Scheduling with opting out: Improving upon Random Priority[*]

Hervé Crès[y] and Hervé Moulin[z]

February 28, 2000

Abstract

In a scheduling problem where agents can opt out, we show that the familiar Random Priority (RP) mechanism can be improved upon by another mechanism dubbed Probabilistic Serial (PS). Both mechanisms are nonmanipulable in a strong sense, but the latter is Pareto superior to the former and serves a larger (expected) number of agents. The PS equilibrium outcome is easier to compute than the RP outcome; on the other hand RP is easier to implement than PS. We show that the improvement of PS over RP is significant but small: at most a couple of percentage points in the relative welfare gain and the relative difference in quantity served. Both gains vanish when the number of agents is large; hence both mechanisms can be used as a proxy of each other.

[*]We are very grateful to two anonymous referees and an associate editor for their critical comments on an earlier version of the paper. We thank seminar and conference participants in Paris, Caen, Jouy-en-Josas, Barcelona, Aix-en-Provence, Montréal and Vancouver. Stimulating conversations with Anna Bogomolnaia, Eric Friedman, Scott Shenker and Rakesh Vohra are gratefully acknowledged. Special thanks to Vianney Dequiedt and Liqan Wang, who developed the numerical computations. In the preparation of this work Moulin was supported by an NSF grant SBR 9809316.

[y]HEC School of Management, 78351 Jouy-en Josas, France; Tel: 33 1 39 67 94 12; Fax: 33 1 39 67 70 85; cres@hec.fr.

[z]Department of Economics, Rice University, 6100 Main Street, Houston TX 77005/1892; moulin@rice.edu.

1 Introduction

1.1 Probabilistic scheduling

The paper addresses a simple scheduling problem. There is a batch arrival of agents, each of whom have a job requiring one unit of time. The planner/manager controls a server processing one job per unit of time. All agents prefer early service but are heterogeneous in their 'type' —i.e., the number of time periods they can afford to wait for service. If an agent anticipates a wait longer than this limit, he immediately opts out and leaves the system. The manager uses a non-price mechanism to schedule the agents; the only information he can use is the type of each agent.

Fix a deterministic priority ordering σ (e.g., alphabetical order) of the agents and, following this ordering, let them successively choose either to stay in line —and be served at the best non-assigned date— or to opt out. The *Priority mechanism* (denoted Prio(σ) throughout the paper) is played in a very simple way: the first agent in the ordering who is scheduled to be processed at a time period beyond his type opts out first. Once he does so, all later-to-be-processed jobs (agents) improve their scheduled time by one period. Then the next agent who, based on the improved schedule, is scheduled to be processed at a time period beyond his type, opts out, and so on. Hence the q-th agent in line faces a wait of k time periods if k ¡ 1 jobs before him chose to stay in line and q ¡ k opted out, and decides to stay or opt out according to his type. Such a deterministic priority mechanism has good incentive compatibility properties (agents have no interest to misreport their type) but is unfair.

A simple and natural way to restore fairness is the *Random Priority* (RP) mechanism: the planner selects *at random and without bias* a certain priority ordering σ of the agents (among the n! possible orderings if there are n agents), then priority mechanism Prio(σ) is played as described in the previous paragraph. A first contribution of the present paper is to provide a recursive algorithm computing the outcome of RP, i.e., the expected assignment of agents to time slots, called the RP equilibrium assignment in the sequel[1].

A second contribution of the paper is to propose another scheduling protocol dubbed *Probabilistic Serial* (PS). It resembles Random Priority closely, in particular shares its properties of incentive compatibility (strategyproofness) and fairness. The advantage of Probabilistic Serial over Random Priority is twofold: (1) from an efficiency point of view PS always improves upon RP welfarewise, in the strong sense of the Pareto ranking (no

[1] Or even in short RP assignment.

agent is worse off and some agents are strictly better off in PS than in RP) and moreover, it always serves a larger expected number of agents than RP; (2) from a computational point of view the outcome of PS, i.e., the expected assignment of agents to time slots (called the PS equilibrium assignment in the sequel), is much easier to compute than that of RP; moreover the two equilibrium assignments are shown to be close to each other and even to converge to one another when the number of agents becomes large.

Probabilistic Serial is implemented in the same way as Random Priority: by selecting a certain priority ordering σ of the agents and then playing this priority mechanism Prio(σ). The only difference is that the probability distribution according to which the ordering of agents is selected is computed by a polynomial algorithm using the agents' reports about their type, in contrast with RP selecting this ordering from the uniform distribution.

Thirdly, we show that the improvement of PS over RP is significant but small: if a gain of a couple of percentage points (in the number of agents served and in the surplus collected) matters, then the additional effort of implementing PS instead of RP is justified. Otherwise, we may in effect take each mechanism as a proxy of the other; this is especially true if many agents are involved.

1.2 Overview of the results and related literature

The model is defined in Section 2, and the Random Priority mechanism is analyzed in Section 3, where we give a recursive algorithm to compute its expected outcome (Proposition 1). The Probabilistic Serial mechanism is defined in Section 4 in a 'backward' fashion: we first define its equilibrium outcome (by means of an easy formula), then we show that it can be implemented by randomly choosing, and playing, a priority mechanism (Proposition 2). Proposition 3 in Section 5 shows that both mechanisms, RP and PS, share the very strong incentive compatibility property known as 'group strategyproofness'. Next, Theorem 1 establishes that the equilibrium outcome of PS is Pareto superior (or indifferent) to that of RP. Section 6 gathers some concluding comments.

A convergence result (Theorem 2) is reported in Appendix A: when the number of agents grows large, the difference between the RP and PS assignments vanishes. Extensive numerical computations are reported in Appendix B. All tedious proofs are gathered in Appendix C.

We discuss now the literature related to our model. In the mathematical economics literature, scheduling is a special case of the random assignment problem, where q objects must be randomly assigned to n agents with heterogeneous preferences over the objects.

The other related stream of literature bears on scheduling and queuing; a good survey is Lawler et al. (1993). Both streams of literature discuss incentive compatibility and fairness, but they give different meanings to these terms.

In the latter, the discussion of incentive compatibility typically relies on tolls (Naor, 1968; Dolan, 1979; Suijs, 1996) or nonlinear prices (Mendelson, 1985; Mendelson and Whang, 1990) whereas cash transfers of any sort are ruled out in our model. An exception, where incentive compatibility is by means of randomization, is the work of Shenker (1995), see also Nagle (1987) and Demers, Keshav and Shenker (1990), focusing on the case where each agent may demand a different number of jobs. Fairness is also dicussed in that work, and interpreted , as is common in the queuing literature, as requiring that no job takes all resource capacity at the expense of other jobs.

In the mathematical economics literature on random assignment, on the other hand, fairness means, at least, that users with identical demands should be treated equally (ex ante), and sometimes is interpreted as the stronger requirement of envy-freeness (no agent prefers –ex ante– the assignment of another agent to his own). The main finding is the impossibility of meeting simultaneously fairness, incentive compatibility and efficiency: Hylland and Zeckhauser (1979), Gale (1981), Zhou (1990), Bogomolnaia and Moulin (1999b). But in the particular context of scheduling with opting out, these three requirements are compatible, and in fact their combination characterize Probabilistic Serial: Bogomolnaia and Moulin (1999a)[2], briefly discussed at the end of Section 6. Finally, Friedman (1994) assumes that individual preferences are dichotomous (namely flat up to a certain 'deadline') and discusses a mechanism similar to, but different from, our Probabilistic Serial.

2 The model

The set I of agents is fixed throughout, $I = \{1, 2, \ldots, n\}$. Each agent i is endowed with a von Neumann–Morgenstern utility function u_i, $u_i \in \mathbf{R}^n$, over the n possible periods or dates at which he could be served. If agent i receives service at date k, his utility is $u_i(k)$: in this case, we say below that he *consumes* date k, or that he is *assigned* date k. We always assume that preferences are monotonic, namely:

$$u_i(k) \geq u_i(k+1), \quad \text{for all } k = 1, \ldots, n - 1 . \tag{1}$$

[2] Note that these two papers were inspired by the present work.

Moreover, the zero of the utility function is interpreted as the utility for the outside option. We say that agent i is of *type* k if

$$u_i(k) > 0 \geq u_i(k+1) . \tag{2}$$

Here k varies from zero (if $u_i(1) \leq 0$) to n (if $u_i(n) > 0$). The type tells us how long the agent is willing to wait before exercising his outside option. In the strategic analysis we assume that, faced with the choice between opting out and consuming the k-th date, he opts out whenever indifferent ($u_i(k) = 0$). This altruistic tiebreaking rule will simplify the strategic analysis without any real loss of generality.

We denote by $\mathcal{U}$ the set of utility functions (the subset of $\mathbf{R}^n$ defined by (1)) and by $\mathcal{U}_k$, $k = 1, \ldots, n$, the subset of utility functions of type k (defined by (2)). To a profile of utility functions $U \in \mathcal{U}^n$, we associate a profile of types $T = (\mathcal{I}_0, \ldots, \mathcal{I}_n)$ which keeps track of the type of every agent. Thus $(\mathcal{I}_k)_{0 \leq k \leq n}$ is the partition of $\mathcal{I}$ defined by $i \in \mathcal{I}_k \iff u_i \in \mathcal{U}_k$. Because our two mechanisms, Random Priority and Probabilistic Serial, use only the profile of types to compute the (random) assignment of dates to agents, we often omit the underlying profile of von Neumann–Morgenstern utility functions; however the latter are key to the strategic and welfare analysis, and the primitive constituents of our model.

We denote by Δ the set of (random) assignments over the n dates and the outside option. An element $z \in \Delta$ is written as a nonnegative n-vector $z = (z_1, \ldots, z_n)$ such that $\sum_k z_k \leq 1$. The outside option is deliberately omitted from this notation; its probability is $1 - \sum_k z_k$.

An $\mathcal{I}$-assignment $(z_i)_{1 \leq i \leq n}$ specifies a random assignment to each agent. It will be convenient to write such an assignment in matrix form, $Z = [z_{ik}]$, where the entry z_{ik} is the probability that agent i consumes date k: the row index i runs over $\mathcal{I}$ and the column index k over $K = \{1, \ldots, n\}$.

