

Markovian Assignment Rules*

Francis Bloch
Ecole Polytechnique

and

David Cantala
El Colegio de Mexico

Abstract

We analyze dynamic assignment problems where agents successively receive different objects (positions, offices, etc.). A finite set of n vertically differentiated indivisible objects are assigned to n agents who live n periods. At each period, a new agent enters society, and the oldest agent retires, leaving his object to be reassigned. We define independent assignment rules (where the assignment of an object to an agent is independent of the way other objects are allocated to other agents), efficient assignment rules (where there does not exist another assignment rule with larger expected surplus), and fair assignment rules (where agents experiencing the same circumstances have identical histories in the long run). When agents are homogenous, we characterize efficient, independent and fair rules as generalizations of the seniority rule. When agents draw their types at random, we prove that independence and efficiency are incompatible, and that efficient and fair rules only exist when there are two types of agents. We characterize two simple rules (type-rank and type-seniority) which satisfy both efficiency and fairness criteria in dichotomous settings.

JEL Classification Numbers: C78, D73, M51

Keywords: dynamic assignment, finite Markov chains, seniority, promotion rules

*We thank Olivier Compte, Geoffroy DeClippel, Matt Jackson, Philippe Jehiel, Jean-François Laslier, Hervé Moulin, Nicolas Houy, Ütku Ünver, William Thomson and seminar participants at different institutions for their comments. This project was started when Francis Bloch visited the Department of Economics at Brown University. Financial support from CONACYT is gratefully acknowledged.

1 Introduction

This paper considers assignment rules to allocate a fixed number of vertically differentiated objects to an overlapping population of economic agents. These allocation problems arise frequently in organizations where exit and entry of new members lead to a reassignment of fixed resources like positions, offices, housing, priority rank, etc.. For example, how should one allocate the office of a faculty member who retires or leaves an academic department? How should one post career diplomats across embassies whose ranking by importance is commonly accepted? How should one transfer high level bureaucrats or high school teachers across regions in a centralized bureaucracy? How should one allocate campus housing among generations of students? How should one choose the priority sequence in serial dictatorship systems like the one used by airline pilots to select routes and schedules every month?

While all these assignment problems share a common structure, the allocation rules used in practice vary widely. In order to limit moving costs, department chairs will often choose to assign the offices of departing faculty members to newcomers, a rule that we will label *the replacement rule*. Career diplomats often climb up the hierarchy of diplomatic posts by being assigned to an embassy which is one step up the ladder, a rule that we will label *the rank rule*. Slots in campus dormitories are usually allocated according to random serial dictatorship: an ordering of students is picked at random, and students choose sequentially following this order. When all students agree on the ranking of the dormitories, this allocation method is equivalent to the *uniform rule*, where every agent has the same probability of receiving the object. Finally, airline pilots select routes and schedules according to a *seniority rule*, where the pilots' priority ranks are determined by the seniority list.² High level bureaucrats and high school teachers are transferred according to complex priority systems, taking into account seniority, on-the-job seniority and merit.³

In this paper, our objective is to characterize *efficient* and *fair* assign-

²The pilots' seniority lists is of paramount importance in many airlines. For example, the planned merger between Delta and Northwestern failed because pilots could not agree on the merger of the seniority lists of the two companies. See "NWA pilots set merger conditions," *Minneapolis Star Tribune*, January 18, 2008.

³See the illuminating account of the transfer of members of the Indian Administrative Service in Iyer and Mani (2008). The transfer of high school teachers in France, a concrete assignment problem which motivated one of us to undertake this study, is explained in detail in Appendix A. We note that in this example, agents have heterogeneous preferences over regions. Hence our model cannot be directly applied, but it casts some light on the assignment problem.

ment rules in allocation problems with overlapping generations of agents. We first consider a model with homogeneous agents. In this environment, all assignments generate the same value and are equally efficient, and we focus attention on fair rules which guarantee that in the long run all agents experience the same history. We prove that both the rank and the seniority rules are fair, but the replacement and uniform rules are not. We characterize a one-parameter family of rules (which are convex combinations of rank and seniority rules) as the unique fair assignment rules for which the assignment of an object to an agent is independent of the way other objects are assigned to other agents. We also analyze a model with heterogeneous players with different productivity levels. In this setting, optimal matchings are assortative and assign higher ranked objects to more productive agents. However, individual rationality constrains the reassignment of objects as agents cannot be forced to accept an object of lower rank than the one they currently hold. Taking this constrain into account, we characterize efficient allocation rules as lexicographic assignments where objects are first allocated according to the agents' productivity levels. If there are more than three levels of productivity, these efficient rules cannot be fair, and generate trajectories where different agents who were born and lived in societies characterized by the same distribution of agents' productivity levels, experience different histories. The incompatibility between efficiency and fairness does not arise in dichotomic societies where the type-rank and type-seniority rules, which allocate goods lexicographically, first taking into account an agent's productivity, and then selecting according to the rank or seniority rules among agents of the same productivity, are both efficient and fair.

Our characterization of efficient and fair assignment rules relies on the analysis of the dynamical system of assignments among overlapping agents. Because agents can always choose to keep the object they currently hold, the current assignment constrains the allocation of objects in future periods, generating the dynamics of the model. Each assignment rule gives rise to a finite Markov chain over the set of allocations. When the Markov chain is convergent, (or quasi-convergent in a model where the drawing of agents' productivity levels creates an exogenous source of randomness in the system), the dynamical system will eventually reach a steady state, at which point all agents will experience the exact same history. We thus interpret assignment rules which generate convergent and quasi-convergent rules as fair allocation rules. In the simpler case of homogeneous agents, we also investigate other dynamical properties of the system, providing sufficient conditions for assignment rules to be ergodic (generating a long run behavior which is independent of the initial conditions) or irreducible (achieving a long run

distribution where all assignments have positive probability). While some of our results are general and apply to all assignment rules, others are obtained for the special class of independent assignment rules, where the allocation of an object to an agent is independent of the way other objects are allocated to the other agents.

The intuition underlying our characterization of efficient and fair Markovian assignment rules can easily be grasped. When agents are homogeneous, the only candidate for a steady state of the dynamical system is a monotonic allocation, where older agents receive objects of higher rank. This monotonic allocation is an absorbing state of the Markov chain for the seniority and rank rules, but not for the uniform or replacement rules. Furthermore, it is easy to check that, starting from any initial condition, the rank and seniority rules lead to this monotonic allocation in a small number of steps, so that the Markov chain on assignments generated by these rules is convergent. The rank rule also satisfies the property of independence, and we show that the family of independent rules generating convergent Markov chains is a slight generalization of the rank rule, which allocates any object but the highest according to the rank rule, and allocates the highest object through a convex combination of the rank and seniority rules. When agents are heterogeneous, the incompatibility of fairness and efficiency is explained by the fact that allocations must be path dependent. Efficiency and individual rationality typically constrain the assignments along different trajectories of the dynamical system, and may generate different assignments for the same distribution of agent's productivity levels. When productivity levels can only take on two values, this path dependence disappears. Any efficient assignment partitions the set of objects into two classes, assigning higher objects to high productivity agents and lower objects to low productivity agents. When objects inside each class are allocated according to the seniority or rank rules, the dynamical system converges to assignments which respect the two-class structure, and allocate objects inside each class in a monotonic way.

In its vast majority, the literature on matching and allocation rules only considers static assignment problems (See Roth and Sotomayor (1990) and Thomson (2007) for exhaustive accounts of the literature on matching and allocation models respectively). In matching markets where agents enter/exit the market, Blum et al. (1997) and Cantala (2004) study the re-stabilization process triggered by such disruptions. The individual rationality condition, which is central to our analysis, was introduced by Abdulkadiroglu and Sönmez (1999) in the context of house allocations and incorporated to senior matching markets in Cantala and Sanchez (2008). Moulin and Stong (2002) and (2003) analyze a problem of allocation of balls to different urns which

bears some resemblance to the problem of allocation of objects to agents with different ages.

Very recently, there have been various attempts to embed assignment problems in dynamic contexts.⁴ Ünver (2008) analyzes kidney exchange in a dynamic framework, where pairs of kidney donors-recipients enter and exit the pool of matchable agents according to a random Poisson process. In a context where side-payments are allowed, Parkes and Singh (2003), Athey and Segal (2007), Bergemann and Valimäki (2006) and Gershkov and Moldovanu (2008a, 2008b) study dynamic assignment problems, where agents enter sequentially, and participate in a Vickrey-Clarke-Groves revelation mechanism which determines transfers and good allocations. They show that Vickrey-Clarke-Groves mechanisms and optimal stopping rules can be combined to obtain efficient dynamic mechanisms. In these models, objects can only be assigned once at the time of entry. Abdulkadiroglu and Loerscher (2007) analyze a dynamic version of the school choice problem, explicitly allowing for reallocation of students to schools and for a dynamical linkage between the different allocations. In their model, individual rationality constraints are not present, and the set of agents is fixed. Finally, Kurino (2008) analyzes a dynamic model of house allocation with overlapping agents. His starting point is very close to ours – he studies dynamic reallocation of campus housing with overlapping generations of students – but his modeling strategy is very different. He bases his analysis on a dynamic extension of Abdulkadiroglu and Sönmez (1999)’s study of housing allocation with existing tenants, and checks whether the rules analyzed by Abdulkadiroglu and Sönmez (1999) in a static setting keep the same efficiency and incentive properties in a dynamic environment.

The rest of the paper is organized as follows. We present an illustrative example with three agents and three objects in Section 2. The general model, notations and definitions are introduced in Section 3. Section 4 is devoted to the analysis of the model with homogeneous agents, and Section 5 considers the model with heterogeneous agents. Section 6 contains our conclusions and directions for future work.

⁴There is also an older literature in operations research and management which has studied dynamical control of matching processes. See for example Talluri and Van Rysin (2004) for an introduction to the literature on dynamic pricing and revenue management and Bartholomew (1982) and Nilakantan and Ragavhendra (2005) for an account of the literature on manpower planning.

2 An illustrative example

We introduce the main concepts and results of our analysis with an illustrative example. Suppose that there are three objects to be allocated, $j = 1, 2, 3$. Agents are homogeneous and have the same ranking for the objects, preferring 3 to 2 to 1. Each agent lives for three periods and any period t , $i = 1, 2, 3$ denotes the ages of the agents in the system. An assignment is a bijection between the set of agents and the set of objects. There are six possible assignments:

$$\begin{aligned}\mu_1 & : (1, 2, 3) \\ \mu_2 & : (1, 3, 2) \\ \mu_3 & : (2, 1, 3) \\ \mu_4 & : (2, 3, 1) \\ \mu_5 & : (3, 1, 2) \\ \mu_6 & : (3, 2, 1).\end{aligned}$$

In assignment μ_1 (the “monotonic assignment”), the youngest agent gets object 1, the middle agent object 2 and the oldest agent object 3. Assignment μ_6 is the reverse assignment, where the youngest agent gets the best object and the oldest agent the worst. The other assignments can be described in a similar way.

