Centro de Estudios Económicos
El Colegio de México, A.C.
www.colmex.mx

Serie documentos de trabajo

MARKOVIAN ASSIGNMENT RULES

Francis Bloch
Ecole Polytechnique

David Cantala
El Colegio de México

DOCUMENTO DE TRABAJO
Núm. XVIII - 2010
.

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 de Clippel, Philippe Jéhiel, Jean-François Laslier, Hervé Moulin and Utku Ünver for constructive discussions about the paper. We are also grateful to seminar participants at Auckland, Columbia, CORE, Harvard/MIT, Paris School of Economics, Sevilla, Stanford GSB and Yale as well as audiences at conferences in Caen, Stony Brook and Maastricht (SED 2009) for their comments. Francis Bloch thanks the Department of Economics at Brown University for its hospitality when the project was started, and David Cantala thanks CONACYT for financial support.

1 Introduction

How should one allocate offices to faculty members and graduate students in academic departments, dormitory rooms to students, subsidized housing to low-income families? How should one post career diplomats to embassies, officers to military bases, teachers to high schools? All these assignment problems share a common structure: they involve a fixed number of durable objects (positions, offices, housing units) which are successively assigned to generations of agents. Because the number of objects is fixed, reassignments typically only occur when one agent (or a cohort of agents) leaves the organization. For example, offices and embassy posts are reassigned when faculty members or diplomats leave or retire ; dormitory rooms become vacant when students leave the university.[2] At this point, assignment mechanisms might generate a cascade of reassignments, as agents who receive a new object relinquish their old object which in turn is assigned to some other agent.

In practice, the rules which are used to solve these assignment problems vary widely.[3] Oftentimes, the rules involve a priority structure based on seniority, merit and the history of past assignments. Many rules also respect an individual rationality constraint, as agents cannot be forced to accept an object worse than the one they currently hold.[4] Finally, the rules often balance the benefits of reassigning objects with the moving costs of reallocations. Our objective in this paper is to provide an axiomatic study of assignment rules for durable objects to successive generations of agents, in order to better understand the importance of seniority, randomness and moving costs in the design of dynamic allocation rules.

The dynamic assignment problems we consider differ from standard static assignment problems in three important dimensions. First, the set of agents varies over time, as agents enter and exit society. Second, as agents live for

[2]Reassignments could also occur even if no agent leaves the organization. However, we will focus on situations where agents cannot jointly benefit from reallocating the objects among themselves, so that reassignment opportunities only arise when an agent leaves the organization.

[3]Kurino (2008) discusses the rules used by American universities and colleges to assign students to dormitories. The assignment of teachers to French high schools which spurred our interest in Markovian assignment rules is documented on the website of the French Ministry of Education, http://www.education.gouv.fr. Casual discussions at different academic departments show that the rules used to allocate offices to faculty members vary widely from uniform random assignments to fixed, seniority-based priority rules.

[4]This individual rationality constraint, which was introduced by Abdulkadiroglu and Sonmez (1999) in the context of house allocation mechanisms, is almost always imposed in mechanisms to allocate student dormitories, as noted by Kurino (2008). This constraint is also binding in the mechanism allocating teachers to high schools in France.

more than one period, they evaluate entire sequences of assignments over their lifetime rather than a single assignment. Third, by the individual rationality constraint, the assignment at date t constrains the assignment at date $t+1$, creating a dynamic dependence in the assignment problem. These features introduce new difficulties, and lead us to define new dynamic axioms over assignment rules and to apply techniques borrowed from the study of finite Markov chains in order to study the long run properties of assignment rules. In order to deal with this dynamic dimension of assignment problems, we make a number of simplifying assumptions. First, we suppose that objects are vertically ranked and that all agents share the same preferences.[5] Second, we restrict attention to Markovian assignments, which do not depend on history or calendar time, but only on the current assignment of objects to agents. Finally, we suppose that utilities are additive separable over different periods. Under these simplifying assumptions, we are able to obtain sharp characterizations of assignment rules satisfying various dynamic properties.