The planner/manager can choose at random the order in which agents are offered service. Hence Z is feasible if and only if it is a convex combination of deterministic priority assignments. In order to make this statement precise, we introduce some notation.

Let $\overline{\mathcal{S}}$ denote the set of sequences $\overline{\sigma} = (i_1, \ldots, i_q)$ of q distinct elements in $\mathcal{I}$, for some q, $0 \leq q \leq n$ (with the convention that the empty sequence corresponds to $q = 0$). Let $\mathcal{S}$ be the subset of $\overline{\mathcal{S}}$ comprising the sequences of length exactly n: thus $\mathcal{S}$ is identified with the set of priority orderings of $\mathcal{I}$. To each $\overline{\sigma} \in \overline{\mathcal{S}}$ we associate the following *truncated permutation matrix* $P_{\overline{\sigma}}$:

$$P_{\overline{\sigma}} = [z_{ik}], \qquad z_{ik} = 1 \text{ if } k = 1, \ldots, q \text{ and } i = i_k; \qquad z_{ik} = 0 \text{ for all other } i, k .$$

The matrix $P_{\overline{\sigma}}$ is the deterministic assignment resulting from serving the agents in $\overline{\sigma}$, in that order, and ignoring the others.

To each $\sigma \in \mathcal{S}$ we associate the *Priority mechanism* Prio(σ) : following the ordering σ , the agents are successively offered the best non-assigned date, and choose either to stay in line for service or to opt out. If σ : $(\sigma(1), \ldots, \sigma(n))$, agent $\sigma(q)$ is offered date k if exactly $k - 1$ agents among $\sigma(1), \ldots, \sigma(q - 1)$ accepted the offer.

Given a profile of types T and a priority ordering $\sigma \in \mathcal{S}$, we denote by $\overline{\sigma}$ the (possibly shorter) sequence in $\overline{\mathcal{S}}$ of the agents successively served in equilibrium, when the priority ordering is σ. That is, an agent of type k accepts any date not later than k and refuses any later date. We call $\overline{\sigma}$ the equilibrium sequence associated with the ordering σ. For instance consider $(I_0, I_1, I_3, I_4) = (\{4\}, \{1, 3\}, \{5\}, \{2, 6\})$, then

$$\sigma = (3, 2, 5, 4, 1, 6) \implies \overline{\sigma} = (3, 2, 6) ,$$
$$\sigma = (4, 6, 1, 3, 5, 2) \implies \overline{\sigma} = (6, 5, 2) .$$

Given the profile T and an ordering $\sigma \in \mathcal{S}$, we denote by Prio$(\sigma, T) = P_{\overline{\sigma}}$ the truncated permutation matrix resulting from the corresponding equilibrium sequence.

Definition 1 *An assignment matrix Z is feasible at the profile of types T if and only if Z is a convex combination of the matrices Prio(σ, T), $\sigma \in \mathcal{S}$. We let $\mathcal{F}$ be the set of feasible assignment matrices.*

Remark: Clearly, and assignment matrix $Z \in \mathcal{F}$ is substochastic:

$$z_{ik} \geq 0 \text{ for all } i, k \; ; \quad \sum_{i \in \mathcal{I}} z_{ik} \leq 1 \text{ and } \sum_{k=1}^{n} z_{ik} \leq 1 \text{ for all } i, k \; . \tag{3}$$

This follows straightforwardly from the fact that each matrix $P_{\overline{\sigma}}$ is substochastic. Conversely, these inequalities are not sufficient to characterize $\mathcal{F}$. Baiou and Balinski (1998), Theorem 4, offer a conjecture about a characterization of the related set of convex combinations of the matrices $P_{\overline{\sigma}}$, $\overline{\sigma} \in \overline{\mathcal{S}}$ (independently of agents' types).

3 The RP mechanism and equilibrium assignment

As mentioned in the introduction, the description of the RP mechanism is very simple, but that of its equilibrium assignment is not.

Definition 2 *The Random Priority mechanism selects at random and without bias a priority ordering of* I*, namely and element* σ *in* S*. Then the agents play the priority mechanism Prio(*σ*). The equilibrium*[3] *assignment corresponding to the profile of types* T *is denoted RP(*T*):*

$$RP(T) = \frac{1}{n!} \sum_{\sigma \in \mathcal{S}} Prio(\sigma, T)\ .$$

We will refer to the above matrix as the RP assignment at T*.*

Let $M_q = \bigcup_{k \geq q} I_k$ denote the set of agents of type at least q. As long as the RP mechanism is assigning the dates 1 to q, all agents in M_q behave in exactly the same way. Consider then the assignment of date q. The probability that an agent in $I \setminus M_q$ consumes it is zero, whereas all agents in M_q have an equal probability to consume it. Denoting m_q the cardinality of M_q, the latter probability equals β_q / m_q where β_q is the probability that date q is assigned at all. Therefore the RP assignment at T is $Z^{RP} = [z_{ik}]$ where for all k, $1 \leq k \leq n$:

$$\begin{aligned} z_{ik} &= \frac{\beta_k}{m_k} \quad \text{if } i \in M_k \\ &= 0 \qquad \text{if } i \notin M_k \end{aligned} \tag{4}$$

Computing the RP assignment boils down to computing the sequence $(\beta_k)_{k \geq 1}$. It is useful to introduce the 'threshold' quantity[4] q_e associated with a profile of types T:

$$q_e \text{ is the largest quantity } q \text{ such that } q \leq m_q\ . \tag{5}$$

Observation 1 $\beta_q = 1$ *whenever* $q \leq q_e$*. Conversely, whenever* $q > q_e$*, we have* $\beta_q < 1$ *. Moreover the sequence* $(\beta_q)_{q \geq 1}$ *is nonincreasing.*

Proof: For the first assertion, after the first $q - 1$ dates are assigned, some agents in M_q are still not served (because $q \leq q_e \Longrightarrow m_q > q - 1$), hence the claim by induction. The second assertion is because with positive probability the first m_q agents drawn by RP are precisely those in M_q; they all accept to consume date k with $k \leq q - 1$; since

[3] The concept of strategic equilibrium used here is that of a dominant strategy; for an agent of type q, it is a dominant strategy to accept any date k, k ≤ q, and refuse any other one. Alternatively, we may describe our equilibrium as the unique strong equilibrium of the game. We omit the unimportant details, and refer the interested reader to Cres and Moulin (1999), where the strategic discussion is more detailed.

[4] In the slightly different context of Cres and Moulin (1999), it is the 'efficient' quantity, in the economic meaning of the word; it is also the quantity for which the type curve crosses the diagonal.

$q > q_e \Longrightarrow m_q \le q - 1$, all agents in $\mathcal{M}_q$ are served in this case before date q is offered and there is no one left to consume that date.

Unfortunately, it is not possible to give a simple formula for β_q. Our first result provides a recursive algorithm.

Proposition 1 *Given a profile of types $T = (\mathcal{I}_0, \dots, \mathcal{I}_n)$ with corresponding cardinalities $(n_0, \dots, n_n)$ we denote by Q the largest integer q such that $\mathcal{I}_q$ is nonempty, and assume $Q > 1$ (the case $Q = 0$ is trivial). Then the probability β_q that the q-th date be assigned is given by:*

$$\beta_q = \sum_{r=0}^{Q-q} a_{Q,r} \qquad \text{for all } q, 1 \le q \le Q; \qquad \beta_q = 0 \text{ if } q > Q \tag{6}$$

where the double sequence $(a_{q,r})_{0 \le q,r \le Q}$ is computed by the initial conditions:

$$a_{0,0} = 1;\ a_{0,r} = 0 \text{ for } 1 \le r \le Q,$$

and the recursive formulas:

1. *if $r \le m_{q+1}$,*

$$a_{q,r} = \binom{m_q}{n_q}^{-1} \sum_{j=0}^{n_q} \binom{r+j}{j} \binom{m_q - (r+j)}{n_q - j} a_{q-1,r-1+j}; \tag{7}$$

2. *if $r > m_{q+1}$,*

$$a_{q,r} = a_{q-1,r-1+n_q}. \tag{8}$$

Proof: See Appendix C. The formulas are not intuitive. The quantity $a_{q,r}$ is the probability that among the q first periods, $q - r$ of them are assigned to agents of type 1 to q, and r or less to agents of type $q+1$ to Q. (Note that the set of type 0 agents plays no role in the computations.)

The main interest of the proposition is to allow numerical computations and to prove the convergence result in Appendix A. We give several examples in the next section and Appendix B.

4 The PS mechanism and equilibrium assignment

As mentioned in the introduction, the description of the PS equilibrium assignment is very simple, but that of the mechanism to implement it requires to run a polynomial algorithm. An intuitive definition of the PS assignment is by the following algorithm: think of each date as a mass one probability; allocate the dates sequentially starting from the best dates with equal share to all interested agents. Therefore agent i of type q ($i \in I_q$) gets a $1/m_k$ probabilistic share on date k for $k = 1, \ldots, q$. He will accumulate the probabilistic shares of all the dates, starting from the best one, until one of two things happen: he has accumulated a probability one of service; or he has accumulated the shares of all dates he prefers to opting out (i.e., dates 1, 2, ... , k if he is of type k). Formally, for all $q \geq 1$, for all $i \in I_q$:

$$\begin{cases} z_{ik} = \dfrac{1}{m_k} & \text{if } \displaystyle\sum_{1}^{k} \frac{1}{m_h} \leq 1 \,, \\ z_{ik} = 0 & \text{if } \displaystyle 1 \leq \sum_{1}^{k-1} \frac{1}{m_h} \,, \\ z_{ik} = 1 - \displaystyle\sum_{1}^{k-1} \frac{1}{m_h} & \text{if } \displaystyle\sum_{1}^{k-1} \frac{1}{m_h} < 1 < \sum_{1}^{k} \frac{1}{m_h} \,. \end{cases}$$

The key to the above formula is the critical integer, if any, at which the sum $\sum 1/m_h$ passes 1. We define q^* to be the largest integer such that $q \leq n$ and $\displaystyle\sum_{1}^{q} \frac{1}{m_h} \leq 1$. With the convention $1/0 = \infty$, we see that q^* cannot exceed Q (the largest q such that I_q is nonempty). We set $\epsilon = 1 - \displaystyle\sum_{1}^{q^*} \frac{1}{m_h}$, so that $0 \leq \epsilon < 1$.

