Assigning agents to a line

by

Jens Leth Hougaard,

Juan D. Moreno-Ternero

and

Lars Peter Østerdal

Discussion Papers on Business and Economics
No. 11/2012

FURTHER INFORMATION
Department of Business and Economics
Faculty of Social Sciences
University of Southern Denmark
Campusvej 55
DK-5230 Odense M
Denmark

Tel.: +45 6550 3271
Fax: +45 6550 3237
E-mail: lho@sam.sdu.dk
http://www.sdu.dk/ivoe

ISBN 978-87-91657-64-1

Assigning agents to a line*

Jens Leth Hougaard[†], Juan D. Moreno-Ternero[‡], Lars Peter Østerdal[§¶]

July 4, 2012

Abstract

We consider the problem of assigning agents to a facility, represented by slots on a line, where only one agent can be served at a time. There is a finite number of agents, and each one wants to be served as close as possible to his preferred slot. We first consider *deterministic* assignment of agents to slots. We characterize (Pareto) *efficiency* in such setting and provide an algorithm for testing if a given deterministic assignment is efficient. We also characterize *utilitarianism* (minimization of the total gap between preferred and assigned slots) and provide a quick algorithm for testing if a given deterministic assignment is utilitarian. We then consider *probabilistic* assignment of agents to slots. In such framework, we characterize, making use of the previous algorithms, a method which is *ordinally efficient* and *utilitarian*.

JEL numbers: *C78, D61, D63.*

Keywords: *Random assignment, ordinal efficiency, ex post efficiency, congested facility, utilitarianism.*

---

*We thank Anna Bogomolnaia and Herve Moulin for helpful comments and suggestions. Financial support from the Danish Council for Strategic Research, the Spanish Ministry of Science and Innovation (ECO2011-22919) as well as from the Andalusian Department of Economy, Innovation and Science (SEJ-4154, SEJ-5980) via the "FEDER operational program for Andalusia 2007-2013" is gratefully acknowledged.

[†]University of Copenhagen.

[‡]Universidad Pablo de Olavide, and CORE, Université catholique de Louvain.

[§]University of Southern Denmark

[¶]Address for correspondence: Lars Peter Østerdal, Department of Business and Economics, University of Southern Denmark, Campusvej 55, 5230 Odense M. Denmark. E-mail: lpro@sam.sdu.dk.

1 Introduction

Imagine four applicants are shortlisted for a job interview. Each interview takes an hour. Three applicants prefer to go at 1pm and one applicant prefers to go at 2pm. Every applicant would like an interview as close as possible to his preferred time slot. How should we assign time slots to the four applicants given their preferences?

We analyze in this paper assignment problems like the one above. More precisely, we consider situations where there is a finite number of agents, each with a preferred slot, caring only about the gap between their assigned slot and their preferred slot. Due to the geometric interpretation of the problem, we refer to such a situation as assigning agents to a line. Our model will have potential applications within a wide range of matching problems in which agents are served sequentially and only limited preference information (e.g., the preferred serving time) is available from the agents. For instance, assigning drivers to congested facilities, personnel to work shifts, tennis players to courts, students to oral exams etc.

Assignment problems can be solved deterministically or probabilistically (i.e., randomizing over deterministic allocations) and we consider both cases.

First we characterize Pareto efficiency of a deterministic allocation. In particular, we show that if it is possible to make a Pareto improvement by switching slots between n agents then it is also possible to make a Pareto improvement between 2 agents. Based on this, we propose an algorithm to determine whether a given (deterministic) allocation is Pareto efficient. Despite the rather elementary (although non-trivial) nature of the problem, we have been, somewhat surprisingly, unable to locate a direct characterization and algorithm in the literature.

Next, we characterize when a given (deterministic) allocation is *utilitarian* in the sense that the total gap between agents' preferred slot and their allocated slot is minimized. We propose an algorithm building on this key characterization result stating that an allocation is *constrained utilitarian* (i.e., utilitarian subject to the constraint that only occupied slots can be used) if and only if there does not exist a total gap-reducing *pairwise* swap of slots between any two agents. We use this result to focus on common shifts of blocks of assignments and whether they give rise to reductions in the total gap. As such, we obtain a direct characterization of utilitarianism and an algorithm for checking it, which is tailor-suited to the problem, as opposed to checking it by first solving the classical *linear sum assignment problem*, identifying costs with gaps, and then seeing if the total gap is the same as in the allocation at hand.

Deterministic assignment methods are usually criticized on equity grounds. If the objects to be assigned are indivisible then all deterministic methods will violate any fundamental notion of fairness. A usual course of action in normative economics to deal with such problem is to resort to randomization. Probabilistic assignment methods, which are commonly used in practice, arise when randomizing over deterministic allocations. The first probabilistic assignment method that comes to mind is the so-called *random priority* (RP) method, which assigns objects based on a random ordering of the agents. RP obeys in our setting a fundamental principle of fairness known as *equal treatment of equals*, as well as the basic incentive-compatibility principle of *strategy-proofness*. It is, however, in conflict with efficiency (and, in the case of our setting, even with weak forms of it). An alternative to RP is the so-called *probabilistic serial* (PS) method, which assigns probability distributions over objects by means of a natural constructive argument (the so-called "eating algorithm") in which each indivisible object is considered as a divisible object of probability shares. PS is (ordinally) efficient and envy-free, although only weakly strategy-proof. As we shall show later, it is not utilitarian in our context (in fact, a utilitarian solution cannot be envy-free).

We suggest in this paper a modified version of RP satisfying equal treatment of equals which is not only (ordinally) efficient but also utilitarian. Intuitively, the new rule behaves as RP by ordering agents randomly and then assigning them to slots sequentially, but with the important modification of doing so in a way that will ensure that the allocation following each step in the algorithm will be utilitarian among those agents already assigned. The construction of this new rule is partly based on our algorithms for checking efficiency and utilitarianism of deterministic assignments.

The rest of the paper is organized as follows. In Section 2, we summarize the related literature to this paper on assignment problems. In Section 3, we set the preliminaries of our model. In Section 4, we deal with the deterministic assignments and present our algorithms to check efficiency and utilitarianism. In Section 5, we move to probabilistic assignments, showing some logical relations among axioms and introducing our proposed (efficient and utilitarian) rule for this context. For a smooth passage, proofs of lemmas can be found in an appendix.

2 Related literature

Assignment problems have long been analyzed in the operations research literature, mostly taking a classical combinatorial optimization approach (see, for instance, Burkard et al., (2012) and the literature cited therein). In the economics literature, however, the interest on assignment problems has mostly relied on issues of efficiency, incentive compatibility and fairness.[1] For instance, Hylland and Zeckhauser (1979) deal, in an early an influential contribution, with the general problem of achieving efficient allocations in matching problems in which individuals must be allocated to positions with limited capacities, and propose an algorithm for it based on market-clearing prices, which also guarantees that assignments are *envy-free*, i.e., no agent would prefer the assignment of some other agent. Their algorithm, however, is not *strategy-proof*, i.e., it does not always give agents the incentive to report their true preferences. It turns out, as shown by Zhou (1990), that no rule in such setting satisfies strategy-proofness, (ex-ante) efficiency, and an even weaker notion of fairness than no-envy. Zhou (1990) also discussed the so-called *random priority* (RP) solution (known as *serial dictatorship* by Abdulkadiroglu and Sönmez (1998), among others) which had played the role of a *folk solution* for a long time, although it had not been considered in the economics literature before.

Bogomolnaia and Moulin (2001) can be considered as the seminal work that sparked the recent interest within the economic literature on assignment problems. They restrict attention to strict preferences and ordinal mechanisms (i.e., mechanisms that rely only on ordinal rankings of objects) and introduce the notion of *ordinal efficiency*.[2] They characterize, building on the contribution of Cres and Moulin (2001) in a related setting, all ordinally efficient assignments, which include the so-called *probabilistic serial solution* (PS) as a central element (but not random priority), and show that ordinal efficiency is incompatible with strategy-proofness, and equal treatment of equals.[3] They also show that PS is envy-free but not strategy-proof, whereas the opposite holds for RP.

PS is actually considered as a prominent method to solve assignment problems. Besides being popularized by Bogomolnaia and Moulin (2001), it has been characterized, in related settings to theirs, by Bogomolnaia and Moulin (2002, 2004), Katta and Sethuraman (2006), Kesten (2009), Manea (2009), Yilmaz (2009), Che and Kojima (2010), Kojima and Manea (2010), Hashimoto and Hirata (2011), Heo (2011), and Bogolmonaia and Heo (2012), among others.

[1] Another concept that might lie at the intersection of both literatures is the so-called notion of stability, which has also received considerable attention since Gale and Shapley (1962) and their analysis of the so-called marriage problem. See also Roth and Sotomayor (1990) and the literature cited therein.

[2] As we shall define later, a probabilistic assignment is ordinally efficient if it is not stochastically dominated with respect to individual preferences over certain objects.

[3] Abdulkadiroglu and Sönmez (2003) provide an alternative characterization of ordinal efficiency. See also McLennan (2002) and Manea (2008) for the relationship between ordinal efficiency and ex-ante efficiency.

In this paper, we study a slightly different version of the model considered by Bogomolnaia and Moulin (2001) and their followers. Not only we allow for weak preferences, as in Katta and Sethuraman (2006), but we actually assume that preferences are symmetric with respect to a *peak* (*target*). As such, our model is reminiscent of the so-called division problem with single-peaked preferences, initiated by Sprumont (1991).[4] It also touches some recurrent topics in queueing problems (e.g., Dolan, 1978).

3 Preliminaries

Imagine a facility with a fixed service capacity that can serve one agent at each *slot*. We use the term "slot" which can refer to both a point in time (i.e., a time slot) or a location (i.e., a physical slot) arranged on a line. In the latter case, we could imagine that the situation is one where the agents are showing up at a particular location (slot) before a given deadline, and then they are going to be served as close as possible to the slot they arrived at.

The set of slots is identified by the set of integers. It is assumed that slots are equidistant (in time, or in space).

Agents are labeled by letters $A, B, \ldots$, with generic elements i and j. Each agent i has a preferred slot t_i which we refer to as the agent's *target*. We label the agents so that $t_A \leq t_B \leq \ldots$

A *problem of assigning agents to a line* (in short, a *problem*), consists of a finite number of agents each having a target. In general, a problem can be represented by an *agent-target* table, and be depicted graphically accordingly, as in following example.

Example 1: Two agents wish to be served at slot 3, one agent wishes to be served at slot 4, and eight agents wish to be served at slot 6.

Example 1: Table

Agent	A	B	C	D	E	F	G	H	I	J	K
Target	3	3	4	6	6	6	6	6	6	6	6

[4]Kasajima (2012) is actually a recent instance of a study of probabilistic assignment when agents have single-peaked preferences.