At any period, the oldest agent leaves the system and a new agent enters. The object held by the oldest agent, $\mu_i(3)$ is then reassigned to one of the agents in the system, according to the fixed assignment rule. This assignment may trigger a cascade of reallocations: if the object is assigned to agent 2 or agent 1, object $\mu_i(2)$ or $\mu_i(1)$ must in turn be reallocated. The cascade of reassignments ends when the entering agent receives an object. By individual rationality, agents cannot be forced to accept an object worse than the object they currently hold. This implies that object j , can only be assigned to an agent who holds an object $j' < j$. The entering agent can receive any of the objects $j = 1, 2, 3$.

For any allocation rule, this sequence of reallocations generates a Markov chain over the set of assignments $\mathcal{M} = \{\mu_1, \mu_2, \mu_3, \mu_4, \mu_5, \mu_6\}$. We now analyze the properties of this Markov chain for four natural assignment rules.

2.1 The seniority rule

The seniority rule assigns any object to the oldest agent who can receive it. The following example shows how one constructs a Markov chain for this assignment rule. Suppose that the current assignment is the monotonic assignment μ_1 . At the next period, object 3 is reassigned to the oldest agent who holds object 2. In turn, object 2 is reassigned to the middle agent, who holds object 1, and object 1 is finally reallocated to the entering agent. This shows that, when the current state is μ_1 , the next state will also be μ_1 with probability 1. The transition probability matrix of the Markov chain generated by the seniority rule is given by:

$$P = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}$$

The transitions between states are represented in Figure 1. The Markov chain generated by the seniority rule is convergent: the system converges to the steady state μ_1 in at most two steps.

2.2 The rank rule

The rank rule assigns object j to the agent who holds object $j - 1$. The transition probability matrix is given by:

$$P = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}$$

Figure 2 shows the transitions between states for the Markov chain generated by the rank rule. The dynamical properties of the Markov chain generated by the rank rule are very similar to those of the Markov chain generated by the seniority rule. The system converges to the steady state μ_1 in at most two steps.⁵

⁵The only difference between the two dynamical systems is that it takes two steps to go from μ_3 to μ_1 with the rank rule, and only one step with the seniority rule.

2.3 The uniform rule

The uniform rule assigns object 3 with equal probability to all three agents, object 2 with equal probability to the entering agent and the agent holding object 1, and object 1 to the entering agent. We compute the transition probability matrix of the uniform rule as:

$$P = \begin{bmatrix} \frac{1}{6} & \frac{1}{3} & \frac{1}{6} & 0 & \frac{1}{3} & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} & 0 & 0 & 0 \\ \frac{1}{3} & \frac{1}{6} & 0 & \frac{1}{6} & 0 & \frac{1}{3} \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}$$

The transitions between states are given in Figure 3. One can easily check that there is a path from any assignment to any assignment. Hence the Markov chain is irreducible, and in the long run, the system converges to an invariant distribution whose support is the entire set \mathcal{M} .⁶ We compute the invariant distribution as:

$$p_1 = \frac{27}{84} \simeq 0.32, p_2 = \frac{21}{84} \simeq 0.25, p_3 = \frac{15}{84} \simeq 0.18, p_4 = \frac{7}{84} \simeq 0.08, \\ p_5 = \frac{9}{84} \simeq 0.10, p_6 = \frac{5}{84} \simeq 0.06.$$

The invariant distribution puts the highest weight on the monotonic assignment μ_1 and the lowest weight on the reverse assignment μ_6 .

2.4 The replacement rule

The replacement rule assigns any object to the entering agent, and generates a Markov chain represented by the following transition probability matrix:

$$P = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}$$

⁶Any irreducible finite Markov chain admits a unique invariant distribution, which can be computed by solving the equation: $\pi P = \pi$ in π . (See e.g. Isaacson and Masden (1976), Theorem III.2.2 p. 69)

The Markov chain generated by the replacement rule results in two cycles, as illustrated in Figure 4. This Markov chain is not ergodic: the long run behavior of the system depends on the initial conditions.

2.5 Discussion

The study of the four natural assignment rules shows that two rules are convergent (the seniority and rank rules), one is irreducible (the uniform rule) and one is nonergodic (the replacement rule). In the long run, the rank and seniority rule guarantee that every agent in the system will experience the same history, receiving object 1 when young, then object 2 and object 3 when old. By contrast, neither the uniform rule nor the replacement rule guarantee that two agents born at different times will experience the same history in the long run. When objects are allocated uniformly (a mechanism which may at first glance appear to be “fair” but in fact results in inequitable outcomes), with positive probability, some agents may enjoy object 3 for the three periods, whereas other agents will keep object 1 for their entire life. The same inequitable outcomes arise with the replacement rule, where two agents born at successive dates will experience very different histories, as agents keep the same object throughout their entire life.

3 The Model

3.1 Markovian assignment rules

3.1.1 Agents and goods

We consider the allocation of n vertically differentiated indivisible goods to overlapping generations of agents. Time is discrete and runs as $t = 1, 2, \dots$. The set of goods is denoted J , and $j = 1, 2, \dots, n$ indexes the qualities of the goods.

Agents live for exactly n periods, and enter and exit society one by one. Hence, at any period t , there are exactly n agents present in the society. We let I denote the set of agents, and $i = 1, 2, \dots, n$ be the index of an agent’s age or cohort.

We consider two models: one where agents are *homogeneous* and are only indexed by their age i , and one where agents are *heterogeneous*, draw independently a productivity k in the set $K = \{1, 2, \dots, m\}$ according to the probability distribution $q(k)$. In heterogeneous societies, agents are characterized by the pair (i, k) . Alternatively, we can define the mapping θ , from

I to K , assigning a productivity to agents according to their age in every society. When agents are heterogeneous, we assume that the surplus of a match, $\sigma(j, k)$ is an increasing, supermodular function of agent's types and goods' qualities:

$$\begin{aligned}\sigma(j', k) &> \sigma(j, k) \text{ for all } j' > j, \\ \sigma(j, k') &> \sigma(j, k) \text{ for all } k' > k,\end{aligned}$$

$$\text{If } k' \geq k \text{ and } j' > j, \text{ then } \sigma(k', j') + \sigma(k, j) \geq \sigma(k', j) + \sigma(k, j'),$$

with strict inequality when $k' > k$. Hence, total surplus is maximized when the matching is *assortative*, assigning higher objects to agents of higher productivity.

3.1.2 Assignments

Definition 1 An assignment μ is a bijection from I to J assigning to every agent i the object $j = \mu(i)$. Given that I and J have the same cardinality, an assignment can be identified with a permutation over the finite set $\{1, 2, \dots, n\}$.

The set of all assignments is denoted \mathcal{M} . A state s in the economy is defined both by an assignment μ and a type profile θ , $s = (\mu, \theta)$. Hence, for each type profile, there are $n!$ states in a society of n agents, each corresponding to a different permutation over the finite set $\{1, 2, \dots, n\}$.

We now consider the dynamical structure of assignments. Let μ^{t-1} denote the assignment of the economy at date $t - 1$. At the end of period $t - 1$, the oldest agent alive in period $t - 1$, n , relinquishes object $\mu^{t-1}(n)$ and leaves the society. A new agent, 1, enters at t with no object – by convention we denote the null object by 0 – and agents $i = 2, 3, \dots, n$ retain the objects they were assigned in the previous period, $\mu^{t-1}(i - 1)$.

Definition 2 A truncated assignment ν is an injection from I to $J \cup \{0\}$, such that

- For all agents $i \neq 1$, $\nu(i) \in J$,
- $\nu(1) = 0$

A truncated assignment specifies the goods held by all agents but the entering agent in society, and assigns the null object to the entering agent. Clearly, for any truncated assignment ν , there is a unique object j which is unassigned, such that $\nu^{-1}(j) = \emptyset$. At the beginning of period t , the truncated assignment verifies $\nu(i) = \mu^{t-1}(i-1)$ for agents $i = 2, \dots, n$ and $\nu(1) = 0$ with object $j = \mu^{t-1}(n)$ unassigned. It will be assigned to one of the current members of society according to an assignment rule α . This, in turn, will free a new object to be reassigned, etc.. The cascade of re-assignments will end when the entering agent is assigned an object. We consider well-defined assignment rules, namely rules such that the re-assignment process does not cycle.

3.1.3 Assignment rules

We impose two restrictions on assignment rules. First, we suppose that agents cannot be forced to give back their current object, and can only receive an object which is better than the one they currently hold.⁷ Second, we focus on assignment rules which only depend on the current truncated assignment of objects to agents, and not on the entire history of assignments.⁸ Assignment rules which possess this Markovian property are called *Markovian assignment rules*. Formally, we define:

Definition 3 *A Markovian assignment rule is a collection of probability distributions $\alpha_j(\nu, \theta, i)$ in $[0, 1]^n$, describing the probability that agent i receives object j , given the truncated assignment ν (such that $\nu^{-1}(j) = \emptyset$) and the type profile θ . Moreover, $\sum_{i|\nu(i) < j} \alpha_j(\nu, \theta, i) = 1$.*

3.1.4 Examples of Markovian assignment rules

We now describe examples of Markovian assignment rules. The four assignment rules for homogeneous agents discussed in the previous Section can now formally be defined:⁹

⁷This assumption is of course an extreme simplification. It captures the idea that demotions and promotions are asymmetric, and that it is easier to promote than to demote, or that demotion incurs high costs.

⁸Notice that some common rules are *not* Markovian. For example, rules based on job seniority require information about the number of periods during which agent i has owned object j , an information which cannot be recovered from the current assignment ν .

⁹When agents are homogeneous, we omit the type profile θ in order to simplify notations.

The *seniority rule* assigns object j to the oldest agent with an object smaller than j , $\alpha_j(\nu, i) = 1$ if and only if $i = \max\{k | \nu(k) < j\}$.

The *rank rule* assigns object j to the agent who currently owns object $j - 1$, $\alpha_j(\nu, i) = 1$ if and only if $\nu(i) = j - 1$.

The *uniform rule* assigns object j to all agents who own objects smaller than j with equal probability, $\alpha_j(\nu, i) = \frac{1}{|\{k | \nu(k) < j\}|}$ for all i such that $\nu(i) < j$.

The *replacement rule* assigns object j to the entering agent, $\alpha_j(\nu, i) = 1$ if and only if $i = 1$.

When agents are heterogeneous, an important class of assignment rules are *type-lexicographic rules* which assign objects using productivity as the first criterion.

Definition 4 A Markovian assignment rule α is type-lexicographic if, for any j, ν and θ , $\alpha_j(\nu, \theta, i) > 0 \Rightarrow \theta(i) \geq \theta(k) \forall k, \nu(k) < j$.

The type-seniority and type-rank are two examples of type-lexicographic rules, which assign objects among agents with the same productivity according to seniority and rank respectively:

The *type-seniority rule* is defined by $\alpha_j(\nu, \theta, i) = 1$ if $\theta(i) \geq \theta(k)$ for all k such that $\nu(k) < j$ and $i > l$ for all l such that $\theta(l) = \theta(i)$ and $\nu(l) < j$.

The *type-rank rule* is defined by $\alpha_j(\nu, \theta, i) = 1$ if $\theta(i) \geq \theta(k)$ for all k such that $\nu(k) < j$ and $\nu(i) > \nu(l)$ for all l such that $\theta(l) = \theta(i)$ and $\nu(l) < j$.

