference relation lying within distance m (according to metric $d(\cdot, \cdot)$) from his true preference relation.

***e*-Within-*m* Strategy-proofness:** For each $R \in \mathcal{D}^N$, each $i \in N$, and each $R'_i \in \mathcal{D}$ such that $d(R_i, R'_i) \leq m$, $\varphi_i(R)\, R_i^e\, \varphi_i(R'_i, R_{-i})$.

Clearly, *e-adjacent strategy-proofness* and *e-strategy-proofness* are special cases of *e-within-m strategy-proofness*, and the proof exploits this observation. We show that for each $e \in \{sd, dl, ul\}$ and each $m \in \mathbb{N}$, *e-within-m strategy-proofness* implies *e-within-$(m+1)$ strategy-proofness* on the respective preference domains. Then it follows that *e-adjacent strategy-proofness* implies *e-strategy-proofness*. We first give the arguments for the *dl*- and *ul*-extensions because the argument for the *sd*-extension is a combination of the two.

Lemma A.1. *Let $\mathcal{D}$ be a preference domain satisfying the POF direct path property. For each $m \in \mathbb{N}$, if a rule is dl-*within-m strategy-proof *on $\mathcal{D}^N$, then it is dl-*within-$(m+1)$ strategy-proof *on $\mathcal{D}^N$.*

Proof. Let $\mathcal{D}$ be as in the lemma. Let φ be a *dl-within-m strategy-proof* rule defined on $\mathcal{D}^N$. Let $R \in \mathcal{D}^N$ and $i \in N$. Let $R'_i \in \mathcal{D}$ be such that $d(R_i, R'_i) \leq m + 1$. Let $\pi_i \equiv \varphi_i(R)$ and $\pi'_i \equiv \varphi_i(R_i, R_{-i})$. If $d(R_i, R'_i) \leq m$, then by *dl-within-m strategy-proofness*, $\pi_i \, R^{dl}_i \, \pi'_i$. Thus, assume, henceforth, that $d(R_i, R'_i) = m + 1$. Suppose, without loss of generality, that R_i is such that $1 \, P_i \, 2 \, P_i \, \cdots \, P_i \, n$ and that R'_i is such that $k_1 \, P'_i \, k_2 \, P'_i \, \cdots \, P'_i \, k_n$. We distinguish two cases.

Case 1: $k_1 \neq 1$.

Since $\mathcal{D}$ satisfies the POF direct path property, there is $\hat{R}_i \in \mathcal{D}$ such that (i) $d(R_i, \hat{R}_i) = 1$; and (ii) the ranking of object k_1 is higher according to $\hat{R}_i$ than according to R_i. Let $\hat{\pi}_i \equiv \varphi_i(\hat{R}_i, R_{-i})$. Since $m \geq 1$, by *dl-within-m strategy-proofness*, $\pi_i \, R^{dl}_i \, \hat{\pi}_i$ and $\hat{\pi}_i \, \hat{R}^{dl}_i \, \pi_i$, so that

$$\begin{aligned} &\text{either} \quad \text{(i) } \pi_i = \hat{\pi}_i \\ &\text{or} \quad \text{(ii) for each } \ell \in \{1, \cdots, k_1 - 2\}, \; \pi_{i\ell} = \hat{\pi}_{i\ell}; \\ &\qquad \pi_{i,k_1-1} > \hat{\pi}_{i,k_1-1}; \text{ and } \pi_{ik_1} < \hat{\pi}_{ik_1}. \end{aligned} \tag{5}$$

Further, since $d(\hat{R}_i, R'_i) = m$, again by *dl-within-m strategy-proofness*, $\hat{\pi}_i \, \hat{R}^{dl}_i \, \pi'_i$ and $\pi'_i \, (R'_i)^{dl} \, \hat{\pi}_i$.