Definition 3 *Given a profile of types $T = (I_0, \ldots, I_n)$, define the sequence $(\gamma_q)_q$ as*

$$\begin{cases} \gamma_q = 1 & \text{for} \quad 1 \leq q \leq q^* \,, \\ \gamma_{q^*+1} = \epsilon \cdot m_{q^*+1} & \\ \gamma_q = 0 & \text{for} \quad q^* + 2 \leq q \leq n \,. \end{cases}$$

Interpret γ_q as the probability that, in the PS assignment, the q-th date be assigned [5]*. Then the PS assignment $Z^{PS} = [z_{ik}]$ is defined by:*

$$\begin{aligned} z_{ik} &= \frac{\gamma_k}{m_k} && \text{if } i \in M_k \\ &= 0 && \text{if } i \notin M_k \end{aligned} \tag{9}$$

[5] Note that the last part of the formula disapears if $q^* = n$.

Given the similar formula (4) defining the assignment Z^{RP}, comparing the two assignments amounts to comparing the vectors (β_q) and (γ_q), namely the probabilities that the q-th date be assigned by the two mechanisms.

For q not larger than the threshold quantity q_e (see (5)), the q-th date is consumed for sure in both assignments: $\beta_q = \gamma_q = 1$. This was shown for β_q in Observation 1. As for γ_q, it results from $q_e \leq q^*$, which itself follows from the implication:

$$\left\{ \text{for } q \text{ such that } q \leq q_e : \ \frac{1}{m_q} \leq \frac{1}{m_{q_e}} \leq \frac{1}{q_e} \right\} \Longrightarrow \sum_{1}^{q_e} \frac{1}{m_h} \leq 1 .$$

Beyond q_e, the probabilities β_q and γ_q may or may not coincide, as demonstrated by two examples illustrating Definitions 2 and 3. In the examples, we describe the profile simply by listing the cardinality $n_1, \ldots, n_Q$ of the subsets $N_1, \ldots, N_Q$.

Example 1: $(n_1, n_2, n_3) = (1, 2, 2)$

We have five agents, and the threshold quantity is $q_e = 2$ (because $m_2 = 4$ and $m_3 = 2$). In RP, the probabilities of allocating the three dates are[6]: $\beta_1 = 1$, $\beta_2 = 1$, $\beta_3 = \frac{13}{15}$. Turning to PS, here $q^* = 3$ (because $1/5 + 1/4 + 1/2 < 1$), therefore $\gamma_1 = \gamma_2 = \gamma_3 = 1$. Hence the RP and PS assignments:

$$Z^{RP} = \begin{pmatrix} \frac{1}{5} & 0 & 0 \\ \frac{1}{5} & \frac{1}{4} & 0 \\ \frac{1}{5} & \frac{1}{4} & 0 \\ \frac{1}{5} & \frac{1}{4} & \frac{13}{30} \\ \frac{1}{5} & \frac{1}{4} & \frac{13}{30} \end{pmatrix} \quad \text{and} \quad Z^{PS} = \begin{pmatrix} \frac{1}{5} & 0 & 0 \\ \frac{1}{5} & \frac{1}{4} & 0 \\ \frac{1}{5} & \frac{1}{4} & 0 \\ \frac{1}{5} & \frac{1}{4} & \frac{1}{2} \\ \frac{1}{5} & \frac{1}{4} & \frac{1}{2} \end{pmatrix} .$$

Example 2: $(n_1, n_2, n_3, n_4) = (1, 1, 1, 1)$

There are four agents, and $q_e = 2$ (because $m_2 = 3$ and $m_3 = 2$). The probability that RP assigns the four dates are[7]: $\beta_1 = 1$, $\beta_2 = 1$, $\beta_3 = \frac{3}{4}$, $\beta_4 = \frac{1}{24}$. As for PS, $q^* = 2$ (because $1/4 + 1/3 \leq 1 < 1/4 + 1/3 + 1/2$), and $\gamma_1 = \gamma_2 = 1$, $\gamma_3 = \frac{5}{6}$ and $\gamma_4 = 0$. Hence

[6] Indeed, the third date is not assigned if: (1) either the first two agents in line are those in N_3 — which has probability $\frac{2}{5}\frac{1}{4}$; (2) or the first and third agents in line are those in N_3, and the second the one in N_1 — which has probability $\frac{2}{5}\frac{1}{4}\frac{1}{3}$. Hence a total of $\frac{2}{15}$.

[7] For instance, the fourth date is assigned only if the priority ordering is by increasing type: proba $\frac{1}{24}$.

the resulting RP and PS assignments:

$$Z^{RP} = \begin{pmatrix} \frac{1}{4} & 0 & 0 & 0 \\ \frac{1}{4} & \frac{1}{3} & 0 & 0 \\ \frac{1}{4} & \frac{1}{3} & \frac{3}{8} & 0 \\ \frac{1}{5} & \frac{1}{4} & \frac{3}{8} & \frac{1}{24} \end{pmatrix} \quad \text{and} \quad Z^{PS} = \begin{pmatrix} \frac{1}{4} & 0 & 0 & 0 \\ \frac{1}{4} & \frac{1}{3} & 0 & 0 \\ \frac{1}{4} & \frac{1}{3} & \frac{5}{12} & 0 \\ \frac{1}{5} & \frac{1}{4} & \frac{5}{12} & 0 \end{pmatrix}.$$

Table 1 in Appendix B provides ten more profiles, where the number of agents goes up to 36, for which the vectors (β_q) and (γ_q) are explicitely computed.

We turn to the definition of the PS mechanism: this amounts to describe PS as a convex combination of deterministic priority assignments (see Definition 1). Given a profile of types and the corresponding PS assignment Z^{PS}, we must show the existence of a probability distribution π over $\mathcal{S}$ such that, when an order σ is drawn in $\mathcal{S}$ according to π, and Prio(σ) is played, the resulting assignment is Z^{PS}. With our matrix notation, this reads:

$$Z^{PS} = \sum_{\sigma \in \mathcal{S}} \pi_\sigma \cdot P_{\overline{\sigma}} \,. \tag{10}$$

Proposition 2 *Given a profile of types $T = (\ell_0, \ldots, \ell_n)$ there exists a probability distribution π over $\mathcal{S}$ (computed by a polynomial algorithm) such that (10) holds. (Moreover each sequence $\overline{\sigma}$ is of length q^* or $q^* + 1$.)*

Proof: See Appendix C. The proof is very similar to the standard problem of representing a bistochastic matrix as a convex combination of permutation matrices (Birkhoff-von Neumann theorem[8]). In fact our proof consists simply of providing a bistochastic 'cover' of the PS matrix, and then invoking the Birkhoff-von Neumann theorem.

For Example 1, a probability distribution π satisfying Proposition 2 obtains as follows: (1) draw the first agent in line (among the five agents) with uniform probability; (2a) if the type 1 agent is drawn first, go on as in RP; (2b) if a type 2 agent is drawn first, draw *both* type 3 agents (in random order, with equal chance on both orders); (2c) if a type 3 agent is drawn first, draw *one* type 2 agent (with equal proba on both agents), then draw the other type 3 agent.

[8] Well-known to be solvable in polynomial time; a polynomial algorithm for turning any feasible solution to a linear program into a convex combination of the extreme points of the feasible region may be found in Bazaraa et al. (1990).

For example 2, we can choose π as indicated below, where we write the (unique) agent of type i as i and omit all agents who decline (i.e., we report $\overline{\sigma}$ instead of σ):

priority ordering	probability	priority ordering	probability
1 3 4	1/8	3 2 4	1/6
1 4 3	1/8	4 2 3	1/6
2 3 4	1/8	3 4	1/12
2 4 3	1/8	4 3	1/12

With respect to implementation, the main difference between RP and PS is that the former uses the uniform distribution over all priority orderings, whereas the latter chooses a distribution only *after* eliciting the type of every agent. Thus the PS mechanism requires to process more information than RP, and to compute a polynomial algorithm. We ask now if the PS mechanism could be open to strategic manipulation: what if an agent finds it profitable to misreport his type? Fortunately, such manipulations do not pose any more problem to the PS mechanism than they do to the RP mechanism as we shall see in the next section.

5 Incentive compatibility and welfare comparison

5.1 Incentive compatibility properties

To each profile of types $T = (l_0, \ldots, l_n)$, the RP mechanism and the PS mechanism associate a (probabilistic) assignment $Z \in \mathcal{F}$ (given respectively by (4) and (9)). This defines two mechanisms: $T \to Z$ and we now show that they both are *strategyproof*: it is never profitable for any agent to misreport his type in the hope of receiving an assignment improving his utility.

In fact, we show a stronger property: these mechanisms are both *group strategyproof*, namely a joint deviation by any coalition of agents either leaves the utilities of all agents in the coalition unchanged, or strictly decreases the utility of at least one of them. The group strategyproofness property is one of the strongest incentive compatibility requirements: its well known informational and normative implications are discussed, e.g., by Barbera (1995), Moulin (1996).

To formally define the property 'group strategyproofness', we fix a mechanism f. Consider an arbitrary nonempty subset J of $\mathcal{I}$ and two profiles $U^t, t = 0$ or 1, in $\mathcal{U}^{\mathcal{I}}$ that may differ along the J-coordinates: for all $i \in \mathcal{I} \setminus J : u_i^0 = u_i^1$. We denote by T^t the profile of types corresponding the the utility profile U^t and by z^t the assignment resulting from f. The GSP property requires the following:

$$\text{if } \{\text{for all } i \in J, u_i^0 \cdot z_i^1 \geq u_i^0 \cdot z_i^0\} \text{ then } \{\text{for all } i \in J, u_i^0 \cdot z_i^1 = u_i^0 \cdot z_i^0\}. \tag{11}$$

Proposition 3 *The RP assignment and the PS assignment both define a group strategyproof mechanism.*

Proof: See Appendix C.

This result states the incentive compatibility of our two mechanism, viewed as 'revelation mechanisms': that is, the manager asks each agent to report his type and enforces the equilibrium assignment at the reported profile.

