Markovian Assignment Rules*

Francis Bloch
Ecole Polytechnique and Brown University

and

David Cantala
Colegio de Mexico

August 2008

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. 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), and analyze the dynamic properties of the Markov chains generated by assignment rules. When agents are homogenous, 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.

JEL Classification Numbers: C78, D73, M51

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

*We thank Matt Jackson, Jean-François Laslier, Hervé Moulin, William Thomson and seminar participants at various institutions for useful discussions on dynamic assignment rules.

1 Introduction

Economic models of matching and assignment are essentially static, and only consider assignments one at a time. In Gale and Shapley (1962)'s marriage problem, divorce is not allowed and men and women are married forever ; in the roommate problem, college students are assigned to dorms every year, independently of their previous year's assignment ; in the assignment problem (Gale and Shapley (1971)), workers and firms negotiate their contract irrespective of their history, and in the school choice problem (Abdulkadiroglu and Sonmez (2003)), pupils are admitted to schools independently of their previous schooling history.[2] Models of assignment of indivisible goods, like Shapley and Scarf (1974)'s house allocation problem also focus on a single, static, assignment. The history is entirely captured by the current ownership or tenancy structure (Abdulkadiroglu and Sonmez (1999)), and successive reassignments of houses are not considered.[3]

However, there exist situations where assignment rules are dynamic, and successive assignments cannot be analyzed separately. For example, consider the assignment of positions to civil servants in centralized systems. In France, teachers who want to transfer are reassigned to high schools according to a complex priority system, which takes into account current position, seniority and on the job seniority, as illustrated in Table 1.[4] Regions in France are not equally attractive, as shown in Figure 1 which gives the minimal number of points needed to transfer into different regions for high school teachers of English in March 2008. In India, officers of the Indian Administrative Service are also reassigned according to their seniority, current position, career history and rank at the entrance exam to the IAS (Iyer and Mani (2008)). More generally, successive job assignments inside organizations are decided according to complex rules putting weight on seniority, performance and career history.

Priority systems based on seniority seem to be prevalent in many different settings. Airline pilots and flight attendants get to choose their flight assignments according to seniority.[5] Assignment of subsidized housing to potential

[2]For a survey of models of two-sided matching, see the excellent monograph by Roth and Sotomayor (1990).

[3]See Thomson (2007) for an exhaustive account of the literature on the allocation of indivisible goods.

[4]The exact priority rules are published every year by the French Ministry of Education. For the 2008 rules, see for example www.ac-rouen.fr/rectorat/profession_mutation/inter_bareme.pdf.

[5]This emphasis on seniority explains the difficulties in merging pilots' seniority lists when two airline companies contemplate a merger. Northwest Airlines and Delta have

tenants often gives priority to agents who have the highest seniority on the waiting list. In many industries, seniority rules govern priorities for layoffs and promotions (see for example the historical account given by Lee (2004)).

In other situations, assignment rules do not favor the agents with highest seniority. For example, in order to minimize moving costs, offices of retiring employees are likely to be reassigned to newcomers. This rule favors agents with the lowest seniority, and we will term it the "replacement rule". Alternatively, offices and positions can be reassigned at random. Random assignments based on uniform distributions, termed "uniform rules" are used in a wide range of contexts and deserve a special study.

Dynamic assignment rules differ from static assignment rules in two important ways. First, in a dynamic setting, agents are not only characterized by their current preferences, but also by their history (past assignments and past preferences). Assignment rules can use this information, and condition the allocation of objects to characteristics summarizing the agent's history, like seniority or on the job seniority. Second, in a dynamic setting, the set of agents to whom objects are allocated is not constant. Agents enter and leave the pool of agents to whom objects are allocated. For example, every year, some civil servants retire while others are recruited. Given that the number of positions is fixed, retiring civil servants free their positions, which can be reallocated to other agents, etc.., until all positions are filled. Similarly, if the number of offices is fixed, the allocation of offices in organizations depends on the flow of agents entering and leaving the organization. Agents leaving the organization free their offices, which can be reallocated sequentially until the agents entering the organization are assigned an office.

In this paper, we analyze dynamic assignment rules to allocate a fixed number of objects (offices or positions) to a dynamic set of agents entering and leaving society. In order to simplify the analysis, we assume that all agents have the same preferences over the objects, which can thus be ranked according to their desirability.[6]. We identify a *state* of the society with an assignment of objects to agents, where agents are distinguished by their seniority. A *Markovian assignment rule* specifies how objects are allocated to agents according to the current assignment of objects in society.