If $\hat{\pi}_i = \pi'_i$, then by $\pi_i \, R^{dl}_i \, \hat{\pi}_i$, $\pi_i \, R^{dl}_i \, \pi'_i$. Thus, assume, henceforth, that $\hat{\pi}_i \neq \pi'_i$, so that $\hat{\pi}_i \, \hat{P}^{dl}_i \, \pi'_i$. Now there are four subcases.

Case 1.1: *There is $\ell^* \in \{1, \cdots, k_1 - 2\}$ such that (i) for each $\ell \in \{1, \cdots, \ell^* - 1\}$, $\hat{\pi}_{i\ell} = \pi'_{i\ell}$; and (ii) $\hat{\pi}_{i\ell^*} > \pi'_{i\ell^*}$.*

Case 1.2: *For each $\ell \in \{1, \cdots, k_1 - 2\}$, $\hat{\pi}_{i\ell} = \pi'_{i\ell}$ and $\hat{\pi}_{ik_1} > \pi'_{ik_1}$.*

Case 1.3: *For each $\ell \in \{1, \cdots, k_1 - 2, k_1\}$, $\hat{\pi}_{i\ell} = \pi'_{i\ell}$ and $\hat{\pi}_{i,k_1-1} > \pi'_{i,k_1-1}$.*

Case 1.4: *There is $\ell^* \in \{k_1 + 1, \cdots, n\}$ such that (i) for each $\ell \in \{1, \cdots, \ell^* - 1\}$, $\hat{\pi}_{i\ell} = \pi'_{i\ell}$; and (ii) $\hat{\pi}_{i\ell^*} > \pi'_{i\ell^*}$.*

In Cases 1.1, 1.3, and 1.4, Statement (5) implies that $\pi_i \, P^{dl}_i \, \pi'_i$. In Case 1.2, it follows that $\hat{\pi}_i \, (P'_i)^{dl} \, \pi'_i$, in contradiction to $\pi'_i \, (R'_i)^{dl} \, \hat{\pi}_i$. (Note that if $k_1 = 2$, then Case 1.1 does not apply.) In sum, $\pi_i \, R^{dl}_i \, \pi'_i$, as desired.

Case 2: *There is $h \in \{2, \cdots, n\}$ such that (i) for each $\tilde{h} \in \{1, \cdots, h - 1\}$, $k_{\tilde{h}} = \tilde{h}$; and (ii) $k_h \neq h$.*

The argument is essentially the same as that in Case 1; only minor changes in notation are needed. The proof is omitted. □

Next, we prove a similar lemma for the *ul*-extension.

Lemma A.2. *Let $\mathcal{D}$ be a preference domain satisfying the LOF direct path property. For each $m \in \mathbb{N}$, if a rule is ul-*within-m strategy-proof *on $\mathcal{D}^N$, then it is ul-*within-$(m+1)$ strategy-proof *on $\mathcal{D}^N$.*

Proof. Let $\mathcal{D}$ be as in the lemma. Let φ be a *ul-within-m strategy-proof* rule defined on $\mathcal{D}^N$. Let $R \in \mathcal{D}^N$ and $i \in N$. Let $R'_i \in \mathcal{D}$ be such that $d(R_i, R'_i) \leq m+1$. Let $\pi_i \equiv \varphi_i(R)$ and $\pi'_i \equiv \varphi_i(R_i, R_{-i})$. If $d(R_i, R'_i) \leq m$, then by *ul-within-m strategy-proofness*, $\pi_i \, R_i^{ul} \, \pi'_i$. Thus, assume, henceforth, that $d(R_i, R'_i) = m+1$. Suppose, without loss of generality, that R_i is such that $1 \, P_i \, 2 \, P_i \, \cdots \, P_i \, n$ and that R'_i is such that $k_1 \, P'_i \, k_2 \, P'_i \, \cdots \, P'_i \, k_n$. We distinguish two cases.

Case 1: $k_n \neq n$.