5.2 Welfare comparison of the RP and PS mechanisms

With two or three agents, the RP and PS assignments coincide. This fact is easily checked from Definition 3 and by computing directly the RP assignment. For problems involving four agents or more, these assignments may be different as illustrated by Examples 1 and 2 in the previous section. Note first that the expected number of agents served is higher in the PS assignment: in the first example, $\sum_q \beta_q = 2.87 < \sum_q \gamma_q = 3$; in the second, $\sum_q \beta_q = 2.79 < \sum_q \gamma_q = 2.83$. Moreover, in Example 1, an agent in I_3 strictly prefers his PS assignment to his RP assignment (because he gets a higher proba of consuming period 3, ceteris paribus); on the other hand, agents in I_1 and I_2 are indifferent between RP and PS. The situation is similar in Example 2: the first two agents are indifferent whereas the last two strictly prefer PS to RP (e.g., the proba that the type 4 agent consumes date 4 in RP is 'transfered' in PS to his consumption of date 3 —as 3/8+1/24=5/12). These observations generalize.

Theorem 1 *Given is a profile of types $T = (I_0, \ldots, I_n)$ with threshold quantity q_e.*

1. *For every agent, the probability that he be served at all (that he does not opt out) is not smaller in the PS than in the RP assignment.*

2. *In both assignments, the expected number of agents served is between q_e and $2q_e$[9]:*

$$q_e \leq q^{RP} \leq q^{PS} \leq 2q_e. \tag{12}$$

[9] We note than the upper bound on q^{PS} in statement 2 above is tight. That is, for any q, there exists a profile of types where q is the threshold quantity and $q^{RP} = q^{PS}$ is arbitrarily close to $2q_e$ (see proof).

3. *The PS assignment is Pareto superior to the RP assignment, or they are welfare equivalent:*

$$\text{for all } i \in I, \qquad u_i\left(z_i^{PS}\right) \geq u_i\left(z_i^{RP}\right).$$

4. *Every agent of type at most q_e gets the same assignment in RP and PS:*

$$\text{for all } i \in I \setminus M_{q_e+1}, \qquad z_i^{PS} = z_i^{RP}.$$

5. *Assume $q^* \geq q_e + 1$. Then an agent of type at least $q_e + 1$ with strictly monotonic preferences, strictly prefers his PS assignment to his RP assignment.*

Proof: See Appendix C.

Theorem 1 demonstrates the unambiguous welfare advantage of the PS mechanism over the RP one (statement 3). Moreover, it says that the agents with high types strictly prefer PS to RP, whereas the agents with low types are indifferent (statements 4 and 5). Finally, we learn that PS (in expectation) serves any agent more often (statement 1), but does not serve more than twice the 'efficient' (in the economic meaning of the word) number of agents (statement 2).

The literature on scheduling and queuing regards statement 1 as an argument in favor of PS: it means that its *failure rate* is smaller (Mendelsson and Wang (1990), Gelenbe and Mitrani (1980), Lawler et al. (1993)). However, in another interpretation of our model inspired by the tragedy of the commons (the joint exploitation of a decreasing returns technology), one important normative goal is to reduce the level of production (see Crès and Moulin (1999) and references therein); in that context, statements 1 and 2 argue in favor of RP over PS.

In Appendix B, we explain numerically the gap between RP and PS (quantitywise and welfarewise) and we show it is small (we conjecture in never exceeds 8.33 % quantitywise, and is usually a couple of percentage points on both dimensions). On top of that we establish in Appendix A a result of asymptotic equivalence of the two mechanisms when the number of agents, n, tends toward ∞.

6 Concluding comments

1. An important feature of our two mechanims RP and PS is that they only use the ordinal information on types in the computation of the final outcome. This

makes for very simple mechanisms and allows the very strong incentive compatibility property described in Proposition 3; on the other hand, it rules out the possibility to take advantage of cardinal information. For instance, if an agent has dichotomous preferences such as $u_i(k) = 1$ for $k = 1, 2, 3, 4$ and $u_i(k) = -1$ otherwise, it is inefficient to serve him with positive probability in date 1 if there is at least one other agent with strictly monotonic preferences. Yet the RP and PS mechanisms cannot use such information.

2. A more general perspective on our model comes from the mechanism design approach. A general mechanism elicits from the agents the profile of utility functions U, then determines the random assignment Z under the feasibility constraint described in Definition 1. Bogomolnaia and Moulin (1999a) apply this viewpoint to our model, and restrict attention to the class of random assignment mechanisms that only elicit such ordinal information. They offer a concept of efficiency adapted to this informational structure (called 'ordinal efficiency') and they characterize PS by the combination of its properties of efficiency, fairness, and incentive compatibility. More precisely, they show that:

    (a) PS is the only such *mechanism* that satisfy the combination of *(i)* ordinal efficiency, *(ii)* strategyproofness and *(iii)* equal treatment of equals (two agents sending identical reports receive the same random assignment);

    (b) it is also the only such *assignment* that satisfy the combination of *(i)* ordinal efficiency and *(ii)* envy-freeness (no agent prefers the random assignment of another agent to his own).

References

Baiou, M. and Balinski, M. (1998). Polytopes of trucated permutations: a conjecture, Mimeo, Laboratoire d'Économétrie, École Polytechnique, Paris.

Barbera, S. (1995). Notes on strategyproof social choice functions, in *Social Choice Re-examined*, Arrow, K., Sen, A. and Suzumura, K., Eds, Macmillan, London.

Bazaraa, M., Jarvia, J. and Sherali, H. (1990). *Linear programming in Network Flows*, 2nd ed., New York: Wiley.

Bogomolnaia, A. and Moulin, H. (1999a). A simple random assignment problem with a unique solution, Mimeo, Duke University.

Bogomolnaia, A. and Moulin, H. (1999b). A new solution to the random assignment problem, forthcoming *Journal of Economic Theory.*

Crès, H. and Moulin, H. (1998). Random Priority : a Probabilistic Resolution of the Tragedy of the Commons, *Cahiers de recherche, HEC*, CR662/1999.

Demers, A., Keshav, S. and Shenker, S. (1990). Analysis and simulation of a fair queuing algorithm, *International Research and Experience*, 1: 3-26.

Dolan, R. (1979). Incentive mechanisms for priority queuing problems, *Bell Journal of Economics*, 421-436.

Friedman, E. (1994). An incentive compatible scheduling algorithm, Mimeo, Duke University.

Gale, D. (1987). College course assignments and optimal lotteries, Mimeo, University of California, Berkeley.

Gelenbe, E. and Mitrani, I. (1980). *Analysis and synthesis of computer systems models*, New York: Academic Press.

Hardy, G., Littlewood, J.E. and Polya, G. (1934). *Inequalities*, Cambridge, Cambridge University Press.

Hylland, A. and Zeckhauser, R. (1979). The efficient allocation of individuals to positions, *Journal of Political Economy*, 91: 293-313.

Lawler, E.L., Lenstra, J.K., Rinnooy Kan, A.H.G. and Shmoys, D.B. (1993). Sequencing and Scheduling: Algorithms and Complexity, in *Logistics of Production and Inventory*, Graves, S.C., Rinnooy Kan, A.H.G. and Zipkin, P.H., Eds. Amsterdam: North Holland Press.

Mendelson, H. (1985). Pricing computer services: Queuing effects, *Communications*, ACM 28: 312-321.

Mendelson, H. and Whang, S. (1990). Optimal incentive compatible pricing for the M/M/1 queue, *Operations Research*, 38: 870-883.

Moulin, H. (1996). Procedural versus Endstate Justice: An implementation view point, forthcoming in *Justice, Political Liberalism and Utilitarianism*, Salles, M. and Weymark, J., Eds. Proceedings of a conference in honor of John Harsanyi and John Rawls.

Naor, P. (1969). The regulation of queue size by levying tolls, *Econometrica*, 37: 15-24.

Nagle, J. (1987). On packet switches with infinite storage, *IEEE Transactions Communications*, 35: 435-438.

Shenker, S. (1977). Making greed work in network: A game theoretic analysis of switch service discipline, *IEEE/ACM Transactions on Networking*, 3: 819-831.

Suijs, J. (1996). On incentive compatibility and budget balancedness in public decision making, *Economic Design*, 2: 193-209.

Zhou, S. (1990). On a conjecture by Gale about one-sided matching problems, *Journal of Economic Theory*, 52:123-135.

Appendix A: Asymptotic equivalence between the RP and PS mechanisms

Denote L the set of probability measures over the unit interval $[0,1]$. An element $\lambda \in L$ generates a distribution function δ of agents across types:

$$\forall x \in [0,1], \qquad \delta(x) = \lambda([x,1]).$$

To avoid unnecessary technical difficulties, we assume that λ is measurable (absolutely continuous) with respect to the Lebesgue measure and that its density is nonzero in the neighborhood of 1. In particular δ is continuous, nonincreasing and $\delta(0) = 1$; $\delta(1) = 0$.

We fix the measure λ throughout this appendix. Denote a profile of types as $T = (n_0, \ldots, n_Q)$ where Q is the largest nonempty type and n_q is the number of agents of type q. Thus $Q \leq n = \sum_0^Q n_q$ and $n_Q > 0$. To avoid 'empty economies', we assume $Q \geq 1$. For any pair $(n, Q) \in N_+^2$, one can define a profile of types $T^{(n,Q)} = (n_0, \ldots, n_Q)$ for all q, $0 \leq q \leq Q$, through:

$$m^{(n,Q)} = \left\lfloor n\delta\left(\frac{q}{Q+1}\right)\right\rfloor$$
$$n^{(n,Q)} = \left\lfloor n\delta\left(\frac{q}{Q+1}\right)\right\rfloor - \left\lfloor n\delta\left(\frac{q+1}{Q+1}\right)\right\rfloor \qquad (13)$$

where $\lfloor x \rfloor$ denotes the highest integer smaller than or equal to x. For instance, take $\lambda \equiv 1$, the uniform density over $[0,1]$. It generates the distribution function $\delta(x) = 1 - x$. For $n = (Q+1)k$, $k \in N$, the uniform profile of types $T = (k, \ldots, k)$ obtains.

Notice however that our measurability assumption on λ prevents us from describing 'atomic' profiles of types where a positive fraction of agents have the same type for any n and any Q. For instance, there is no function δ such that eqs. (13) yield an homogeneous profile $T^{(n,Q)}$ for all n,Q.