3.2 Markov processes generated by assignment rules

Given an assignment rule α and a probability distribution q over types, we can compute the transition probability from state $s = (\theta, \mu)$ to state $s' = (\theta', \mu')$:

- First, the conditional probability of type profile θ' given type profile θ is given by:

$$\begin{aligned} q(\theta' | \theta) &= q(k) && \text{if } \theta'(1) = k, \theta'(i) = \theta(i - 1), i = 2, \dots, n, \\ q(\theta' | \theta) &= 0 && \text{otherwise.} \end{aligned}$$

- Second, given the new type profile θ' , and the assignment rule α , construct sequences of reassignments defining the probability of reaching state μ' from state μ , $p(\mu' | \mu)$ as follows:

1) Consider the sequence of agents $i^0 = n + 1, i^1 = \mu'^{-1}(\mu(i^0 - 1)), \dots, i^m = \mu'^{-1}(\mu(i^{m-1} - 1)), \dots, i^M = 1$. This sequence of agents corresponds to the unique sequence of reallocations of goods for which society moves from assignment μ to assignment μ' . First, the good held by the last agent at date t , $\mu(n)$ is assigned to agent $i^1 = \mu'^{-1}(\mu(n))$. Then the good held by agent i^1 at period $t + 1$ (or by agent $i^1 - 1$ at period t) is reallocated to the agent $i^2 = \mu'^{-1}(\mu(i^1 - 1))$, etc. The process continues for a finite number of periods until a good is assigned to agent $i^M = 1$, after which no other good can be reallocated.

2) The probability of reaching μ' from μ is thus simply the probability that the sequence of reallocations of goods between agents i^0, \dots, i^M is realized:

$$p(\mu'|\mu) = \prod_{m=0}^{M-1} \alpha_{\mu(i^{m+1})}(\nu^m, \theta, i^{m+1}) \quad (1)$$

where $\nu^m(i) = \mu(i - 1)$ for $i \neq i^t, t = 1, 2, \dots, m$ and $\nu^m(i) = \mu'(i)$ for $i = i^t, t = 1, 2, \dots, m$.

Hence the transition probability from state $s = (\theta, \mu)$ to state $s' = (\theta', \mu')$ is $q(\theta'|\theta)p(\mu'|\mu)$.

It is interesting to note that while any assignment rule α generates a unique Markov process on the set of assignments and type profiles, the converse statement is not true. One can easily construct Markov processes over the set of states which cannot be generated by Markovian assignment rules, as shown by the following example.¹⁰

Example 1 Consider three homogeneous agents, and consider a Markov process over states such that $p(\mu_1|\mu_1) = 1$ and $p(\mu_3|\mu_2) = 1$. Let the truncated assignment be given by $\nu(1) = 0, \nu(2) = 1$ and $\nu(3) = 3$. Because $p(\mu_1|\mu_1) = 1$, we must have $\alpha_2(\nu, 2) = 1$. However, because $p(\mu_3|\mu_2) = 1$ we must also have $\alpha_2(\nu, 1) = 1$, a contradiction.

3.3 Efficiency and Independence

3.3.1 Efficient assignment rules

When agents are heterogeneous, the total surplus varies with the assignments, and different assignment rules result in different total surpluses. We define a notion of *efficiency* of assignment rules, based on the following (static) criterion.

¹⁰We are indebted to Nicolas Houy for this example.

Definition 5 An assignment rule α is efficient if there does not exist another assignment rule α' , such that, for every state $s = (\mu, \theta)$

$$\sum_{\theta'} \sum_{\mu'} \sum_i q(\theta'|\theta) p'(\mu'|\mu) \sigma(\theta'(i), \mu'(i)) \geq \sum_{\theta'} \sum_{\mu'} \sum_i q(\theta'|\theta) p(\mu'|\mu) \sigma(\theta'(i), \mu'(i)),$$

with strict inequality for some state.

Some remarks are in order. First, because the assignment rule α can be made conditional on the type profile θ' , definitions of efficiency ex ante and at the interim stage (after the type profile θ' has been drawn), are equivalent.¹¹ Second, this definition of efficiency is static, and only considers total surplus at the next step of the Markov chain, and not the surplus generated by the two assignment rules α and α' along the entire path of the Markov chain, or in the long run. Third, by using this definition we impose the same constraint on the assignment rules α and α' , and in particular, we do not consider efficiency improving reassignments which would violate the individual rationality condition, namely the fact that an agent holding object j cannot be reassigned an object of value smaller than j . Finally, when agents are homogeneous, any assignment rule is efficient.

3.3.2 Independent assignment rules

A Markovian assignment rule may condition the assignment of object j to agent i on the objects currently held by the other agents (the truncated assignment ν) and the profile. A simple property of Markovian assignment rules is *independence*, stating that the assignment of object j to player i does not depend either on the current assignment of objects held by the other players, nor on their types:

Definition 6 A Markovian assignment rule $\alpha \in \mathcal{M}$ satisfies independence if, for any i, j , ν, ν' such that $\nu(i) = \nu'(i)$, $\nu^{-1}(j) = \emptyset$, and $\theta(i) = \theta'(i)$, $\alpha_j(\nu, \theta, i) = \alpha_j(\nu', \theta', i)$.

The independence property is appealing because it states that an agent's assignment only depends on his characteristics (age, productivity, and object currently held) and not on the characteristics of the other agents. A stronger independence property states that an agent's assignment is also independent of his age:

¹¹This is reminiscent of the equivalence between the definition of Bayesian equilibria using ex ante or interim calculations. See Fudenberg and Tirole (1991) p. 215.

Definition 7 A Markovian assignment rule $\alpha \in \mathcal{M}$ satisfies strong independence if, for any i, j, k, ν, ν' such that $\nu(i) = \nu'(k)$, $\nu^{-1}(j) = \emptyset$, and $\theta(i) = \theta'(k)$, $\alpha_j(\nu, \theta, i) = \alpha_j(\nu', \theta', k)$,

Notice that the rank, uniform and replacement rules are all strongly independent. The seniority rule is not independent, but satisfies a weaker independence property, stating that the assignment $\alpha_j(\nu, i)$ only depends on the truncated assignment of objects for agents who currently hold objects smaller than j and are thus eligible to receive object j . Formally:

Definition 8 A Markovian assignment rule $\alpha \in \mathcal{M}$ satisfies weak independence if, for any i, j, ν, ν' such that $\nu(i) = \nu'(i)$, $\nu(k) = \nu'(k)$ for all k such that $\nu(k) < j$, $\nu^{-1}(j) = \emptyset$, $\theta(i) = \theta'(i)$, and $\theta(k) = \theta'(k)$, for all k such that $\nu(k) < j$, $\alpha_j(\nu, \theta, i) = \alpha_j(\nu', \theta', i)$.

3.4 Fairness and Dynamical Properties of Markovian rules

3.4.1 Fair assignment rules

We propose a definition of fairness based on the comparison of histories experienced by agents born at different times. An assignment rule will be called fair if, in the long run, two agents born at different times in the same circumstances experience the same history. More precisely, we interpret “same circumstances” as meaning that the agents belong to societies with the same distribution of types throughout their lifetime. Formally,

Definition 9 An assignment rule is fair if for any two agents i and i' entering society at dates t and t' , any realization of type profiles such that $\theta^{t+\tau} = \theta^{t'+\tau}$ for $\tau = 0, 1, \dots, n-1$, the assignment rule α generates a deterministic sequence of assignments such that $\mu^{t+\tau}(i) = \mu^{t'+\tau}(i')$ for $\tau = 0, 1, \dots, n-1$.

Our definition of fairness captures different aspects. First, it excludes probabilistic rules because they generate different outcomes for identical agents. Second, it guarantees that agents receive the same sequence of assignments along their lifetime when they live in the same circumstances. Hence, this definition compares deterministic trajectories rather than expected payoffs. In that sense, this is a stronger notion than a notion based on expected lifetime outcomes. When agents are homogeneous, this definition of fairness collapses into a simple requirement: an assignment rule is fair if it is deterministic and guarantees the same sequence of assignments to all the agents.

3.4.2 Convergent, ergodic, irreducible assignment rules

There is a close connection between the notion of fairness and convergence of the Markov process generated by an assignment rule. In order to illustrate this connection, we first recall the definition of well-known dynamical properties of Markov chains (See Kemeny and Snell (1960) or Isaacson and Madsen (1976)).

Definition 10 *Two states s and s' intercommunicate if there exists a path in the Markov chain from s to s' and a path from s' to s .*

Definition 11 *A set of states C is closed if, for any states $s \in C$, $s' \notin C$, the transition probability between s and s' is zero.*

Definition 12 *A recurrent set is a closed set of states such that all states in the set intercommunicate. If the recurrent set is a singleton, it is called an absorbing state.*

With these definitions in hand, we can define dynamical properties of the Markov chains generated by the assignment rule:

Definition 13 *A Markovian assignment rule α is irreducible if the induced Markov chain is irreducible (the only recurrent set is the entire state set). A Markovian assignment rule α is ergodic if the induced Markov chain is ergodic (has a unique recurrent set).¹² A Markovian assignment rule α is convergent if the induced Markov chain is convergent (admits a unique absorbing state, and any initial assignment converges to the absorbing state).*

when agents are heterogeneous, while the notions of ergodic and irreducible assignment rules are well defined, it is clear that assignment rules *are never convergent*. The random drawing of the type of the entering agent every period introduces a source of randomness in the Markov chain which prevents the existence of absorbing states. However, distinguishing between the two sources of randomness (one linked to the exogenous drawing of the type of the entering agent every period, and one to the dynamics of reassignments), we propose the following notion of *quasi-convergence*

¹²This definition of ergodicity does not agree with the definition given by Isaacson and Masden (1976) who also require all recurrent states to be aperiodic, so that an invariant distribution exists, nor with Kemeny and Snell (1960)'s definition where an ergodic Markov chain is defined by the fact that the only recurrent set is the entire state set. For lack of better terminology, we call ergodic a finite Markov chain such that the long run behavior of the chain (whether it is a cycle or an invariant distribution) is independent of the initial conditions.

Definition 14 *A Markovian assignment rule α is quasi-convergent if the induced Markov chain has a unique recurrent set of n^m states S such that, for any s, s' in S , $\theta(s) \neq \theta(s')$.*

In words, a quasi-convergent Markov chain ultimately settles in a recurrent state, where a single assignment arises for every type profile θ . When there is a unique type, this definition is of course equivalent to convergence to a unique absorbing state. Quasi-convergent and fair allocation rules are related by the following observation:

Remark 1 *A Markovian assignment rule with heterogeneous agents is fair if and only if any recurrent set S of the induced Markov chain contains exactly n^m states such that for any s, s' in S , $\theta(s) \neq \theta(s')$.*

Hence, any quasi-convergent assignment rule is fair; conversely, any fair rule such that the long run trajectory of agents is independent of initial conditions is quasi-convergent. In order to study fair allocation rules, we will thus focus attention on convergent rules in the model with homogeneous agents, and quasi-convergent rules in the model with heterogeneous agents.

4 Markovian assignment rules with homogeneous agents

In this Section, we study Markovian assignment rules when agents are homogeneous. We first characterize independent assignment rules, then analyze the dynamical properties of the Markov chains generated by assignment rules.

4.1 Independent assignment rules

The following Lemma characterizes assignment rules satisfying independence, and highlights the gap between independence and strong independence.