Since $\mathcal{D}$ satisfies the LOF direct path property, there is $\hat{R}_i \in \mathcal{D}$ be such that (i) $d(R_i, \hat{R}_i) = 1$; and (ii) the ranking of object k_n is lower according to $\hat{R}_i$ than according to R_i. Let $\hat{\pi}_i \equiv \varphi_i(\hat{R}_i, R_{-i})$. Since $m \geq 1$, by *ul-within-m strategy-proofness*, $\pi_i \, R_i^{ul} \, \hat{\pi}_i$ and $\hat{\pi}_i \, \hat{R}_i^{ul} \, \pi_i$, so that

$$\begin{array}{lll} \text{either} & \text{(i) } \pi_i = \hat{\pi}_i & \quad (6) \\ \text{or} & \text{(ii) for each } \ell \in \{n, n-1, \cdots, k_n+2\},\ \pi_{i\ell} = \hat{\pi}_{i\ell}; & \\ & \quad \pi_{i,k_n+1} < \hat{\pi}_{i,k_n+1};\ \text{and } \pi_{ik_n} > \hat{\pi}_{ik_n}. & \end{array}$$

Further, since $d(\hat{R}_i, R'_i) = m$, again by *ul-within-m strategy-proofness*, $\hat{\pi}_i \, \hat{R}_i^{ul} \, \pi'_i$ and $\pi'_i \, (R'_i)^{ul} \, \hat{\pi}_i$.

If $\hat{\pi}_i = \pi'_i$, then by $\pi_i \, R_i^{ul} \, \hat{\pi}_i$, $\pi_i \, R_i^{ul} \, \pi'_i$. Thus, assume, henceforth, that $\hat{\pi}_i \neq \pi'_i$, so that $\hat{\pi}_i \, \hat{P}_i^{ul} \, \pi'_i$. Now there are four subcases.

Case 1.1: *There is* $\ell^* \in \{n, n-1, \cdots, k_n+2\}$ *such that (i) for each* $\ell \in \{n, n-1, \cdots, \ell^*+1\}$, $\hat{\pi}_{i\ell} = \pi'_{i\ell}$*; and (ii)* $\hat{\pi}_{i\ell^*} < \pi'_{i\ell^*}$.

Case 1.2: *For each* $\ell \in \{n, n-1, \cdots, k_n+2\}$, $\hat{\pi}_{i\ell} = \pi'_{i\ell}$ *and* $\hat{\pi}_{ik_n} < \pi'_{ik_n}$.

Case 1.3: *For each* $\ell \in \{n, n-1, \cdots, k_n+2, k_n\}$, $\hat{\pi}_{i\ell} = \pi'_{i\ell}$ *and* $\hat{\pi}_{i,k_n+1} < \pi'_{i,k_n+1}$.

Case 1.4: *There is* $\ell^* \in \{k_n-1, k_n-2, \cdots, 1\}$ *such that (i) for each* $\ell \in \{n, n-1, \cdots, \ell^*+1\}$, $\hat{\pi}_{i\ell} = \pi'_{i\ell}$*; and (ii)* $\hat{\pi}_{i\ell^*} < \pi'_{i\ell^*}$.

In Cases 1.1, 1.3, and 1.4, Statement (6) implies that $\pi_i \, P_i^{ul} \, \pi'_i$. In Case 1.2, it follows that $\hat{\pi}_i \, (P'_i)^{ul} \, \pi'_i$, in contradiction to $\pi'_i \, (R'_i)^{ul} \, \hat{\pi}_i$. (Note that if $k_n = n-1$, then Case 1.1 does not apply.) In sum, $\pi_i \, R_i^{ul} \, \pi'_i$, as desired.

Case 2: *There is* $h \in \{n, \cdots, 2\}$ *such that (i) for each* $\tilde{h} \in \{n, n-1, \cdots, h+1\}$, $k_{\tilde{h}} = \tilde{h}$*; and (ii)* $k_h \neq h$.

