Ordinal Efficiency of the Random Priority Rule

Uluç Şengil
Koç University
usengil@ku.edu.tr

July 16, 2019

Abstract

The *random assignment problem* is the task of assigning a number of objects to an equal number of agents. A common solution to this is the *Random Priority Rule* which is *strategy-proof*, *ex-post efficient* and *treats equals equally*. However, the *Random Priority Rule* is not guaranteed to be *ordinally efficient* when there are more than three agents. I provide necessary and sufficient conditions for *ordinal efficiency* for the *Random Priority Rule* by characterizing *ordinal efficiency* in *random assignment problems* with few objects. The results are generalized by provided methods to obtain smaller problems from any *random assignment problem* while preserving *ordinal efficiency* under the *Random Priority Rule*, and where the methods do not yield small enough problems I provide an algorithm that characterizes *ordinal efficiency*.
Keywords: The random assignment problem, the Random Priority rule, ordinal efficiency

Özet

Rastgele tayin problemi belli bir sayıda objenin eşit sayıdaki ekonomik bireye verilmesini işler. Bu problem için yaygınca kullanılan bir çözüm, *Rastgele Öncelik Kuralı*, *streteji-geçirmez*, *ex-post verimlidir* ve *eşite eşit davranır.* Ancak, *Rastgele Öncelik Kuralı* üçten fazla birey olduğu durumlarda *ordinal verimliliği* sağlamayı garanti edemez. *Rastgele tayin problemlerinde ordinal verimliliği* karakterize ederek, az sayıda obje bulunduğu durumlar için *Rastgele Tayin Kuralı*'nın *ordinal verimli* olması için gerekli ve yeterli koşuları öne sürüyorum. Sonrasında bu elde ettiğimiz sonuçları genellemek için methodlar oluşturuyorum ve genellemenin erişemediği durumlar için de *ordinal verimliliği* karakterize eden bir algoritma öne sürüyorum.
Anahtar Kelimeler: Rastgele tayin problemi, Rastgele Öncelik Kuralı, ordinal verimlilik

Acknowlegdements

I want to offer my deepest gratitude to my advisor Özgür Yılmaz; my professors Levent Koçkesen, Mert Kimya, Kamil Yılmaz, İnsan Tunalı, Cem Çakmaklı and Sumru Altuğ; my accomplice Ahmet Gülek; and my dearest friend Sude Ganidağlı for their extensive support during my studies.

1 Introduction

I consider the problem of assigning n objects to n agents where each agent may have differing preferences over objects. This is called the *random assignment problem* or simply referred to hereafter as *problem*. This problem is usually encountered in assignments where fairness is crucial and monetary transfers are not allowed, such as assignments of dorms to college students or public housing to applicants. One solution to this problem is the *Random Priority Rule* (Abdulkadiroğlu and Sönmez (1998)). The rule is applied by randomly selecting an ordering of agents and letting them take turns according to that order. The agent in turn chooses their most preferred object among the remaining ones. The chosen object is then assigned to the agent and removed from the pool of remaining objects. The process ends when all agents have taken turns, or there are no objects remaining. This solution leads to *ex-post efficient* outcomes, where any trade that makes an agent better off will make at least one agent worse off.

Another solution, the *Probabilistic Serial Rule* is favored in some situations as it satisfies *ordinal efficiency*, which is a stronger than *ex-post efficiency* and ensures no trade can make an agent better off while not making some agent worse off, even when agents are allowed to trade the probabilities of obtaining objects. However this solution is not *strategy-proof*, so some agents might be in a position where misreporting their preferences may lead to a better outcome for them. The lack of *strategy-proofness* has an impact on fairness when the agents have varying levels of information (Abdulkadiroglu et al. (2006)). Being *strategy-proof* is a desirable property where *random assignments* are being used to deliver equal opportunities to agents coming from differing backgrounds.

The *Random Priority (RP) Rule* gives the agents a very strong incentive to report their preferences truthfully by being *strategy-proof*. But in cases where efficiency is considered to be important, not satisfying *ordinal efficiency* raises concerns. This can be mitigated by the fact that although the *RP Rule* lacks *ordinal efficiency* when considering all *problems*, the *RP Rule* is *ordinally efficient* in some domains of *random allocation problems*. One example of these domains is the domain with three objects and three agents (Bogomolnaia and Moulin (2001)). By singling out the properties of such domains, I derive more generalized requirements for *ordinal efficiency*, or lack thereof. This leads to the full characterization of *ordinal efficiency* for *problems* involving two objects and any number of agents.