Example 1: Figure

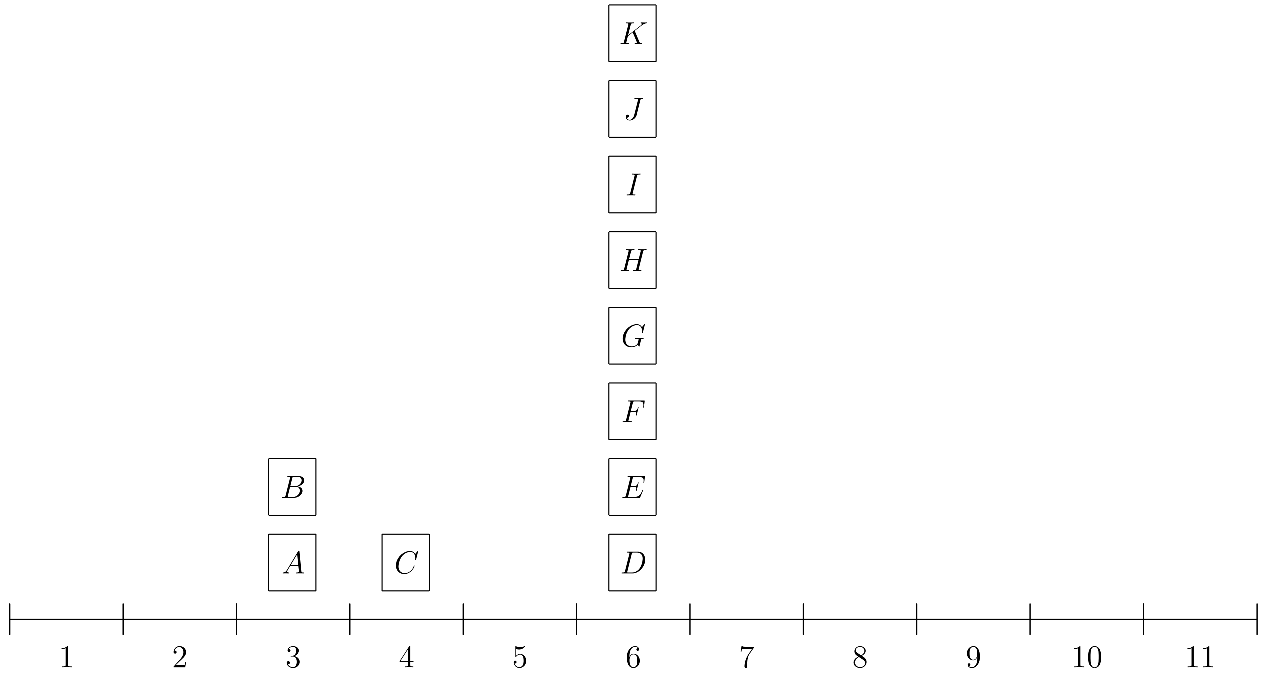

We add two modeling assumptions from the outset. First, *anonymity*, which says that the identity of agents is irrelevant for solving problems. As such, the property can be interpreted as reflecting the principle of impartiality, which excludes ethically irrelevant information (e.g., Moreno-Ternero and Roemer, 2006). Second, *slot invariance*, which says that our methods for solving problems will not depend on how we label the slots. For instance, the problem of Example 1, in which two agents share slot 3 as their target, one other agent has slot 4 as his target, and eight more agents share slot 6 as their target, is solved as the problem in which two agents share slot 9 as their target, one other agent has slot 10 as his target, and eight more agents share slot 12 as their target, except for the corresponding shift in six slots. The two assumptions together permit the use of a concise notation, in which only the number of agents preferring (consecutive) targets are considered. For instance, the problem of Example 1 would be described by the notation $[2, 1, 0, 8]$. Note that 0 is inserted to indicate that there is indeed a slot between the target of one agent and the target of eight agents, which itself is not the target of any agent.

4 Deterministic Assignments

An *allocation* (of a problem) is a deterministic assignment of agents to slots. Denote by x_i the slot assigned to agent i in allocation x. We shall often refer to x_i as agent i's outcome (in allocation x).

For each allocation x of a problem we define the *gap* of agent i as $g_i(x) = |t_i - x_i|$, to be interpreted as the disutility of agent i from being assigned the slot x_i while having target t_i. We assume that only the size of the gap matters, i.e., agents' preferences are symmetric around the target.

Example 1 (cont.): $[2, 1, 0, 8]$. We can display an allocation and the associated gaps by adding additional columns (see below for an example).

agent	target	assignment	gap
A	3	1	2
B	3	2	1
C	4	3	1
D	6	4	2
E	6	5	1
F	6	6	0
G	6	7	1
H	6	8	2
I	6	9	3
J	6	10	4
K	6	11	5

Example 1 (cont.)

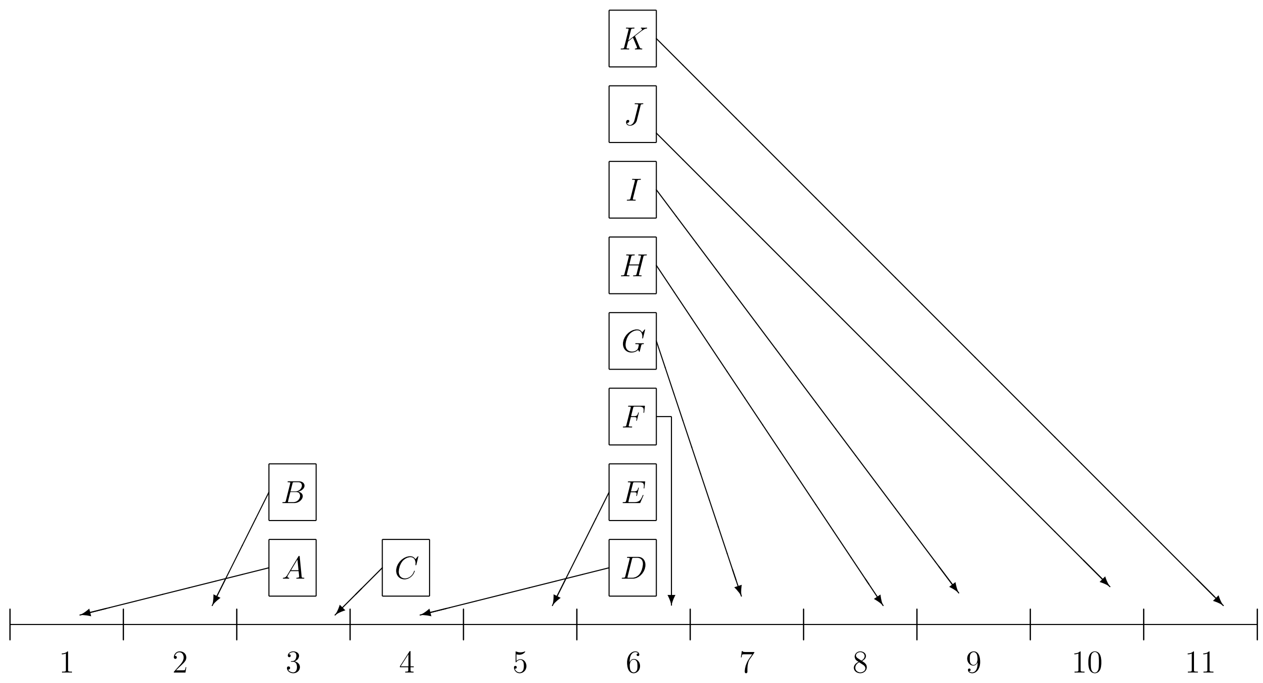

4.1 Efficiency

An allocation x is a *Pareto improvement* of another allocation y, if $g_i(x) \leq g_i(y)$ for each agent i, with at least one strict inequality. An allocation is Pareto efficient if there does not exist a Pareto improvement of it. We present in this section an intuitive algorithm for checking whether a given allocation is Pareto efficient.

The next lemma, which is interesting on its own, provides an intermediate step towards such algorithm. It focusses on reallocations among already occupied slots and states that if a Pareto improvement through some reallocation of outcomes is possible, then it is possible to obtain a Pareto improvement by a bilateral reallocation of outcomes between two agents.[5]

Lemma 1: *Suppose that there exists a Pareto improvement by reallocating outcomes between n agents. Then there exists a Pareto improvement by reallocating the outcomes of two agents.*

[5] This is reminiscent of some results dealing with coalitional manipulation of agents' characteristics in related models of resource allocation, which state that if outcomes can be manipulated, then that can be done with pairwise manipulations (e.g., Ju et al, 2007; Barberá et al., 2010).

Lemma 1 implies that if no Pareto improving pairwise swap of slots can be found, it is impossible to find a Pareto improvement of the current allocation *using the same slots*. It suggests the following procedure for checking efficiency:

Efficiency Check Algorithm

Step 0: Check if any pair of agents can benefit from a swap.

If yes: conclude that the allocation is not efficient; if not: go to Step 1.

Step 1: Check if any unilateral reassignment (to an empty slot) can reduce the gap of an agent.

If yes: conclude that the allocation is not efficient; if not: go to Step 2.

Step 2: Mark with "$*$" all agents that are indifferent between their current outcome and some empty slot. Also mark that associated occupied slot with "$*$". Then, check if any agent prefers a "$*$"-marked slot to his own slot.

If yes: conclude that the allocation is not efficient; if not: go to Step 3.

Step 3: Mark with "$**$" all agents that are indifferent between their current outcome and a "$*$" marked slot. Also mark the associated slot with "$**$". Then, check if any agent prefers a "$**$"-marked slot to his own slot.

If yes: conclude that the allocation is not efficient; if not: go to Step 4.

⋮

Step k: Mark with "$\overbrace{*\cdots*}^{k-1}$" all agents that are indifferent between their current outcome and "$\overbrace{*\cdots*}^{k-2}$" marked slots. Also mark that associated occupied slot with "$\overbrace{*\cdots*}^{k-1}$". Then, check if there is any agent that prefers a "$\overbrace{*\cdots*}^{k-1}$"-marked slot to his own slot.

If yes: conclude that the allocation is not efficient; otherwise go to Step $k+1$.

The algorithm stops when no agents have been "star-marked" in the latest step, or all agents have been "star-marked". In both cases, the assignment is deemed efficient.

The following figure illustrates the Efficiency Check Algorithm for a specific problem: $[2, 0, 2, 1]$. For such problem, in the second step only one star is assigned to slot 3, because only agent A would be indifferent between that slot (assigned to him) and 1, which is empty. No other agent would prefer 3 to their assignment and, thus, we move to the third step. In the third step, two stars are only assigned to slot 5, because only agent C would be indifferent between that slot (assigned to him) and 3, which is marked with one star. Agent E prefers 5 to his assignment and thus we conclude that the assignment is inefficient.