The following two observations are immediate. Consider a fixed Q, and a number of agents n growing infinitely. Then

$$\lim_{n\to\infty}\frac{n_q}{n}=\lambda\left(\left[\frac{q}{Q+1},\frac{q+1}{Q+1}\right]\right);\qquad \text{for all } q, q\le Q:$$

When n and Q tend toward ∞, for any sequence $(q^{(n,Q)})_{(n,Q)}$[10],

$$\text{if }\lim_{(n,Q)\to\infty}\frac{q^{(n,Q)}}{Q}=x\in[0,1];\ \text{ then }\lim_{(n,Q)\to\infty}\frac{m_{q^{(n,Q)}}}{n}=\delta(x):$$

Theorem 2 below describes the asymptotic behavior of the RP and PS assignments, for any sequence (n,Q) where n grows infinitely and where the crowding factor $c=n/Q$ converges, possibly to infinity. The key fact is that these two assignments coincide in the limit. In order to give a precise statement of this property, we establish some preliminary results.

First we compute the limit ratio between the threshold quantity $q_e^{(n,Q)}$ and Q, in terms of the two parameters of the model — the distribution function δ and the (limit) crowding factor c, $1\le c\le\infty$. We claim that for any sequence (n,Q) such that $n\to\infty$ and $n/Q\to c$, the ratio $q_e^{(n,Q)}/Q$ converges; its limit x_e is as follows:

$$\begin{aligned}&\text{if }\lim_{n\to\infty}\frac{n}{Q}=c=\infty;\ \text{ then } x_e=1;\\ &\text{if }\lim_{(n,Q)\to\infty}\frac{n}{Q}=c<\infty;\ \text{ then } x_e \text{ is defined by } c\delta(x_e)=x_e:\end{aligned}\tag{14}$$

The proof is in Appendix C.

In the case $c=\infty$, Property (14) says that the threshold quantity $q_e^{(n,Q)}$ is (almost) equal to Q when n is large. Since we know that the PS and RP assignments coincide for all agents of type at most q_e, this establishes their asymptotic equivalence at once (see Theorem 2 for a precise statement).

We turn to the case of a sequence (n,Q) for which the limit crowding factor is finite.

Lemma 1 Consider a sequence (n,Q) such that $\lim_{n\to\infty} n/Q=c$, $1\le c<\infty$. For any sequence $(q^{(n,Q)})_{(n,Q)}$:

$$\text{if }\lim_{(n,Q)\to\infty}\frac{q^{(n,Q)}}{Q}=x\in[0,1[;\qquad\text{then }\lim_{(n,Q)\to\infty}\sum_{r=1}^{q^{(n,Q)}}\frac{1}{m_r}=\frac{1}{c}\int_0^x\frac{du}{\delta(u)}:$$

Proof: For simplicity, we write q instead of $q^{(n,Q)}$ and m instead of $m^{(n,Q)}$. We compute

$$\sum_{r=1}^{q}\frac{1}{m_r}=\underbrace{\frac{q}{Q}}_{\to x}\underbrace{\frac{Q}{n}}_{\to 1/c}\left[\frac{1}{q}\sum_{r=1}^{q}\frac{1}{\frac{1}{n}\,n\delta\left(\frac{r}{q}\frac{q}{Q+1}\right)}\right]:$$

[10] In the following expression m stands for $m^{(n,Q)}$ as defined by eqs. (13): we drop the superscript, here and below, to lighten the notation.

Moreover, for any integer n and $t \in [0,1]$, one has:

$$\frac{1}{\frac{1}{n}\lfloor n\delta(t)\rfloor} \geq \frac{1}{\frac{1}{n}+\delta(t)} \geq \frac{1}{\delta(t)} - \frac{1}{n\delta(t)^2};$$

from which we get:

$$\frac{1}{q}\sum_{r=1}^{q} \frac{1}{\frac{1}{n}\left\lfloor n\delta\left(\frac{r}{q}\frac{q}{Q+1}\right)\right\rfloor} \geq \frac{1}{q}\sum_{r=1}^{q} \frac{1}{\delta\left(\frac{r}{q}\frac{q}{Q+1}\right)} - \frac{1}{n\delta\left(\frac{r}{q}\frac{q}{Q+1}\right)^2}:$$

The two terms in the right-hand sum converge, respectively, to $\int_0^1 \frac{du}{\delta(xu)}$ and to zero. Hence a lower bound converging to $\frac{1}{c}\int_0^x \frac{du}{\delta(u)}$. A similar computation based on the inequality

$$\frac{1}{\frac{1}{n}\lfloor n\delta(t)\rfloor} \leq \frac{1}{-\frac{1}{n}+\delta(t)} \leq \frac{1}{\delta(t)} + \frac{2}{n\delta(t)^2}$$

delivers an upper bound converging to the same limit.

We can now describe the asymptotic behavior of the integer $q^*(n,Q)$ around which the PS assignment is defined (Definition 3). In fact, the ratio $q^*(n,Q)/Q$ converges, and its limit x^* can be computed by distinguishing two cases:

1. if $\int_0^1 \frac{du}{\delta(u)} = +\infty$ then x^* is the number in $[0,1[$ such that $\int_0^{x^*} \frac{du}{\delta(u)} = c$;
2. if $\int_0^1 \frac{du}{\delta(u)} < +\infty$ then two subcases occur:
   (a) if $c < \int_0^1 \frac{du}{\delta(u)}$ then x^* is the number in $[0,1[$ such that $\int_0^{x^*} \frac{du}{\delta(u)} = c$,
   (b) if $c \geq \int_0^1 \frac{du}{\delta(u)} < +\infty$ then $x^* = 1$.

We are now ready to state the convergence result; i.e., the statement that when n grows infinitely, the two distributions $\beta_q^{(n,Q)}$ and $\gamma_q^{(n,Q)}$ become arbitrarily close.

Theorem 2 Given a probability distribution λ with cumulative δ, and a number c, $1 \leq c \leq \infty$, we define x^* as follows:

$$\begin{cases} x^* = 1 & \text{if} \quad \int_0^1 \frac{du}{\delta(u)} \leq c \\ x^* \in [0,1[\quad \text{solves} & \int_0^{x^*} \frac{dt}{\delta(t)} = c \quad \text{otherwise}. \end{cases}$$

Consider a sequence (n,Q) such that $n \to \infty$ and $\lim_{n\to\infty} \frac{n}{Q} = c, 1 \leq c \leq \infty$; and a sequence of integers $\left(q^{(n,Q)}\right)_{(n,Q)}$ in $\{0,\dots,Q\}^{\mathbb{N}^2}$ such that $\lim_{(n,Q)\to\infty} \frac{q^{(n,Q)}}{Q} = x, 0 \leq x \leq 1$. Then the asymptotic behavior of the probabilities $\left(\beta_{q^{(n,Q)}}^{(n,Q)}\right)$ and $\left(\gamma_{q^{(n,Q)}}^{(n,Q)}\right)$ associated with the profile of types $T^{(n,Q)}$ (see Eqs 13) is as follows:

$$\begin{cases} x < x^* \implies \lim_{n\to\infty} \beta_q^{(n,Q)} = \lim_{n\to\infty} \gamma_q^{(n,Q)} = 1 \,; \\ x > x^* \implies \lim_{n\to\infty} \beta_q^{(n,Q)} = \lim_{n\to\infty} \gamma_q^{(n,Q)} = 0 \end{cases} .$$

Proof: Available from the authors and online.

In other words, the sequences of 'curves' $\left((q/Q\,;\beta_q)_{q=1}^{Q}\right)_{(n,Q)}$ and $\left((q/Q\,;\gamma_q)_{q=1}^{Q}\right)_{(n,Q)}$ converge together pointwise toward the threshold 'function':

$$\Phi_c(x) = \begin{cases} 0 & \text{for } x > x^* \\ [0;1] & \text{for } x = x^* \\ 1 & \text{for } x < x^* \end{cases} .$$

Appendix B: Numerical computations

Proposition 1 gives a recursive algorithm for computing the probabilities β_q, and Definition 3 gives a (much simpler) algorithm for computing γ_q. In the ten representative examples reported in Table 1, we give this probabilities in full, as well as the quantity: $\frac{\Delta q}{q_e} = \frac{q^{PS} - q^{RP}}{q_e} \times 100$, i.e., the difference in expected number of agents served, relative to the threshold number, in percentage. The examples are ordered according to their crowding factor $c = n/Q$.

Profile : $(n_q)_{1\le q\le Q}$	$c = \frac{n}{Q}$	$(\beta_q)_{1\le q\le Q}$ $(\gamma_q)_{1\le q\le Q}$	$\frac{\Delta q}{q_e}$
$(0;2;0;0;0;0;0;2)$	0.5	$(1\ 1\ 0.83\ 0.17\ 0\ 0\ 0\ 0)$ $(1\ 1\ 1\ 0\ 0\ 0\ 0\ 0)$	0%
$(1;1;1;1;1;1;1;1;1;1;1;1)$	1	$(1\ 1\ 1\ 1\ 1\ 1\ 0.98\ 0.69\ 0.17\ 0.01\ 0\ 0)$ $(1\ 1\ 1\ 1\ 1\ 1\ 1\ 0.9\ 0\ 0\ 0\ 0)$	0.8%
$(0;2;0;2;0;2;0;2;0;2;0;2)$	1	$(1\ 1\ 1\ 1\ 1\ 1\ 0.99\ 0.84\ 0.28\ 0.03\ 0\ 0)$ $(1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 0.2\ 0\ 0\ 0)$	1.0%
$(0;0;0;0;0;10;0;2)$	1.5	$(1\ 1\ 1\ 1\ 1\ 1\ 0.77\ 0.23)$ $(1\ 1\ 1\ 1\ 1\ 1\ 1\ 0)$	0%
$(2;2;2;2;2;2;2;2;2;2;2;2)$	2	$(1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 0.96\ 0.61\ 0.1)$ $(1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 0.79\ 0)$	0.5%
$(6;0;0;0;0;6;6;0;0;0;0;6)$	2	$(1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 0.98\ 0.87\ 0.54\ 0.18)$ $(1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 0.58\ 0)$	0.1%
$(0;0;2;0;4;4;6;2)$	2.25	$(1\ 1\ 1\ 1\ 1\ 1\ 1\ 0.77)$ $(1\ 1\ 1\ 1\ 1\ 1\ 1\ 1)$	3.1%
$(6;5;4;3;2;1;1;1;1)$	2.67	$(1\ 1\ 1\ 1\ 1\ 0.99\ 0.7\ 0.13\ 0)$ $(1\ 1\ 1\ 1\ 1\ 1\ 0.89\ 0\ 0)$	1.2%
$(3;3;3;3;3;3;3;3;3;3;3;3)$	3	$(1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 0.99\ 0.67)$ $(1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 0.9)$	2.7%
$(8;7;6;5;4;3;2;1)$	4.5	$(1\ 1\ 1\ 1\ 1\ 1\ 0.94\ 0.24)$ $(1\ 1\ 1\ 1\ 1\ 1\ 1\ 0.22)$	0.7%