The argument is essentially the same as that in Case 1; only minor changes in notation are needed. The proof is omitted. □

Finally, we prove a lemma for the *sd*-extension.

Lemma A.3. *Let* $\mathcal{D}$ *be a preference domain satisfying the PLOF direct path property. For each* $m \in \mathbb{N}$*, if a rule is sd-*within-m strategy-proof *on* $\mathcal{D}^N$*, then it is sd-*within-$(m+1)$ strategy-proof *on* $\mathcal{D}^N$.

Proof. Let $\mathcal{D}$ be as in the lemma. Let φ be an *sd-within-m strategy-proof* rule defined on $\mathcal{D}^N$. Let $R \in \mathcal{D}^N$ and $i \in N$. Let $R'_i \in \mathcal{D}$ be such that $d(R_i, R'_i) \leq m + 1$. Let $\pi_i \equiv \varphi_i(R)$ and $\pi'_i \equiv \varphi_i(R_i, R_{-i})$. If $d(R_i, R'_i) \leq m$, then by *sd-within-m strategy-proofness*, $\pi_i \, R_i^{sd} \, \pi'_i$. Thus, assume, henceforth, that $d(R_i, R'_i) = m + 1$. Suppose, without loss of generality, that R_i is such that $1 \, P_i \, 2 \, P_i \, \cdots \, P_i \, n$ and that R'_i is such that $k_1 \, P'_i \, k_2 \, P'_i \, \cdots \, P'_i \, k_n$. Since $\mathcal{D}$ satisfies the PLOF direct path property, there are two cases.

Case 1: *There is* $\hat{R}_i \in \mathcal{D}$ *be such that (i)* $d(R_i, \hat{R}_i) = 1$*; and (ii) the ranking of object* $k^*(R'_i|R_i)$ *is higher according to* $\hat{R}_i$ *than according to* R_i.

Adapt the proof of Lemma A.1.

Case 2: *There is* $\hat{R}_i \in \mathcal{D}$ *be such that (i)* $d(R_i, \hat{R}_i) = 1$*; and (ii) the ranking of object* $k_*(R'_i|R_i)$ *is lower according to* $\hat{R}_i$ *than according to* R_i.

Adapt the proof of Lemma A.2.

In each of the above cases, the argument becomes more involved, but the idea of the proof is the same. We omit the details. □

B Appendix: Adjacent Lowest Welfare for Inverse Reports

In this appendix, we introduce an axiom that weakens *e-lowest welfare for inverse reports* in much the same way as *e-adjacent strategy proofness* weakens *e-strategy proofness*, and show that for each $e \in \{sd, dl, ul\}$, it is sufficient for *e-lowest welfare for inverse reports*.

Fix the announcement of all but agent i, say. Assume that agent i's true preference relation is R_i, and let $R'_i \in \mathcal{R}(A)$ be adjacent to R_i^{-1}. The following axiom requires that the lottery that agent i receives by reporting R'_i be at least as desirable, according to R_i^e, as the lottery that he receives by reporting R_i^{-1}.

***e*-Adjacent Lowest Welfare for Inverse Reports:** For each $R \in \mathcal{R}(A)^N$, each $i \in N$, and each $R'_i \in \mathcal{R}(A)$ such that R'_i is adjacent to R_i^{-1}, $\varphi_i(R'_i, R_{-i}) \, R_i^e \, \varphi_i(R_i^{-1}, R_{-i})$.

The first observation is that a connection similar to that between *e-strategy proofness* and *e-lowest welfare for inverse reports* holds between *e-adjacent strategy proofness* and *e-adjacent lowest welfare for inverse reports*. The proof of the following proposition is analogous to that of Proposition 4.4 and is omitted.

Proposition B.1. *Let* e *be an extension. Then* e-adjacent strategy-proofness *is equivalent to* e^d-adjacent lowest welfare for inverse reports.