In the first part of the paper, we consider homogeneous societies formed of identical agents. All assignments are equally efficient, and we concentrate our analysis on equity. Our main result is easily illustrated in a simple example where a single object is assigned to agents living two periods. When the object becomes available, there are only two ways to assign it: either to the old or to the young agent. If the object is assigned to the old agent (the seniority rule), every agent will obtain it for one period when he is old. If the object is assigned to the young agent (the replacement rule), agents who are born every other period will keep the object for two periods, whereas other agents will never get it. Hence, the seniority rule is fair in the sense that it guarantees the same sequence of assignments to every agent irrespective of his date of birth. On the other hand, the replacement rule, while very inequitable, minimizes reassignment costs. This intuition can be generalized to a society with an arbitrary number of agents and objects. We characterize a family of rules (convex combinations of the seniority rule and the rule which assigns object j to the agent who owns object $j-1$) as the only rules which satisfy two properties of independence (the assignment of object j to agent i is independent of the current assignment of objects to other agents) and fairness (in the long run, every agent is guaranteed to obtain the same sequence of assignments). However, these rules imply that all objects are reassigned every period, and result in high reallocation

[5]This assumption is reasonable in some contexts, like the assignment of diplomats to embassies or teachers to high schools, where most diplomats and teachers have a common preference over different positions. It is also likely to hold in the assignment of offices or objects whose characteristics result in a clear vertical ranking.

costs. By contrast, the replacement rule minimizes reassignment costs. We also study two other dynamic properties of the Markov chain generated by assignment rules and provide sufficient conditions for ergodicity (in the long run, assignments of agents are independent of the initial conditions) and irreducibility (all assignments can be reached with positive probability in the long run).

In the second part of the paper, we analyze heterogeneous societies where agents' productivities are drawn at random at the time they are born. We assume that the surpluses generated by a match between an agent and an object is supermodular, so that an efficient assignment involves assortative matching, with agents of higher productivity receiving objects of higher rank. Our results for heterogeneous societies are mostly negative, showing incompatibility between the Markovian structure of the assignment rule and efficiency. In particular, we prove two impossibility results concerning the existence of assignment rules which are independent and efficient and of assignment rules which are efficient and guarantee that two agents born at different times in identical societies receive the same sequence of assignments.[6] This last result highlights the strength of the requirement that assignment sequences be independent of history in heterogeneous societies, and indicates that some level of inequity must be tolerated to reach efficient assignments.

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 Sonmez (1999) in the context of house allocations. (See also Sonmez and Ünver (2008) and 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.

In independent work, Kurino (2008) has formulated a similar model of dynamic assignment of objects to overlapping generations of agents. On the one hand, Kurino (2008) considers a more general framework, allowing agents to have nonseparable utilities over sequences of assignments and to have different preferences. He also supposes that groups of agents (rather than

[6] This last impossibility result is obtained when the number of possible productivities is greater or equal to three. In dichotomous societies where agents can only draw low or high productivities, efficient and fair assignment rules exist.

single individuals) enter and leave society at every period. On the other hand, instead of considering abstract families of assignment rules, Kurino (2008) focuses on a small number of special rules, and discusses the properties that these rules satisfy. The two approaches are complementary, addressing the same problem from very different angles, raising different issues and giving different solutions.

The paper is also related to a broader literature embedding assignment problems in dynamic contexts.[7] Ünver (2010) 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. Abdulkadiroglu and Loertscher (2007) analyze a two-period dynamic house allocation problem where agents who choose a bad house in the first period get priority for the good house in the second, and show that this dynamical linkage between allocations improves welfare. In a very different context where side-payments are allowed, Athey and Segal (2007), Bergemann and Välimäki (2010) and Gershkov and Moldovanu (2009a, 2009b), Parkes and Singh (2003), 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.

The rest of the paper is organized as follows. We present the model and illustrate it with an overlapping generations model where agents live for three periods in Section 2. Section 3 is devoted to the analysis of the model with homogeneous agents, and Section 4 considers the model with heterogeneous agents. In Section 5, we discuss two extensions of the analysis and outline some directions for future research.