Lemma 1 *If a Markovian rule α satisfies independence, then for any $j < n$, ν, ν' and i, k such that $\nu(i) = \nu'(k)$, $\alpha_j(\nu, i) = \alpha_j(\nu', k)$. Furthermore, for any ν, ν' such that $\nu(i) = \nu'(j)$, $\nu(j) = \nu'(i)$, $\alpha_n(\nu, i) + \alpha_n(\nu, j) = \alpha(\nu', i) + \alpha_n(\nu', j)$.*

Lemma 1 shows that if a Markovian assignment rule satisfies independence, the assignment of any object $j < n$ is strongly independent, and fully

determined by the probabilities $\alpha_j(k)$ of assigning object j to an agent currently holding object $k < j$. However, this property *does not hold for the assignment of the highest object*, n . For the assignment of the last object, the only constraint imposed by independence is that, for any two assignments which only differ in the positions of i and j , the total probability assigned to agents i and j be constant. As the following simple example shows, there exist assignment rules satisfying independence which allocate object n with different probabilities to two agents of different ages holding the same object.

Example 2 *Let $n = 3$. Consider the assignment of object 3 and the two truncated assignments $\nu(2) = 1, \nu(3) = 2, \nu'(2) = 2, \nu'(3) = 1$. Independence puts no restriction on the assignment rule α_3 , as there is no agent i for which $\nu(i) = \nu'(i)$. Now, we must have: $\alpha_3(\nu, 1) + \alpha_3(\nu, 2) + \alpha_3(\nu, 3) = 1 = \alpha_3(\nu', 1) + \alpha_3(\nu', 2) + \alpha_3(\nu', 3)$. This implies that the assignment rules satisfying independence are characterized by three numbers, $\alpha_3(\nu, 1), \alpha_3(\nu, 2)$ and $\alpha_3(\nu', 2)$, but it does not imply that $\alpha_3(\nu, 2) = \alpha_3(\nu', 3)$ nor $\alpha_3(\nu', 2) = \alpha_3(\nu, 3)$.*

4.2 Convergent Markovian assignment rules

We first characterize convergent assignment rules. Recall that an agent is never reassigned an object of lower value than the object she currently holds. Hence, for any $i = 1, \dots, n - 1$, $\mu^{t+1}(i + 1) \geq \mu^t(i)$. If an assignment μ is an absorbing state, we must have

$$\mu(i + 1) = \mu^t(i + 1) = \mu^{t+1}(i + 1) \geq \mu^t(i) = \mu(i). \quad (2)$$

Hence, at an absorbing state, the assignment must be *monotone*, assigning higher objects to older agents. The only monotone assignment is the identity assignment ι for which $\iota(i) = i$ for all $i = 1, \dots, n$. Hence, *the only candidate absorbing state is the identity assignment ι .*

Proposition 1 *Both the seniority and rank assignment rules are convergent.*

Proposition 1 shows that both the seniority and rank rules are convergent and that the absorbing state is reached in at most n periods. Furthermore, a careful inspection of the proof of the Proposition reveals that any Markovian assignment rule which can be written as a convex combination of the rank and seniority rule, is also convergent. However, the seniority and rank rules (and their convex combinations) are not the only convergent rules. A complete

characterization of convergent assignment rules is difficult, because the condition guaranteeing that the identity assignment is absorbing only pins down the assignment rule for the truncated assignments $\tilde{\nu}^j$, where $\tilde{\nu}^j(i) = i - 1$ for $i \leq j$ and $\tilde{\nu}^j(i) = i$ for $i > j$, but does not impose any conditions for other assignments. When assignments are independent of the assignments of the other agents, progress can be made and the next Theorem characterizes the one-parameter family of independent convergent rules.

Theorem 1 *An assignment rule α is independent and convergent if and only if $\alpha_j(j - 1) = 1$ for all $j < n$, $\alpha_n(\nu, n) = 1$ if $\nu(n) = n - 1$, and there exists $\lambda \in [0, 1]$ such that $\alpha_n(\nu, n) = \lambda$ and $\alpha_n(\nu, \nu^{-1}(n - 1)) = 1 - \lambda$ if $\nu(n) \neq n - 1$.*

Theorem 1 characterizes the family of independent and convergent assignment rules as rules which allocate any object $j < n$ according to the rank rule, and allocate object n according to a convex combination of the rank and seniority rules. If, in addition, we require the assignment rule to be strongly independent, if $\alpha_n(\nu, n) = 1$ when $\nu(n) = n - 1$, we must have $\alpha_n(n - 1) = 1$, so that:

Corollary 1 *The only strongly independent, convergent assignment rule is the rank rule.*

4.3 Ergodic assignment rules

We first recall some definitions of special permutations.

Definition 15 *A permutation from a set of n elements to itself is a cycle, denoted κ , if $\pi(i) = \pi(i + 1)$ for all $i = 1, 2, \dots, n$ and $\pi(n) = 1$.*

Definition 16 *A permutation from a set of n elements to itself is an (i, j) transposition, denoted $\tau_{i,j}$, if $\pi(i) = j$, $\pi(j) = i$ and $\pi(k) = k$ for all $k \neq i, j$. For a shorthand, we will denote any $(1, i)$ transposition as τ_i .*

Using these definitions, we can decompose the evolution of the Markov chain as a composition of cycles and transpositions. Consider an initial state μ at period t , and succession of reassignments i^0, \dots, i^M . The state μ' at period $t+1$ is obtained by (i) first applying a cycle, which lets object $\mu(n)$ be assigned to the entering agent, (ii) then applying a transposition between agent 1 and agent i^1 , assigning object $\mu(n)$ to i^1 , (iii), then applying a transposition between agent 1 and agent i^2 , assigning object $\mu(i^1)$ to agent i^2 , etc... Hence, we may write:

$$\mu' = \mu \circ \kappa \circ \tau_{i^1} \dots \circ \tau_{i^m} \circ \dots \circ \tau_1,$$

where it is understood that τ_1 , the identity permutation, is just added for the sake of completeness, and to signal the end of the process of composition of permutations.

We are now ready to provide a simple characterization of ergodic assignment rules based on the accessibility of an assignment where the highest object is assigned to the oldest player.

Theorem 2 *An assignment rule α is ergodic if and only if there exists an assignment μ' with $\mu'(n) = n$ such that, for all assignments μ with $\mu(n) = n$, the permutation $\mu^{-1}(\mu')$ can be decomposed into a sequence of permutations, $\mu^{-1}(\mu'^1 \circ \dots \circ \pi^m \circ \dots \pi^M)$ such that either π^m is a cycle or a $(1, i)$ transposition and, if it is a $(1, i)$ transposition, $\alpha_{(\mu \circ \dots \circ \pi^{m-1})^{-1}(1)}(\nu^{m-1}, i) > 0$, where $\nu^{m-1}(j) = (\mu \circ \dots \circ \pi^{m-1})(j)$ for all $j = 2, \dots, n$.*

Theorem 2 is based on the simple observation that any recurrent set must contain an assignment for which $\mu(n) = n$, so that in order to check ergodicity, one only needs to check that there exists an assignment assigning the highest object to the oldest agent which can be reached from any assignment assigning the highest object to the oldest agent. This condition is always violated for the *replacement rule*, for which the set of states can be decomposed into n cycles, each cycle containing a single assignment such that $\mu(n) = n$, and for which there is no path between the cycles.

Proposition 2 does not pin down a simple condition guaranteeing the existence of a path in the Markov chain from an assignment μ to an assignment μ' with $\mu(n) = \mu'(n) = n$. A simple sufficient condition is that any object i is assigned with positive probability to an agent of age i holding object $i - 1$:

Corollary 2 *Suppose that $\alpha_i(i, \nu) > 0$ whenever $\nu(i) = i - 1$, then the assignment rule α is ergodic.*

Corollary 2 generalizes our result on the convergence of the rank and seniority rules, by showing that any assignment rule which assigns object i to agent i when he holds object $i - 1$ with positive probability (a condition satisfied both by the rank and seniority rule) must be ergodic. Furthermore, if the condition of Corollary 2 is satisfied, then it is possible to reach the identity assignment ι from itself, so that the period of the recurrent state ι is equal to one. As all states in a recurrent set must have the same period (Isaacson and Masden (1976), Theorem II.2.2 p.54), all states in the unique recurrent

set are aperiodic. Hence, the Markov chain is ergodic in the stronger sense of Isaacson and Masden (1976), and admits a unique invariant distribution.

The sufficient condition identified in Corollary 2 is not necessary. As the following four player example shows, a Markovian assignment rule may be ergodic even when it allows some "gaps" (situations where the probability of assigning object j to the agent holding object $j - 1$ is equal to zero).

Example 3 *Let $n = 4$. Consider the strongly independent assignment rule $\alpha_4(3) = 1, \alpha_3(1) = 1, \alpha_2(1) = 1, \alpha_1(0) = 1$ ¹³.*

Let all states such that $\mu(4) = 4$ be ordered as in Subsection 2. In addition, define the states:

$$\begin{aligned}\mu_7 &: (1, 3, 4, 2) \\ \mu_8 &: (1, 2, 4, 3) \\ \mu_9 &: (1, 4, 3, 2) \\ \mu_{10} &: (1, 4, 2, 3) \quad .\end{aligned}$$

Figure 5 illustrates the transitions between these states and shows that there exists a path leading to the identity matching from any other state, proving that the assignment rule is ergodic.

4.4 Irreducible assignment rules

In this Subsection, we characterize irreducible assignment rules, generating irreducible finite Markov chains, where any state can be reached from any other state.

Theorem 3 *An assignment rule α is irreducible if and only if*
(i) For all j , all truncated assignments ν of objects in $J \setminus j$, $\alpha_j(\nu, 1) > 0$ and
(ii) For all assignments μ, μ' such that $\mu(n) = \mu'(n) = n$, the permutation $\mu^{-1}(\mu')$ can be decomposed into a sequence of permutations, $\mu^{-1}(\mu^1 \circ \dots \circ \pi^m \circ \dots \pi^M$ such that either π^m is a cycle or a $(1, i)$ transposition, if π^m is a $(1, i)$ transposition, both π^{m-1} and π^{m+1} are cycles, and $\alpha_{(\mu \circ \dots \circ \pi^{m-1})^{-1}(1)}(\nu^{m-1}, i) > 0$, where $\nu^{m-1}(j) = (\mu \circ \dots \circ \pi^{m-1})(j)$ for all $j = 2, \dots, n$.

Theorem 3 provides a characterization of irreducible assignment rules which relies on two conditions: (i) assumes that replacement (the allocation of any object to the entering agent) occurs with positive probability at

¹³For strongly independent assignment rules, we simplify notations and let $\alpha_j(i)$ denote the probability that the agent who currently holds object i receives object j .

all states, (ii) assumes that any two assignments which allocate the highest object to the oldest agent are related through a sequence of elementary permutations, with cycles and $(1, i)$ transpositions such that any transposition in the sequence is followed by a cycle.

At first glance, condition (ii) may appear to be a mere rephrasing of the irreducibility condition – guaranteeing that any state can be reached from any state. However, condition (ii) is weaker than the irreducibility condition, as it only applies to a set of states of cardinality $(n - 1)!$ rather than $n!$ Condition (ii) also focusses attention on a special sequence of “elementary permutations” rather than arbitrary assignments. When condition (i) is satisfied, any path from a state μ to a state μ' can be generated through elementary permutations. Hence, in the direction of sufficiency, requiring that the states can be reached through elementary permutations is not more demanding than requiring that the states can be reached through any arbitrary reassignment. In the direction of necessity, checking that there is no elementary permutations leading from one state to another is easier than checking that states cannot be reached through any reassignment. Furthermore, the description of elementary permutations will serve as a building block for the analysis of irreducible assignment rules satisfying independence.