Illustration of the Efficiency Check Algorithm for the problem $[2, 0, 2, 1]$

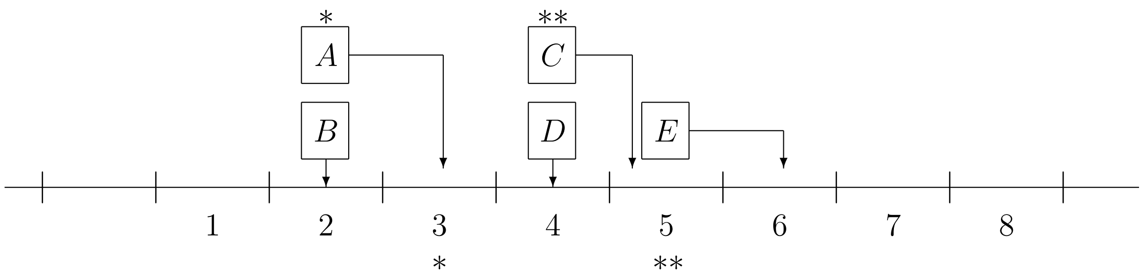

Theorem 1: *The Efficiency Check Algorithm works: it finishes in a finite number of steps and concludes that an allocation is Pareto efficient if and only if this is in fact the case.*

Proof: Note that any agent (and any slot) is "star-marked" at most once. Thus, the algorithm finishes in a finite number of steps. If at some step m some agent i prefers a "star-marked" slot k to his own, a Pareto improvement is possible through the change in outcomes given by letting agent i get his preferred position, letting the agent who currently has agent i's preferred slot change to his indifferent outcome, etc., until the last agent (with a single star) in the sequence moves to an empty slot.

It remains to verify that if a Pareto improvement exists, then it is detected in one of the steps.

If a Pareto improvement exists by reallocating outcomes among n agents, by Lemma 1, there is also a Pareto-improving pairwise swap. Thus, Step 1 yields a "yes" if a Pareto-improvement exists from outcomes that are already used.

If a Pareto improvement exists that involves several previously unused slots, then there exists also a Pareto improvement involving a single previously unused slot. Indeed, we can decompose the Pareto improvement into a sequence of Pareto improvements that each makes use of a single previously unused slot. Thus, it remains to verify that Step 2 or some of the following steps detects a Pareto improvement involving a single previously unused slot. But this follows by the construction of the algorithm. Q.E.D.

As we shall see it in the next section dealing with probabilistic assignment, this procedure is also instrumental to check if a given lottery is *ex-post efficient*, by checking efficiency of each possible realization.

4.2 Utilitarianism

An allocation x for a given problem is *utilitarian* if it minimizes the *aggregate gap* of the agents, $\sum_i g_i(x)$. Similarly, an allocation x for a given problem is *constrained utilitarian* if it minimizes the aggregate gap of the agents, subject to the constraint that only slots assigned in x can be used. As shown by the following example, efficiency does not even imply constrained utilitarianism.

Example 2: $[6, 0, 3]$

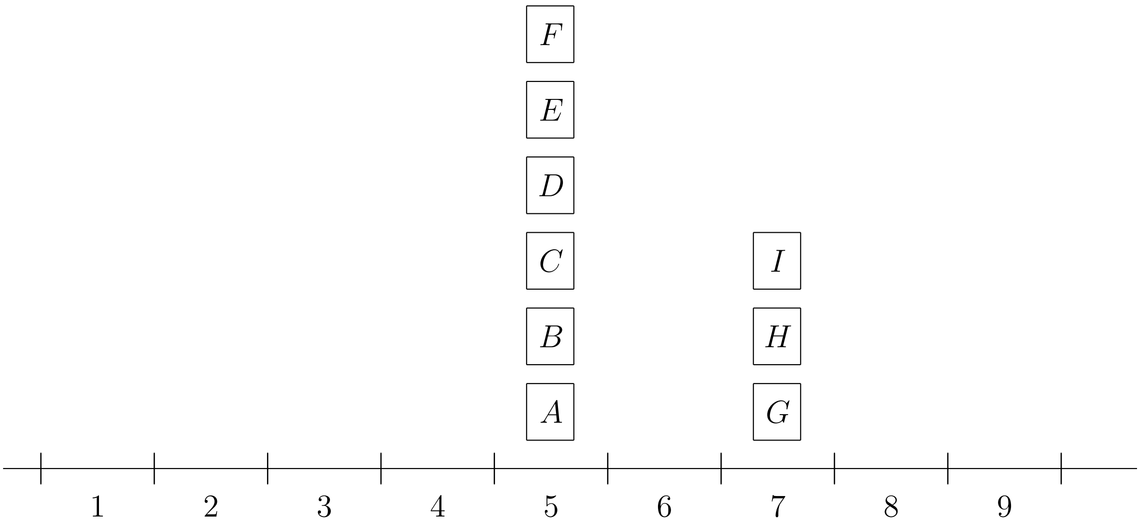

Consider the following two allocations for the problem above:

Example 2 (allocation x)

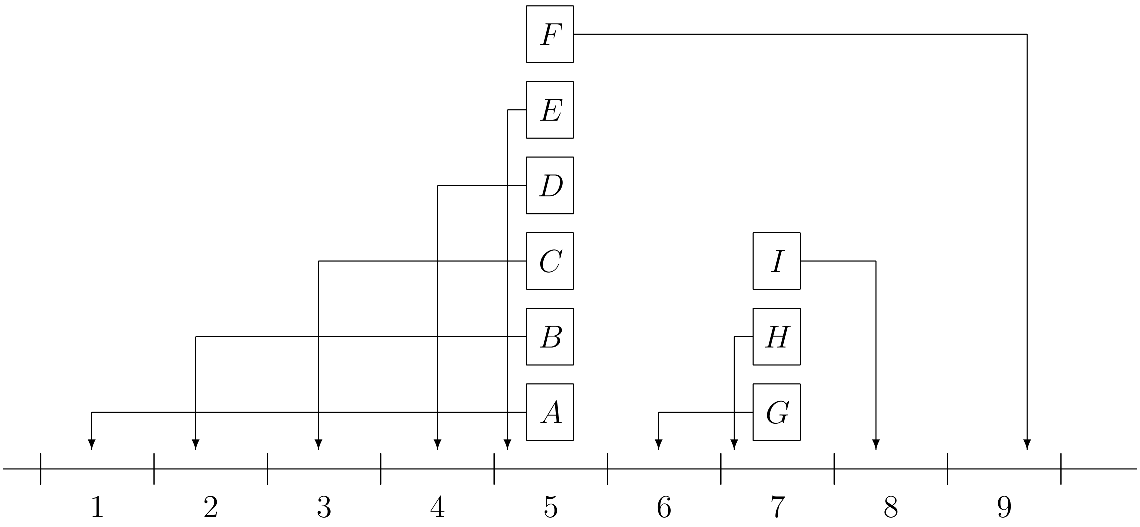

Example 2 (allocation y)

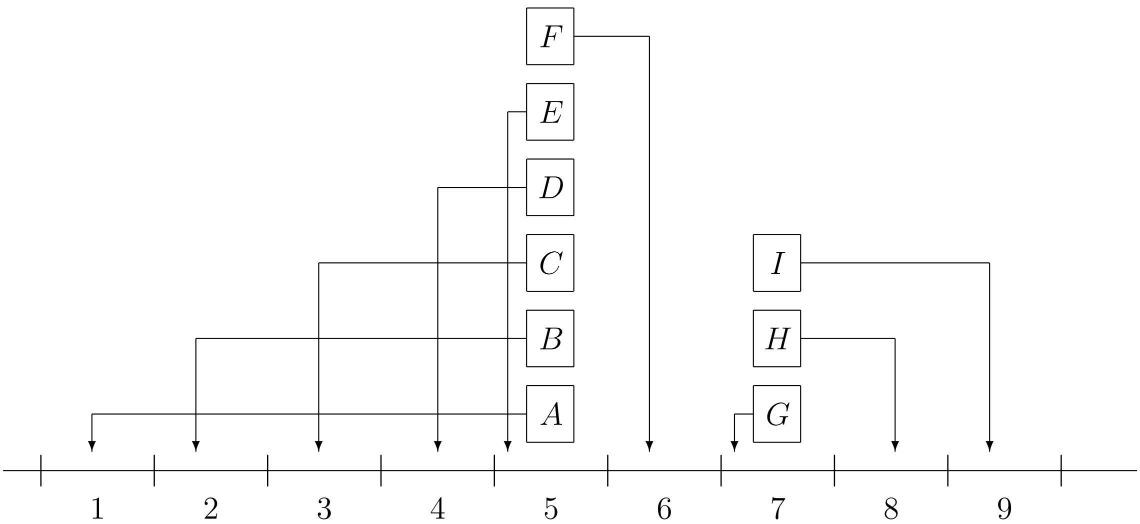

We have $\sum_i g_i(x)$ =16 and $\sum_i g_i(y)$=14 . However both assignments are efficient.

In what follows, we often make use of the following straightforward observations:[6]

[6]These observations follow from applications of general results for assignment problems satisfying the so-called *Monge property*, see Burkard et al. 2012 , p. 150.

For a given allocation x, any other allocation using the same slots can be obtained by a sequence of pairwise swaps (of outcomes). In particular, a constrained utilitarian allocation can be obtained from x by sequentially conducting pairwise swaps where each swap is ordering outcomes like the targets. Note also that each such swap is aggregate gap-reducing or gap-neutral.

Lemma 2: *Let x be an allocation where assigned slots are ordered like targets. Then, the following statements hold:*

1. *x is constrained utilitarian (and thus any pairwise swap of outcomes is either aggregate gap-neutral or aggregate gap-increasing).*
2. *Any constrained utilitarian allocation using the same slots as in x can be obtained from x by a sequence of aggregate gap-neutral pairwise swaps.*

Lemma 3: *A given allocation y is constrained utilitarian if and only if there does not exist a pairwise aggregate gap-reducing swap of outcomes between two agents.*

In order to introduce our next result, we need the following definition. Given an allocation x, a *connected segment* of outcomes is a set of occupied slots with empty neighbor slots in both ends.

The Utilitarian Check Algorithm:

Step 1: Order all outcomes like the targets through some sequence of pairwise swaps (if necessary). If any such pairwise swap is aggregate gap-reducing then conclude that the allocation is not utilitarian. Otherwise, go to Step 2.

Step 2: Given the target ordered allocation produced in Step 1, check all connected segments using the following procedure:

Step 2.A: Label slots in the m-slot connected segment from left to right by $a, b, c, \ldots, m$, and let $A, B, C \ldots, M$ be the agents with outcomes $a, b, c, \ldots, m$ respectively.

Start counting from the left: Define c_i recursively as follows,

$c_1 = -1$ if $t_A \geq a$ and $c_1 = 1$ otherwise

$c_2 = c_1 - 1$ if $t_B \geq b$ and $c_2 = c_1 + 1$ otherwise

$\vdots$

If $c_i > 0$, for some $i \in \{a, \ldots, m\}$, then conclude that the allocation is not utilitarian. Otherwise, go to Step $2.B$.

Step 2.B: Perform a procedure similar to that in Step $2.A$ but this time counting from the right to the left. Formally, define d_i recursively as follows,

$d_1 = -1$ if $t_m \leq m$ and $d_1 = 1$ otherwise

$d_2 = c_1 - 1$ if $t_{m-1} \leq m - 1$ and $d_2 = d_1 + 1$ otherwise.

$\vdots$

If $d_i > 0$, for some $i \in \{a, \ldots, m\}$, then conclude that the allocation is not utilitarian. Otherwise, go to Step 3.

Step 3: Repeat the process described in Step 2 for any other (non-visited) connected segment.