Following that, I further generalize our findings by devising methods to obtain *problems* with fewer objects while not affecting the *ordinal efficiency* of the *RP Rule*. A domain of *problems* that lends itself to this is where the preferences of the agents is *tiered*, ie there are groups of objects where all agents have consistent preferences between the groups, but have differing preferences inside the groups. However this does not cover all possible preferences, so there is the need of developing more general methods. Thus I devise other methods for obtaining *problems* with fewer objects such as *deconstruction* and *reduction* which can be applied to almost all preferences. I show that both methods preserve *ordinal efficiency* and, thus, it is possible to characterize the notion for any *problem* at hand.

Related Literature

Abdulkadiroğlu and Sönmez (1998) define the *Random Priority Rule* and show that the rule is *ex-post efficient*. Another solution to the *random assignment problem*, the *Probabilistic Serial (PS) Rule* is proposed by Bogomolnaia and Moulin (2001). The new solution lacks *strategy-proofness* but satisfies the stronger notion of *ordinal efficiency*. The *Random Priority Rule* is also shown to satisfy *ordinal efficiency* when the number of objects and agents are fewer than four. The two rules have been shown to be asymptotically equivalent when objects have multiple copies and the ratio of agents to types of objects increases (Che and Kojima (2010)).

Ordinal efficiency in *random assignments* is studied in Abdulkadiroğlu and Sönmez (2003) by using the *ex-post efficiency* of *lotteries*. Restrictions on *preference domains* has been studied for inducing *strategy-proofness* on the *PS Rule* (Liu (2017), Cho (2016)). *Tiered preference structures* have also been studied regarding *strategy-proofness* in the *PS Rule* (Liu and Zeng (2017)).

2 The Model

Let O be a finite set of objects and I a finite set of agents with $|O| = |I|$. Fix the set of agents and objects. For each agent $i \in I$, let $\succ_i$ be a preference relation which is complete, antisymmetric and transitive. Let $\succ = (\succ_i)_{i \in I}$ be a preference profile. An agent i **ranks** an object a k'th, or $r_i(a) = k$ if the number of objects o_n such that $o_n \succ_i a$ equals $k-1$. A **random assignment problem** (or simply **problem**) is a triple $(O, I, \succ)$.

An **ex-post assignment**, or **assignment** is an injective function $\mu : I \leftarrow O$. A **random consumption** is a probability distribution over objects. A **random assignment** is denoted by $Q = \left[q_{ia}\right]$, where $q_{ia} \in [0,1]$ is the probability of agent i receiving object a and the i'th row Q_i is agent i's *random consumption*. A **rule** assigns a *random assignment* to each **problem**.

For a *random assignment* $Q = \left[q_{ia}\right]$ and a preference profile $\succ$, Define the following binary relation:

$$\tau(Q, \succ) = \{(a,b) \in O \times O \mid \exists i \in I \text{ such that } a \succ_i b, q_{ib} > 0\} \tag{1}$$

For a pair of objects $a, b \in O$, $\tau(Q, \succ)$ contains an $a-b$ **cycle**, or $\tau(Q\ \succ)$ is $a-b$ **cyclic** if $(a,b), (b,a) \in \tau(Q, \succ)$. $\tau(Q, \succ)$ is $a-b$ **acyclic** if it is not $a-b$ *cyclic*. $\tau(Q, \succ)$ is a-**cyclic** if for some b, $\tau(Q, \succ)$ is $a-b$ *cyclic* and $\tau(Q, \succ)$ is **cyclic** if for some a, $\tau(Q, \succ)$ is a-*cyclic*. $\tau(Q, \succ)$ is **acyclic** if it is not *cyclic*.

A *random assignment* Q is **ordinally efficient** for a *problem* $(O, I, \succ)$ if $\tau(Q, \succ)$ is *acyclic*. An *assignment* is **ex-post efficient** for a *problem* $(O, I, \succ)$ if $\tau(Q, \succ)$ is *acyclic*. A *rule* is *ex-post efficient* if for each *problem*, the resulting *assignment* μ is *ex-post efficient* and it is *ordinally efficient* if for each problem, the resulting *random assignment* Q is *ordinally efficient*. A *rule* is **strategy-proof** if truth telling is a dominant strategy for the agents.