Theorem 4 *For any independent assignment rule α , consider the graph $G(\alpha)$ defined over the nodes $\{1, 2, \dots, n - 1\}$ by $g_{i,j} = 1$ if and only if either $\alpha_j(i) > 0$ or $\alpha_i(j) > 0$. Any independent Markovian assignment rule α such that $\alpha_j(0) > 0$ for all $j \geq 1$, and for which the graph $G(\alpha)$ is connected is irreducible.*

Theorem 4 provides a simple sufficient condition to check whether an independent assignment rule is irreducible. This condition is satisfied when the set of states for which transitions occur with positive probability is rich enough. For example, it is always satisfied for the uniform assignment rule where $\alpha_j(i) > 0$ for all $i \leq j$, or when the probability of assigning object j to an agent holding $j - 1$ is positive, $\alpha_j(j - 1) > 0$ (in which case the graph $G(\alpha)$ is a connected line), or if the probability of assigning object j to the agent holding object 1 is positive for all j , $\alpha_j(1) > 0$ (in which case the graph $G(\alpha)$ is a connected star with 1 as the hub).

However, as shown by the following example, the condition is not necessary. There exist irreducible assignment rules for which the graph $G(\alpha)$ is not connected.

Example 4 *Let $n = 4$. Consider the strongly independent assignment rule, $\alpha_1(0) = 1, \alpha_2(0) = 1, \alpha_3(0) = \alpha_3(1) = \frac{1}{2}, \alpha_4(0) = \alpha_4(1), \alpha_4(2) = \alpha_4(3) = \frac{1}{4}$.*

In this Example, the graph $G(\alpha)$ only contains the link $(1, 3)$ and is not connected. However, all assignments with $\mu(n) = 4$ intercommunicate, as illustrated in Figure 6, which uses the same ordering of three player assignments as that used in Subsection 2.

5 Markovian assignment rules with heterogeneous agents

In this Section, we allow for heterogeneity across agents. We first prove that independence and efficiency are incompatible, then study rules satisfying efficiency and quasi-convergence.

5.1 Independent and efficient assignment rules

We first observe that independence places very strong restrictions on assignment rules with heterogeneous agents.

Lemma 2 *Let α be an independent assignment rule among heterogeneous agents. Then, for any $\theta, \theta' \in \Theta$, any $j \in J$, $\nu \in \mathcal{T}$ and $i \in I$, $\alpha_j(\nu, \theta, i) = \alpha_j(\nu, \theta', i)$.*

With heterogeneous players, independence thus limits the set of rules to those rules which do not depend on agents' types and satisfy independence for homogeneous players (e.g. the rank or uniform rules, which do not take into account players' types). Our next result characterizes efficient Markovian assignment rules, and shows that they exhibit a simple lexicographic structure.

Theorem 5 *An assignment rule α is efficient if and only if it is type-lexicographic.*

We combine Theorem 5 and Lemma 2 to show that efficiency and independence are incompatible when agents are heterogeneous. Theorem 5 shows that when there exists at least two types of agents, there exist states where an efficient rule must allocate objects according to agents' types. Hence, efficient rules cannot satisfy the necessary condition for independence derived in Lemma 2.

Corollary 3 *Suppose that the set K contains at least two types. Then there is no assignment rule satisfying independence and efficiency.*

5.2 Efficient and fair assignment rules

We finally characterize rules satisfying both efficiency and quasi-convergence. When $K = 2$, any lexicographic rule which first assigns objects to agents of high type, and then selects among them according to the seniority or rank rule satisfies both efficiency and quasi-convergence. When $K > 2$, we show that there does not exist any rule satisfying both efficiency and quasi-convergence.

Theorem 6 *Suppose that the set K contains at least three types. Then there is no assignment rule satisfying efficiency and quasi-convergence.*

The proof of Theorem 6 relies on a simple argument, which is worth reproducing here. Consider states where only three types L, M, H can be drawn. Because the assignment rule is quasi-convergent, at states where all agents have the same type, the only candidate assignment is the identity assignment (See Subsection 4.2). Hence, if type (L, L, L) is realized, $\mu(i) = i$. Similarly, if type (M, M, M) is realized, $\mu(i) = i$. Now consider successive changes in the type profiles, from (L, L, L) to (M, L, L) , (H, M, L) and (H, H, M) ; and from (M, M, M) to (H, M, M) and (H, H, M) . Applying Theorem 5, an efficient rule is type lexicographic, and objects must be assigned according to agents' types. However, these assignments result in two different assignments according to the path of changes in the type profile. Starting from (L, L, L) , one reaches the assignment $\mu(1) = 1, \mu(2) = 2, \mu(3) = 3$, and starting from (M, M, M) , the assignment $\mu'(1) = 2, \mu'(2) = 3, \mu'(3) = 1$. Hence, the rule cannot be quasi-convergent at state (H, H, M) .

The incompatibility between quasi-convergence and efficiency hinges on the fact that agents cannot be forced to accept objects lower than the ones they currently hold. This creates a path-dependence which prevents the emergence of quasi-convergent rules, when efficiency is satisfied.¹⁴ The intuition underlying Theorem 6 relies on the existence of at least three types. With only two types, efficiency and quasi-convergence can be satisfied simultaneously.

Theorem 7 *Suppose that $|K| = 2$. The type-rank and type-seniority rules are both efficient and quasi-convergent.*

Theorem 7 indicates that, if one can separate the set of types in dichotomous categories, there exist assignment rules satisfying both criteria of intergenerational equity (time invariance) and efficiency (static efficiency). The

¹⁴If efficiency is not required, quasi-convergent rules exist. For example, the rank and seniority rules satisfy both quasi-convergence and independence, but not efficiency.

type-rank and type-seniority rules stand out as simple rules which should be used to allocate objects in a dichotomous world.

To understand why these rules satisfy efficiency and quasi-convergence, consider the set S' of states $s = (\mu, \theta)$ such that μ allocates objects according to a lexicographic criterion, first using an agent's type, and then his seniority. (Formally, for all i, j , $\mu(i) > \mu(j) \Rightarrow \theta(i) > \theta(j)$ or $\theta(i) = \theta(j), i > j$.) The type-seniority and type-rank rules have the property that the set S' is a closed set (from any state in S' , all transitions lead to another state in S').¹⁵ Furthermore, because the probability of any type is positive, there exists a path between any two states in S' , which is then a recurrent set. Because the state $s^H = (\iota, (H, H\dots H))$ belongs to the set S' , and there exists a path under the type-seniority and type-rank rules from any state $s = (\mu, (H, H, \dots, H))$ to s^H , and a path from any state to a state where all types are high, there exists a path from any state s to state s^H in S' , showing that S' is the unique recurrent set of the Markov chain.

The previous argument also highlights why a complete characterization of efficient and quasi-convergent rules for two types may be difficult. The argument shows that the transitions are only pinned down for a small number of states (states in S' and states where all agents have high types), and transitions among other states can be arbitrary. Nevertheless, the important conclusion is that efficient and quasi-convergent assignment rules exist for two types, and the type-rank and type-seniority rules emerge as simple, useful rules to apply in dichotomous settings.

6 Conclusion

In this paper, we analyze dynamic assignment problems where agents successively receive different objects (positions, offices, etc.). A finite set of n vertically differentiated indivisible objects are assigned to n agents who live n periods. At each period, a new agent enters society, and the oldest agent retires, leaving his object to be reassigned. A Markovian assignment rule specifies the probability that agents receive objects, and generates a finite Markov chain over the set of assignments. We define independent assignment rules (where the assignment of an object to an agent is independent of the objects currently held by the other agents), efficient assignment rules (where there does not exist another assignment rule with larger expected surplus), fair assignment rules (where two agents living in equal circumstances expe-

¹⁵This is the step in the argument which cannot be generalized to more than two types, as shown by the heuristic argument illustrating Theorem 6.

rience the same history) and analyze the dynamic properties of the Markov chains generated by assignment rules. When agents are homogeneous, we characterize independent convergent assignment rules, and provide sufficient conditions for irreducibility and ergodicity. When agents draw at random their types, we prove that independence and efficiency are incompatible, and assignment and quasi-convergent rules only exist for two types. We characterize two simple rules (type-rank and type-seniority) which satisfy both equity and efficiency criteria in dichotomous settings.

The results of our analysis rely on the assumption that agents cannot be forced to accept objects lower than the one they currently hold. This individual rationality requirement creates a history dependence in the process of allocation, and lies at the root of the dynamic aspects of our model. We believe that individual rationality is likely to be a necessary requirement in many assignment processes, but if it is not, history-independent efficient and equitable assignment rules could easily be devised.

While our analysis represents a first step in the understanding of dynamic assignment processes, it is based on a number of simplifying assumptions that we hope to relax in future research. First, we have assumed that all agents have the same preferences over the objects. Allowing for diversity in preferences would open an entire new set of questions on stability of assignment rules. Second, we have supposed that agents' types are perfectly observable. Relaxing this assumption would lead us to study incentive properties of assignment rules and move closer to the literature on dynamic mechanism design. Finally, we have considered a model where agents enter and exit one-by-one. If agents were to enter and exit in groups, our analysis would remain valid if the assignment rule of groups of agents could be decomposed in successive individual assignment rules. However, there exist other group assignment rules which do not satisfy this decomposability property and deserve further study.

7 References

1. Abdulkadiroglu A. and S. Loerscher (2007) “Dynamic house allocations,” mimeo., Duke University and University of Melbourne.
2. Abdulkadiroglu, A. and T. Sönmez (1999) “House allocation with existing tenants,” *Journal of Economic Theory* **88**, 233-260.
3. Athey, S. and I. Segal (2007) “An efficient dynamic mechanism,” mimeo., Stanford University.
4. Bartholomew, J. (1982) *Stochastic Models for Social Processes*, New York: Wiley.
5. Bergemann, D. and J. Valimäki (2006) “Efficient dynamic auctions,” Cowles foundation discussion paper 1584, Yale University.
6. Blum, Y. Roth, A.E. and Rothblum, U.G., (1997) “Vacancy chains and equilibration in senior-level labor markets,” *Journal of Economic theory* **103**, 429-443.
7. Cantala, D. (2004) “Restabilizing matching markets at senior level,” *Games and Economic Behavior* **48**, 1-17.
8. Cantala, D. and Sanchez, F. (2008) “Welfare and stability in senior matching markets,” *International Journal of Game Theory* **36**, 369-392.
9. Fudenberg, D. and J. Tirole (1991) *Game Theory*, Cambridge, MA: MIT Press.
10. Gershkov, A. and B. Moldovanu (2008a) “Learning about the future and dynamic efficiency,” forthcoming, *American Economic Review*.
11. Gershkov, A. and B. Moldovanu (2008b) “Dynamic revenue maximization with heterogeneous objects: A mechanism design approach,” mimeo., Bonn University.
12. Isaacson, D. and R. Masden (1976) *Markov Chains: Theory and Applications*, New York: John Wiley.
13. Iyer, L. and A. Mani (2008) “Traveling agents? Political change and bureaucratic turnover in India,” mimeo., Harvard Business School and University of Warwick.