[7]There 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 The model

2.1 Agents, objects and assignment rules

2.1.1 Agents and objects

We consider an overlapping generations model where every agent lives for n periods. At each date $t = 0, ...,$ one agent enters and one agent exists society.[8] Hence, every agent is characterized by his date of birth t.

There are n durable objects, indexed by $j = 1, 2, ...n$ which are successively reassigned to agents in society. [9] Every agent thus receives a sequence of n assignments over his lifetime, $\mathbf{j} = (j_1, ..., j_n)$. We suppose that agents have additively separable utilities,

$$U_t = \sum_{\tau=1}^{n} u_t(j_\tau).$$

Furthermore, we suppose that all agents have common preferences over the objects, and that objects are ordered in such a way that higher objects generate higher utility, so that

$$u_t(j) = u(j),$$

where $u(j) > u(j')$ if and only if $j > j'$. Finally, we concentrate on *Markovian* assignment rules, which are independent of calendar time t, so that an agent's assignment only depends on his *age* and not on his *date of birth*. As a consequence, we can identify any agent in society by his age $i = 1, ..., n$ rather than his date of birth t.

2.1.2 Assignments and assignment rules

An assignment μ is a one-to-one bijective mapping between the set of agents $I = \{1, 2, ..., n\}$ and the set of objects $J = \{1, 2, ..., n\}$.[10] We denote by M the set of all assignments.

[8] We make this assumption to simplify notations. However, as indicated in Section 5 our analysis carries over to a model where groups of agents enter and exit society at any date t.

[9] Again, the assumption that there is an equal number of objects and agents is made for convenience; but the analysis could easily be extended to societies where the number of agents and objects are different, as discussed in Section 5.

[10] Alternatively, the matching μ can be represented as a permutation on the set $\{1, 2, ..., n\}$.

In our dynamic model, we are interested in mechanisms which, starting from a given assignment μ at date t generate a probability distribution over new assignments at date $t+1$. General Markovian assignment rules would specify transition probability matrices over M. However, most rules used in practice are not formulated in terms of general Markov chains, which allow for simultaneous reallocation of objects, but as rules allocating single objects. In this paper, we adopt this view, and consider a more restrictive setting, where an assignment rule operates object by object rather than over the entire set of objects simultaneously. Hence, a Markovian assignment rule determines, for each object j and given the assignment of objects in $J \setminus j$ to the agents in I, the probability with which object j is assigned to an agent of age i. We incorporate the individual rationality constraint in this definition by assuming that object j is only assigned with positive probability to agents who hold an object worse than j.

Formally, we define a *truncated assignment* ν as a one-to-one injective map from $I \setminus \{1\}$ to J, assigning to any agent but the youngest agent, an object in J. A *Markovian assignment rule* α is a collection of probability distributions over I, $\alpha_j(\nu)$ for any $j \in J$, any ν such that j is not assigned in ν, such that $\sum_{i|\nu(i)<j} \alpha_j(\nu, i) = 1$.

2.1.3 Four natural assignment rules

We describe four natural ways to assign objects in our dynamic setting.

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 $i = 1$.

2.1.4 Markov chains generated by assignment rules

Given an assignment rule α, we compute the transition probability matrix of the Markov chain over assignments generated by α. To this end, note that for any two assignments μ and μ', there is a unique sequence of agents

$i^0 = n+1, i^1 = \mu'^{-1}(\mu(i^0-1)), ..., i^m = \mu'^{-1}(\mu(i^{m-1}-1)), ..., i^M = 1$, such that the reassignment of objects following the sequence leads from μ to 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 the probability that the sequence of reallocations of goods between agents $i^0, ..., i^M$ is realized:

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

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

We conclude by noting that assignment rules operating over single objects are *less general* than assignment rules allowing for simultaneous reallocations. As the following example shows, there exist Markov chains over $\mathcal{M}$ that cannot be generated by single object assignment rules.