A *rule* commonly applied to *random assignment problems* is the **Random Priority (RP) Rule** (Abdulkadiroğlu and Sönmez (1998)). This rule satisfies *ex-post efficiency* and *strategy-proofness*. I explain the rule in greater detail in the next section.

2.1 The Random Priority Rule

For a problem $(O, I, \succ)$, the *RP Rule* leads to an *assignment* by the following steps:

1. Draw an ordering of the agents in O in a uniformly random way.
2. Give turns to agents sequentially by following the drawn ordering.
3. In the first turn, the first agent picks their first *ranked* object. The chosen object is removed.
4. In the following turns, the agent who has the turn picks the first *ranked* object among the remaining ones.
5. The sequence ends when either objects or the order runs out.

In order to illustrate the *rule*, consider following example: Let there be four objects a, b, c, d and agents $1, 2, 3, 4$ with the following *preference domain*:

$$\begin{aligned} &a \succ_1 b \succ_1 c \succ_1 d \\ &a \succ_2 b \succ_2 c \succ_2 d \\ &b \succ_3 a \succ_3 d \succ_3 c \\ &b \succ_4 a \succ_4 d \succ_4 c \end{aligned} \tag{2}$$

For the algorithm, first we obtain an ordering of the agents. Let the ordering to be $(2, 3, 1, 4)$. Then the *rule* resolves as follows:

- Agent 2's turn: they pick their first ranked object a.
- Agent 3's turn: they pick their first ranked object b as it is still available.
- Agent 1's turn: Their first and second ranked objects a, b are removed but their third ranked object c is available, thus they pick c.
- Agent 4's turn: Their first and second ranked objects are removed, so they pick their third ranked object d.

The result is the following *assignment*: Agents $1, 2, 3, 4$ obtain c, a, b, d respectively. Notice that this *assignment* satisfies *ex-post efficiency*. For the RP Rule, the *random assignment* is obtained by exhausting the possible orderings that can be picked in the initial step and giving each ordering an equal probability. For our example *problem*, the RP Rule gives the following *random assignment* Q:

$$\begin{bmatrix} 5/12 & 1/12 & 5/12 & 1/12 \\ 5/12 & 1/12 & 5/12 & 1/12 \\ 1/12 & 5/12 & 1/12 & 5/12 \\ 1/12 & 5/12 & 1/12 & 5/12 \end{bmatrix} \tag{3}$$

Agents 1 and 3 contribute (a, b) to $\tau(Q, \succ)$ and agents 2 and 4 contribute (b, a). See that the *random assignment* is thus not *ordinally efficient*. The RP Rule does not satisfy *ordinal efficiency* in general, but it assigns some *problems* to *ordinally efficient random assignments*.

The RP Rule is also *strategy-proof* for each *problem* (Abdulkadiroğlu and Sönmez (1998)).

3 Characterizations of Efficiency

Our goal is to characterize the *preference domain* for which the RP Rule always gives an *ordinally efficient random assignment*. Let us fix a *problem* $(O, I, \succ)$ and the *random assignment* Q given by the RP Rule for $(O, I, \succ)$. A *preference domain* is *ordinally efficient* if for each *preference profile* in the preference domain, Q is *ordinally efficient*.

A trivial *ordinally efficient preference domain* is such that for each $i, j \in I$, $\succ_i = \succ_j$. The *preference domain* consists of a single *preference profile* up to all permutations of objects. The following result provides a slight extension to this *domain*:

Proposition 1. *Let $a, b \in O$ and define $A = \{i \mid a \succ_i b\}$, $B = \{i \mid b \succ_i a\} = I - A$. Then $\tau(Q, \succ)$ is not* a-b cyclic *if $|A| \leq 1$ or $|B| \leq 1$.*

Proof. Without loss of generality, assume $|B| \geq |A|$. If $|A| = 0$, then for each $i \in I$, $b \succ_i a$ and $\tau(Q, \succ)$ is trivially $a - b$ *acyclic*.