---

recently faced this problem, as pilot associations of the two airlines have so far been unable to agree on a common seniority list. See "NWA pilots set merger conditions," *Minneapolis Star Tribune*, January 18, 2008.

[6]This is of course a very strong simplifying assumption, but it can often be justified. For example, civil servants often have congruent preferences over positions, ranking the importance of various positions in the same way. Similarly, most agents will agree on the ranking of offices according to their sizes and locations.

We focus attention in this paper on the dynamic properties of the finite Markov chains generated by different assignment rules. We study which assignment rules are *convergent* (every assignment leads to a unique absorbing state), *ergodic* (the long run behavior of the chain is independent of the initial assignment) and *irreducible* (all assignments occur with positive probability in the invariant distribution). We also define different notions of *independence*, specifying how the assignment of object j to agent i depends on the objects currently held by the other agents. Finally, we consider a static notion of *efficiency* requiring that there does not exist an alternative assignment rule which generates an expected total surplus at least as large (and sometimes strictly larger) at every state.

We discuss four specific Markovian assignment rules. The seniority rule allocates object j to the oldest agent who holds an object smaller than j. The rank rule allocates object j to the agent holding object $j-1$. The uniform rule allocates object j with equal probability to all agents who currently hold objects smaller than j. Finally, the replacement rule allocates the object of the agent leaving the society to the agent entering society.

Our main result in that part of the paper shows that any convergent rule satisfying independence must be a weighted combination of the rank and seniority rules. Convergence in our setting can be understood as a condition of equity across cohorts. In a convergent Markov chain, the absorbing state will eventually be reached, and agents entering society at different dates will experience the exact same history. Hence, our analysis gives support to the rank and seniority rules as rules satisfying both a natural notion of independence and a condition of equity.

Ergodic and irreducible rules cannot be characterized as easily. However, we show that any rule which allocated object j to the agent currently holding object $j-1$ with positive probability is ergodic. On the other hand, rules which do not allow for transitions across two different assignments μ and μ' such that $\mu(n) = \mu'(n) = n$, like the replacement rule, are not ergodic, and different initial conditions will lead to different long run behaviors of the Markov chain.

We also provide a sufficient condition for irreducibility of independent chains, in terms of the graph of intercommunicability of different states. This condition states that, for an independent assignment rule, any object is reassigned to the entering agent with positive probability and the undirected graph formed by all pairs (i, j) of agents such that an agent holding object i receives object j with positive probability, is connected. This condition will always be satisfied when the probability of allocating object j to the entering agent and the agent holding object $j-1$ is positive, and also when the probability of allocating object j to both the entering agent and the agent holding object 1 is positive.

In the second part of the paper, we explore the properties of assignment rules where agents are heterogeneous, and draw at random a type before entering society. We first show that independence becomes a very strong condition when agents are heterogeneous, as it implies that assignments are independent of the agent's type. Assuming that the surplus is a supermodular function of the quality of the object and the agent's ordered type, efficiency requires an assortative matching between objects and types. We first characterize efficient allocation rules, and show that they must exhibit a lexicographic structure, first assigning objects to agents of higher types. This characterization also shows that there exists an incompatibility between independence and efficiency.

When agents' types are drawn at random, the evolution of states is governed by two simultaneous forces: the exogenous draw of type profiles and the endogenous evolution of assignments. We define quasi-convergent assignment rules as rules for which the unique recurrent set only admits one assignment per type profile, thereby generalizing convergent rules. This generalization captures the same equity consideration as before, as any two agents born at different points in time the same societies will experience the same history if the assignment rule is quasi-convergent. When agents only belong to two types, quasi-convergent and efficient allocation rules exist and the type-rank and type-seniority rules, where objects are allocated according to a a lexicographic ordering putting type first and rank or seniority second, stand out as simple, efficient and equitable rules. With more than two types of agents, we obtain an impossibility result: there is no allocation rule which satisfies simultaneously efficiency and quasi-convergence.

1.1 Related literature

We situate our paper with respect to the existing literature in economics and operations research. As noted above, economic models of assignment do not allow for the type of dynamic changes in the population and objects that we consider here. The papers which are more closely related to ours in this literature are papers by Moulin and Stong (2002) and (2003) which consider axiomatic rules to allocate balls entering in succession into multicolor urns. As in our model, the state of the urns varies over time, and the allocation depends on the current state and influences the transition across states. However, in all other respects, the problems we study are very different.