Next, we establish the equivalence of *e-adjacent lowest welfare for inverse reports* and *e-lowest welfare for inverse reports* for each $e \in \{sd, dl, ul\}$.

Proposition B.2. *For each $e \in \{sd, dl, ul\}$, e-*adjacent lowest welfare for inverse reports *is equivalent to* e-lowest welfare for inverse reports.

Proof. We only prove the statement for the *dl*-extension; the statements for the *sd*- and *ul*-extensions can be proved similarly. First, by Proposition B.1, *dl-adjacent lowest welfare for inverse reports* is equivalent to *ul-adjacent strategy-proofness*, which, by Theorem 4.2, is equivalent to *ul-strategy-proofness*. But by Proposition 4.8, *ul-strategy-proofness* is equivalent to *dl-lowest welfare for inverse reports.* □

C Appendix: *dl*-strategy-proofness of the Serial Rule

In this appendix, we prove part (1) of Theorem 4.7. By Theorem 4.2, it suffices to show that the serial rule S is *dl-adjacent strategy-proof*. Let $R \in \mathcal{R}(A)^N$ and $i \in N$. Let $R'_i \in \mathcal{R}(A)$ be adjacent to R_i. Without loss of generality, assume that $1\, P_i\, 2\, P_i \cdots P_i\, n$. Let $k \in A$ be such that $(k+1)\, P'_i\, k$. Let $\pi \equiv S(R)$ and $\pi' \equiv S(R'_i, R_{-i})$. Suppose, on the contrary, that $\pi'_i\, P^{dl}_i\, \pi_i$.

Consider the simultaneous consumption algorithm applied to the economy R. We use the following notation throughout the proof.[22] For each $(\ell, t) \in A \times \mathbb{R}_+$, let $\boldsymbol{N(\ell, t)}$ be the set of agents who consume object ℓ at time t; i.e., for each $\ell \in A$ and each $t \in \mathbb{R}_+$ such that $t \in [t^{s-1}, t^s)$ for some $s \in \mathbb{N}$, $N(\ell, t) \equiv N^*(R, A^{s-1}, \ell)$. Note that $N(\ell, t)$ may be empty. Also, for each $\ell \in A$, let $\boldsymbol{t(\ell)}$ be the time at which object ℓ is exhausted; i.e., $t(\ell) \equiv \sup\{t \in \mathbb{R}_+ : N(\ell, t) \neq \emptyset\}$. Let $\boldsymbol{t_0} \equiv \max_{\ell \in \{1, \cdots, k-1\}} t(\ell)$ if $k \neq 1$; and $t_0 \equiv 0$ otherwise. Define $N'(\ell, t)$, $t'(\ell)$, and t'_0 similarly for the economy (R'_i, R_{-i}).

It is easy to see that for each $\ell \in \{1, \cdots, k-1\}$, $\pi_{i\ell} = \pi'_{i\ell}$ and $t_0 = t'_0$. Now we proceed in four steps.

Step 1: $t(k) \leq t'(k)$.

In the algorithm that determines π, agent i consumes object k during the interval $[t_0, t(k))$, so that $\pi_{ik} = t(k) - t_0$. On the other hand, in the algorithm that determines π', agent i consumes object k during a subinterval of $[t_0, t'(k))$. Thus, if $t(k) > t'(k)$,

$$\pi'_{ik} \leq t'(k) - t_0 < t(k) - t_0 = \pi_{ik},$$

contradicting that $\pi'_i\, P^{dl}_i\, \pi_i$.

Step 2: *For each $t \in [t_0, t(k))$, $N(k, t) \backslash \{i\} = N'(k, t) \backslash \{i\}$.*

[22] Some of the notations are borrowed from Bogomolnaia and Moulin (2001).