Let $|A| = 1$ and i be the only agent with the preference $a \succ_i b$. Thus, for $j \neq i$, $b \succ_j a$. If i chooses first or a is available when it is i's turn, they will never choose b instead of a. If a is chosen before the turn gets to i, since $b \succ_j a$ for all agents $j \neq i$, b has to be chosen before i's turn. Thus, i cannot take b. Since $q_{ib} = 0$, only (b, a) can be in $\tau(Q, \succ)$. Thus, $\tau(Q, \succ)$ cannot be $a - b$ cyclic. □

This proposition naturally extends to an alternative proof of the *ordinal efficiency* of the RP Rule when $|I| \leq 3$. Since the condition in Proposition 1 is satisfied in this case, the *RP Rule* is *ordinally efficient*.

Another result is obtained by applying the same steps in the previous allocation:

Proposition 2. *Let a be an object and i be an agent. If q_{ia} is positive, then for each object $b \succ_i a$, there exists a unique agent j_b with $b \succ_{j_b} a$.*

Proof. Follows simply through the proof of Proposition 1. □

To further examine *ordinal efficiency*, construct a minimal *preference domain* which violates it: Let a, b, c, d be objects with $a \neq b, c \neq a, d \neq b$ (we allow $c = b$ and $d = a$). Then the following preference profile induces an $a - b$ cyclic $\tau(Q, \succ)$:

$$\begin{aligned} & a \succ_1 b \succ_1 \dots \\ & b \succ_2 a \succ_2 \dots \\ & a \succ_3 c \succ_3 \dots \\ & b \succ_4 d \succ_4 \dots \end{aligned} \tag{4}$$

$\tau(Q, \succ)$ has an $a - b$ cycle if the following conditions are satisfied:

1. Agents i_1 and i_2 disagree on their preferences over objects a and b, while there are no other objects c that are preferred between those two (ie, $a \succ_i c \succ_i b$) for both of these agents.
2. For each object $x \in \{a, b\}$, there are two agents ranking object x first. Thus, each of these agents is assigned to their second ranked object with positive probability.

This can be stated as follows:

Remark 1. *$\tau(Q, \succ)$ is $a-b$ cyclic if both a and b are ranked first by at least two agents and for both a and b, at least one of the agents who ranks one first ranks the other second.*

This is a very specific *preference domain* that is not *ordinally efficient*. But with further generalization methods, we can apply this result to some *profiles* where the structure is not immediately apparent. I introduce such a method in the following section.

4 Tiered Preference Profiles

To characterize larger domains, refer to the following definition.

Definition 1. *A* preference profile $\succ$ *is* ***tiered*** *if there exists an ordered partition $T = (T_1, T_2, ..., T_n)$ of O such that for each agent i and object $a \in T_m$, $b \in T_l$, $m < l$ implies $a \succ_i b$. Each $T_i \in T$ is a* ***tier***.

Note that every *preference profile* is *tiered* with a single *tier* trivially. When a *preference profile* has n *tiers*, it is a member of a *preference domain* containing $n!$ *profiles* where each *profile* has the same *radom assignment* under the *RP Rule*.

The aforementioned *preference domain* can be constructed starting from any *tiered preference profile* it includes. To make it easier to deal with subsets of O, define the following:

Definition 2. *Let $'O \subseteq O$, $I' \subseteq I$, $\succ(I', O')$ be a preference profile such that for each $a, b \in O'$, $a \succ_i b$ if and only if $a \succ(I', O')_i b$. Then $\succ'$ is a* ***projection of*** *$\succ$ to the set O'.*

Proposition 3. *Let $\succ, \succ'$ be two* tiered preference profiles *with* tiers *T and T'. If T and T' are the same partitions of O with a different order and for each $T_m \in T, T'$, $\succ(T_m) = \succ'(T_m)$, then the* random assignments *the RP Rule gives for the problems $(O, I, \succ)$ and $(O, I, \succ')$ are equal.*

Proof. This proof shows that there exists a bijection from the *ex-post assignments* of the problem $(O, I, \succ)$ to the *assignments* of the *problem* $(O, I, \succ')$.

In any ordering $I_\#$ picked for the RP Rule in $(O, I, \succ)$, the first $|T_1|$ picks will be from T_1. Then the next $|T_2|$ picks will be from T_2 and so on due to the *tiered profile*.