Example 1 *Let $n = 3$. We represent an assignment as a triple (i, j, k), where $i = \mu(1), j = \mu(2)$ and $k = \mu(3)$. We focus on the three assignments $\mu_1 = (1,2,3), \mu_2 = (1,3,2), \mu_3 = (2,1,3)$. Consider the Markov process $p(\mu_1|\mu_1) = 1$ and $p(\mu_3|\mu_2) = 1$. Let $\alpha_2(\nu)$ be the assignment rule allocating object 2 for the truncated assignment $\nu(2) = 1, \nu(3) = 3$. 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.*

2.2 An illustrative example

We analyze assignment rules in a three agent society. An assignment is given by a triple (i, j, k), where $i = \mu(1), j = \mu(2)$ and $k = \mu(3)$. There are six possible assignments:

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

We analyze the properties of the Markov chains generated by four natural assignment rules.

2.2.1 The seniority rule

The seniority rule assigns an object to the oldest agent eligible to receive it. In order to understand how the Markov chain generated by the seniority rule is constructed, suppose that the current assignment is μ_1. At the next period, object 3 is reassigned to agent 3 who holds object 2. In turn, object 2 is reassigned to agent 2, who holds object 1, and object 1 is finally allocated to the entering agent. This shows that the Markov chain generated by the seniority rules moves from μ_1 to μ_1 with probability 1. The transition probability matrix is given by:

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

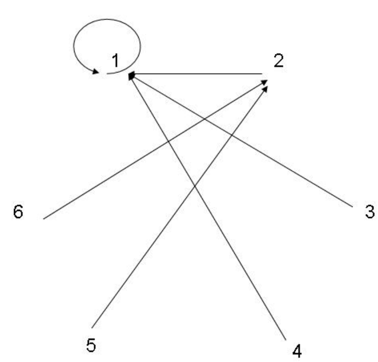

Figure 1: Seniority assignment for three agents

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

2.2.2 The rank rule

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

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

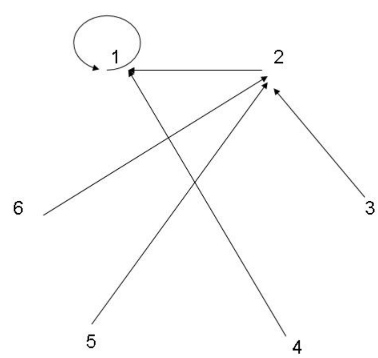

Figure 2: Rank assignment for three agents

Figure 2 shows the transitions between states for the Markov chain generated by the rank rule. The dynamic 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 absorbing state μ_1 in at most two steps.[11]

[11] The 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.

2.2.3 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:

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

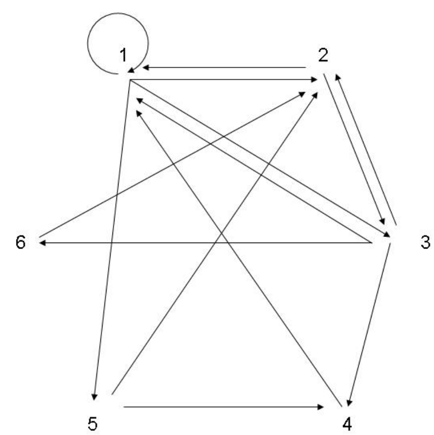

Figure 3: Uniform assignment for three agents

The transitions between states are given in Figure 3. All feasible transitions (which satisfy the individual rationality condition) are present on the graph. 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 the unique invariant distribution $(p_1, ..., p_6)$ over states $(\mu_1, \mu_2, ..., \mu_6)$:

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

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

2.2.4 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:

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

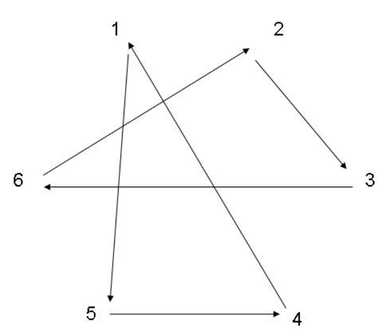

Figure 4: Replacement assignment for three agents

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.

3 Properties of Assignment Rules in Homogeneous Societies

In this Section, we characterize assignment rules in homogeneous societies by imposing axioms of independence, convergence, ergodicity and irreducibility.