$\vdots$

Step k: If no further connected segments exist, and the process did not terminate in any of the previous steps, conclude the allocation is utilitarian.

Illustration of the Utilitarianism Check Algorithm for the problem $[4, 0, 3]$

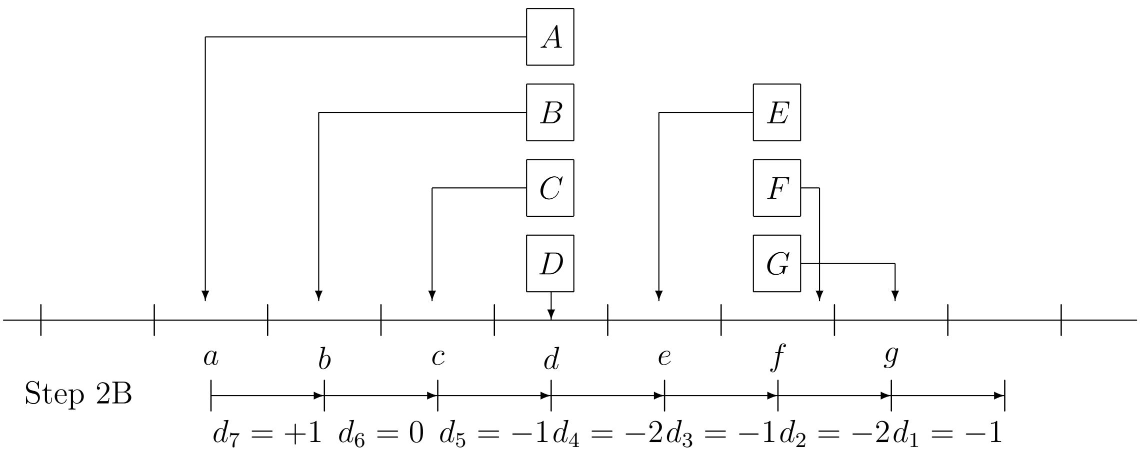

Theorem 2: *The Utilitarian Check Algorithm works: it finishes in a finite number of steps and concludes that an allocation is utilitarian if and only if this is in fact the case.*

Proof: Suppose a given allocation x has survived the algorithm and denote by y the corresponding allocation after ordering outcomes from x as targets by a sequence of pairwise swaps (cf. Step 1).

Our aim is to show that y, and hence x, is utilitarian.

First we observe that, for y, all agents with outcomes in a connected segment have targets within the segment. Indeed, if some agent has his target outside the segment there would be an aggregate gap-reducing shift involving this agent and all agents, either before, or after, in the segment, towards the target. This is not the case, as y survived the algorithm.

Now, by contradiction, assume that there exists an alternative allocation z that reduces the aggregate gap compared to y. If z involves several unused slots compared to y, then there also exists an aggregate gap-reducing allocation z' involving a single unused slot. Indeed, we can go from y to z by a number of sequences of reallocations, where each such sequence involves the use of a single slot unused at y.

Without loss of generality, we can assume that the agent being reallocated to an unused slot moves to a neighbor slot to his segment. Indeed, this can be assumed because in any other case the agent that moves to the new slot could benefit from a unilateral change to such an unused neighbor slot of his own segment, which is closer to his target.

Moreover, without loss of generality, we can further assume that all agents involved in the sequence of reallocations leading from y to z' have targets within the same segment. Otherwise, we could obtain at least the same aggregate gap reduction by eliminating agents with targets outside the segment from the sequence.

Now, order the outcomes in z' like the targets (as in y) by a sequence of pairwise swaps and denote this allocation z''. Note that, as all such swaps are either aggregate gap-neutral or aggregate gap-reducing the allocation obtained is also aggregate gap-reducing compared to y. However, it then follows that y would not survive the algorithm, as the allocation z'' corresponds to a shift of outcomes in the relevant connected segment of y; a contradiction. Q.E.D.

5 Probabilistic Assignments

We now move to focus on random (probabilistic) assignments, partially building onto the analysis of deterministic assignments from the previous section. Formally, a *probabilistic assignment* is a specification of marginal probabilities for each agent over slots. Our convention regarding probabilistic assignments will be to disregard *inactive slots*, i.e., slots used with zero probability.

Furthermore, if the number of *active slots* (i.e., slots assigned to some agents with a positive probability) is greater than the number of agents, we shall introduce *dummy agents* to match the total number of slots. In doing so, a probabilistic assignment will always be represented by a bistochastic matrix.[7]

For a given problem, a *lottery* is a probability distribution over deterministic allocations. It is clear that a lottery induces a probabilistic assignment. Conversely, by the classical Birkhoff-von Neumann decomposition theorem (e.g., Budish et al., 2011), any probabilistic assignment is consistent with some lottery, although, typically, it is not uniquely determined.

For any given probabilistic assignment, let p_i^j be the probability that agent i obtains a gap j. A vector $p_i = (p_i^0, p_i^1, ..., p_i^m)$ is a marginal distribution over gaps if $p_i^j \geq 0$ and $\sum p_i^j = 1$. A marginal distribution $p_i = (p_i^0, p_i^1, ...)$ is (first order) *stochastically dominated* by a marginal distribution $s_i = (s_i^0, s_i^1, ..., s_i^n)$ if $\sum_{h=1}^k p_i^k \leq \sum_{h=1}^k s_i^k$ for each $k \leq \max\{m, n\}$, with the convention that $p_i^h = 0$ if $m < h$ and $s_i^h = 0$ if $n < h$.[8] We say that stochastic dominance is *strict* if the inequalities hold with at least one strict inequality.

A probabilistic assignment P is stochastically dominated by S if, for each agent, the marginal distribution in S stochastically dominates that in P.

In this paper, we endorse the approach in the seminal contribution of Bogomolnaia and Moulin (2001) and assume that each agent cares only about his marginal distribution over outcomes/slots. Thus, agents are indifferent between all lotteries consistent with a given probabilistic assignment, and, therefore, we shall focus on probabilistic assignments.[9]

We also add three basic assumptions over probabilistic assignments that our analysis endorses. First, *equal treatment of equals*, which says that agents with the same target should get the same probability distribution over slots. Second, *slot invariance*, as described in Section 2. Third, *symmetry*, which says that if we "flip" the problem from left to right, the corresponding solution flips too.

We also assume that the only relevant information of an agent is his probability distribution over gaps. In other words, if two probability distributions over slots have the same induced

[7]Note that a deterministic assignment is, formally speaking, a one-to-one correspondence between the set of agents and the set of objects. Thus, it could also be convenient to think of a deterministic assignment as a $0-1$ matrix, with rows indexed by agents and columns indexed by objects, and containing exactly one 1 in each row and each column.

[8]Equivalently, p_i is first order stochastically dominated by s_i if it is possible to obtain p_i from s_i by moving probability mass from larger to smaller gaps.

[9]See Budish et al., (2012) for further discussion on this issue.

probability distribution over gaps, then they are equally desirable for the agent. Hence an agent is indifferent between being served k slots *before* or k slots *after* his target slot. Note that this assumption is much weaker than assuming that only the expected gap matters for an agent.

5.1 Properties of probabilistic assignments

We now consider some notions of probabilistic assignments. Most of them are standard in the literature.

First, we deal with efficiency. A lottery is *ex-post efficient* if any allocation that occurs with positive probability is Pareto efficient. A probabilistic assignment is *ordinally efficient* if it is not stochastically dominated by any other probabilistic assignment.

We then move to utilitarianism. A lottery is utilitarian if every allocation that occurs with positive probability is utilitarian. There are two natural ways of defining that a probabilistic assignment is utilitarian. On the one hand, a probabilistic assignment is *universally utilitarian* if every lottery consistent with it is utilitarian. A probabilistic assignment is *potentially utilitarian* if some lottery consistent with it is utilitarian.

We then consider a classical notion of fairness. A probabilistic assignment is *envy-free* if the probability distribution over gaps for an agent i stochastically dominates the probability distribution over gaps induced from obtaining instead the distribution over outcomes that any other agent j gets. *It is weakly envy-free* if the probability distribution over gaps for an agent i is not strictly stochastically dominated by the probability distribution over gaps induced from obtaining instead the distribution over outcomes that any other agent j gets.

A *solution* is a mapping from problems to probabilistic assignments. We say that a solution satisfies one of the above properties if each probabilistic assignment it gives rise to satisfies that property.

Finally, we introduce a strategic notion. A solution is *strategy-proof* if the probability distribution over gaps for an agent i stochastically dominates the probability distribution over gaps induced from misrepresenting his target. It is *weakly strategy-proof* if the probability distribution over gaps for an agent i is not strictly stochastically dominated by the probability distribution over gaps induced from misrepresenting his target.

5.2 Reasonable probabilistic assignments

We argue in this section that even in extremely simple settings, several reasonable probabilistic allocations might exist. More precisely, consider the following example involving four agents

Example 3: [3,1].

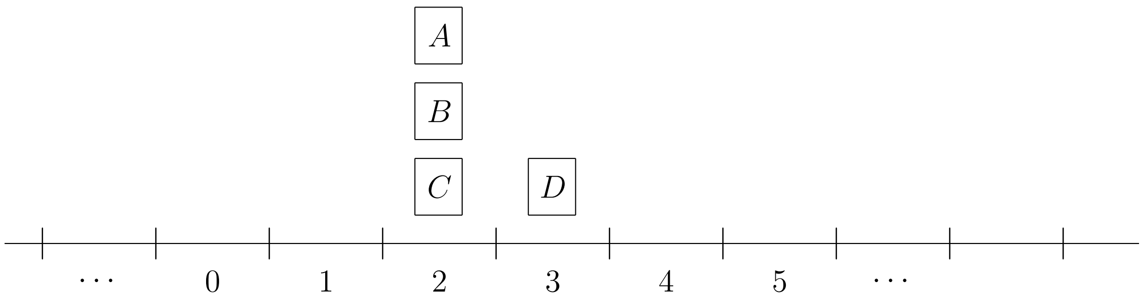

The relevant slots (i.e., those that could be considered in an ex post efficient allocation) are the slots 0 to 4. Slot 0 is, however, never needed: for any allocation where an agent is assigned to slot 0, we can find another allocation involving only slots 1-4 and in which no agent is worse off. In what follows, we therefore focus on (probabilistic) assignments for this problem assigning agents to slots 1-4 by means of a lottery.[10]

We start with a *first-priority-first-served* probabilistic assignment, which gives priority to assigning a slot to those who have it as target. That is, the solution below is the only ordinally efficient solution which holds the property that if a slot is a first priority for someone, then it would always go to someone who has it as first priority.[11]

FIRST	1	2	3	4
A	$\frac{1}{3}$	$\frac{1}{3}$	0	$\frac{1}{3}$
B	$\frac{1}{3}$	$\frac{1}{3}$	0	$\frac{1}{3}$
C	$\frac{1}{3}$	$\frac{1}{3}$	0	$\frac{1}{3}$
D	0	0	1	0

The following is a *minmax-gap* probabilistic assignment that minimizes the maximal possible gap.

[10]Note, however, that in more complex problems there are generally more active slots (i.e., slots that are occupied for at least some ex post efficient allocation) than agents.