Let $c_m = |T_m|$ and divide the ordering $I_\#$ to n tiers:

$$\begin{array}{cccc} \text{from } T_1 & \text{from } T_2 & \dots & \text{from } T_n \\ p_1^1, p_1^2, \dots, p_1^{c_1}, & p_2^1, p_2^2, \dots, p_2^{c_2}, & \dots, & p_n^1, p_n^2, \dots, p_n^{c_n} \end{array} \tag{5}$$

Here p_m^r corresponds to the r'th pick from the m'th *tier*.

For $I_\#$ see that the k'th pick p_k corresponds to

$$p_k = p_m^r, \text{ where } j = r + \sum_{j=1}^{m-1} c_j \tag{6}$$

This means for a set of *tiers* $\mathcal{T}_m$ with

$$\mathcal{T}_m = \{T_o \mid \forall i \in I, a \in T_o, b \in T_m \implies a \succ_i b\} \tag{7}$$

there exists the following relation:

$$p_k = p_m^r, \text{ where } k = r + \sum_{T_o \in \mathcal{T}_m} |T| \tag{8}$$

This fact leads to a natural mapping $P^* : \{1, ..., n\} \to \{1, ..., n\}$ from the picks of $(O, I, \succ)$ to the picks of $(O, I, \succ')$. For each $p_k \in I_\#$, let

$$P^*(p_k) = p_m^r, \text{ where } k = r + \sum_{T_o \in \mathcal{T}'_m} |T| \tag{9}$$

given that

$$\mathcal{T}'_m = \{T_o \mid \forall i \in I, a \in T_o, b \in T_m \implies a \succ'_i b\} \tag{10}$$

The ordering $I'_\# = \{P^*(p_k)\}$ results in the same *assignment* for the problem $(O, I, \succ')$ that $I_\#$ results in $(O, I, \succ)$.

The final thing to do is to show P^* is a bijection. Because P^* maps a finite domain to itself, being an injection implies being a bijection. Let α, β be permutations over tier numbers $(1 \ldots m)$ and assume that $P^*(\alpha) = P^*(\beta)$. By the definition of P^*, $\alpha = \beta$. Both the changes in orders within a *tier* and in orders between *tiers* changes the resulting order.

Thus, P^* is injective and therefore bijective. Since each ordering in $(O, I, \succ)$ is mapped to one ordering in $(O, I, \succ')$, their *random assignments* are equal to one another. □

By this proof, each *tier* is shown to be independent from each other under the RP Rule. This leads to the following result.

Corollary 1. *Let* $\succ$ *be a* tiered preference profile *with* tiers $T = \{T_1, T_2, ..., T_n\}$ *and* $Q^m = \left[q_{ia}^m\right]$ *be the* random assignment *given by the RP Rule for the* problem $(T_m, I, \succ(T_m))$. *Then,*

$$Q = \sum_{m=1}^{n} Q^m = \left[\sum_{m=1}^{n} q_{ia}^m\right] \tag{11}$$

As *ordinal efficiency* of a *preference profile* only depends on the *random assignment*, the equivalence of the *random assignments* imply the equivalence of ordinal efficiency. Thus, *problems* can be divided to decrease the number of objects in our problem thanks to the *tier structure*:

Corollary 2. *Let* $\succ$ *be a* tiered preference profile *with* tiers $T = \{T_1, T_2, ..., T_n\}$. Q *is ordinally efficient if and only if each* random assignment Q^m *given by the RP Rule for* $(T_m, I, \succ(T_m))$ *is ordinally efficient.*

Using *tier structure*, Remark 1 can be generalized to a larger *preference domain*:

Proposition 4. *Let* $\succ$ *be a* tiered preference profile *with* tiers $T = \{T_1, ..., T_n\}$. Q *is not ordinally efficient if for any* $T_m \in T$*, the requirements of Remark 1 is satisfied.*

Proof. Follows directly from Corollary 2 and Remark 1. □

This condition is clearly sufficient for ordinal inefficiency, and it is also necessary for problems with few objects:

Proposition 5. *Let* $\succ$ *be* tiered *and* T_m *be a* tier *with* $|T_m| \leq 2$.