To the best of our knowledge, the literature in operations research and computer science on queuing and assignment has not considered the problem we study in this paper.[7] Typically, problems of queuing in operations research do not model preferences of agents over the servers, do not identify agents who enter repeatedly the system and whose utility is defined over a sequence of assignments, and do not impose, as we do, individual rationality conditions, stating that an agent cannot be assigned a lower object than the one he currently holds.

In management, the literature on Markov manpower systems also analyzes the allocation of agents to vertically differentiated positions as a Markov process.[8] But the focus of the literature is very different from our approach. Studies of manpower planning investigate how planners can achieve a fixed target (in terms of sizes of different grades in the hierarchy of the organization) by controlling promotion and recruitment flows. By contrast, our study is based on the structure and properties of allocation rules which determine the reassignment of objects at every period.

In personnel and labor economics, seniority rules for promotion and layoffs have been analyzed, both theoretically and empirically. (See Lazear (1995)). Theoretical models emphasize seniority promotion rules as a way to provide incentives to worker to acquire firm-specific human capital (Carmichael (1983)). Empirical studies of internal labor markets, like Chiappori, Salanié and Valentin (1999)'s study of promotions and careers in a large firm in France, provide a more complex and nuanced view of the effects of seniority, human capital acquisition and innate abilities on career histories.

1.2 Structure of the paper

The rest of the paper is organized as follows. We introduce the model in Section 2 and illustrate our concepts by undergoing an exhaustive analysis of the case of three agents. In Section 3, we study the dynamic properties of the model with homogeneous agents, characterizing convergent, ergodic and irreducible assignment rules. Section 4 is devoted to the model with heterogeneous agents, and contains our impossibility results and the characterization of efficient, quasi-convergent rules for dichotomic types. We conclude and give directions for future research in Section 5.

[7] See for example, Kleinrock (1975) and (1976) for a simple exposition of classical assignment models in queuing theory and computer science.

[8] See the books by Bartholomew (1982) and Vajda (1978) for a review of the early contributions, and the recent paper by Nilakantan and Ragavhandra (2005) for an account of recent work.

2 The Model

2.1 Markovian assignment rules

We consider a society I of n agents indexed by their age, $i = 1, 2, ..., n$, and n vertically differentiated indivisible goods, indexed by $j = 1, 2, ..., n$ in the set J . Time is discrete and runs as $t = 1, 2, ...$ Each agent lives for exactly n periods, and at each date t, one agent dies and leaves the society whereas another agent enters the society.

Definition 1 *An* assignment μ *is a mapping 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* $\{1, 2, ...n\}$.

A *state* of the society is totally described by the assignment of objects to agents. Hence there are $n!$ states in a society of n agents, each corresponding to a different permutation over the finite set $\{1, 2, ...n\}$.

We now consider the dynamical structure of the society. Let μ^t denote the state of the economy at date t. At the beginning of period $t + 1$, the oldest agent alive in period t relinquishes object $\mu^t(n)$ and leaves the society. A new agent enters with no object – by convention we denote the null object by 0 – and agents $i = 2, 3, ..., n$ retain the objects they were assigned in the previous period, $\mu^t(i-1)$. A new assignment will then be chosen, to allocate object $\mu^t(n)$ to one of the current members of society. The assignment of object $\mu^t(n)$ will in turn, free a new object to be reassigned, etc.. The cascade of assignments will end when the new agent is assigned an object.

Definition 2 *A* truncated assignment ν*, given some object* j *is a mapping from* $I \setminus \{1\}$ *to* $J \setminus \{j\}$*, reflecting the assignment of the objects in* $J \setminus \{j\}$ *to the* $n - 1$ *oldest agents in the society.*

We focus on assignment rules which only depend on the current truncated assignment ν of objects to agents, and not on the entire history of assignments in the society. Assignment rules which possess this Markovian property are called *Markovian assignment rules.* An assignment rule allocates any object j to those agents who currently possess an object $k < j$. *We suppose that agents cannot be forced to give back their current object. The assignment rule must satisfy an individual rationality condition, and cannot assign to agent* i *an object* $j < \nu(i)$*.* Formally, we define:

Definition 3 *A* Markovian assignment rule *is a collection of vectors* $\alpha_j(\nu)$ *in* $\Re^n$ *for* $j = 1, 2, ...n$ *satisfying:* $\alpha_j(\nu, i) \geq 0$ *for all* i *and* $\sum_{i,\nu(i)<j} \alpha_j(\nu, i) = 1$.

The number $\alpha_j(\nu, i)$ denotes the probability that agent i receives object j given the truncated assignment ν.