3.1 Independent Assignment Rules

We first state our independence axiom:

Definition 1 *A Markovian assignment rule α satisfies* independence *if, for any i, j, 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 current assignment) and not on the characteristics of the other agents. Independent assignment rules are simple to implement, have low informational requirements, and capture the idea that an agent should not be held responsible for the way objects are assigned to the other agents. A stronger independence property states that an agent's assignment only depends on the object he currently holds and not on his age:

Definition 2 *A Markovian assignment rule α satisfies* strong independence *if, for any i, i', any j and any ν, ν' such that $\nu(i) = \nu'(i')$, $\alpha_j(\nu, i) = \alpha_j(\nu', i')$.*

The rank, uniform and replacement rules are all strongly independent. The seniority rule is not independent.[12] As the following Lemma shows, the gap between independent and strongly independent rules is small.

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

Lemma 1 shows that if a Markovian assignment rule satisfies independence, the assignment of any object $j < n$ is strongly independent. However, this property *does not hold for the assignment of the highest object*, n. For

[12] The seniority rule satisfies a weaker independence property, namely that the assignment of object j to an agent only depends on the assignment of agents who are eligible to receive object j.

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.

This result relies on a simple algebraic argument. Consider the assignment of some object j under two truncated assignments $\tilde{\nu}$ and $\tilde{\nu'}$ which only differ in the assignment of objects to some pair of agents, i, i'. In other words, $\tilde{\nu}(i) = \tilde{\nu'}(i') < j$, $\tilde{\nu}(i') = \tilde{\nu'}(i')$ and $\tilde{\nu}(k) = \tilde{\nu'}(k)$ for all $k \neq i, i'$. By independence, the assignment of object j to an agent $k \neq i, i'$ must be identical under $\tilde{\nu}$ and $\tilde{\nu'}$, so that $\alpha_j(\tilde{\nu}, k) = \alpha_j(\tilde{\nu'}, k)$. Now, because $\alpha_j(\tilde{\nu}, \cdot)$ and $\alpha_j(\tilde{\nu'}, \cdot)$ are probability distributions,

$$\sum_k \alpha_j(\tilde{\nu}, k) + \alpha_j(\tilde{\nu}, i) + \alpha_j(\tilde{\nu}, i') = 1 = \sum_k \alpha_j(\tilde{\nu'}, k) + \alpha_j(\tilde{\nu'}, i) + \alpha_j(\tilde{\nu'}, i'),$$

implying that $\alpha_j(\tilde{\nu}, i) + \alpha_j(\tilde{\nu}, i') = \alpha_j(\tilde{\nu'}, i) + \alpha_j(\tilde{\nu'}, i')$.

At this point, we distinguish between the allocation of the best object n and of any other object $j < n$. If one assigns $j < n$, the argument above applies to any assignment such that $\tilde{\nu}(i') = \tilde{\nu'}(i) > j$, so that $\alpha_j(\tilde{\nu}, i') = \alpha_j(\tilde{\nu'}, i) = 0$, resulting in $\alpha_j(\tilde{\nu}, i) = \alpha_j(\tilde{\nu'}, i')$. To finish the argument, we need to show that this equality holds for all assignments ν and ν' such that $\nu(i) = \nu'(i')$. This last step derives from the independence of the assignment rule: as $\nu(i) = \tilde{\nu(i)}$ and $\nu'(i') = \tilde{\nu'}(i')$, by independence, $\alpha_j(\nu, i) = \alpha_j(\tilde{\nu}, i) = \alpha_j(\tilde{\nu'}, i') = \alpha_j(\nu', i')$.

We finally note that strongly independent assignment rules have a very simple representation: for any object j, the assignment rule specifies a probability distribution over the finite set $\{1, 2, ..., j-1\}$, where $\alpha_j(k)$ describes the probability that object j is assigned to the agent who holds object k.

3.2 Dynamic properties of assignment rules

We recall the definitions of well-known properties of finite Markov chains (See Kemeny and Snell (1960) or Isaacson and Madsen (1976)).