Then the random assignment Q^m *given by the RP Rule for* $(T_m, I, \succ(T_m))$ *is ordinally efficient if* $\succ(T_m)$ *doesn't satisfy the conditions in Remark 1.*

Proof. Let $A = \{i \mid a \succ_i b\}$, $B = \{i \mid b \succ_i a\} = I - A$ and assume $|B| \leq |A|$ without loss of generalization. If the *tier* has two objects and does not satisfy Remark 1's conditions, then there we will have $|A| \leq 1$. This is shown to be efficient in Proposition 1. □

5 Reduction Operation on Preference Profiles

By Proposition 3, we have obtained a way to decrease the number of *objects* in some *preference profiles*. This is important because we can fully determine *ordinal efficiency* by decreasing the number of objects to two. Separating to *tiers* is useful, but it requires a special *preference domain* structure. Here we introduce a method to decrease the number of objects in almost any *preference domain*.

Consider the following slight variation to the *RP Rule*: Instead of determining an ordering of the agents first, we randomly select an agent at the start of each turn. This variation still produces the same assignments with the same probabilities, thus it is equivalent to the *RP Rule*.

Using this variation, after the first turn we'll again be left with a *problem* that lacks one agent and their first *ranked* object by the recursive nature of this rule. this will again be a *problem*. To establish the new *problem's* relation with the initial *problem*, we define it as an *immediate subproblem*:

Definition 3. *If* $r_i(a) = 1$*, the following* problem *is an* immediate subproblem *of* $(O, I, \succ)$*:*

$$D_{ia}(O, I, \succ) = (O - \{a\}, I - \{i\}, \succ(O - \{a\}, I - \{i\})) \tag{12}$$

For a problem*; the* immediate subproblems*, the* immediate subproblems *of its* immediate subproblems *and so on are the* problem's subproblems*.*

The set of all immediate subproblems *of a* problem $(O, I, \succ)$ *is a* decomposition

$$\mathcal{D}(O, I, \succ) = \{D_{ia} \mid i \in I, a \in O, r_i(a) = 1\}$$

Thus, there are $|I|$ *immediate subproblems* of a *problem* with each *immediate subproblem* having $|I|-1$ agents and $|O| - 1$ objects.

With the RP Rule, every agent has an equal chance of acting first and taking their first ranked object. Using this fact, we write the *random assignment* of a *problem* as a sum of the *random assignments* of its *immediate subproblems*.

Remark 2. *Let* $\mathcal{D}$ *be the* decomposition *of* $(O, I, \succ)$*,* Q_{ia} *the* random assignment *given by the RP Rule for each* $D_{ia}(O, I, \succ) \in \mathcal{D}$ *and* Q_f *be a matrix with* $q^f_{ia} = {}^1\!/_{|I|}$ *if* $r_i(a) = 1$ *and* 0 *otherwise. Then,*

$$Q = Q_f + \frac{1}{|I|} \sum_{i \in I} Q_{ia} \tag{13}$$

Being able to remove any object from the *problem* is useful due to Proposition 5, but with each *decomposition* the number of *subproblems* increases combinatorially.

We can mitigate this by assigning objects to agents simultaneously. For an object a, if we have two agents i and j who rank a first, the *decomposition* includes both $D_{ia}(O, I, \succ)$ and $D_{ja}(O, I, \succ)$. This means that both i and j can be assigned an object that is not a. By merging such similar *subproblems*, we can avoid the combinatorial increase in numbers.

We propose *reduction* as a method of merging such *subproblems*. Unlike *decomposition*, we will need to keep track of the objects we've removed and the agents who *ranked* the object first, because we want to be able to obtain the *subproblems* by solely using the *reduced* form of the *problem*.

Before we define it, as an example consider the following preference profile $\succ$:

$$\begin{aligned} &a \succ_1 b \succ_1 c \succ_1 d \\ &a \succ_2 b \succ_2 c \succ_2 d \\ &b \succ_3 a \succ_3 d \succ_3 c \\ &b \succ_4 a \succ_4 d \succ_4 c \end{aligned} \tag{14}$$

If a is assigned first, then only agents 1 or 2 may have taken it. This corresponds to the *immediate subproblems* obtained by applying D_{1a} or D_{2a} with the following *preference profiles*:

$$\begin{aligned} &b \succ_2 c \succ_2 d \qquad && b \succ_1 c \succ_1 d \\ &b \succ_3 d \succ_3 c && b \succ_3 d \succ_3 c \\ &b \succ_4 d \succ_4 c && b \succ_4 d \succ_4 c \end{aligned} \tag{15}$$