[11]Note that, as we have more slots than agents, and therefore each agent must be assigned to some slot (but not vice versa) a probabilistic assignment in our framework is a specification of marginal probabilities for each agent over slots (that sum to 1) and marginal probabilities for each slot over agents (that sum to at most 1).

MINMAX	1	2	3	4
A	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$	0
B	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$	0
C	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$	0
D	0	0	0	1

A third assignment equalizing the expected gap (and minimizing the maximal expected gap) among agents could also be considered.

EQUAL	1	2	3	4
A	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{4}$	$\frac{1}{12}$
B	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{4}$	$\frac{1}{12}$
C	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{4}$	$\frac{1}{12}$
D	0	0	$\frac{1}{4}$	$\frac{3}{4}$

Finally, we consider a probabilistic assignment obtained from a central mechanism in the literature, the so-called *EPS algorithm* (e.g., Bogolmonaia and Moulin, 2001; Katta and Sethuraman, 2006). Such algorithm (adapted to our framework) works as follows.[12]

1. Give each agent a probability share of his most preferred slot, as large as possible, given the constraint that each agent must receive the same probability share. (Thus, the slots that are the most popular targets will be completely shared). Remove those slots that have been completely shared.

2. Give each agent a probability share of his most preferred slots among those not removed, as large as possible, given the constraint that each agent must receive the same total probability share. Remove slots that have been completely shared.

⋮

The algorithm stops when all agents have obtained marginal probabilities over slots (that sum to 1).[13]

It is not difficult to show that the EPS algorithm yields the following probabilistic assignment for the example above:

[12]We refer to Katta and Sethuraman (2006) for a more formal treatment.

[13]In general, each step can be handled by solving an appropriately defined max flow problem, as described by Katta and Sethuraman (2006).

EPS	1	2	3	4
A	$\frac{4}{12}$	$\frac{4}{12}$	$\frac{1}{12}$	$\frac{3}{12}$
B	$\frac{4}{12}$	$\frac{4}{12}$	$\frac{1}{12}$	$\frac{3}{12}$
C	$\frac{4}{12}$	$\frac{4}{12}$	$\frac{1}{12}$	$\frac{3}{12}$
D	0	0	$\frac{9}{12}$	$\frac{3}{12}$

Remark. In each of the above four assignments, the average expected gap is $\frac{3}{4}$, which happens to be the minimum possible value of it for the problem being considered.

5.3 Logical relations among axioms

The above remark raises the question of whether this property holds more generally for all ordinally efficient solutions (in the context of our model). That is, are all ordinally efficient solutions utilitarian? The following example provides a negative answer for such a question.

Example 4: [4,0,3].

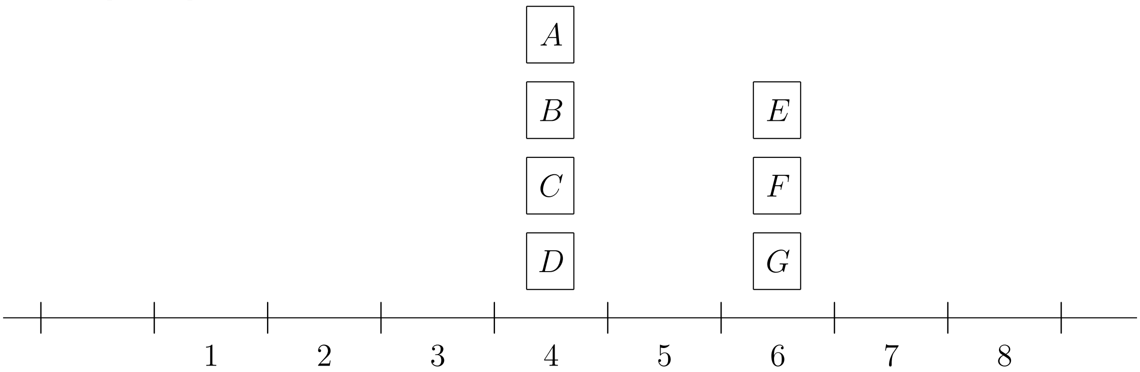

The EPS algorithm yields the following probabilistic assignment for this example:

agent / slot	1	2	3	4	5	6	7	8
A	3/84	1/4	1/4	1/4	9/42			
B	3/84	1/4	1/4	1/4	9/42			
C	3/84	1/4	1/4	1/4	9/42			
D	3/84	1/4	1/4	1/4	9/42			
E					1/21	1/3	1/3	6/21
F					1/21	1/3	1/3	6/21
G					1/21	1/3	1/3	6/21

Hence, with probability 1/7, an allocation occurs that does not minimize aggregate gap. More precisely, the canonical ordinally efficient solution - the EPS solution - is not utilitarian, i.e., it does not minimize the average gap (over all agents) for each possible realization.

The above discussion is summarized in one of the statements of the upcoming proposition which refers to further connections in our setting between the notions of efficiency and utilitarianism defined above.

Proposition 1. The following statements hold in our framework:

- Universal and potential utilitarianism are equivalent.
- Utilitarianism implies ordinal efficiency, but ordinal efficiency does not imply utilitarianism.
- Ordinal efficiency implies ex post efficiency, but ex post efficiency does not imply ordinal efficiency.
- Utilitarianism is incompatible with no-envy.

Proof. In order to prove the first statement, suppose, by contradiction, that there exists a problem and a probabilistic assignment for it, which is *potentially utilitarian*, but not *universally utilitarian.* Then, there is a lottery consistent with the probabilistic assignment such that, with positive probability, an allocation occurs which is not utilitarian. Thus, the expected total gap is then higher than for the utilitarian lottery. However, note that, for each agent, the expected individual gap is the same for all lotteries (and the expected total gap is equal to the sum of individual expected gaps). Thus, all lotteries consistent with a given probabilistic assignment yield the same expected total gap. This represents a contradiction. This proves the first statement. Thus, in what follows, we refer to either potential or universal utilitarianism simply as utilitarianism.

The second part of the second statement has been shown above. As for the first part of it, suppose, by contradiction, that the average expected gap is minimized in an allocation, but that the allocation is not ordinally efficient. If so, some probability mass can be moved towards strictly preferred alternatives for at least one agent. Thus, expected gap cannot be minimized, a contradiction.

The first part of the third statement is straightforward. The second part of it is shown by means of the following example

Example 5: $[3, 3]$

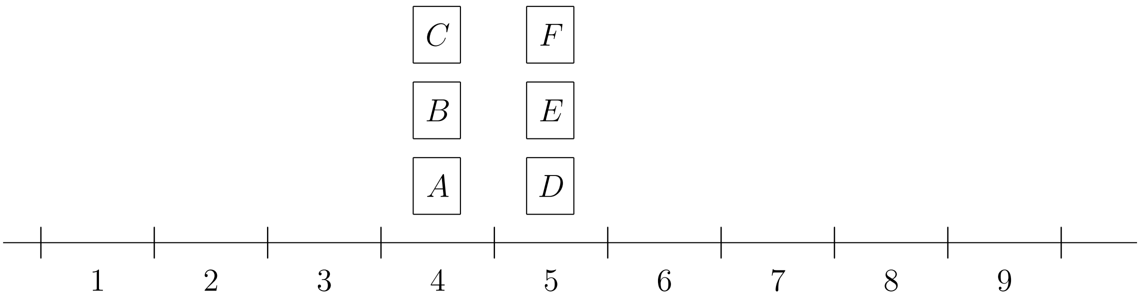

agent	target	assignment 1	assignment 2
A	4	3	1
B	4	4	2
C	4	5	3
D	5	6	4
E	5	7	5
F	5	8	6

Example 5 (assignment 1)

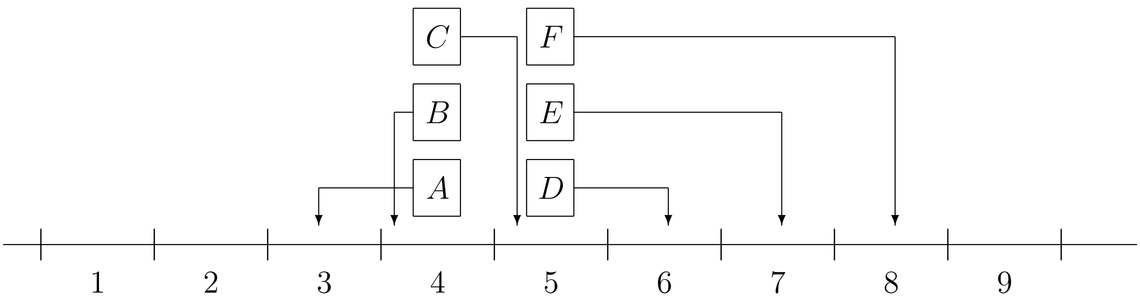

Example 5 (assignment 2)

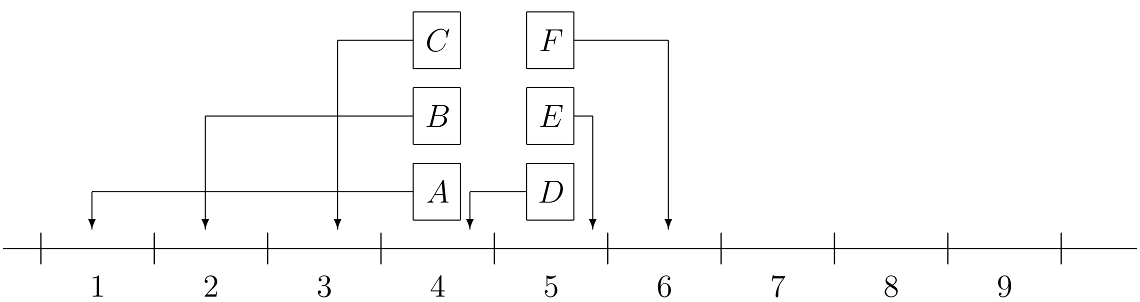

Note that both assignment 1 and assignment 2 are efficient. Moreover, any lottery (satisfying equal treatment of equals) that gives positive probability to both assignments 1 and 2, fails to be ordinally efficient, as agent C gets slot 5 with positive probability and agent D gets slot 4 with positive probability; hence it is possible to swap probabilities between these two agents for slots 4 and 5 making both agents better of ex ante (while keeping the other agents unaffected). This concludes the proof of the third statement of Proposition 1.

Finally, the fourth statement is proved by means of the following example:

Example 6: $[3, 1, 3]$

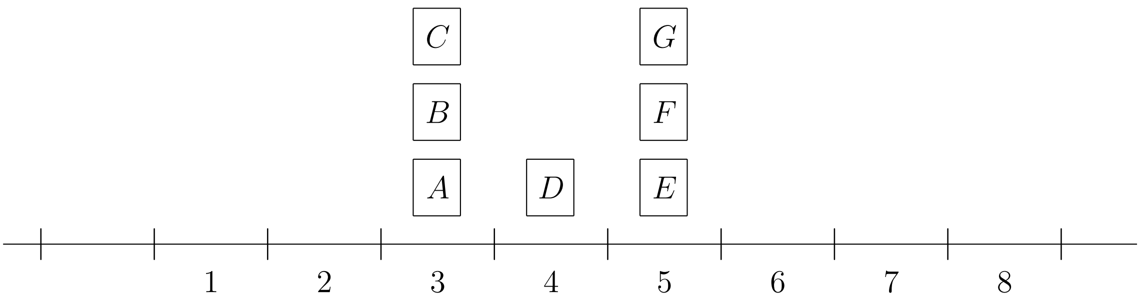

Example 6 (cont.)