We only show that for each $t \in [t_0, t(k))$, $N(k,t)\backslash\{i\} \subseteq N'(k,t)\backslash\{i\}$; the reverse inclusion can be proved similarly. Suppose, on the contrary, that there are $\hat{t} \in [t_0, t(k))$, $j \in N\backslash\{i\}$, and $\hat{\ell} \in A\backslash\{k\}$ such that $j \in N(k,\hat{t}) \cap N'(\hat{\ell},\hat{t})$. We proceed in three steps.

Step 2.1: *Let* $B \equiv \{\ell \in A\backslash\{k\} : t(\ell) < t'(\ell)\}$ *and* $h \equiv \arg\min_{\ell \in B} t(\ell)$. *Then* $B \neq \emptyset$ *and* $t(h) < t(k)$.

Since $\hat{t} < t(k) \leq t'(k)$ and $j \in N'(\hat{\ell},\hat{t})$, $\hat{\ell}\, P_j\, k$ and $\hat{t} < t'(\hat{\ell})$. Also, since $j \in N(k,\hat{t})$, $t(\hat{\ell}) \leq \hat{t}$. Thus, $\hat{\ell} \in B$, so that $B \neq \emptyset$. Further, since $t(h) \leq t(\hat{\ell}) \leq \hat{t} < t(k)$, $t(h) < t(k)$.

Step 2.2: *There are* $\bar{t} \in [t_0, t(h))$ *and* $j' \in N$ *such that* $j' \in N(h,\bar{t}) \cap N'(h,\bar{t})^c$.

Suppose, on the contrary, that for each $t \in [t_0, t(h))$, $N(h,t) \subseteq N'(h,t)$. Since $t(h) < t'(h)$,

$$\begin{aligned} 1 &= \int_0^{t'(h)} |N'(h,t)|\, dt \\ &= \int_0^{t_0} |N'(h,t)|\, dt + \int_{t_0}^{t(h)} |N'(h,t)|\, dt + \int_{t(h)}^{t'(h)} |N'(h,t)|\, dt \\ &> \int_0^{t_0} |N(h,t)|\, dt + \int_{t_0}^{t(h)} |N(h,t)|\, dt \\ &= \int_0^{t(h)} |N(h,t)|\, dt \\ &= 1, \end{aligned}$$

where the inequality follows from the fact that there is $\tilde{t} \in [t(h), t'(h))$ such that for each $t \in [\tilde{t}, t'(h))$, $N'(h,t) \neq \emptyset$. This is a contradiction.

Step 2.3: *Let* $h' \in A$ *be such that* $j' \in N'(h',\bar{t})$. *Then* $h' \in B$ *and* $t(h') < t(h)$, *contradicting our choice of* h.

By Steps 2.1 and 2.2, $\bar{t} < t(h) < t(k)$. Also, in the algorithm that determines π, agent i consumes object k during $[t_0, t(k))$. Thus, $j' \neq i$. Since $\bar{t} < t(h) < t'(h)$ and $j' \in N'(h',\bar{t})$, $h'\, P_{j'}\, h$ and $\bar{t} < t'(h')$. Since $j' \in N(h,\bar{t})$, the fact that $h'\, P_{j'}\, h$ implies that $t(h') \leq \bar{t}$. Thus, $h' \in B$. Moreover, $t(h') \leq \bar{t} < t(h)$.

Step 3: *(i)* $t(k) = t'(k)$*; (ii)* $\pi_{ik} = \pi'_{ik}$*; and (iii) in the algorithm that determines* π'*, agent* i *consumes object* k *during the interval* $[t_0, t'(k)) = [t_0, t(k))$.