Table 1

Table 2 reports the maximal values of $\Delta q/q_e$ when the type of each agent is at most 8 and the number n of agents varies between 4 and 25 (all bounds imposed for computational tractability). For each value of n, we report also one maximizing profile.

n	maximizing profile : $(n_q)_{q=1}^8$	$\Delta q/q_e$	n	maximizing profile : $(n_q)_{q=1}^8$	$\Delta q/q_e$
4	$(0;2;2;0;0;0;0;0)$	8.33%	11	$(1;3;4;0;3;0;0;0)$	7.94%
5	$(2;1;2;0;0;0;0;0)$	6.67%	12	$(4;0;5;0;3;0;0;0)$	7.94%
6	$(3;1;2;0;0;0;0;0)$	8.33%	13	$(4;2;4;0;3;0;0;0)$	7.81%
7	$(4;1;2;0;0;0;0;0)$	7.14%	15	$(7;1;4;0;3;0;0;0)$	7.79%
8	$(0;4;2;2;0;0;0;0)$	6.75%	18	$(9;3;3;0;3;0;0;0)$	7.71%
9	$(0;0;6;0;3;0;0;0)$	7.94%	20	$(13;0;4;0;3;0;0;0)$	7.52%
10	$(1;1;5;0;3;0;0;0)$	7.84%	25	$(17;2;3;0;3;0;0;0)$	7.46%

Table 2

The constraint $q \leq 8$ appears to be nonbinding, suggesting that the above values are the absolute maxima of $\Delta q/q_e$ for any profile with the corresponding number of agents. The decline in the maximal value for $n \geq 12$, together with the convergence result reported in Appendix A, suggest the conjecture that 8.33% is the absolute upper bound on all profiles.

We turn to an evaluation of the surplus gain. We assume here that utility functions are quasi-linear: given an increasing sequence of costs c_q, $c_0 = 0$ and $c_q < c_{q+1}$, interpreted as the cost of waiting q periods, the utility function of agent i takes the form: $u_i(q) = u_i - c_q$ for all $i \in I$, and all q, $1 \leq q \leq n$. (Thus agent i is of type q if and only if $c_q < u_i \leq c_{q+1}$.)

Denote by α_q^t the probability that date q is assigned, with $t = 0$ for RP and $t = 1$ for PS (so that $\alpha_q^0 = \beta_q$ and $\alpha_q^1 = \gamma_q$). The surplus S^t at a profile $T = (I_0; \ldots; I_n)$ is

$$S^t = \sum_{q=1}^{n} \alpha_q^t \left(\frac{1}{m_q} \sum_{i \in M_q} u_i - c_q \right):$$

We now evaluate the surplus gain from RP to PS relative to the efficient surplus S_e, corresponding to the assignment of the q_e first periods to the q_e agents of highest type (see eq. (5)). In Table 3 we report the maximal values of the relative surplus gain $\frac{\Delta S}{S_e} = \frac{S^1 - S^0}{S_e} \times 100$ (in %) when we specify linearly increasing costs $c_q = q$, and assume that the willingness to wait of all agents of a given type q are evenly spread in the interval $]c_q; c_{q+1}]$. Naturally, we maintain the constraint $Q \leq 8$.

n	maximizing profile: $(n_q)_{q=1}^8$	$\Delta S/S_e$	n	maximizing profile: $(n_q)_{q=1}^8$	$\Delta S/S_e$
4	(0;2;0;2;0;0;0;0)	2.78%	11	(5;2;2;0;2;0;0;0)	2.23%
5	(2;1;0;2;0;0;0;0)	2.22%	12	(6;2;2;0;2;0;0;0)	2.31%
6	(3;1;0;2;0;0;0;0)	2.78%	13	(7;2;2;0;2;0;0;0)	2.23%
7	(4;1;0;2;0;0;0;0)	2.38%	15	(5;7;1;0;2;0;0;0)	2.14%
8	(0;4;2;0;2;0;0;0)	2.34%	18	(13;1;2;0;2;0;0;0)	2.16%
9	(2;3;2;0;2;0;0;0)	2.29%	20	(15;1;2;0;2;0;0;0)	2.21%
10	(5;0;3;0;2;0;0;0)	2.29%	25	(20;1;2;0;2;0;0;0)	2.08%

Table 3

Appendix C: Proofs

Proof of Proposition 1

Step 1. Let $A_{q;r}$ denote the set of sequences of the n agents such that, in the equilibrium of the corresponding priority mechanism, exactly $q - r$ among the q first dates are assigned to agents in $I \backslash M_{q+1}$, and r or less to agents in M_{q+1}. Let $a_{q;r}$ denote the probability of $A_{q;r}$ (namely $\# A_{q;r} = n! a_{q;r}$). Then

$a_{Q,Q-q}$ is the probability that exactly q dates are eventually assigned. Clearly β_q equals the probability that at least q dates at least are assigned, so equation (6) follows.

Step 2: Three simple cases. We must prove that the probabilities $a_{q,r}$ are given by the recursive formulas (7) and (8). Consider first the case $r > q$. Formula (7) gives (by induction) $a_{q,r} = 0$, as desired.

Next consider (q, r) such that $n_q = 0$. Then equation (7) reads: $a_{q,r} = a_{q-1,r-1}$. Indeed, all binomial numbers in the formula are equal to 1 since the only value j can take is zero. On the other hand $a_{q,r}$ is the proba that $q - r$ of the q first dates be assigned to agents of types 1 to q. Since there are no agent of type q and since agents of types strictly smaller than q cannot consume the q-th date, the probability of this event is the same as the probability that among the $q-1$ first dates, $q - r [= (q-1) - (r-1)]$ be assigned to agents of types 1 to $q-1$, i.e., $a_{q-1,r-1}$.

Third, consider (q, r) such that $q \geq r > m_{q+1}$ and prove equation (8). Set $r = m_{q+1} + k \leq q$, with $k \geq 1$, then $a_{q,r}$ is the probability that, after assignment of the q first dates, $q - m_{q+1} - k$ be assigned to agents of types 1 to q; but there are only m_{q+1} agents of types $q+1$ to Q. Therefore ($k \geq 1$) the q-th date is not assigned. For such orderings, the last round, i.e., the assignment of a q-th date, can just be cancelled: all agents who were going to accept it have already been proposed a earlier date, and are served. Then $a_{q,m_{q+1}+k}$ is just the proba that at the preceding round $(q - m_{q+1} - k) - n_q$ dates have been served to agents of types smaller or equal to $q-1$, i.e., is equal to $a_{q-1,\underbrace{m_{q+1}+n_q+k-1}_{r-1+n_q}}$.

Step 3: It remains to consider a pair (q, r) such that $n_q > 0$ and $r \leq \min\{q, m_{q+1}\}$ and prove equation (7), giving $a_{q,r}$ as a function of $a_{q-1,\cdot}$:

For all j such that $0 \leq j \leq \min\{n_q, q-r\}$, let $A_{q,r}(j)$ be the subset of $A_{q,r}$ such that exactly j dates are assigned to agents of type q, and let $a_{q,r}(j)$ be its probability. We show now:

$$a_{q,r}(j) = \binom{m_q}{n_q}^{-1} \binom{r+j}{j} \binom{m_q-(r+j)}{n_q-j} a_{q-1,r-1+j}. \tag{15}$$

This establishes formula (7) in view of two facts:

$$a_{q,r} = \sum_{j=0}^{\min\{n_q,q-r\}} a_{q,r}(j) \qquad \text{and} \qquad \{j > q-r \Rightarrow a_{q-1,r-1+j} = 0\}.$$

(The latter fact follows from $q - 1 < r - 1 + j$ and Step 2.) In order to establish (15), we note that $A_{q,r}(j)$ is a subset of $A_{q-1,r-1+j}$, and we compute the conditional probability of $A_{q,r}(j)$ given $A_{q-1,r-1+j}$.

Consider a sequence in $A_{q-1,r-1+j}$. With respect to the assignment of the first $q-1$ dates, agents in M_q are equivalent (from a probabilistic point of view): they all accept any date offered. To make the final argument intuitive, in the considered ordering, we put a bar on top of each agent who ends up being served. Thus our ordering is in $A_{q-1,r-1+j}$ if and only if, upon reading the ordering from left to right up to the q-th bar, there are $q - r - j$ bars under which one sees an agent in $I \setminus M_q$. The key point is that this is independent of the relative ordering of agents in M_q under the remaining $r+j$ (or less) bars. Making this ordering precise, i.e., distinguishing agents in I_q from agents in M_{q+1}, is exactly what we need to deduce $a_{q;\cdot}$ from $a_{q-1;\cdot}$.

Note that, for an ordering in $A_{q;r}(j)$, there are (up to the q-th bar) exactly j bars under which is an agent of I_q and exactly r bars under which is an agent of M_{q+1} (the latter claim because $r \leq m_{q+1}$). To compute the desired conditional probability, we fix the $q - r - j$ bars (up to the q-th bar) on top of an agent in $I \setminus M_q$, and we compute the probability that the remaining $r+j$ bars are filled with exactly j agents in I_q and r agents in M_{q+1}. This goes first by choosing the j bars, out of those $r+j$, which are placed on top of agents of type q: there are $\binom{r+j}{j}$ ways to do so. And then one has to fill them with agents of type q, and there are, for the first bar, n_q such agents out of the m_q of type q and more, $n_q - 1$ choices out of $m_q - 1$ for the second, and so on, down until $n_q - j + 1$ choices out of $m_q - j + 1$ for the j-th bar; then one has to fill the r other bars: m_{q+1} choices out of the $m_q - j$ for the first one, $m_{q+1} - 1$ choices out of $m_q - j - 1$ for the second, and so on, down until $m_{q+1} - r + 1$ choices out of the $m_q - j - r + 1$ for the last one. Hence the chances that the ordering has the required feature is:

$$\frac{n_q \cdots (n_q - j + 1)\, m_{q+1} \cdots (m_{q+1} - r + 1)}{m_q(m_q - 1) \cdots (m_q - j - r + 1)} = \frac{\binom{m_q - (r+j)}{n_q - j}}{\binom{m_q}{n_q}}:$$

This concludes the proof of equation (15), and of Proposition 1.