2.2 Examples of Markovian assignment rules

We now describe four different Markovian assignment rules which have a simple interpretation.
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 $i = 1$.

Notice that some common rules are *not* Markovian. For example, rules based on on the 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.

2.3 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). A simple property of Markovian assignment rules is *independence*, stating that the assignment of object j to player i does not depend on the current assignment of objects held by the other players:

Definition 4 *A Markovian assignment rule α satisfies* independence *if and only if, for any j, for any i, for any ν, ν' such that $\nu(i) = \nu'(i)$, $\alpha_j(\nu, i) = \alpha_j(\nu', i)$.*

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

Definition 5 *A Markovian assignment rule α satisfies* strong independence *if and only if, for any j, for any i, k, for any ν, ν' such that $\nu(i) = \nu'(k)$, $\alpha_j(\nu, i) = \alpha_j(\nu', k)$.*

If an assignment rule is strongly independent, it is fully characterized by the probability of assigning object j to an agent holding object $k < j$. Abusing notations, we will denote this probability by $\alpha_j(k)$. 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 6 *A Markovian assignment rule α satisfies* weak independence *if and only if, for any j, for any i, for any ν, ν' such that $\nu(k) = \nu'(k)$ for all $k < j$, $\alpha_j(\nu, i) = \alpha_j(\nu', i)$.*

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 1 *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 $\alpha_3(\nu', 2) = \alpha_3(\nu, 3)$.*

2.4 Convergent, irreducible and ergodic assignment rules

Starting with any assignment μ^0, any Markovian assignment rule α generates a finite Markov chain over the set of assignments. More precisely, we can define the probability of reaching state μ' from state μ, $p(\mu'|\mu)$ as follows:

Consider the sequence of agents $i^0 = n+1, i^1 = \mu'^{-1}(\mu(i^0-1)), ..., i^m = \mu'^{-1}(\mu(i^{m-1}-1)), ..., i^M = 1$. 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.

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

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

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

Having defined the Markov chain over assignments, we now consider the dynamic properties of this chain, and relate it to the Markovian assignment rules. The following definitions are borrowed from classical books on finite Markov chains (Kemeny and Snell (1960), Isaacson and Madsen (1976)).

Definition 7 *Two states i and j* intercommunicate *if there exists a path in the Markov chain from i to j and a path from j to i.*

Definition 8 *A set of states C is closed if, for any states $i \in C$, $k \notin C$, the transition probability between i and k is zero.*

Definition 9 *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.

Definition 10 *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).*

Definition 11 *A Markovian assignment rule α is* irreducible *if the induced Markov chain is irreducible (the only recurrent set is the entire state set).*

Definition 12 *A Markovian assignment rule α is* ergodic *if the induced Markov chain is ergodic (has a unique recurrent set).*[9]

We finally define a dynamic notion of equity, based on the idea that two agents born at different dates must be treated identically in the long run.

Definition 13 *A Markovian assignment rule α is* time invariant *if there exists $T > 0$ such that, any two agents born at $t, t' > T$ experience the same history, $\mu^{t+\tau}(i) = \mu^{t'+\tau}(i')$ for $\tau = 0, 1, ..., n-1$, and agents i and i' entering society at dates t and t'.*

The following Lemma is immediate and does not require a proof:

Lemma 2 *A Markovian assignment rule α is time invariant if and only if any recurrent set of the induced Markov chain is an absorbing state.*

2.5 Markovian assignment rules among three agents

In this Section, we completely characterize the Markovian assignment rules among three agents. If $n = 3$, there are six possible assignments defined by:

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

Given the constraint that $\sum_{i\nu(i)\leq j} \alpha_j(i,\nu) = 1$, not all transitions among states can occur with positive probability, and the transition matrix must display the following zero entries:

[9] This 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.

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

Alternatively, Figure 1 below illustrates the transitions between states in the Markov process induced by an assignment rule putting positive probabilities on all feasible transitions:

We now examine in turn the four assignment rules described above.

2.5.1 The seniority rule

The seniority rule is represented by the transition matrix:

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

The transitions between states are represented in Figure 2, which shows that the seniority rule is in fact convergent.

2.5.2 The rank rule

The rank rule is represented by the transition matrix:

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

The transitions between states are represented in Figure 3, which shows that the rank rule is convergent.