Definition 3 *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 4 *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 5 *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 notions in hand, we can define dynamic properties of the Markov chains generated by assignment rules:

Definition 6 *A Markovian assignment rule α is* irreducible *if the induced Markov chain is irreducible (the only recurrent set is the entire state space). A Markovian assignment rule α is* ergodic *if the induced Markov chain is ergodic (has a unique recurrent set).*[13] *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).*

Why do we want to study the dynamical properties of Markovian assignment rules? We claim that each of the three properties highlighted above has important implications on the study of assignment processes. If an assignment rule is irreducible, we know that all assignments μ in $\mathcal{M}$ will be visited with positive probability, and can compute the unique invariant distribution as the left eigenvector of the transition probability matrix satisfying $\sum_i \pi_i = 1$ (Isaacson and Masden (1976), Theorem III.2.2 p. 69). If an assignment rule is ergodic, we know that the Markov chain will settle in a unique long run behavior irrespective of the initial conditions. If an assignment rule is convergent, we know that in the long run, all agents will experience the same sequence of assignments in their lifetime. Convergence is thus related to the following strong notion of fairness.

Definition 7 *An assignment rule is* fair *if for any two agents i and i' entering society at dates t and t', the assignment rule α generates a deterministic sequence of assignments such that $\mu^{t+\tau}(i) = \mu^{t'+\tau}(i')$ for $\tau = 0, 1, ..., n-1$.*

Clearly, any convergent assignment rule is fair, and if a fair assignment rule has a unique recurrent set, it must be convergent.

[13]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)for whom a Markov chain is "ergodic" if the recurrent set is the entire state space. 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.

3.3 Convergent Assignment Rules

3.4 Convergent Markovian assignment rules

We start by characterizing convergent assignment rules. Because an agent is never reassigned an object of lower value than the one he currently holds, the only candidate for an absorbing state is the identity assignment ι, $\iota(i) = i$ for all $i = 1, ...n$. It is easy to check that the identity assignment is an absorbing state for the seniority and rank rules, and that, starting from any other assignment, the identity assignment is reached in a finite number of steps.

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

The seniority and rank rules 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. However, we can characterize fully 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.*

3.5 Ergodic assignment rules

There is no simple characterization of ergodic assignment rules. The following condition is a sufficient condition for ergodicity of the Markov chain generated by an assignment rule.

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

In other words, if the assignment rule assigns object j with positive probability to the agent who holds object $j-1$, then the assignment rule is ergodic. The proof of Proposition 2 is based on the following observation. Under the condition of the Proposition, there exists a path from any assignment μ to the identity assignment ι, showing that the recurrent set is unique.

The proof of Proposition 2 generalizes the argument used to prove that the identity assignment ι can be reached from any other assignment using the rank or seniority rules. Both the rank and seniority rules satisfy the condition of Proposition 2, as $\alpha_j(j,\nu) = 1$ whenever $\nu(j) = j-1$. Furthermore, if the condition of Corollary 2 is satisfied, 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), the recurrent set does not admit any cycle and the Markov chain 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 $\alpha_4(3) = 1, \alpha_3(1) = 1, \alpha_2(1) = 1, \alpha_1(0) = 1$[14].*

Let all states such that $\mu(4) = 4$ be ordered as in Subsection 2.2. 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.

[14]For 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.

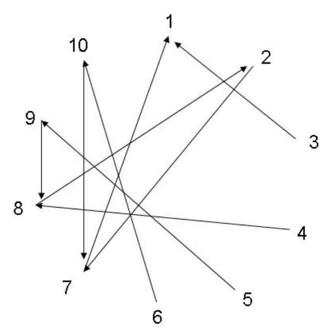

Figure 5: TRANSITIONS BETWEEN STATES FOR EXAMPLE 2

3.6 Irreducible assignment rules

The characterization of assignment rules generating irreducible Markov chains is difficult and we provide instead a sufficient condition for irreducibility.

Theorem 2 *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 2 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 2.2.