14. Kemeny, J. and L. Snell (1960) *Finite Markov Chains*, New York: Van Nostrand.
15. Kurino, M. (2008) "House allocation with overlapping agents: A mechanism design approach," mimeo., University of Pittsburgh.
16. Moulin, H. and R. Stong (2002) "Fair queuing and other probabilistic allocation methods," *Mathematics of Operations Research* **27**, 1-31.
17. Moulin, H. and R. Stong (2003) "Filling a multicolor urn: An axiomatic analysis," *Games and Economic Behavior* **45**, 242-269.
18. Nilakantan, K. and B.G. Ragavendhra (2005) "Control aspects in proportionality Markov manpower systems," *Applied Mathematical Modelling* **29**, 85-116.
19. Parkes, D. C. and S. Singh (2003) "An MDP approach to online mechanism design," *Proceedings of the 17th Annual Conference on Neural Information Processing Systems (NIPS 03)*
20. Roth, A. and M. Sotomayor (1990) *Two-sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Cambridge: Cambridge University Press.
21. Talluri, K. and G. van Ryzin (2004) *The theory and practice of revenue management*, Springer: New York.
22. Thomson, W. (2007) "Fair allocation rules," mimeo., University of Rochester.
23. Ünver, Ü. (2007) "Dynamic exchange mechanisms and an application to kidney pairwise donation," mimeo., Boston College.

8 Appendix A: The transfer of high school teachers in France

High school teachers in France are allocated to different schools by a centralized mechanism administered by the Ministry of Education. The assignment rule by which high school teachers are matched to vacant positions is a very good example of a dynamic assignment rule.¹⁶

Every year in November, high school teachers who wish to be transferred are asked to declare their intentions to the Ministry of Education. There are two different phases in the allocation. In the first phase, in February, teachers are transferred across regions (called "académies"). In the second phase, in May, teachers are assigned to high schools in every académie.

Teachers are asked to submit their preferences for the two phases. In the first phase, they must rank the different regions in which they want to be transferred ; in the second phase, they must rank twenty high schools in the region to which they have been assigned.¹⁷ Positions are then allocated to teachers following a serial dictatorship rule, where the priority ranking of teachers is determined by their accumulated "points". These points are computed taking into account seniority, on-the-job-seniority and personal circumstances according to a complex formula given in Table 1.

Figure 7 displays the threshold number of points needed to enter specific académies for high school teachers of English in February 2008. The Ministry of Education also publishes every year the number of points needed to enter specific high schools in the second phase of the assignment mechanism.

¹⁶The complete details of the procedure can be found on the website of the French Ministry of Education, <http://www.education.gouv.fr>.

¹⁷Teachers who don't want to be transferred across regions only participate in the second phase.

9 Appendix B: Proofs

Proof of Lemma 1: Consider two assignments ν and ν' such that $\nu(i) = \nu'(k), \nu'(i) = \nu(k) = n$ and $\nu(l) = \nu'(l)$ for all $l \neq i, k$. For any $j < n$, $\alpha_j(\nu', i) = \alpha_j(\nu, k) = 0$. For any j , by independence, $\alpha_j(\nu, l) = \alpha_j(\nu', l)$. Now, as $\sum_{i, \nu(i) < j} \alpha_j(\nu, i) = \sum_{i, \nu'(i) < j} \alpha_j(\nu', i)$, we conclude that $\alpha_j(\nu, i) = \alpha_j(\nu', k)$.

Next, consider two assignments ν, ν' such that $\nu(i) = \nu'(j), \nu(j) = \nu'(i)$. As $\sum_k \alpha_n(\nu, k) = \sum_k \alpha_n(\nu', k)$, and, by independence, $\alpha_n(\nu, k) = \alpha_n(\nu', k)$ for all $k \neq i, j$, so that $\alpha_n(\nu, i) + \alpha_n(\nu, j) = \alpha_n(\nu', i) + \alpha_n(\nu', j)$.

Proof of Proposition 1: We first check that the identity assignment is indeed an absorbing state. A necessary and sufficient condition for this to occur is that:

$$\prod_j \alpha_j(\tilde{\nu}^j, j) = 1, \quad (3)$$

where $\tilde{\nu}^j(i) = i - 1$ for $i \leq j$ and $\tilde{\nu}^j(i) = i$ for $i > j$.

Both the seniority and rank assignment rules satisfy this condition, as j is at the same time the oldest agent eligible to receive object j and the agent with the highest ranked object in the matching $\tilde{\nu}^j$.

Next we show that starting from any initial state μ , there exists a time t at which the Markov chain is absorbed into the identity assignment ι .

In the rank rule, if $\mu(n) = k$, all objects $j = 1, 2, \dots, k$ are reassigned to the agents sequentially. In particular, at period 1, object 1 will be reassigned to the entering agent so that $\mu^1(1) = 1$. At period 2, object 2 is reassigned to agent 2 (who currently holds object 1) and object 1 is reassigned to the entering agent, so that $\mu^2(1) = 1$ and $\mu^2(2) = 2$. Following this argument, it is easy to see that $\mu^n = \iota$.

In the seniority rule, notice that the entering agent can never receive any object but object 1. Similarly if $\nu(1) = 1$, $\alpha_k(\nu, 1) = 0$ for all $k > 2$, and more generally if $\nu(i) = i$, $\alpha_k(\nu, i) = 0$ for all $k > i + 1$.

The preceding argument shows that, starting from any μ at period 0, in the seniority rule $\mu^1(1) = 1$. Furthermore, at period 2, object 1 must be reassigned to the entering agent so that $\mu^2(1) = 1, \mu^2(2) = 2$. We thus conclude that $\mu^n = \iota$, namely, the Markov chain is absorbed into the identity assignment in at most n periods.

Proof of Theorem 1: By Proposition 1, the rank rule and the seniority rules are convergent, so that the rule α , which is a convex combination of the seniority and rank rules, is also convergent.

Next suppose that the rule α satisfies independence and is convergent. Because it is convergent, the identity assignment is an absorbing state, so that

$$\alpha_j(\tilde{\nu}_j, j) = 1.$$

By independence, from Lemma 1, $\alpha_j(j-1) = \alpha_j(\tilde{\nu}_j, j) = 1$ for all $j < n$. Furthermore, by independence again, from Lemma 1, for any two assignments ν, ν' which only differ in the position of two agents, the total probability of assigning object n to the two agents is constant. As $\alpha_n(\tilde{\nu}_n, n) + \alpha_n(\tilde{\nu}_n, k) = 1$ for all $k < n$, we conclude that, for all ν ,

$$\alpha_n(\nu, n) + \alpha_n(\nu, \nu^{-1}(n-1)) = 1.$$

Next construct two different truncated assignments ν and ν' such that $\nu^{-1}(n) = i, \nu'^{-1}(n) = j$ and $\nu^{-1}(n-1) = \nu'^{-1}(n-1) = k$. By independence, $\alpha_n(\nu, \nu^{-1}(n-1)) = \alpha_n(\nu', \nu'^{-1}(n-1))$ so that $\alpha_n(\nu, n) = \alpha_n(\nu', n)$. Applying independence again, for any ν, ν' such that $\nu(n) = i < n-1$ and $\nu'(n) = j < n-1$, we have:

$$\alpha_n(\nu, n) = \alpha_n(\nu', n) = \lambda,$$

so that

$$\alpha_n(\nu, \nu^{-1}(n-1)) = 1 - \lambda$$

for any ν such that $\nu(n) \neq n-1$, establishing the result.

Proof of Theorem 2: Suppose first that the condition holds. Because object n is reassigned at least every n periods, and can only be reassigned when $\mu(n) = n$, any recurrent set must contain an assignment for which $\mu(n) = n$. Suppose by contradiction that there are two recurrent sets, each containing an assignment where $\mu(n) = n$, denoted μ_1 and μ_2 . If the condition holds, there exists an assignment μ' with $\mu'(n) = n$ which can be reached from both μ_1 and μ_2 , contradicting the fact that μ_1 and μ_2 belong to two distinct recurrent sets.

Conversely, suppose that the condition is violated and that there is a single recurrent set. There must exist one assignment μ' with $\mu'(n) = n$ in the recurrent set. However, if the condition is violated there exists another assignment μ with $\mu(n) = n$ such that there is no path in the Markov chain from μ to μ' , contradicting the fact that there is a single recurrent set.

Proof of Corollary 2: We will show the existence of a path to the identity assignment ι . Because object 1 is reassigned at least every n periods, there exists a time t at which $\mu^t(1) = 1$. Then, as $\alpha_2(2, \nu^t) > 0$, we can construct a path where $\mu^{t+1}(1) = 1, \mu^{t+1}(2) = 2$. Repeating the argument, we eventually reach the identity assignment.

Proof of Theorem 3: (Sufficiency) Consider two assignments μ and μ' . We will exhibit a path from μ to μ' . First observe that because $\alpha_j(\nu, 1) > 0$ for all j , by successively applying cycles κ , one eventually reaches an assignment μ^0 for which $\mu^0(n) = n$. Similarly, there exists an assignment μ^1 for which $\mu^1(n) = n$, and such that, by successively applying cycles, one reaches assignment μ' from μ^1 .

By condition (ii), there exists a sequence of permutations from μ^0 to μ^1 such that, at each step, with positive probability, the object held by the last player is assigned to player i , and the object held by player i to the entering player. Hence, one can construct a path between μ^0 and μ^1 , concluding the sufficiency part of the proof.

(Necessity) Suppose first that condition (i) is violated, i.e. there exists j and a truncated assignment ν of objects in $J \setminus j$ such that $\alpha_j(\nu, 1) = 0$. Consider the assignment μ such that $\mu(1) = j, \mu(i) = \nu(i)$ for $i = 2, \dots, n$. For this assignment to be reached, it must be that object j is assigned to the entering player with positive probability when all other players hold the objects given by the truncated assignment ν . Hence, if $\alpha_j(\nu, 1) = 0$, assignment μ can *never be reached from any other state*, contradicting the fact that the Markov chain is irreducible.

Next suppose that condition (ii) is violated.

We will first show that any reassignment from a matching μ to a matching μ' , $\pi = \mu^{-1}(\mu')$ can be decomposed into a sequence of cycles and $(1 - i)$ transpositions such that, if π^m is a $(1, i)$ transposition, then π^{m-1} and π^{m+1} are cycles. Let $i^0 = n + 1, i^1 = \mu'^{-1}(\mu(i^0 - 1)), \dots, i^m = \mu'^{-1}(\mu(i^{m-1} - 1)), \dots, i^M = 1$ be the sequence of reassignments from μ to μ' .

Consider the following sequence of permutations:

- First assign object $\mu(n)$ to agent i^1 and object $\mu(i^1 - 1)$ to the entering agent (apply $\kappa \circ \tau_{i^1}$)
- Then, during $n - 1$ periods, assign the object held by the last agent to the entering agent (apply κ^{n-1})

After this first cycle of n permutations, we have that, for all $j \neq i^1 - 1, n$, $(\pi^1 \circ \dots \circ \pi^n)(j) = j$. For $i^1 - 1$, we have $(\pi^1 \circ \dots \circ \pi^n)(i^1) = \pi^1(\pi^2 \circ \dots \circ \pi^n)(i^1)$.