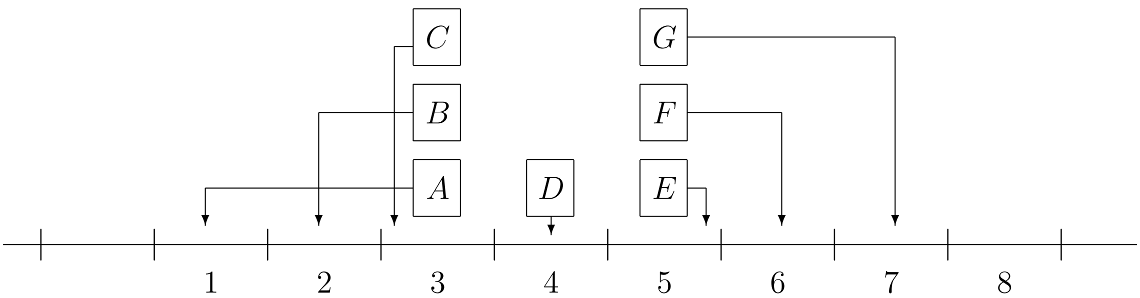

In this problem, a utilitarian solution assigns agents A, B, and C to outcomes 1, 2 and 3, respectively, agent D to 4, and agents E, F and G to 5, 6 and 7, respectively. Note that the other agents do not face a probability distribution of outcomes that first-order (stochastically) dominates agent D's assignment. Thus, the solution fails to be envy-free. Q.E.D.

5.4 A utilitarian solution

Motivated by the observation that the EPS solution (the canonical ordinally efficient solution) fails to be utilitarian, we introduce in this last section a solution which is not only ordinally efficient but also utilitarian.

A simple and intuitive way of assigning agents to slots probabilistically is by means of the random priority (RP) solution: First, pick an agent randomly and assign him to his most preferred slot. Then, pick another agent randomly (among the remaining agents) and assign him to a preferred slot (among those still free), and so on until all agents have been assigned to slots. Note that, in our model, the random priority (RP) solution fails to be even ex post efficient (think for instance of the example [2,1]). However, we will provide a modification which is not only ex post efficient but also utilitarian (and hence also ordinally efficient). Intuitively, the *modified random priority* solution behaves as RP by ordering agents randomly and then assigning them to slots sequentially, but with the important modification that it does so in a way that will ensure that the allocation following each step in the algorithm will be utilitarian among those agents already assigned.

More precisely, the modified RP works as follows: Order agents randomly. Assign the first agent his target. Assign the second agent a most preferred free slot; if there are two such slots pick one of them with equal probability, ... Suppose that $k-1$ agents have been assigned, and consider now the next agent k in the order. If his target is free, assign him the target and go to the next agent. If his target is not free, we consider two allocations constructed as follows:

(a) Assign k to the slot furthest to the right that is occupied by some agent i with target to the left of k and who is assigned to an outcome to the right of his own target (if no such outcome exists, assign k to the first free slot to the left of his target, and go to (b)). Assign agent i to the slot furthest to the right that is occupied by some agent j with target to the left of i and who is assigned to an outcome to the right of his own target (if no such outcome exists, assign j to the first free slot to the left of his target, and go to (b)), etc. until some agent has been assigned the first free slot to the left.

(b) We now consider the symmetric counterpart construction involving agents having target to the right of k's target that stops when an agent has been assigned the first free slot to the right of k's target.

Now, if the two allocations in (a) and (b) respectively give rise to the same aggregate gap among assigned agents $1, 2, ..., k$, choose one of them with equal probability. Otherwise, choose the one with lowest aggregate gap among assigned agents $1, 2, ..., k$.

Formally,

Modified Random Priority algorithm: Let N denote the set of all agents in a given problem.

Step 1: Pick an agent randomly with equal probability among N, and denote him agent A_1. Let $x^1_{A_1} = t_{A_1}$.

$\vdots$

Step k: Suppose that the agents $A_1, ..., A_{k-1}$, have already been chosen randomly in steps 1,..., $k-1$ and assigned to slots $x^{k-1}_{A_1}, ...x^{k-1}_{A_{k-1}}$. Pick an agent randomly from the set $N\backslash\{A_1, ..., A_{k-1}\}$ with equal probability, and denote him A_k.

Let α be the unoccupied slot closest to t_{A_k} for which $\alpha \leq t_{A_k}$ and let γ be the unoccupied slot closest to t_{A_k} for which $t_{A_k} \leq \gamma$.

If $\alpha = \gamma$ let $x^k_{A_k} = t_{A_k}$, and $x^k_i = x^{k-1}_i$ for each $i \in \{A_1, ..., A_{k-1}\}$ and go to step $k+1$.

Otherwise, let $z_{A_k} = \max\{z \mid z = \alpha, \text{ or } x^{k-1}_i = z, t_i < t_{A_k}, t_i < x^{k-1}_i, \text{ for some } i \in \{A_1, ..., A_{k-1}\}\}$. Let l_1 denote the agent in $\{A_1, ..., A_{k-1}\}$ for which $x^{k-1}_{l_1} = z_{A_k}$. Let $z_{l_1} = \max\{z|z = \alpha, \text{ or } x^{k-1}_i = z, t_i < t_{l_1}, t_i < x^{k-1}_i, \text{ for some } i \in \{A_1, ..., A_{k-1}\}\backslash\{l_1\}\}$. Let l_2 denote the agent for which $x^{k-1}_{l_2} = z_{l_1}$, define $z_{l_2} = \max\{z|z = \alpha, \text{ or } x^{k-1}_i = z, t_i < t_{l_1}, t_i < x^{k-1}_i, \text{for some } i \in \{A_1, ..., A_{k-1}\}\backslash\{l_1, l_2\}\}$ etc., until we have defined agents, $l_1, \ldots, l_r$ such that $t_{A_k} > t_{l_1} > t_{l_2} > \ldots$, with $z_{l_r} = \alpha$. Now, define

$$g(l_h) = \begin{cases} z_{l_{h-1}} - z_{l_h}, & \text{if } t_{l_h} \leq z_{l_h} \\ (t_{l_h} - z_{l_h}) - (z_{l_{h-1}} - t_{l_h}), & \text{if } t_{l_h} > z_{l_h}, \end{cases}$$

with the convention that $z_{l_0} = z_{A_k}$. Intuitively, $g(z_{l_h})$ measures the gain (possibly negative) that l_h gets from being reallocated from $z_{l_{h-1}}$ to z_{l_h}. Define $G = \sum_{h=1,...,r} g(l_h)$, with the convention that $G = 0$ if $z_{A_k} = \alpha$.

As a natural mirror-image of the above definitions, we also define the following. Let $q_{A_k} = \min\{z|z = \gamma, \text{or } x^{k-1}_i = z, t_i > t_{A_k}, t_i > x^{k-1}_i, \text{ for some } i \in \{A_1, ..., A_{k-1}\}\}$. Let m_1 denote the agent in $\{A_1, ..., A_{k-1}\}$ for which $x^{k-1}_{m_1} = q_{A_k}$. Let $q_{m_1} = \min\{z|z = \gamma, \text{ or } x^{k-1}_i = z, t_i > t_{m_1}, t_i > x^{k-1}_i, \text{ for some } i \in \{A_1, ..., A_{k-1}\}\backslash\{m_1\}\}$. Let m_2 denote the agent for which $x^{k-1}_{m_2} = q_{m_1}$, define $q_{m_2} = \min\{z|z = \gamma, \text{ or } x^{k-1}_i = z, t_i > t_{m_1}, t_i > x^{k-1}_i, \text{for some } i \in \{A_1, ..., A_{k-1}\}\backslash\{m_1, m_2\}\}$ etc., until we have defined agents, $m_1, ..., m_s$ such that $t_{A_k} < t_{m_1} < t_{m_2} < ..$, with $q_{m_s} = \gamma$. Now, define

$$g(m_h) = \begin{cases} q_{m_h} - q_{m_{h-1}}, & \text{if } z_{m_h} \leq t_{m_h} \\ (q_{m_h} - t_{m_h}) - (q_{m_h} - q_{m_{h-1}}), & \text{if } t_{m_h} < z_{m_h}, \end{cases}$$

with the convention that $q_{m_0} = q_{A_k}$. Intuitively, $g(m_h)$ measures the gain (possibly negative) that m_h gets from being reallocated from $q_{l_{h-1}}$ to q_{l_h}. Define $H = \sum_{h=1,...,s} g(m_h)$, with the convention that $H = 0$ if $z_{A_k} = \gamma$.

Now, we consider three cases:

(1) If $|z_{A_k} - t_{A_k}| - G > |z_{A_k} - t_{A_k}| - H$, let $x^k_{A_k} = z_{A_k}$ and $x^k_{l_h} = z_{l_h}$, for each $h = 1, ...r$, and $x^k_i = x^{k-1}_i$ for $i \in \{A_1, ..., A_{k-1}\} \backslash \{l_1, ..., l_r\}$.

(2) If $|z_{A_k} - t_{A_k}| - G < |z_{A_k} - t_{A_k}| - H$, let $x^k_{A_k} = q_{A_k}$ and $x^k_{l_h} = q_{m_h}$, for each $h = 1, ...s$, and $x^k_i = x^{k-1}_i$ for $i \in \{A_1, ..., A_{k-1}\} \backslash \{m_1, ..., m_s\}$.

(3) If $|z_{A_k} - t_{A_k}| - G = |z_{A_k} - t_{A_k}| - H$ choose either the allocation in (1) or (2) with equal probability.

Go to step $k + 1$.

Stop the algorithm when all agents have been assigned to an outcome.

The next figure illustrates a situation where $k = 14$, i.e., 13 agents have been assigned to slots by way of the algorithm. A new agent, H, with target at 11, has to be assigned to a slot. There are two options, either assign him to target 11 and push two agents to the left, or assign him to slot 12 and push an agent to the right. The arrows below the lines show how they are going to be pushed according to the algorithm and the associated gain/loss from doing so. Clearly, the outcome in this example will be that the agent is assigned to slot 11.