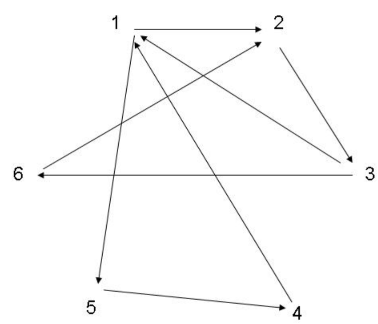

Figure 6: Transitions between states for Example 3

4 Properties of assignment rules in heterogeneous societies

4.1 Heterogeneous societies

4.1.1 Agents and matchings

In heterogeneous societies, agents draw independently a productivity k in the set $K = \{1, 2, ..., m\}$ according to the probability distribution $q(k)$. This productivity is drawn by the agents at birth and remains constant over their lifetime. Agents are now characterized by a *pair* (i, k) describing their age and productivity. In heterogeneous societies, the value of the match between an object and an agent is a function $u(j, k)$ which depends both on the quality of the object and the agent's productivity. We assume that agent's types are ordered so that higher productivity agents generate larger surpluses:

$$u(j, k') > u(j, k) \text{ for all } k' > k,$$

and that the function $u(j, k)$ is supermodular:

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

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

In heterogeneous societies, we define a mapping $\theta : I \to K$, assigning to each agent a productivity type, and let $\mathcal{T}$ define the set of all those mappings. Any $\theta \in \mathcal{T}$ defines a type profile for agents in the society. A state is now described by a pair (μ, θ) in $\mathcal{M} \times \mathcal{T}$, and a Markovian assignment rule defines assignment probabilities which depend both on the truncated assignment ν and on the type profile θ: $\alpha_j(\nu, \theta, i)$.

4.1.2 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 criterion. Let $\delta \in [0, 1)$ denote the discount factor of the social planner.

Definition 8 *An assignment rule α is efficient if it maximizes:*

$$\mathbf{E}_{k_0,\ldots} \sum_{t=0}^{\infty} \delta^t \sum_{i=1}^{n} u_i(\mu_t(i), \theta_t(i)).$$

We thus adopt the expected discounted sum of utilities as our efficiency criterion. Notice that, when $\delta = 0$, this criterion becomes equivalent to the myopic efficiency criterion, where the assignment rule maximizes the sum of expected surpluses at every period. However, by the individual rationality constraint, the assignment at period t imposes constraints on the assignment of period $t + 1$, so that the dynamic efficiency criterion we adopt typically differs from the myopic criterion when δ is positive.

4.1.3 Quasi-convergent assignment rules

When agents are heterogeneous, an assignment rule generates a Markov chain over the set $\mathcal{M} \times \mathcal{T}$. As types are randomly drawn every period, this Markov chain does not admit absorbing states, and it is clear that assignment rules *are never convergent*. However, distinguishing between the two sources of randomness (one linked to the exogenous draw 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 9 *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-convergence is also related to the following strong notion of fairness:

Definition 10 *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 $\tau = 0, 1, ..., n-1$, the assignment rule α generates a deterministic sequence of assignments such that $\mu^{t+\tau}(i) = \mu^{t'+\tau}(i')$ for $\tau = 0, 1, ..., n-1$.*

This fairness notion extends the definition of fairness in homogeneous societies, by requiring that two agents born at different times but who live throughout their entire lifetime in the same societies experience the same sequence of assignment. As in the case of homogeneous societies, it is easy to see that fairness and quasi-convergence are equivalent notions when the Markov chain admits a unique recurrent set.

4.1.4 Type-lexicographic assignment rules

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

Definition 11 *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 $\nu(l) < j$.

4.2 Properties of assignment rules in heterogeneous societies

4.2.1 Independent assignment rules in heterogeneous societies

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 θ, θ', any j, ν and i, $\alpha_j(\nu, \theta, i) = \alpha_j(\nu, \theta', i)$.*

Lemma 2 is easily proved. Consider two type profiles $\theta, \tilde{\theta}$ such that $\theta_k = \tilde{\theta}_k$ for all $k \neq i$ and $\theta_i \neq \tilde{\theta}_i$. For any j and any ν, $\sum_{l|\nu(l)<j} \alpha_j(\nu, \theta, l) = \sum_{l|\nu(l)<j} \alpha_j(\nu, \tilde{\theta}, l) = 1$. By independence, $\alpha_j(\nu, \theta, k) = \alpha_j(\nu, \tilde{\theta}, k)$ for any $k \neq i$, so that $\alpha_j(\nu, \theta, i) = \alpha_j(\nu, \tilde{\theta}, i)$. Applying independence again, this argument shows that for any type profile θ' such that $\theta'(i) = \tilde{\theta}(i)$,$\alpha_j(\nu, \theta, i) = \alpha_j(\nu, \theta', i)$.