$\pi^n)(i^1 - 1)) = \pi^1(i^1) = n$. For n , we have $(\pi^1 \circ \dots \circ \pi^n)(n) = \pi^1(\pi^2 \circ \dots \circ \pi^n)(n) = \pi^1(1) = i^1 - 1$. Hence, after the first cycle of n permutations, we have $\mu^n(j) = \mu(j)$ for all $j \neq i^1 - 1, n$, $\mu^n(i^1 - 1) = \mu(n) = \mu^{i^1}$ and $\mu^n(n) = \mu(i^1 - 1)$.

In the second cycle, we allocate object $\mu(i^1 - 1)$ to agent i^2 , by applying a cycle followed by a transposition τ_{i^2} , and then apply a cycle during $n - 1$ periods, κ^{n-1} . The process is repeated until we assign object $\mu(i^{M-1})$ to the entering agent.

Next, if condition (i) is satisfied, it must be that all cycles can be generated with positive probability. Furthermore, if state μ' can be reached from state μ through some reassignment, we must have $\alpha_{\mu(i^m-1)}(\nu^m, i^{m+1}) > 0$ for all m , so that the transpositions τ_{i^m} occur with positive probability. Hence, if state μ' can be reached from state μ by some arbitrary reassignment, it must also be reached through a sequence of permutations alternating cycles and transpositions τ_{i^m} as in the statement of the Proposition.

Hence, if for any sequence of permutations alternating cycles and transpositions, state μ' cannot be reached from state μ , we conclude that there is no path in the Markov chain from μ to μ' , and the Markov chain induced by the assignment rule is not irreducible.

Proof of Theorem 4: Given Theorem 3, we need to show that there exists a path from any assignment μ such that $\mu(n) = n$ to any assignment μ' such that $\mu'(n) = n$. The mapping $\mu'^{-1} \circ \mu$ is a permutation over $\{1, \dots, n\}$ which leaves the last element invariant. As any permutation can be decomposed into a sequence of transpositions, the induced permutation over the elements in $\{1, \dots, n-1\}$ can be decomposed onto a sequence of transpositions $\tau_{i^m, i^{m+1}}$, $m = 1, \dots, M-1$, where $1 \leq i^m \leq n-1$ for all m .

$$\mu' = \mu \circ \tau_{i^1, i^2} \circ \dots \circ \tau_{i^m, i^{m+1}} \circ \dots \circ \tau_{i^{M-1}, i^M}.$$

We now show that there exists a path in the Markov chain corresponding to this sequence of transpositions. We first consider the first transposition, τ_{i^1, i^2} . Let $j^1 = \mu(i^1)$ and $j^2 = \mu(i^2)$. Suppose without loss of generality that $j^1 \geq j^2$. Because the graph $G(\alpha)$ is connected, there exists a sequence $j_1 = j^1, \dots, j_q, \dots, j_Q = j^2$ such that $\alpha_{j_q}(j_{q+1}) > 0$ for all $q = 1, \dots, Q-1$. Let $i_q = \mu^{-1}(j_q)$ be the agent holding good j_q in μ . We can decompose the transposition τ_{i^1, i^2} as:

$$\begin{aligned} \tau_{i^1, i^2} &= \tau_{i_1, i_2} \circ \dots \circ \tau_{i_{Q-2}, i_{Q-1}} \circ \tau_{i_{Q-1}, i_Q} \circ \tau_{i_{Q-1}, i_{Q-2}} \circ \dots \circ \tau_{i_2, i_1} \\ &\equiv \chi \end{aligned}$$

To check this equality, notice that, for any i not included in the sequence i^q , $\tau_{i^1, i^2}(i) = i = \chi(i)$. Furthermore,

$$\begin{aligned}\chi(i^1) &= \tau_{i_1, i_2} \circ \dots \tau_{i_{Q-2}, i_{Q-1}} \circ \tau_{i_{Q-1}, i_Q} \circ \tau_{i_{Q-1}, i_{Q-2}} \circ \dots \tau_{i_2, i_1}(i_1) \\ &= \tau_{i_1, i_2} \circ \dots \tau_{i_{Q-2}, i_{Q-1}}(i_Q) \\ &= i_Q \\ &= i^2.\end{aligned}$$

Similarly,

$$\begin{aligned}\chi(i^2) &= \tau_{i_1, i_2} \circ \dots \tau_{i_{Q-2}, i_{Q-1}} \circ \tau_{i_{Q-1}, i_Q} \circ \tau_{i_{Q-1}, i_{Q-2}} \circ \dots \tau_{i_2, i_1}(i_Q) \\ &= \tau_{i_1, i_2} \circ \dots \tau_{i_{Q-2}, i_{Q-1}}(i_{Q-1}) \\ &= i_1 \\ &= i^1.\end{aligned}$$

Finally, for i_q , $q \neq 1, Q$,

$$\begin{aligned}\chi(i^q) &= \tau_{i_1, i_2} \circ \dots \tau_{i_{Q-2}, i_{Q-1}} \circ \tau_{i_{Q-1}, i_Q} \circ \tau_{i_{Q-1}, i_{Q-2}} \circ \dots \tau_{i_2, i_1}(i_Q) \\ &= \tau_{i_1, i_2} \circ \dots \tau_{i_{q-1}, i_q} \circ \dots \circ \tau_{i_{Q-1}, i_Q} \circ \dots \tau_{i_{q-1}, i_q}(i_q) \\ &= \tau_{i_1, i_2} \circ \dots \tau_{i_{q-1}, i_q} \circ \dots \circ \tau_{i_{Q-1}, i_Q} \circ \dots \tau_{i_q, i_{q+1}}(i_{q-1}) \\ &= \tau_{i_1, i_2} \circ \dots \tau_{i_{q-1}, i_q}(i_{q-1}) \\ &= \tau_{i_1, i_2} \circ \dots \tau_{i_{q-2}, i_{q-1}}(i_q) \\ &= i_q.\end{aligned}$$

We now construct a path from μ to $\mu \circ \tau_{i_1, i_2}$. We first apply cycle κ for $n - i_1 + 1$ periods, so that $\mu \circ \kappa^{n-i_1+1}(1) = j_1$.

If $i_2 \leq i_1 - 1$, then $j_2 = \mu \circ \kappa^{n-i_1+1}(i_2 + n - i_1 + 1)$, and, because $\alpha_{j_1}(j_2) > 0$, we can apply the transposition $\tau_{i_2+n-i_1+1}$, to obtain $\mu \circ \kappa^{n-i_1+1} \circ \tau_{i_2+n-i_1+1}(1) = j_2$ and $\mu \circ \kappa^{n-i_1+1} \circ \tau_{i_2+n-i_1+1}(i_2 + n - i_1 + 1) = j_1$. Applying the cycle κ again for $i_1 + 1$ periods, we finally have: $\mu \circ \kappa^{n-i_1+1} \circ \tau_{i_2+n-i_1+1} \circ \kappa^{i_1-1}(i_1) = j_2$, $\mu \circ \kappa^{n-i_1+1} \circ \tau_{i_2+n-i_1+1} \circ \kappa^{i_1-1}(i_2) = j_1$ and $\mu \circ \kappa^{n-i_1+1} \circ \tau_{i_2+n-i_1+1} \circ \kappa^{i_1-1}(i) = \mu(i)$ for all $i \neq i_1, i_2$.

If $i_2 \geq i_1 + 1$, then $j_2 = \mu \circ \kappa^{n-i_1+1}(i_2 - i_1 + 1)$, and we now apply the transposition $\tau_{i_2-i_1+1}$ followed by $i_1 - 1$ cycles to finally obtain $\mu \circ \tau_{i_1, i_2}$.

A similar construction can be applied to construct a path from μ to any composition of μ with a sequence of transpositions, concluding the proof of the Theorem.

Proof of Lemma 2: Consider two type profiles θ, θ' such that $\theta_k = \theta'_k$ for all $k \neq i$ and $\theta_i \neq \theta'_i$. For any j and any ν , $\sum_{i|\nu(i) < j} \alpha_j(\nu, \theta, i) =$

$\sum_{i|\nu(i)<j} \alpha_j(\nu, \theta', i) = 1$. By independence, $\alpha_j(\nu, \theta, k) = \alpha_j(\nu, \theta', k)$ for any $k \neq i$, so that $\alpha_j(\nu, \theta, i) = \alpha_j(\nu, \theta', i)$ for any ν, j . By independence again, for any θ'' such that $\theta''(i) = \theta'(i)$, $\alpha_j(\nu, \theta', i) = \alpha_j(\nu, \theta'', i)$, concluding the proof of the Lemma.

Proof of Theorem 5: We first establish the following claim.

Claim 1 *Let α and α' be two rules which are type-lexicographic for all reassignments following the assignment of good j at ν and ν' , where $\{k|\nu(k) < j\} = \{k|\nu'(k) < j\}$. For any assignments μ and μ' generated by α and α' , $\sum_{k|\nu(k)<j} \sigma(\theta(k), \mu(k)) = \sum_{k|\nu'(k)<j} \sigma(\theta(k), \mu'(k))$.*

To prove the claim, consider the change in aggregate surplus (from the original truncated assignment ν), generated by the reassignment μ :

$$\Delta(\mu, \nu) = \sigma(\theta(i^1), j) + \sum_{m=2}^{M-1} (\sigma(\theta(i^{m+1}), \nu(i^m)) - \sigma(\theta(i^m), \nu(i^m))).$$

For any type k , let $i(k)$ denote the agent i for which $\theta(i) = k$ and $\nu(i) \leq \nu(j)$ for any j such that $\theta(j) = k$. In words, $i(k)$ denotes the agent of type k with the lowest object in ν . Any type-lexicographic assignment rule generates a reassignment sequence i^1, \dots, i^M such that $\theta(i^m) \geq \theta(i^{m+1})$ for all m . Furthermore, for any type k such that $\theta(1) \leq k$, agent $i(k)$ must belong to the sequence. (Suppose not, then there exist k, m such that $\theta(i^m) \geq k, \theta(i^{m+1}) < k$ and $\nu(i^m) > \nu(i(k))$, contradicting the fact that the assignment rule is type-lexicographic for the assignment of object $\nu(i^m)$.)

Hence, for any type-lexicographic rule, the reassignment sequence can be decomposed into intervals I_1, I_2, \dots, I_K , where K is the number of types k such that $\theta(1) \leq k$ and $\{i|\theta(i) = k, \nu(i) < j\} \neq \emptyset$, any two agents in the same interval has the same type, and agents $i(k)$ are the last agents in the intervals I_k . This implies that for any μ generated by a type-lexicographic assignment rule,

$$\Delta(\mu, \nu) = \sigma(\theta(i^1), j) + \sum_{k|\theta(1) \leq k, \{i|\theta(i)=k, \nu(i) < j\} \neq \emptyset} \sigma(k, \nu(i(k+1))) - \sigma(k+1, \nu(i(k+1)))$$

which shows that any two type-lexicographic rules generate the same aggregate surplus.