2.5.3 The uniform rule

The uniform rule is represented by the transition matrix:

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

The transitions between states are identical to those of Figure 1. The Markov chain is irreducible and the invariant distribution[10] is given by

$$p_1 = \frac{36}{127} \simeq 0.28, p_2 = \frac{28}{127} \simeq 0.22, p_3 = \frac{30}{127} \simeq 0.24, p_4 = \frac{11}{127} \simeq 0.08,$$
$$p_5 = \frac{12}{127} \simeq 0.09; p_6 = \frac{10}{127} \simeq 0.07.$$

2.5.4 The replacement rule

The replacement rule is represented by the transition matrix:

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

The Markov chain is not ergodic. As shown in Figure 4, there are two cyclical recurrent sets of period 3.

[10] Any 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)

3 Dynamic properties of Markovian assignment rules

3.1 Convergent Markovian assignement rules

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

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

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 $i = 1, ...n$. Hence, *the only candidate absorbing state is the identity assignment* ι. This observation also shows that an assignment rule is time invariant if and only if it is convergent.

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.*

3.2 Ergodic assignment rules

We first recall some definitions of special permutations.

Definition 14 *A permutation from a set of n elements to itself is a* cycle, *denoted κ, if $\pi(i) = \pi(i+1)$ for all $i = 1, 2, .., n$ and $\pi(n) = 1$.*

Definition 15 *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 $i^0, ..., i^M$. 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 apply a transposition between agent 1 and agent i^1, assigning object $\mu(n)$ to i^1, (iii), then apply a transposition between agent 1 and agent i^2, assigning object $\mu(i^1)$ to agent i^2, etc... Hence,we may write:

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

where it is understood that τ_1, the identity permutation, is just added for the sake of completeness, and to show that the composition of permutation ends.

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, $\mu^{-1}(\mu') = \pi^1 \circ ... \circ \pi^m \circ ...\pi^M$ such that either π^m is a cycle or a $(1, i)$ transposition and, if it is a $(1, i)$ transposition, $\alpha_{(\mu \circ ... \circ \pi^{m-1})^{-1}(1)}(\nu^{m-1}, i) > 0$, where $\nu^{m-1}(j) = (\mu \circ ... \circ \pi^{m-1})(j)$ for all $j = 2, ..., n$.*

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

Let all states such that $\mu(4) = 4$ be ordered as in Subsection 2.5. In addition, define the states:

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

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.

3.3 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, $\mu^{-1}(\mu') = \pi^1 \circ ... \circ \pi^m \circ ... \pi^M$ 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 $\alpha_{(\mu \circ ... \circ \pi^{m-1})^{-1}(1)}(\nu^{m-1}, i) > 0$, where $\nu^{m-1}(j) = (\mu \circ ... \circ \pi^{m-1})(j)$ for all $j = 2, ..., n$.

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 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 though 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 $\{1, 2.., n-1\}$ 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 3 *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.

4 Markovian assignment rules among heterogeneous agents

In this Section, we extend the model by allowing for heterogeneity across agents. More precisely, we suppose that agents independently draw types (or abilities) which affect the value of the surplus formed in any matching. Assuming that surpluses are supermodular functions of objects and types, efficient assignments require to assign higher objects to agents with higher types. Of course, this requirement conflicts with the use of simple seniority and rank rules, and the object of the analysis is to characterize richer classes of rules, which take into account agent's types as well as their histories.

4.1 Assignment rules with heterogeneous agents

Let $K = \{1, 2, ..., m\}$ be a finite ordered set of types indexed by k. At every period, the type of the entering agent is drawn according to an independent draw of a finite probability distribution $q(k)$. The set of objects and types are ordered in such a way that the surplus obtained by matching an agent of type k with an object j, $\sigma(k, j)$ is *strictly supermodular*: If $k' \geq k$ and $j' > j$,

$$\sigma(k', j') + \sigma(k, j) \geq \sigma(k', j) + \sigma(k, j'),$$

with strict inequality when $k' > k$. Hence, total surplus in society will be maximized by assigning objects of higher quality to agents with higher types.

A *state* s is now defined both by an assignment μ and a type profile θ, $s = (\mu, \theta)$ where

- The assignment μ is a one-to-one mapping from the set I of agents to the set J of objects,
- The type profile θ is a mapping from the set I of agents to the set K of types.