Proof of Proposition 2

We fix the profile of types $(I_1, \ldots, I_Q)$ and proceed first to prove Proposition 2 in the case $Q \leq n$: Steps 1, 2 and 3. The case $Q > n$ is discussed in Step 4.

Step 1: Doubly stochastic cover of $\mathcal{P}^{PS}$. Denote by $N = I \times \{1, \ldots, n\}$ the set of entries of an assignment matrix. We denote the matrix $\mathcal{P}^{PS}$ as $[p^*_{ik}]$ and define three subsets of N:

$$\begin{aligned} F &= \{(i,k) \in N \mid p^*_{ik} > 0 \text{ and } k \leq q^*\} \\ A_1 &= \{(i,k) \in N \mid k \leq q^* \text{ and } i \notin M_k\} \\ A_2 &= \{(i,k) \in N \mid k \geq q^* + 1 \text{ and } i \in M_{q^*+1}\} \end{aligned}$$

Note that A_i may be empty, but that F is not. Figure 1 illustrates these 3 sets in the case of the six agents profile $(2,1,0,1,2)$ by conveniently assigning agents to rows by increasing types (the agents in I_1 are assigned to the top n_1 rows, etc.). The remaining set of indices, namely $B = N \setminus (F \bigcup A_1 \bigcup A_2)$ is then rectangular (on Figure 1 it occupies a top right corner of $\mathcal{P}^{PS}$):

$$(i,k) \in B \iff i \in I \setminus M_{q^*+1} \text{ and } q^* + 1 \leq k \leq n:$$

In this step, we prove the existence of a doubly stochastic matrix $P = [p_{ik}]$ that differs from $\mathcal{P}^{PS}$ only in

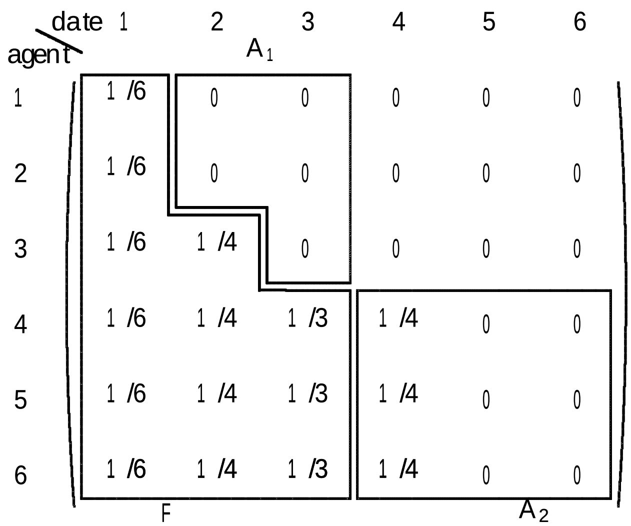

Figure 1

B, that is to say:

$$p_{ik} = p^*_{ik} \quad \text{on} \quad F \bigcup A_1 \bigcup A_2 \, ;$$
$$p_{ik} \geq 0 \quad \text{for all } i;k\,; \qquad \sum_{i \in I} p_{ik} = 1 \quad \text{for all } k\,; \qquad \sum_{k \in K} p_{ik} = 1 \quad \text{for all } i:$$

First we take care of the case $q^* = n$. Then B is empty, so we cannot alter Z^{PS}, but the latter is already doubly stochastic is this case. Indeed, $q^* = n$ implies $Q = n$ (because $q^* \leq Q$ and we assume $Q \leq n$ until Step 4); and the inequality: $\sum_{k=1}^{n} \frac{1}{m_k} \leq 1$, combined with $\frac{1}{m_k} \geq \frac{1}{n}$ for all k, gives $m_k = n$ for all k, $1 \leq k \leq n = Q$. This means that the profile of types is $I_k = \emptyset$ for $k < n$ and $I = I_n$, in which case all entries of Z^{PS} are $1/n$.

Let us now assume $q^* < n$, so that B is nonempty (if $M_{q^*+1} = I$, then $q^* = n$ by definition of q^*). We distinguish two cases for defining P on B.

Case 1: $M_{q^*+1} = \emptyset$ (namely $q^* = Q$). We set for all q, $1 \leq q \leq q^*$, for all $i \in I_q$, all k, $q^*+1 \leq k \leq n$: $p_{ik} = \frac{1-\mu_q}{n-q^*} \geq 0$, with the notation $\mu_q = \sum_{1}^{q} \frac{1}{m_k}$. In the i-row, for $i \in I_q$, we have $\sum_{1}^{q^*} p^*_{ik} = \mu_q$, hence $\sum_{1}^{n} p_{ik} = 1$. Check that the sum is 1 in the k columns $q^*+1 \leq k \leq n$:

$$\sum_{i \in I} p_{ik} = \sum_{q=1}^{q^*} n_q \frac{1-\mu_q}{n-q^*} = \frac{1}{n-q^*} \left(n - \sum_{q=1}^{q^*} n_q \mu_q \right) = 1\,;$$

where the last two equations follow from the definition of μ_q and $m_{q^*+1} = 0$.

Case 2: $M_{q^*+1} \neq \emptyset$. Here we choose differently the entries in column q^*+1 versus columns $q^*+2, \ldots, n$ (if any). Construct first the entries of column q^*+1: for all q, $1 \leq q \leq q^*$, for all $i \in I_q$: $p_{i;q^*+1} = \lambda(1-\mu_q) \geq 0$

where μ is adjusted to make this column sum to 1:

$$1 = \sum_{q=1}^{q^*} \sum_{i \in I_q} \mu(1-\theta_q) + \sum_{i \in M_{q^*+1}} p^*_{i;q^*+1}$$
$$= \mu \sum_{q=1}^{q^*} n_q(1-\theta_q) + \varepsilon m_{q^*+1} = \mu(n - q^* - \varepsilon m_{q^*+1}) + \varepsilon m_{q^*+1}$$

(we omit the details of the straightforward computation). Observe that by definition of q^*,

$$\frac{1}{m_{q^*+1}} + \theta_{q^*} > 1 \Rightarrow 1 > \varepsilon m_{q^*+1}$$

therefore $0 < \mu = \frac{1-\varepsilon m_{q^*+1}}{n - q^* - \varepsilon m_{q^*+1}} \leq 1$ and the definition of the (q^*+1)-column is complete. Now we define the columns q^*+2 to n. If $q^*+1 = n$, there is nothing to do and indeed the matrix P is doubly stochastic already, because $\mu = 1$. Assume now $n \geq q^*+2$ and define:

$$\text{for all } q, 1 \leq q \leq q^*, \text{ for all } i \in I_q, \text{ all } k, q^*+2 \leq k \leq n : p_{ik} = (1-\mu)\frac{1-\theta_q}{n-q^*-1} \geq 0 :$$

Check first that the i-row sums to 1, for all $i \in I \setminus M_{q^*+1}$:

$$\sum_1^{q^*} p^*_{ik} + p_{i;q^*+1} + \sum_{q^*+2}^{n} p_{ik} = \theta_q + \mu(1-\theta_q) + (1-\mu)(1-\theta_q) = 1 :$$

Next the definition of μ yields $(1-\mu)\frac{1-\theta_q}{n-q^*-1} = \frac{1-\theta_q}{n-q^*-\varepsilon m_{q^*+1}}$ and the verification that for all k, $q^*+2 \leq k \leq n$, the k column sums to 1 is now straightforward.

Step 3: Applying Birkhoff's theorem. A doubly stochastic matrix is the convex combination of permutation matrices, hence there is a probability distribution π on S such that: $P = \sum_{\sigma \in S} \pi_\sigma P_\sigma$. Note that the P matrix is zero in A_1 and in A_2 (except for the q^*+1 column); hence, for any permutation σ such that $\pi_\sigma > 0$, the matrix P_σ is also zero in these two subsets of N; in other words, we have:

$$\text{for all } k; 1 \leq k \leq q^* \quad : \sigma(k) = i \Rightarrow i \in M_k$$
$$\text{for all } k; q^*+2 \leq k \leq n \quad : \sigma(k) = i \Rightarrow i \notin M_{q^*+1}$$

This implies at once that in the equilibrium assignment corresponding to the priority ordering σ, the first q^* agents chose to buy and the last $n-q^*-1$ declined. In particular, $\bar{\sigma} = (i_1; \dots; i_{q^*})$ or $(i_1; \dots; i_{q^*}; i_{q^*+1})$. Moreover $P_{\bar{\sigma}} = \Delta(P_\sigma)$ where Δ is the (linear) operator cancelling the B-entries in a (given) matrix and leaving all other entries intact. We then get $Z^{PS} = \Delta(P) = \sum_{\sigma \in S} \pi_\sigma \Delta(P_\sigma) = \sum_{\sigma \in S} \pi_\sigma P_{\bar{\sigma}}$ and the proof of Proposition 2 is complete in the case $Q \leq n$.

Step 3: The case $n < Q$. Consider a profile of types $(I_1; \dots; I_Q)$ with $n < Q$. Distinguish two cases. If $q^* < n$, then the new profile $(I_1; \dots; I_{n-1}; M_n)$ (that is, the set of agents M_n in the initial profile equals the set of agents of type n in the new profile) yields precisely the same PS assignment to every agent (because the two profiles generate the same sets of agents of type 1 to q^*+1). But in the second profile, we can apply the above argument because the number of types is not larger that of agents, and we are home.

If, on the other hand, $q^* = n$, then an argument given in Step 2 shows that $I_1 = \dots = I_{n-1} = \emptyset$, $M_n = I$ and the matrix Z^{PS} is uniform so the desired statement holds.