Illustration of the Modified RP algorithm

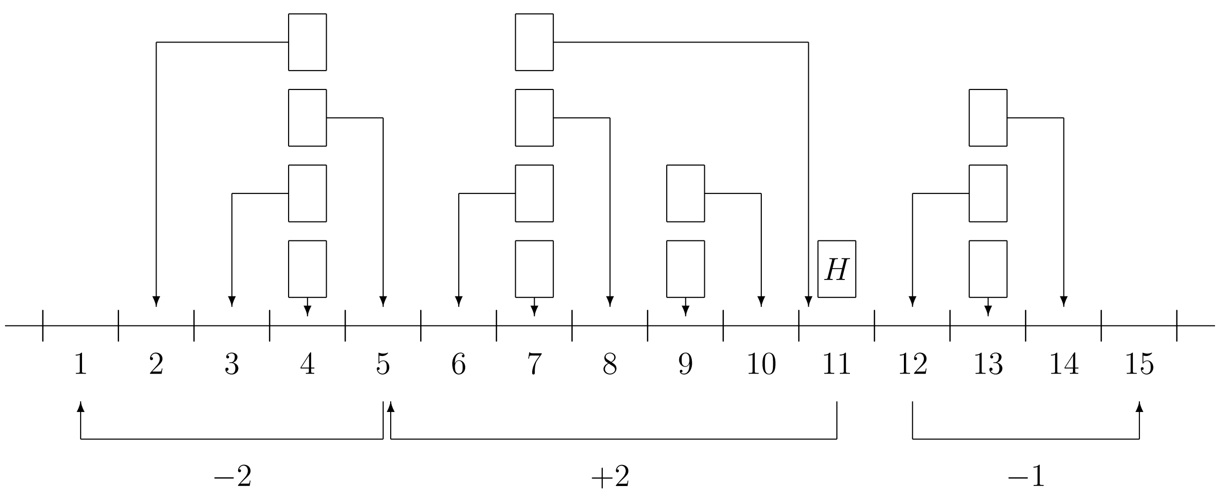

It is clear from the construction of the modified RP algorithm that it satisfies equal treatment of equals, left-right symmetry and slot-invariance. As desired, it is also utilitarian:

Theorem 3: *The Modified RP solution is utilitarian.*

Proof: We need to show that any given allocation $x = (x_1, \ldots, x_n)$ obtained from applying the Modified RP algorithm is utilitarian.

Note first that allocation $x^1_{A_1}(= t_{A_1})$ is utilitarian for the single-agent problem consisting of one agent A_1 with target t_{A_1}.

Next, suppose that $x^{k-1}_{A_1}, ..., x^{k-1}_{A_{k-1}}$ obtained in Step $k-1$ is utilitarian with respect to the set of agents $\{A_1, ...A_{k-1}\}$ (with targets $t_{A_1}, ..., t_{A_{k-1}}$ respectively). We aim to show that then $x^k_{A_1}, ..., x^k_{A_k}$ obtained in Step k is utilitarian with respect to the set of agents $\{A_1, ...A_k\}$ (with targets $t_{A_1}, ..., t_{A_k}$ respectively).

For this, we first argue that $x^k_{A_1}, ..., x^k_{A_k}$ is constrained utilitarian. Indeed, suppose that $x^k_{A_1}, ..., x^k_{A_k}$ fails to be constrained utilitarian. Then by application of Lemma 3 there are agents i, j in $\{A_1, ...A_k\}$ such that $t_i < t_j$, $x_j < x_i, t_i < x_i$ and $x_j < t_j$. However, by construction of step k in the algorithm this cannot happen. Thus, $x^k_{A_1}, ..., x^k_{A_k}$ is constrained utilitarian.

Next, we observe that if a utilitarian allocation can be obtained from using the slots already occupied in $x^{k-1}_{A_1}, ..., x^{k-1}_{A_{k-1}}$ and then either the slot α or the slot γ, the allocation $x^{k-1}_{A_1}, ..., x^k_{A_k}$ is utilitarian. For this note that if a utilitarian allocation can be obtain from using the slots already occupied in $x^{k-1}_{A_1}, ..., x^{k-1}_{A_{k-1}}$ and the slot α, we have $|z_{A_k} - t_{A_k}| - G > |z_{A_k} - t_{A_k}| - H$ and as $x^k_{A_1}, ..., x^k_{A_k}$ is constrained utilitarian, it is also utilitarian. Likewise, if a utilitarian allocation can be obtain from using the slots already occupied in $x^{k-1}_{A_1}, ..., x^{k-1}_{A_{k-1}}$ and the slot γ, we have $|z_{A_k} - t_{A_k}| - G < |z_{A_k} - t_{A_k}| - H$ and as $x^k_{A_1}, ..., x^k_{A_k}$ and is constrained utilitarian, it is also utilitarian. Finally, if a utilitarian allocation can be obtain from using the slots already occupied in $x^{k-1}_{A_1}, ..., x^{k-1}_{A_{k-1}}$ and either the slot α or γ, we have $|z_{A_k} - t_{A_k}| - G = |z_{A_k} - t_{A_k}| - H$ and as $x^k_{A_1}, ..., x^k_{A_k}$ and is constrained utilitarian, it is also utilitarian.

It remains to argue that a utilitarian allocation can be obtained from using the slots already occupied in $x^{k-1}_{A_1}, ..., x^{k-1}_{A_{k-1}}$ and then either the slot α or the slot γ.

For this, let $y_{A_1}, ..., y_{A_k}$ be a utilitarian allocation and consider two cases:

Case 1. y_{A_k} is not one of the slots already occupied in $x^{k-1}_{A_1}, ..., x^{k-1}_{A_{k-1}}$. Obviously, we then have that $(x^{k-1}_{A_1}, ..., x^{k-1}_{A_{k-1}}, y_A)$ is a utilitarian, i.e., a utilitarian allocation can be obtained from using the slots already occupied in $x^{k-1}_{A_1}, ..., x^{k-1}_{A_{k-1}}$ and then some other unoccupied slot. But then clearly this slot is either α or γ and we are done.

Case 2. y_{A_k} is one of the slots already occupied in $x^{k-1}_{A_1}, ..., x^{k-1}_{A_{k-1}}$. Then, we can obtain $y_{A_1}, ..., y_{A_k}$ from $x^{k-1}_{A_1}, ..., x^{k-1}_{A_{k-1}}$ by a sequence of interchanges as follows: The first interchange assigns agent A_k to y_{A_k}, then agent i for which $x^{k-1}_i = y_{A_k}$ gets y_i instead etc, until some agents j gets a slot y_j that is unoccupied at $x^{k-1}_{A_1}, ..., x^{k-1}_{A_{k-1}}$. Next, if the final allocation obtained is not

$y_{A_1}, ..., y_{A_k}$, pick an agent h which does not have y_h yet, move the agent k who currently has that slot to the slot y_k, etc., and continue with such sequences of interchanges until the final allocation is $y_{A_1}, ..., y_{A_k}$ obtained. Note that all interchanges except for the first sequence could have been carried out given the allocation $x^{k-1}_{A_1}, ..., x^{k-1}_{A_{k-1}}$ with the same gain in total gap. This contradicts that $x^{k-1}_{A_1}, ..., x^{k-1}_{A_{k-1}}$ is utilitarian, and we are done. Q.E.D.

Example 4: (continued). The modified RP solution gives the following probabilistic assignment:

agent / slot	1	2	3	4	5	6	7	8
A		1/4	1/4	1/4	1/4			
B		1/4	1/4	1/4	1/4			
C		1/4	1/4	1/4	1/4			
D		1/4	1/4	1/4	1/4			
E						1/3	1/3	1/3
F						1/3	1/3	1/3
G						1/3	1/3	1/3

Hence, it never happens that an allocation occurs that does not minimize aggregate gap.

6 Appendix. Proof of the lemmata

Proof of Lemma 1: Let x be an assignment for a problem involving $n \geq 3$. Suppose that these agents can obtain a Pareto improvement by reallocating outcomes internally and leading to an assignment y. We claim that it is possible to obtain a Pareto improvement by reallocating slots among at most $n-1$ agents. The statement of the lemma then follows by an induction argument.

To prove that claim, we consider four (mutually exclusive and comprehensive) cases. For ease of exposition, let A be an agent such that $t_A \leq t_i$ for each other agent i in the group.

1. Suppose that $t_A \leq x_A < x_i$ for each $i \neq A$. Then there exists a Pareto improvement among the $n-1$ agents (excluding agent A).

Suppose first that no other agent shares target with A.[14] Then, A's gap increases if reassigned to any other outcome x_i, $i \neq A$. Thus, any Pareto improvement involving the n agents

[14] It is straightforward to adapt the argument to cases where more than one agent share the leftmost target.

assigns A to his original outcome x_A, and therefore it is possible to obtain a Pareto improvement among the $n-1$ other agents.

2. Suppose that $x_A < t_A < x_i$ for each $i \neq A$. Then there exists a Pareto improvement among the $n-1$ agents (excluding agent A).

Suppose first that no other agent shares target with A. Let i be the agent for which $y_i = x_A$. We can assume that $i \neq A$ (as if $A = i$ it is clear that a Pareto improvement can be found from x involving the remaining agents only).

Now, let z denote the allocation obtained from taking y and then switching the outcomes of A and i, i.e., $z_i = y_A$, $z_A = y_i = x_A$ and $z_j = y_j$ otherwise. As $t_A < t_i$, $t_A < y_A$ and $y_A < x_i$ then i's gap decreases going from y to z and also weakly decreases going from x to y and thereby also going from x to z. As we also have that $z_A = x_A$, z is a Pareto improvement of x. Thus, as z involves only a reallocation of outcomes for agents other than A (relative to x), we have shown the statement.

Assume now that more than one agent share the (leftmost) target with A. Let i be such that $y_i = x_A$. Clearly, we can assume that such an agent exists as, otherwise, a Pareto improvement can be made by reallocating outcomes among the agents not having target leftmost among all agents (and we are done). Now, the same previous argument applies.

3. Suppose that $t_A < x_A$ and that there exists an agent i for which $x_i < x_A$. Then there exists an agent $j \neq A$ such that A and j can obtain a Pareto improving pairwise swap in x.

Suppose first that no other agent shares target with A. Let $k = |t_A - x_A|$. As $t_A - k \leq y_A < x_A$, there exists $j \neq A$ such that $t_A - k \leq x_j < x_A$ and $x_A \leq y_j$.[15]

We consider two sub-cases:

3a. *Suppose that $x_A \leq t_j$.*

Then we have $|t_A - x_j| \leq |t_A - x_A|$ and $|t_j - x_j| < |t_j - x_A|$. Thus, A and j can obtain a Pareto improving pairwise swap in x.

3b. Suppose that $t_j < x_A$.

First notice that $|t_A - x_j| \leq |t_A - x_A|$ and as $t_A < t_i < x_A$ we have $|t_i - x_A| \leq |t_i - y_j|$. Consequently, $|t_j - x_A| \leq |t_j - x_j|$ (as y is a Pareto improvement of x). As $t_A \neq t_j$ either A or j is strictly better off at the other's outcome, and thus A and j can obtain a Pareto improving pairwise swap.

[15]Note that $y_j < t_A - k$ is not possible as we have $t_A < t_j$ implying that j would be worse off at y.