Instead of dealing with these separately, we can merge them to a single preference profile. To keep track of which *subproblems* we've merged, we use a *merge set*:

Definition 4. *For $a \in O$, the* merge set *of a collection of* immediate subproblems *which lack the same object a, ie $\{D_{ia}(O, I, \succ) \mid r_i(a) = 1\}$ is a set containing a tuple of the object a and the set of agents who rank a first, ie $\{(a, \{i \mid r_i(a) = 1\})\}$.*

With a *merge set* $\{(a, \{1, 2\})\}$, the merged profile becomes:

$$\begin{aligned} &b \succ_1 c \succ_1 d \\ &b \succ_2 c \succ_2 d \\ &b \succ_3 d \succ_3 c \\ &b \succ_4 d \succ_4 c \end{aligned} \tag{16}$$

This gives the *preference profile* of the *reduction* of $(O, I, \succ)$ with the object a.

Definition 5. *Let I_a be the set of agents with $r_i(a) = 1$, $|I_a| > 2$ and $\{(a, I_a)\}$ the* merge set *of $\{D_{ia}(O, I, \succ) \mid i \in I_a\}$. Then the* reduced problem *obtained by removing the object a from $(O, I, \succ)$ is a quadruple $(O - \{a\}, I, \succ(O - \{a\}), \{(a, I_a)\})$, also denoted by $R_a(O, I, \succ)$.*

The *merge set* intuitively defines how to obtain the merged *subproblems* and subsequently which *ex-post assignments* are possible. We say that the *reduced problem induces* those *subproblems*.

Definition 6. *Let $(O', I, \succ(O'), \mathcal{I}_*)$ be a* reduced problem, *$a \in O'$ and $I_a = \{i \mid r_i(a) = 1\}$. $(O', I, \succ(O'), \mathcal{I}_*)$ is* reducible *by an object a if for each $i \in I$, $R_a(O', I, \succ(O'), \mathcal{I}_*)$* induces *a problem $(O' - \{a\}, I', \succ(O' - \{a\}, I'))$ with $I' \subseteq I$ and $i \in I'$.*

Reduction also be applied to *reduced problems*. To *reduce* a *reduced problem* which induces *subproblems* $D_{ia}(O, I, \succ)$, we consider the *decomposition* of each $D_{ia}(O, I, \succ)$ and pick the *subproblems* that lack b. This gives us the collection $\{D_{jb}D_{ia}(O, I, \succ)\}$ and subsequently the *merge set* $\{(a, I_a), (b, I_b)\}$ where $I_a = \{i \mid r_i(a) = 1\}$ and $I_b = \{i \mid r_i(b) = 1 \text{ with respect to } \succ(O')\}$.

Reducing a *reduced problem* simply involves removing the related object from the *reduced problem* and obtaining the related *merge set*.

In our example, the only object that can be assigned in the next turn is b, as all agents *rank* b first. We add them to the *merge set*, which becomes $\{(a, \{1, 2\}), (b, \{1, 2, 3, 4\})\}$. The preference profile is reduced to:

$$\begin{aligned} c &\succ_1 d \\ c &\succ_2 d \\ d &\succ_3 c \\ d &\succ_4 c \end{aligned} \tag{17}$$

This profile only contains two objects and thus can be characterized as *ordinally efficient* or not by Proposition 5. This twice *reduced problem* corresponds to six *subproblems* of $(O, I, \succ)$ and the ratio of *subproblems* to *reduced problems* will increase with each *reduction*.

Note that in the *merge set*, for each agent i there should exist an *ex-post assignment* where agent i is not assigned any object, so that the agent i is included in the *induced subproblem*. Thus not every *reduced problem* is further *reducible*:

To avoid inspecting a large number of *subproblems*, we use the following equivalent condition on *reducibility*:

Definition 7. *Let* $(O', I, \succ(O'), \mathcal{I}_*)$ *be a* reduced problem *obtained by repeatedly* reducing *a problem when it was possible,* O_r *the set of reduced objects,* $a \in O'$*, and* I_a *the set of agents who rank* a *first.*

Then, $(O', I, \succ(O'), \mathcal{I}_*)$ *is* reducible *by* a *if* a *is* multi first ranked *and for each agent* i*, the resulting* reduced problem *induces a* subproblem *that includes* i*.*