Going back to the proof of the Theorem, consider an assignment rule α which is not type-lexicographic. Then there exists j, ν and θ such that

the rule assigns object j with positive probability to some agent i and there exists i' such that $\nu(i') < j$ and $\theta(i') > \theta(i)$. If this occurs in more than one instance, consider a situation where, for all ν' such that $\nu'(i) = \nu(i)$ when $\nu(i) \geq j$, all $j' < j$, $\alpha_{j'}(\nu', \theta, i) > 0 \Rightarrow \theta(i) = \max_{k|\nu(k) < j'} \theta(k)$. In words, for all reassignments following the assignment of good j when the truncated assignment is ν , the assignment rule prescribes to assign the object to an eligible agent with the highest type, and the assignment of good j to agent i is the last instance where the good is assigned with positive probability to an eligible agent of lower type. Furthermore, pick i' such that $\theta(i') = \max_{k|\nu(k) < j} \theta(k)$.

Consider a new assignment rule α' which only differs from α in the assignment of j to i and i' at ν and θ and is defined by $\alpha'_j(\nu, \theta, i') = \alpha_j(\nu, \theta, i') + \alpha_j(\nu, \theta, i)$, $\alpha'_j(\nu, \theta, i) = 0$. We will show that the assignment rule α' generates a higher expected surplus than rule α .

Consider any two sequences of reassignments $i^1 = i, \dots, i^M = 1$ and $i^1 = i'^{M'} = 1$. Suppose first that $\nu(i) > \nu(i')$. Because α is type lexicographic at ν' where $\nu'(i) = \nu(i)$ for all i such that $\nu(i) > j$ and $\nu^1 = j$, $\theta(i^2) = \max_{k|\nu(k) < \nu(i)} \theta(k) = \theta(i')$. By Claim 1, assignment rule α generates the same surplus as a rule which assigns object $\nu(i^1)$ to i' at ν' . Hence we compute the aggregate surplus at μ as:

$$\begin{aligned} S(\mu) &= \sigma(\theta(i), j) + \sigma(\theta(i'), \nu(i)) + \sum_{k|\nu(k) > \nu(i'), k \neq i} \sigma(\theta(k), \nu(k)) \\ &+ \sum_{k|\nu(k) < \nu(i')} \sigma(\theta(k), \mu(k)). \end{aligned}$$

Next, consider the aggregate surplus generated by the reassignment μ' :

$$S(\mu') = \sigma(\theta(i'), j) + \sum_{k|\nu(k) > \nu(i')} \sigma(\theta(k), \nu(k)) + \sum_{k|\nu(k) < \nu(i')} \sigma(\theta(k), \mu'(k)).$$

By Claim 1, $\sum_{k|\nu(k) < \nu(i')} \sigma(\theta(k), \mu'(k)) = \sum_{k|\nu(k) < \nu(i')} \sigma(\theta(k), \mu(k))$, so that

$$S(\mu') - S(\mu) = \sigma(\theta(i'), j) + \sigma(\theta(i), \nu(i)) - \sigma(\theta(i'), \nu(i)) - \sigma(\theta(i), j) > 0,$$

where the last inequality holds by strict supermodularity of the surplus function.

Next suppose that $\nu(i') > \nu(i)$. Let $\tilde{m} = \max m | \nu(i'^m) > \nu(i)$. Because μ' is generated by a type-lexicographic rule, the argument used in the proof of Claim 1 shows that $\theta(i'^{\tilde{m}}) > \theta(i) = \theta(i'^{\tilde{m}+1})$. By Claim 1, the aggregate surplus generated by μ' is equal to the surplus generated when object $\nu(i'^{\tilde{m}})$ is assigned to i . Hence,

$$\begin{aligned} S(\mu') &= \sum_{k | \nu(k) \geq \nu(i'^{\tilde{m}})} \sigma(\theta(k), \mu'^{\tilde{m}}) \\ &\quad + \sigma_{k | \nu(i) < \nu(k) < \nu(i'^{\tilde{m}})} \sigma(\theta(k), \nu(k)) + \sum_{k | \nu(k) < \nu(i)} \sigma(\theta(k), \mu'(k)). \end{aligned}$$

Similarly, we compute

$$S(\mu) = \sigma(\theta(i), j) + \sum_{k | \nu(k) > \nu(i)} \sigma(\theta(k), \nu(k)) + \sum_{k | \nu(k) < \nu(i)} \sigma(\theta(k), \mu(k)).$$

By Claim 1, $\sum_{k | \nu(k) < \nu(i)} \sigma(\theta(k), \mu'(k)) = \sum_{k | \nu(k) < \nu(i)} \sigma(\theta(k), \mu(k))$, so that

$$\begin{aligned} S(\mu') - S(\mu) &= \sigma(\theta(i'), j) + \sum_{m=2}^{\tilde{m}} (\sigma(\theta(i^m), \nu(i^{m-1})) - \sigma(\theta(i^m), \nu(i^m))) \\ &\quad + \sigma(\theta(i), \nu(i'^{\tilde{m}})) - \sigma(\theta(i), j) - \sigma(\theta(i'), \nu(i')). \end{aligned}$$

Next recall that $\theta(i') \geq \theta(i'2) \dots \geq \theta(i'^{\tilde{m}}) > \theta(i)$ and $j > \nu(i'^2) > \dots > \nu(i)$. Hence, strict supermodularity implies

$$\begin{aligned} \sigma(\theta(i'), j) + \sigma(\theta(i), \nu(i')) &> \sigma(\theta(i'), \nu(i')) + \sigma(\theta(i), j), \\ \sigma(\theta(i), \nu(i'^{\tilde{m}})) + \sigma(\theta(i'^{\tilde{m}}), \nu(i')) &> \sigma(\theta(i), \nu(i'^{\tilde{m}}), \nu(i'^{\tilde{m}})), \\ \sigma(\theta(i'^m), \nu(i'^{m-1})) + \sigma(\theta(i'^{m-1}), \nu(i'^{m-1}), \nu(i'^{m-1})) & \\ &\quad + \sigma(\theta(i'^m), \nu(i')) \forall m = 3, \dots, \tilde{m} \end{aligned}$$

Summing up these inequalities,

$$\begin{aligned} \sigma(\theta(i'), j) + \sum_{m=2}^{\tilde{m}} (\sigma(\theta(i^m), \nu(i^{m-1})) - \sigma(\theta(i^m), \nu(i^m))) + \sigma(\theta(i), \nu(i'^{\tilde{m}})) \\ - \sigma(\theta(i), j) - \sigma(\theta(i'), \nu(i')) &> 0 \end{aligned}$$

showing that $S(\mu') - S(\mu) \geq 0$. This last step concludes the argument, as the expected surplus generated by α' is always larger than the expected surplus generated by α .

Conversely, consider a type-lexicographic rule α , and suppose that it is inefficient. Then, there exists a state described by a type vector θ , a truncated assignment ν and an object j and an alternative allocation rule α' which generates a higher expected surplus. By the argument above, if the allocation rule α' is not type-lexicographic, there exists a type-lexicographic rule α'' which generates at least as high an expected surplus. But then, by Claim 1, the two assignment rules α and α'' must generate the same aggregate surplus, resulting in a contradiction.

Proof of Theorem 6: Consider a society with at least three types, and consider the realization of type profiles (L, L, \dots, L) and (M, M, \dots, M) . Because the assignment rule is quasi-convergent, it must result in a single assignment at these states, which has to be the identity assignment. Now consider the sequence of changes in type profiles, where first an agent M enters, then $n - 1$ agents H . By efficiency, the assignment rule is type-lexicographic and object n is assigned to the entering agent M , then object $n - 1$ to the first agent H , etc. In the end, the process will result in the identity assignment, where the oldest agent (agent M) is assigned object n , and any H agent i receives object i . Consider instead a sequence of changes in type profiles from (M, M, \dots) where $n - 1$ H agents enter in sequence. By efficiency, objects will be assigned sequentially to all the entering H agents, resulting in a final assignment where the oldest agent (agent M) holds object 1 and any agent i of type H holds object $i + 1$. Because the rule results in two different assignments for the same type profile, it cannot be quasi-convergent.

Proof of Theorem 7: Consider the set $S' = \{s = (\mu, \theta) | \forall i, j, \mu(i) > \mu(j) \Rightarrow \theta(i) > \theta(j) \text{ or } \theta(i) = \theta(j) \text{ and } i > j\}$. This is the set of states where objects are allocated according to a lexicographic criterion, using first an agent's type and then her seniority. Notice that, for all θ , there is a unique state $s = (\mu, \theta)$ in S' . Furthermore, when $\theta = (H, H, \dots, H)$, $\mu(i) > \mu(j) \Leftrightarrow i > j$, so the assignment μ such that $(\mu, \theta) \in S'$ is the identity assignment ν .

Now, consider a state $s = (\mu, \theta)$ in S' . We will characterize the transitions induced by the type-rank and type seniority rules. Let \prec denote the type-lexicographic ordering among agents at θ' , the type profile obtained after agent n has left and a new agent has entered. Let m be the rank of the new agent in that ordering. We need to distinguish between different cases. First, suppose that $\theta(n) = H$ so that $\mu(n) = n$. Then objects $n, n - 1, \dots, n - m$

will be reassigned sequentially to all agents according to the lexicographic ordering. The new assignment μ' then respects the type-seniority ordering. If, on the other hand $\theta(n) = L$, and $\mu(n) = k$, object k will either be assigned to the entering agent (if he has a high type), or to the oldest agents of low type, inducing a chain of reassignments of objects $1, 2, \dots, k$ among agents of low type. In both cases, the resulting assignment μ also respects the type-seniority ordering. Hence, from any state s in S' , the Markov chain induced by the type-seniority and type-rank rules results in a new state in S' , showing that S' is a closed set.

Next, note that because the probability of high and low types is positive, for any θ, θ' , there exist states $s = (\mu, \theta)$ and $s' = (\mu', \theta')$, with $p(s'|s) > 0$. This shows that all states in S' intercommunicate, and S' is a recurrent set.

Finally, using the argument of Proposition 1, we see that from any state $s' = (\mu, (H, H, H, \dots, H))$, there exists a path to $s^H = (\iota, (H, H, \dots, H))$. For any state s , there exists a path to a state s' where all agents have high type. Hence, there exists a path from any state s to s^H (and then to any state in S'), showing that S' is a unique recurrent set. As it contains exactly one state per type profile, we conclude that the type-rank and type-seniority rules are quasi-convergent.

10 Tables and Figures

Seniority	7 points per "echelon" (on average 3 years of service) + 49 points after 25 years of service
On the job seniority	10 points per year + 25 points every 4 years
Current job	50 points if first assignment If assigned in a "violent" high school + 300 points after 4 years
Family circumstances	150.2 points if spouse is transferred + 75 points per child

TABLE 1: PRIORITY POINTS FOR HIGH SCHOOL TEACHERS IN FRANCE




Figure 1: SENIORITY ASSIGNMENT FOR THREE AGENTS




Figure 2: RANK ASSIGNMENT FOR THREE AGENTS




Figure 3: UNIFORM ASSIGNMENT FOR THREE AGENTS




Figure 4: REPLACEMENT ASSIGNMENT FOR THREE AGENTS




Figure 5: TRANSITIONS BETWEEN STATES FOR EXAMPLE 3




Figure 6: TRANSITIONS BETWEEN STATES FOR EXAMPLE 4



Figure 7: THRESHOLDS FOR TRANSFERS OF ENGLISH TEACHERS 2008