To show (i), suppose, on the contrary, $t(k) \neq t'(k)$. By Step 1, $t(k) < t'(k)$. Note that for each

$t \in [0, t_0)$, $N(k,t) = N'(k,t)$. Thus,

$$\begin{aligned}
1 - \pi_{ik} &= \int_0^{t(k)} |N(k,t)\backslash\{i\}|\, dt \\
&= \int_0^{t_0} |N(k,t)\backslash\{i\}|\, dt + \int_{t_0}^{t(k)} |N(k,t)\backslash\{i\}|\, dt \\
&= \int_0^{t_0} |N'(k,t)\backslash\{i\}|\, dt + \int_{t_0}^{t(k)} |N'(k,t)\backslash\{i\}|\, dt \\
&< \int_0^{t_0} |N'(k,t)\backslash\{i\}|\, dt + \int_{t_0}^{t(k)} |N'(k,t)\backslash\{i\}|\, dt + \int_{t(k)}^{t'(k)} |N'(k,t)\backslash\{i\}|\, dt \\
&= \int_0^{t'(k)} |N'(k,t)\backslash\{i\}|\, dt \\
&= 1 - \pi'_{ik},
\end{aligned}$$

where the third equality follows from Step 2 and the inequality from the fact that there is $\tilde{t} \in [t(k), t'(k))$ such that for each $t \in [\tilde{t}, t'(k))$, $N'(k,t)\backslash\{i\} \neq \emptyset$. Thus, $\pi_{ik} > \pi'_{ik}$, contradicting that $\pi'_i \, P_i^{dl} \, \pi_i$.

To show (ii) and (iii), recall that in the algorithm that determines π', agent i consumes object k during a subinterval of $[t_0, t'(k))$. Since $t(k) = t'(k)$ and $\pi'_{ik} \geq \pi_{ik}$, agent i, in fact, consumes object k during the interval $[t_0, t'(k)) = [t_0, t(k))$. Thus, $\pi'_{ik} = t'(k) - t_0 = \pi_{ik}$.

Step 4: *Concluding.*

If $t(k) \leq t_0$, then object k is not available at time t_0 in each of the two algorithms that determine π and π', respectively. Thus, the two algorithms coincide. If $t(k) > t_0$, then by (iii) in Step 3, $t'(k+1) \leq t_0$ and object $k+1$ is not available at time t_0 in each of the two algorithms. Thus, the two algorithms again coincide. In either case, $\pi = \pi'$, in contradiction to $\pi'_i \, P_i^{dl} \, \pi_i$.

References

[1] Abdulkadiroğlu, A., and T. Sönmez, 2003, "Ordinal efficiency and dominated sets of assignments", *Journal of Economic Theory*, 112, 157-172.

[2] Birkhoff, G., 1946, "Three observations on linear algebra", *Revi. Univ. Nac. Tucuman*, ser A 5, 147–151.

[3] Bogomolnaia, A. and E. J. Heo, 2011, "Probabilistic assignment of objects: Characterizing the serial rule", mimeo.

[4] Bogomolnaia, A., and H. Moulin, 2001, "A new solution to the random assignment problems", *Journal of Economic Theory* 100, 295-328.

[5] Bogomolnaia, A., and H. Moulin, 2002, "A simple random assignment problem with a unique solution", *Economic Theory*, 19, 623-635.

[6] Budish, E., Y.-K. Che, F. Kojima, and P. Milgrom, 2012, "Designing random allocation mechanisms: Theory and applications", mimeo.

[7] Carroll, G., 2012, "When are local incentive constraints sufficient?", *Econometrica*, 80, 661-686.

[8] Che, Y.-K., and F. Kojima, 2010, "Asymptotic equivalence of probabilistic serial and random priority mechanisms", *Econometrica*, 78, 1625-1672.

[9] Chipman, J. S., 1960, "The foundations of utility", *Econometrica*, 28, 193-224.

[10] Cho, W. J., 2012, "On the extension of preferences over objects to preferences over lotteries", mimeo.

[11] Foley, D., 1967, "Resource allocation and the public sector", Yale Economic Essays, 7, 45-98.

[12] Gibbard, A., 1973, "Manipulation of voting schemes: A general result", *Econometrica*, 41, 587-601.

[13] Gibbard, A., 1977, "Manipulation of schemes that mix voting with chance", *Econometrica*, 45,665-681.