If we apply *reduction* to a *reduced problem* with *multi first ranked* objects but all of its *induced subproblems* lack some agent i, that agent will not be able to obtain the objects that are not possibly assigned to them in the *merge set*, or not ranked first. This gives way to the following result:

Proposition 6. *For a problem* $(O, I, \succ)$ *with* $i \in I$ *and* $a \in O$*,* $q_{ia} > 0$ *if and only if there exists a* reduced problem $(O', I, \succ(O'), \mathcal{I}_*)$ *with* $r_i(a) = 1$*.*

Proof. Let $q_{ia} > 0$. Then there exists an *ex-post assignment* where i obtains a. Let the ordering of the agents leading to this result be $(j_1, ..., j_n, i, ...)$ with j_k obtaining b_k. If $r_i(a) = 1$, then we're done. Else, i *ranks* one of b_k first. It is also clear that each j_k *rank* one of b_k first. Since there are n of b_k and $n + 1$ agents who *rank* one of b_k first, at least one of b_k should be *multi first ranked*. Thus, we can reduce the *problem* and eliminate one of b_k.

For the subsequent *reductions*, assume that the *reduced problem* is not *reducible*. So there exists a collection of objects O_b that less than $|O_b| + 1$ agents can only obtain objects from O_b (withholding further reductions). Even without those agents, the number of the remaining agents is still more than the number of remaining objects. Thus, the *reduced problem* must be further *reducible* when there is b_k with $r_i(b_k) = 1$. By repeatedly *reducing*, we find a *reduced problem* with $r_i(a) = 1$.

Assume there exists a *reduced problem* with $r_i(a) = 1$. Following the reduction sequence $b_1, ...b_n$, create an ordering $(j_1, ..., j_n, i, ...)$ with $j_1 \in I_{b_1}$. Applying the *RP Rule* with this ordering results in an *ex-post assignment* where i obtains a. Thus, $q_{ia} > 0$. □

With this result, we can construct $\tau(Q, \succ)$ of a problem $(O, I, \succ)$ using *reduction* without exact knowledge of Q as $\tau(Q, \succ)$ only depends on whether each $q_{ia} > 0$ or not.

Definition 8. *We apply the* Reduction Algorithm *as follows:*

1. *Start with a* problem $(O, I, \succ)$ *and* $Q^+ = 0$.
2. *If the* problem *is not reducible, set* $q^+_{ia} = 1$ *for all* i, a *with* $r_i(a) = 1$ *and end the algorithm.*
3. *Obtain* reduced problems *by applying* reduction *for each* multiple first ranked *object. Set* $q^+_{ia} = 1$ *if* $i \in I_a$ *when applying* R_a.
4. *For each* reduced problem, *iterate the algorithm from step 2.*

Due Proposition 6, this algorithm fully identifies $\tau(Q, \succ)$ and *ordinal efficiency.*

References

[1] Atila Abdulkadiroğlu and Tayfun Sönmez. "Ordinal efficiency and dominated sets of assignments". *Journal of Economic Theory* 112.1 (2003), pp. 157–172.

[2] Atila Abdulkadiroğlu and Tayfun Sönmez. "Random serial dictatorship and the core from random endowments in house allocation problems". *Econometrica* 66.3 (1998), pp. 689–701.

[3] Atila Abdulkadiroglu et al. *Changing the Boston school choice mechanism.* Tech. rep. National Bureau of Economic Research, 2006.

[4] Anna Bogomolnaia and Hervé Moulin. "A new solution to the random assignment problem". *Journal of Economic theory* 100.2 (2001), pp. 295–328.

[5] Yeon-Koo Che and Fuhito Kojima. "Asymptotic equivalence of probabilistic serial and random priority mechanisms". *Econometrica* 78.5 (2010), pp. 1625–1672.

[6] Wonki Jo Cho. "When is the probabilistic serial assignment uniquely efficient and envy-free?" *Journal of Mathematical Economics* 66 (2016), pp. 14–25.

[7] Peng Liu. "Strategy-proofness of the probabilistic serial rule on sequentially dichotomous domains". *Working paper* (2017).

[8] Peng Liu and Huaxia Zeng. "Random assignments on preference domains with a tier structure". *Working paper* (2017).