An assignment rule α is now a collection of mappings $\alpha_j(\nu, \theta, i)$, defining the probability of assigning object j to agent i given the truncated assignment ν and the type profile θ. As before, the assignment rule must satisfy:

$$\sum_{i|\nu(i)<j} \alpha_j(\nu, \theta, i) = 1.$$

Using the additional dimension given by agents' types, we may now define new elementary assignment rules as follows:

Definition 16 *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 $\nu(l) < j$.*

Definition 17 *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 $\nu(l) < j$.*

The type-seniority and type-rank rules use a lexicographic ordering: they first select the set of agents of highest type who may receive the object. If this set contains more than one type, the rule uses a tie-breaking rule (seniority or rank) to allocate the object.[11]

Given an assignment rule α and a probability distribution q over types, we can compute the transition probability from state s to state s':

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

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

- Second, given the new type profile θ', and the assignment rule α, construct sequences of reassignments as in Subsection 3.1 to obtain $p(\mu'|\mu)$.

4.2 Independent assignment rules with heterogeneous agents

As before, we define an assignment rule to be *independent* if the assignment of object j to agent i only depends on the characteristics of agent i: $\alpha_j(\nu, \theta, i) = \alpha_j(\nu', \theta', i)$ whenever $\nu(i) = \nu'(i)$ and $\theta(i) = \theta'(i)$. Similarly, an assignment rule is *weakly independent* if $\alpha_j(\nu, \theta, i) = \alpha_j(\nu', \theta', i)$ whenever $\nu(k) = \nu'(k)$ and $\theta(k) = \theta'(k)$ for all k such that $\nu(k) < j$. Finally, an assignment rule is *strongly independent* if the assignment of object j to agent i only depends on the type and object currently held by agent i: $\alpha_j(\nu, \theta, i) = \alpha_j(\nu', \theta', k)$ whenever $\nu(i) = \nu'(k)$ and $\theta(i) = \theta'(k)$.

The following lemma shows that, if an assignment rule is independent, assignment *cannot depend on an agent's type*

Lemma 3 *Let α be an independent assignment rule among heterogeneous agents. Then, for any θ, θ', any j, ν and i, $\alpha_j(\nu, \theta, i) = \alpha_j(\nu, \theta', i)$.*

With heterogeneous players, independence thus puts a strong restriction on the assignment rule, and 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).

[11] One could similarly define type-uniform or type-replacement rules, where the uniform and replacement rules are used to break tie among agents of the same type.

4.3 Efficient assignment rules

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

Definition 18 *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.[12] 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.

Our next result characterizes efficient Markovian assignment rules, and shows that they exhibit a simple lexicographic structure.

Definition 19 *A Markovian assignment rule α is* type-lexicographic *if, for any j, ν and θ, $\alpha_j(\nu, \theta, i) > 0 \Rightarrow \theta(i) \geq \theta(k) \forall k, \nu(k) < j$.*

In words, in a type-lexicographic rule, objects are assigned in a hierarchical way. Agents with higher type have priority over agents with lower types. Objects are assigned according to a lexicographic criterion, by first considering agents' type, and then applying any second selection criterion.

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

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

Finally, we combine Theorem 5 and Lemma 3 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 3.

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

4.4 Quasi-convergent assignment rules

Finally, while the notions of ergodic and irreducible assignment rules are well defined, it is clear that when agents are heterogeneous, 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*

Definition 20 *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. It is also related to the following extension of the notion of time invariance.

Definition 21 *A Markov chain with heterogeneous agents is* time-invariant *if there exists $T > 0$ such that, any two agents born at $t, t' > T$ facing the same societies, experience the same history. For any two agents i and i' entering society at dates t and t' where $\theta t + \tau = \theta t' + \tau(i')$ for $\tau = 0, 1, ..., n-1$, $\mu^{t+\tau}(i) = \mu^{t'+\tau}(i')$ for $\tau = 0, 1, ..., n-1$.*

Time invariant and quasi-convergent assignment rules are related by the following observation, which is an easy extension of Lemma 2:

Lemma 4 *A Markovian assignment rule with heterogeneous agents is time invariant 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, $\theta(s) \neq \theta(s')$.*

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 3.1). Hence, if type (L, L, L) is realized, $\mu(i) = \iota$. Similarly, if type (M, M, M) is realized, $\mu(i) = \iota$. 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 hedges 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.[13] 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.*

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

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 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').[14] 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 $s^H = (\iota, (H, H...H))$ belongs to the set S', and there exists a path under the type-seniority and type-rank rules from any state $s = (\mu, (H, H, .., H)$ 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.