Proof of Proposition 3

The proof is very simple, as can be checked on a couple of examples, so we only sketch it. The proof is identical for the RP and PS social choice functions, provided we set α_q to be the probability that the q-th date be assigned, with $\alpha_q = \beta_q$ or $\alpha_q = \gamma_q$ respectively. Let U^t, T^t and z^t be as the premises of the group strategyproof property and let q be the smallest type at which T^0 and T^1 differ:

$$I^0_k = I^1_k;\ \text{ for } k = 0, \dots, q-1;\ I^0_q \neq I^1_q:$$

This implies $M^0_k = M^1_k$ for $k = 0, \dots, q$, hence all agents in M_q have, in z^t, the same probability $p_{ik} = \alpha_k / m_k$ of consuming date k. Moreover, if $\alpha_{q+1} = 0$ (which implies that $\alpha_k = 0$ for all $k \geq q+1$) then $z^0 = z^1$ and we are home. Thus we assume from now on $\alpha_{q+1} > 0$. Distinguish two cases.

If there is an agent i in $I^1_k \setminus I^0_k$, this agent, in T^0, is of type at least $q+1$, hence in z^0 he consumes date $q+1$ with positive probability α_{q+1}/m_{q+1}; whereas in z^1 his consumption reduces to dates $1, \dots, q$: this implies $u^0_i \cdot z^0_i > u^0_i \cdot z^1_i$ (because $u^0_i(q+1) > 0$) and we are done.

If there is an agent i in $I^0_k \setminus I^1_k$, this agent, in T^1, is of type at least $q+1$, hence in z^0 he consumes date $q+1$ with positive probability α_{q+1}/m_{q+1}. But his true type is q, and up to period q z^0_i and z^1_i coincide: therefore $u^0_i \cdot z^0_i > u^0_i \cdot z^1_i$ holds true if $u^0_i(q+1)$ is negative. The only possibility is thus $u^0_i(q+1) = 0$.

Clearly the above two cases exhaust all possibilities. They establish the desired property 11 except perhaps in the case where, between z^0 and z^1, some agents report a type q^1, larger than their true type q^0, and are indifferent to the resulting change of assignment because their utility is zero for all dates q between $q^0 + 1$ and q^1, that they consume with positive probability after misreporting. In turn, this implies that everyone in the deviating coalition T is indifferent between z^0 and z^1, which concludes the proof. We omit the details.

Proof of Theorem 1

Step 1: Notation and preliminary remarks. We write α^t_q, $1 \leq q \leq Q$, the probability that date q is assigned by mechanism t, where $t = 0$ refers to RP and $t = 1$ refers to PS. Similarly, $p^t_q = \alpha^t_q / m_q$ denotes the probability that an agent in M_q consumes date q, and π^t_q is the probability that an agent of type q is served at all, namely: $\pi^t_q = \sum_{k=1}^{q} p^t_q$. We say the sequence $(x_1, \dots, x_T)$ of real numbers stochastically dominates another sequence $(y_1, \dots, y_T)$ if:

$$\sum_{r=1}^{t} x_r \geq \sum_{r=1}^{t} y_r;\qquad \text{for all } t;\ 1 \leq t \leq T: \tag{16}$$

We shall use the following well-known fact (Hardy, Littlewood and Polya (1934)): if $(\nu_1, \ldots, \nu_T)$ is a non-increasing sequence of nonnegative real numbers, and $(x_1, \ldots, x_T)$ stochastically dominates $(y_1, \ldots, y_T)$, then:

$$\sum_{r=1}^{t} \nu_r \cdot x_r \geq \sum_{r=1}^{t} \nu_r \cdot y_r\,; \qquad \text{for all } t;\ 1 \leq t \leq T\,: \tag{17}$$

Step 2: Proof of statement 1. This says that the sequence (p^1_q) stochastically dominates the sequence (p^0_q), and follows at once from the following properties:

$$\begin{array}{ll} \beta_q \leq \gamma_q & \text{for all } q;\ 1 \leq q \leq q^*; \\ \pi^0_q \leq \pi^1_q (=1) & \text{for all } q;\ q^*+1 \leq q \leq Q\,: \end{array}$$

Step 3: Proof of statement 4. We already know that $\alpha^t_q = 1$ for $t = 0, 1$ and all q, , $1 \leq q \leq \tilde{q}$ (see discussion after Definition 3). This implies statement 4, as well as the first inequality in (12).

Step 4: Proof of statement 2. Statement 1 says $\pi^0_q \leq \pi^1_q$ for all q; $1 \leq q \leq Q$. The total expected number of agents served by the mechanism t is: $q^t = \sum_{q=1}^{Q} n_q \pi^t_q$, hence $q^0 \leq q^1$, and the second inequality in (12) is proven.

Finally we show $q^1 \leq 2\tilde{q}$. Let q be any number, $1 \leq q \leq q^*$. Because the sequence m_k is nonincreasing, we have

$$1 \geq \sum_{k=q}^{q^*} \frac{1}{m_k} \geq \frac{1}{m_k}(q^* - q + 1): \tag{18}$$

Assume first q^* is odd, $q^* = 2q - 1$. Applying (18) to q yields $q \leq m_q$, hence $q \leq \tilde{q}$ so that $q^* < 2\tilde{q}$ and $q^1 \leq q^* + 1 \leq 2\tilde{q}$ as desired. Next suppose q^* is even, $q^* = 2q - 2$ and apply (18) to q. Note that the left-hand inequality must be strict: an equality would imply $q = 1$ (by definition of q^*) and $q^* = 0$. Therefore (18) implies:

$$q^* - q + 1 < m_q \iff q \leq m_q \iff q \leq \tilde{q} \implies q^* + 1 = 2q - 1 < 2\tilde{q}\,;$$

and we are done.

Step 5: Proof of statement 3. Consider an agent i of type q: his utility at the assignment z^t is $u_i \cdot z^t_i = \sum_{k=1}^{q} u_i(k) \cdot p^t_k$. By definition of type q, we have $u_i(k) \geq 0$ for $k = 1, \ldots, q$, and by Step 2, the sequence (p^1_k) stochastically dominates (p^0_k); this applies as well to their truncated versions where k runs from 1 to q. Hence the desired inequality $u_i \cdot z^0_i \leq u_i \cdot z^0_i$ (Step 1).

Step 6: Proof of statement 5. Assume $q^* \geq \tilde{q} + 1$. Consider an agent i of type q^1, $q^1 \geq \tilde{q} + 1$. We noticed earlier that $\beta_{\tilde{q}+1} < 1$. By assumption $q^* \geq \tilde{q} + 1$, we have now $p^0_{\tilde{q}+1} < p^1_{\tilde{q}+1}$. As $p^0_q = p^1_q$ for $q = 1, \ldots, \tilde{q}$, we deduce $\pi^0_{\tilde{q}+1} < \pi^1_{\tilde{q}+1}$. Now we complete the proof by invoking the following variant of the fact at the end of Step 1: if at least one of the inequalities (16) is strict and if the sequence σ is strictly decreasing, then inequality (16) is strict.

Step 7: Proof of the 'tightness' result. Fix q and denote by λ a large integer. Consider the profile of types: $|I_q| = \lambda \cdot q$; $|I_{2q}| = q$; all other I_k are empty. Check $q_e = q$. Next, when λ goes to infinity, the probability that a random ordering of I starts by q agents in I_q goes to one. When this happens, the number of agents served by RP is $2q$. Therefore, the expected value q^{RP} is arbitrarily close to $2q$ as λ grows large.

Proof of Property 14

Consider first the case where $c = \infty$. The function δ being decreasing, for n sufficiently big, one has

$$\frac{m_q}{q} \geq \frac{n}{q}\delta\left(\frac{q}{Q}\right) = \frac{n}{Q}\frac{\delta(q/Q)}{q/Q}. \quad (19)$$

Fix $x < \infty$. Take a sequence $\left(q^{(n,Q)}\right)$ such that $\liminf_{n\to\infty} \frac{q^{(n,Q)}}{Q} = x \in [0;\infty[$. Inequality (19) ensures us that there exists $(n_1;Q_1)$ such that $(n;Q) \geq (n_1;Q_1) \Rightarrow \frac{m_{q^{(n,Q)}}}{q^{(n,Q)}} > 1$. Then $q^{(n,Q)} \leq q_e^{(n,Q)}$. As a consequence, for all $x < \infty$, $\liminf_{n\ \mathrm{rightarrow} \infty} \frac{q_e^{(n,Q)}}{Q} \geq x$. Hence $\lim_{n\to\infty} \frac{q_e^{(n,Q)}}{Q} = \infty$. Consider then the case where $c < \infty$. Since $q_e \leq m_{q_e}$ and $q_e + 1 \leq m_{q_e+1}$, one has $m_{q_e+1} \leq q_e \leq m_{q_e}$, thus

$$1 - \frac{n_{q_e}}{m_{q_e+1}} \leq \frac{q_e}{m_{q_e}} \leq 1. \quad (20)$$

Take a sequence $\left(q^{(n,Q)}\right)$ such that $\lim_{(n,Q)\to\infty} \frac{q^{(n,Q)}}{Q} = x \in [0;\infty[$. One gets straightforwardly:

$$\lim_{(n,Q)\to\infty} \frac{m_{q^{(n,Q)}}}{q^{(n,Q)}} = c\frac{\delta(x)}{x}.$$

As a consequence, $x_e < \infty$. (Otherwise, if $x_e = \infty$, for all $(n;Q)$; $\frac{m_{q_e^{(n,Q)}}}{q_e^{(n,Q)}} \geq 1$ and $\delta(x_e) = 0$, a contradiction.) Moreover, if $x < \infty$ then $\frac{n_{q^{(n,Q)}}}{m_{q^{(n,Q)}}} \to 0$. (Indeed, $\lim_{n\to\infty} n_{q^{(n,Q)}} = -c\delta'(x)$ and $\lim_{n\to\infty} m_{q^{(n,Q)}} = \infty$.) Given that $x_e < \infty$, one then has $\frac{n_{q_e^{(n,Q)}}}{m_{q_e^{(n,Q)}}} \to 0$. Hence, through (20), $\frac{q_e^{(n,Q)}}{m_{q_e^{(n,Q)}}} \to 1$. The result obtains.