It remains to show that the above arguments apply also to cases where more than one agent share a target with A. It is straightforward for the case of 3a. In case of 3b, if $t_A < t_j$, by the same argument as above, we get that A and j can obtain a Pareto improving pairwise swap. Thus, assume that $t_j = t_A$ and that there does not exist an agent j' for which $t_A < t_{j'}$ and $t_A - k \leq x_{j'} \leq x_A,\ x_A \leq y_{j'}$. As $t_A - k \leq x_j \leq x_A,\ x_A \leq y_j$ we must have that $y_j = x_A$ and $x_j = t_A - k$. In particular, *any* agent j'' for which $t_A - k \leq x_{j''} \leq x_A$ and $y_{j''} \geq x_A$ must have $x_{j''} = t_A - k$ and $y_{j''} = x_A$. Thus, y is obtained from x by an interchange of outcomes between two agents with the same target. Consequently, it is possible to obtain a Pareto improvement with fewer than n agents, as desired.

4. Suppose that $x_A < t_A$ and that there exists an agent i for which $x_i \leq t_A$. Then there exists a Pareto improvement of x involving only $n-1$ agents (excluding agent A).

Suppose first that no other agent shares target with A. Take allocation y and let i be the agent for which $y_i = x_A$. Now, consider the allocation z for which $z_A = x_A$, $z_i = y_A$ and $z_j = y_j$ otherwise. As $t_i > t_A$ and $x_A < t_A$ we get that $|t_i - z_i| < |t_i - y_i|$ (and $|t_i - z_i| < |t_i - x_i|$). Thus, z is a Pareto improvement of x that involves only a reallocation of slots between $n-1$ agents (excluding agent A).

Suppose now that agent i (as defined in the previous paragraph) shares target with agent A. Then both A and i are equally well off in z as in y.[16] Thus, again, z is a Pareto improvement of x that involves at most $n-1$ agents (excluding agent A). Q.E.D.

Proof of Lemma 2: Suppose y is a constrained utilitarian allocation using the same slots as in x. Then, by a sequence of pairwise swaps, where each is aggregate gap-neutral, an allocation y' can be obtained where slots are ordered like the targets (indeed if two outcomes are not ordered like the targets a pairwise swap is either gap-neutral or gap-reducing and as y is constrained utilitarian it cannot be gap-reducing). Note that y' and x are then identical allocations up to pairwise swaps between agents with shared targets, and hence the first two statements follow. Q.E.D.

Proof of Lemma 3: Clearly, if y is constrained utilitarian there does not exist an aggregate gap-reducing pairwise swap. Thus, we prove the converse implication.

Suppose that there does not exist an aggregate gap-reducing pairwise swap and, by contradiction, that the allocation is not constrained utilitarian.

[16] If i has his target to the right of A, the same argument as in the previous paragraph applies.

By the observation prior to Lemma 2, there exists a sequence of pairwise swaps, each ordering outcomes according to targets, leading to gap reduction. This sequence only contains aggregate gap-neutral or aggregate gap-reducing pairwise swaps. By assumption this sequence starts with one or more aggregate gap-neutral pairwise swaps. It remains to be shown that whenever such a sequence exists, an aggregate gap-reducing pairwise swap exists from the initial allocation.

First, we claim that no aggregate gap-neutral pairwise swap between two agents A and B can lead to a subsequent aggregate gap-reducing pairwise swap between any of these two agents, and a third agent C, unless such an aggregate gap-reducing pairwise swap was possible from the initial allocation. Indeed, consider three agents A, B, and C with targets t_A, t_B and t_C, where a swap between the outcomes of agents A and B is aggregate gap-neutral. As the case $t_A = t_B$ is trivial, we will assume, without loss of generality, that $t_A < t_B$.

Consider an allocation where A and B have outcome either x_k or x_h, where $x_k < x_h$, and C has outcome x_C. Note that, as we have assumed that a swap between the outcomes of agents A and B is aggregate gap-neutral, we have either $t_B \leq x_h, x_k$ or $x_h, x_k \leq t_A$. As the two cases can be treated in a symmetric way, we assume for the rest of the proof that $t_B \leq x_h, x_k$. We consider four cases:

Case 1. $t_C \leq t_B$. In this case, $x_C \leq t_B$, as otherwise any sequence of pairwise swaps is aggregate gap-neutral. If $t_C \leq t_A$, then either outcomes are ordered like the targets (in which case no sequence of pairwise swaps reduce the aggregate gap), or the swap between A and B orders outcomes like targets (in which case it is not possible to reduce the aggregate gap from a subsequent swap involving C). If $t_A < t_C \leq t_B$, then $t_C \leq x_C \leq t_B$ (otherwise there exists an aggregate gap-reducing pairwise swap between A and C from the beginning). Thus, a swap between A and C is aggregate gap-neutral regardless of whether A has swap with B or not and a swap between C and B can only increase the aggregate gap regardless of a swap between A and B.

Case 2. $t_B < t_C \leq x_k$. If $x_C \geq t_C$ then all swaps are aggregate gap-neutral. If $x_C < t_C$ then C can make a gap-reducing swap with either A or B from the beginning.

Case 3. $x_k < t_C \leq x_h$. If $x_C \geq x_h$ then either outcomes are ordered like the targets or the swap between A and B orders outcomes like targets. If $x_C < x_k$ then there exists an aggregate gap reducing pairwise swap between A and C from the beginning. If $x_k \leq x_C < x_h$ then

i) if $x_k \leq x_C < t_C$ there exists an aggregate gap-reducing swap initially by swapping with the agent assigned x_h, and

ii) if $t_C \geq x_C < t_h$ the allocation is efficient.

Case 4. $t_C > x_h$.Then either there exists an aggregate gap-reducing swap from the beginning or the swap is weakly increasing the gap.

Now, as the sequence of an aggregate gap-neutral pairwise swap followed by an aggregate gap-reducing pairwise swap cannot exist, unless there exists an aggregate gap-reducing pairwise swap from the beginning, we get, by an induction argument, that there cannot exist any finite sequence of aggregate gap-neutral switches followed by an aggregate gap-reducing swap unless there exists an aggregate gap-reducing pairwise swap from the beginning.

Q.E.D.

References

[1] Abdulkadiroglu A, Sönmez T, 1998. Random serial dictatorship and the core from random endowments in house allocation problems. Econometrica 66, 689-701.

[2] Abdulkadiroglu A, Sönmez T, 2003. Ordinal efficiency and dominated sets. Journal of Economic Theory, 112, 157-172.

[3] Alkan, A., G. Demange, and D. Gale (1991). Fair Allocation of Indivisible Goods and Criteria of Justice. Econometrica, 59, 1023-1059.

[4] Athanassoglou, S., and Sethuraman, J. House Allocation with Fractional Endowments. International Journal of Game Theory, (2010), forthcoming.

[5] Bapat RB, Raghavan, TES, 1997. Non-negative Matrices and Applications, Cambridge University Press, New York

[6] S. Barberá, D. Berga, B. Moreno (2010). Individual versus group strategy-proofness: When do they coincide? Journal of Economic Theory 145, 1648-1674

[7] Bogomolnaia, A., and Heo, E.-J., (2012). Probabilistic Assignment of Objects: Characterizing the Serial Rule. Journal of Economic Theory, Forthcoming

[8] Bogomolnaia A, Moulin H., 2001. A new solution to the random assignment problem. Journal of Economic Theory, 100, 295-328.

[9] Bogomolnaia, A., and Moulin, H., (2002). A Simple Random Assignment Problem with a Unique Solution. Economic Theory, 19, 623-635.

[10] Bogomolnaia A, Moulin H., 2004. Random Matching Under Dichotomous Preferences. Econometrica, 72, 257-279.

[11] Budish, E., and E. Cantillon (2010). The Multi-unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard. American Economic Review. Forthcoming.

[12] Budish, E., Y.-K. Che, F. Kojima, and P. Milgrom (2012). Implementing Random Assignments: A Generalization of the Birkhoff-von Neumann Theorem. American Economic Review. Forthcoming.

[13] Burkard R, Dell'Amico M, Martello S, 2012. Assignment Problems - Revised Reprint. Society for Industrial and Applied Mathematics, Philadelphia.

[14] Chambers, C. (2004). Consistency in the Probabilistic Assignment Model. Journal of Mathematical Economics, 40, 953-962.

[15] Che, Y.-K., and F. Kojima (2010). Asymptotic Equivalence of Random Priority and Probabilistic Serial Mechanisms. Econometrica, 78, 1625-1672.

[16] Cres H., and H. Moulin, (2002). Scheduling with opting out: Improving upon random priority. Operations Research, 49, 565-577

[17] Dolan, R., (1978). Incentive mechanisms for priority queueing problems. Bell Journal of Economics, 9, 421-436.

[18] Foley, D. (1967). Resource Allocation and the Public Sector. Yale Economic Essays, 7, 45-98.

[19] Gale, D., and L. S. Shapley (1962). College Admissions and the Stability of Marriage. American Mathematical Monthly, 69, 915.

[20] Gibbard, A., (1978). Straightforwardness of game forms with lotteries as outcomes. Econometrica, 46, 595-602.

[21] Hylland, A., and Zeckhauser, R. (1979). The Efficient Allocation of Individuals to Positions. Journal of Political Economy, 87, 293-314.

[22] Ju, B.-G., Miyagawa E., Sakai T. (2007). Non-manipulable division rules in claim problems and generalizations. Journal of Economic Theory, 132, 1-26.

[23] Kasajima, Y., (2012). Probabilistic assignment of indivisible goods with single-peaked preferences. Social Choice and Welfare. Forthcoming.

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

[25] Kesten, O. (2009). Why Do Popular Mechanisms Lack Efficiency in Random Environments? Journal of Economic Theory, 144, 2209-2226.

[26] Kesten O, Kurino M, Ünver MU, (2011). Fair and Efficient Assignment via the Probabilistic Serial Mechanism. Working Paper.

[27] Kojima, F., (2009). Random Assignment of Multiple Indivisible Objects. Mathematical Social Sciences, 57, 134-142.

[28] Kojima, F., and M. Manea (2010). Incentives in the Probabilistic Serial Mechanism. Journal of Economic Theory, 144, 106-123.

[29] Manea, M. (2008). A Constructive Proof of the Ordinal Efficiency Welfare Theorem. Journal of Economic Theory, 141, 276-281.

[30] Manea, M. (2009). Asymptotic Ordinal Inefficiency of Random Serial Dictatorship. Theoretical Economics, 4, 165-197.

[31] McLennan, A. (2002). Ordinal Efficiency and The Polyhedral Separating Hyperplane Theorem. Journal of Economic Theory, 105, 435-449.

[32] Moreno-Ternero J, Roemer J (2006). Impartiality, solidarity, and priority in the theory of justice. Econometrica, 74, 1419-1427.

[33] A. Roth and M. Sotomayor, (1990). Two Sided Matching: A Study in Game-Theoretic Modeling and Analysis. Econometric Society Monographs, Cambridge University Press, Cambridge.

[34] Sprumont, Y., (1991). The division problem with single-peaked preferences: a characterization of the uniform rule. Econometrica, 59, 509-519.

[35] H. Varian, (1974). Equity, envy and efficiency. Journal of Economic Theory, 9, 63-91.

[36] Yilmaz, O., (2009). Random Assignment Under Weak Preferences. Games and Economic Behavior, 66, 546-558.

[37] Zhou, L., (1990). On a Conjecture by Gale About One-sided Matching Problems. Journal of Economic Theory, 52, 123-135.