[14] Hashimoto, T., D. Hirata, O. Kesten, M. Kurino, M. U. Ünver, 2012, "Two axiomatic approaches to the probabilistic serial mechanism", mimeo.

[15] Hausner, M., 1954, "Multidimensional utilities", in *Decision Processes* (R. M. Thrall, C. H. Coombs, R. L. Davis, Eds.), Wiley, New York.

[16] Heo, E.-J., 2010, "Probabilistic assignment of objects: A Characterization of the Serial Rule", mimeo.

[17] Hylland, A., and R. Zeckhauser, 1979, "The efficient allocation of individuals to positions", *Journal of Political Economy*, 87, 293-314.

[18] Kasajima, Y., 2011, "Probabilistic assignment of indivisible goods with single-peaked preferences", mimeo.

[19] Katta, A.-K., and J. Sethuraman, 2006, “A solution to the random assignment problem on the full preference domain”, *Journal of Economic Theory*, 131, 231-250.

[20] Kesten, O., 2009, “Why do popular mechanisms lack efficiency in random environments?”, *Journal of Economic Theory*, 144, 2209-2226.

[21] Kemeny, J. G., 1959, “Mathematics without Numbers”, *Daedalus*, 577-591.

[22] Kemeny, J. G., and J. L. Snell, 1962, *Mathematical Models in the Social Sciences*, Blaisdell, New York.

[23] Kojima, F., 2009, “Random assignment of multiple indivisible objects”, *Mathematical Social Sciences*, 57, 134-142.

[24] Kojima, F., and M. Manea, 2010, “Incentives in the probabilistic serial mechanism”, *Journal of Economic Theory*, 145, 106-123.

[25] Liu, Q., and M. Pycia, 2012, “Ordinal efficiency, fairness, and incentives in large markets”, mimeo.

[26] Manea, M., 2008, “A constructive proof of the ordinal efficiency welfare theorem”, *Journal of Economic Theory*, 141, 276-281.

[27] Manea, M., 2009, “Asymptotic ordinal inefficiency of random serial dictatorship”, *Theoretical Economics*, 4, 165-197.

[28] McLennan, A., 2002, “Ordinal efficiency and the polyhedral separating hyperplane theorem”, *Journal of Economic Theory*, 105, 435-449.

[29] von Neumann, J., 1953, “A certain zero-sum two-person game equivalent to the optimal assignment problem”, in *Contributions to the Theory of Games*, 2 (H. W. Kuhn and A. W. Tuckers, Eds.), Princeton University Press, Princeton, New Jersey.

[30] von Neumann, J., and O. Morgenstern, 1944, *Theory of Games and Economic Behavior*, 2nd edition, Princeton University Press, Princeton, New Jersey. 3rd edition published in 1953.

[31] Sato, S., 2010, “A sufficient condition for the equivalence of strategy-proofness and nonmanipulability by preferences adjacent to the sincere one”, forthcoming in *Journal of Economic Theory*.

[32] Satterthwaite, M. A., 1975, "Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions", *Journal of Economic Theory*, 10, 187–217.

[33] Sen, A., 2011, "The Gibbard random dictatorship theorem: A generalization and a new proof", *SERIEs*, 2.

[34] Thomson, W., 2011, *Strategy-proof Allocation Rules*, mimeo.

[35] Tinbergen, I., 1953, *Redelijke Inkomensverdeling*, N. V. DeGulden Pers, Haarlem, The Netherlands, 2nd edition.

[36] Yilmaz, O., 2009, "Random assignment under weak preferences", *Games and Economic Behavior*, 66, 546-558.

[37] Yilmaz, O., 2010, "The probabilistic serial mechanism with private endowments", *Games and Economic Behavior*, 69, 475-491.

[38] Zhou, L., 1990, "On a conjecture by Gale about one-sided matching problems", *Journal of Economic Theory*, 52, 123-135.