5 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

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

of the objects currently held by the other agents), efficient assignment rules (where there does not exist another assignment rule with larger expected surplus), and analyze the dynamic properties of the Markov chains generated by assignment rules. When agents are homogenous, 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 and foremost, we have assumed that all agents have the same preferences over objects, and would like to extend our analysis to richer preference settings. Second, we have restricted out attention to a static concept of efficiency, which might not be an adequate welfare criterion for a dynamic assignment process. We wish to be able to define dynamic efficiency notions for assignment rules. Finally, an important aspect of many job assignment processes is that agents strategically choose when to ask for a transfer, depending on existing priority rules and expectations over vacancies. In future work, we hope to study a model where agents endogenously choose when to ask for a reassignment.

6 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:

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

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 $j = 1, 2, ..., k$ 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, is it 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 $\nu^{-1}(n) = i, \nu'^{-1} = j$ and $\nu^{-1}(n-1) = \nu'^{n-1}(n-1) = k$. 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 $i = 2, ..., n$.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 $i^0 = n+1, i^1 = \mu'{-1}(\mu(i^0 - 1)), ..., i^m = \mu'{-1}(\mu(i^{m-1} - 1)), ..., i^M = 1$ 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 $j \neq i^1 - 1, n$, $(\pi^1 \circ ... \circ \pi^n)(j) = j$. For $i^1 - 1$, we have $(\pi^1 \circ ... \circ \pi^n)(i^1) = \pi^1(\pi^2 \circ .. \circ$

$\pi^n)(i^1-1)) = \pi^1(i^1) = n$. For n, we have $(\pi^1 \circ ... \circ \pi^n)(n) = \pi^1(\pi^2 \circ .. \circ \pi^n)(n)) = \pi^1(1) = i^1 - 1$. 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 $\{1, ..., n\}$ which leaves the last element invariant. As any permutation can be decomposed into a sequence of transpositions, the induced permutation over the elements in $\{1, ..., n-1\}$ can be decomposed onto a sequence of transpositions $\tau_{i^m, i^{m+1}}$, $m = 1, ..., M-1$, where $1 \leq i^m \leq n-1$ for all m.

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

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 $j_1 = j^1, ..., j_q, ..., j_Q = j^2$ such that $\alpha_{j_q}(j_{q+1}) > 0$ for all $q = 1, ..., Q-1$. 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:

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

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,

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

Similarly,

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

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

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

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 3: 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 $i^1, ..., i^M$ 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 $I_1, I_2, ..., I_K$, 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 $i^1 = i, ..., i^M = 1$ and $i^1 = i', ..., i'^{M'} = 1$. 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 $\nu'(i^1) = j$, $\theta(i^2) = \max_{k|\nu(k)<\nu(i)} \theta(k) = \theta(i')$. 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:

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

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,

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

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

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

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

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

Summing up these inequalities,

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

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 $(L, L, ..., L)$ and $(M, M...., M)$. 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 $(M, M, ..)$ 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 $\theta = (H, H, ..., H)$, $\mu(i) > \mu(j) \Leftrightarrow i > j$, so the assignment μ such that $(\mu, \theta) \in S'$ is the identity assignment ι.

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 $n, n-1$,..,$n-m$

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 $1, 2, ...k$ 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 $s' = (\mu, (H, H, H, ...H))$, there exists a path to $s^H = (\iota, (H, H, ..., H))$. 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.

7 References

1. Abdulkadiroglu, A. and T. Sonmez (1999) "House allocation with existing tenants," *Journal of Economic Theory* **88**, 233-260.

2. Abdulkadiroglu, A. and T. Sonmez (2003) "School choice: A mechanism design approach," *American Economic Review* **93**, 729-747.

3. Bartholomew, J. (1982) *Stochastic Models for Social Processes*, New York: Wiley.

4. Carmichael, L. (1983) "Firm-specific human capital and promotion ladders," *Bell Journal of Economics* **14**, 251-258.

5. Chiappori, P.A., B. Salanié and J. Valentin (1999) "Early starters versus late beginners," *Journal of Political Economy* **107**, 731-760.

6. Fudenberg, D. and J. Tirole (1991) *Game Theory*, Cambridge, MA: MIT Press.

7. Gale, D. and L. Shapley (1962) "College admissions and the stability of marriage," *American Mathematical Monthly* **69**, 9-14.

8. Isaacson, D. and R. Masden (1976) *Markov Chains: Theory and Applications*, New York: John Wiley.

9. Iyer, L. and A. Mani (2008) "Traveling agents? Political change and bureaucratic turnover in India," mimeo., Harvard Business School and University of Warwick.

10. Kemeny, J. and L. Snell (1960) *Finite Markov Chains*, New York: Van Nostrand.

11. Kleinrock, L. (1975) *Queuing Systems I. Theory*, New York: Wiley.

12. Kleinrock, L. (1976) *Queuing Systems II. Computer Applications*, New York: Wiley.

13. Lazear, E. (1995) *Personnel Economics*, Cambridge, MA: The MIT Press.

14. Lee, S. (2004) "Seniority as an employment norm: The case of layoffs and promotions in the US employment relationship," *Socio-Economic Review* **2**, 65-86.

15. Moulin, H. and R. Stong (2002) "Fair queuing and other probabilistic allocation methods," *Mathematics of Operations Research* **27**, 1-31.

16. Moulin, H. and R. Stong (2003) "Filling a multicolor urn: An axiomatic analysis," *Games and Economic Behavior* **45**, 242-269.

17. Nilakantan, K. and B.G. Ragavendhra (2005) "Control aspects in proportionality Markov manpower systems," *Applied Mathematical Modelling* **29**, 85-116.

18. Roth, A. and M. Sotomayor (1990) *Two-sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Cambridge: Cambridge University Press.

19. Shapley, L. and H. Scarf (1974) "On cores and indivisibilities," *Journal of Mathematical Economics* **1** 23-37.

20. Thomson, W. (2007) "Fair allocation rules," mimeo., University of Rochester.

21. Vajda, S. (1978) *Mathematics of Manpower Planning*, Chichester: John Wiley.

8 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: THRESHOLDS FOR TRANSFERS OF ENGLISH TEACHERS 2008

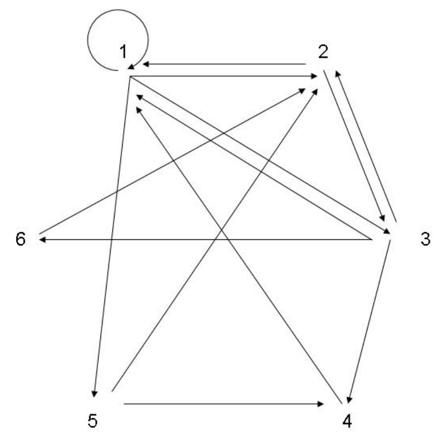

Figure 2: Markov process for three agents

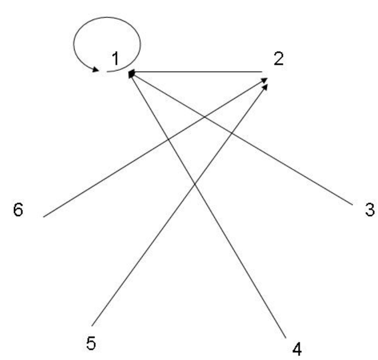

Figure 3: Seniority assignment for three agents

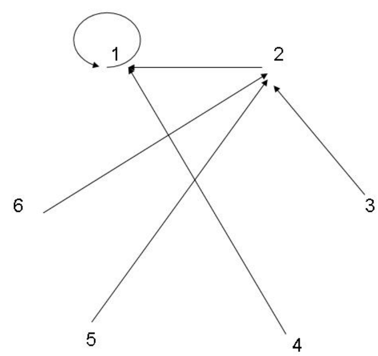

Figure 4: Rank assignment for three agents

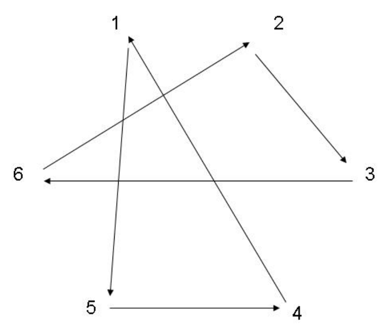

Figure 5: Replacement assignment for three agents

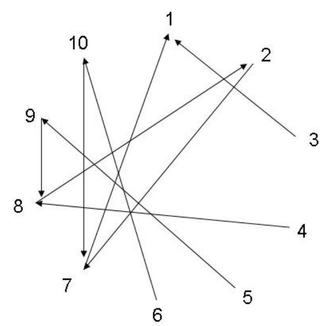

Figure 6: Transitions between states for Example 2

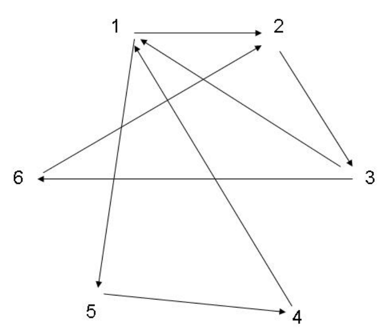

Figure 7: Transitions between states for Example 3