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). Lemma 2 thus shows that in heterogeneous societies, independence and efficiency are incompatible. To illustrate this point, it suffices to consider a two-type society where agents can either be of high or low type. In these dichotomic societies, efficient assignment rules must be type-lexicographic and assign any object to a high type agent when one is eligible to receive it. This will of course stand in contradiction to independence, as it implies that the assignment rule takes into account agents' types.

4.2.2 Efficient and quasi-convergent assignment rules

We now claim that efficiency and quasi-convergence are incompatible requirements for assignment rules in heterogeneous societies.To understand this point, consider a society with three possible types, L, M and H. 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 3.4). Hence, if the type profile $(L, ..., L)$ is realized, $\mu(i) = \iota$. Similarly, if the type profile $(M, ..., M)$ is realized, $\mu(i) = \iota$. Now starting from $(L, ..., L)$ consider successive changes in the type profiles with the sequential entry of one M and $n-1$ H types. When the M agent enters, if the discount factor δ or the probability of high types $p(H)$ is sufficiently close to zero, it is efficient for the social planner to allocate the highest object to the M agent. In the next period, the entering agent H will be assigned object $n-1$, and the agent H entering in the following period object $n-2$. In the end, in society $(H, .., H, M)$, agent M will possess object n, and any agent H of agent i object i.

Suppose instead that $n-1$ H agents enter a society composed only of M agents. By efficiency, when the first H agent enters, he will be assigned object n, the second object $n-1$, etc. In the end, in society $(H, ..., H, M)$, any H agent of age i will be assigned object $i+1$ and the oldest agent of type M object 1. Hence, a young H agent entering a society $(H, ..., H, M)$ will either be assigned object 1 or object 2, depending on past history. This argument shows that the assignment rule is not quasi-convergent and treats differently two identical agents based on past history. This argument establishes the following general result.

Theorem 3 *Suppose that $|K| \geq 3$. There exist probability distributions over types q and/or discount factors δ such that no assignment rule can simultaneously satisfy efficiency and quasi-convergence.*

The incompatibility between quasi-convergence and efficiency is a consequence of the individual rationality constraint, which creates a path dependence that prevents the emergence of quasi-convergent rules when efficiency is satisfied.[15] The intuition underlying Theorem 3 relies on the existence of at least three types. With only two types, efficiency and quasi-convergence can be satisfied simultaneously.

[15] 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 4 *Suppose that $|K| = 2$. The type-rank and type-seniority rules are both efficient and quasi-convergent.*

Theorem 4 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'). 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.[16]

[16]Incidentally, in the assignment of high schools to teachers in France, teachers are actually grouped into two classes (”agrgs” and ”certifis”) reflecting the grades they obtained in the certification exam).

5 Extensions and Conclusions

5.1 Extensions

5.1.1 Group entry, cohort size and numbers of objects

In order to derive the properties of dynamic assignment rules, we have assumed that agents enter sequentially, that there are as many objects as agents, and as many object types as periods in an agent's life. All these assumptions could be relaxed without changing the qualitative results of our analysis. For example, if agents enter in groups and there are more object types than periods in one agent's life, all agents will not experience the same history. However, if one defines the state as the distribution of objects held by agents of the same cohort, our analysis carries over and we can study convergence of the Markov chain generated by assignment rules over the extended state space. Similarly, if there are fewer objects than agents, one needs to redefine the state as the probability that an agent in a given cohort receives an object. If there are more objects than agents, one would restrict attention to the best objects and not assign the worst objects.

5.1.2 Agents' strategic behavior

In the analysis, we have assumed that agents do not behave strategically and always accept the object which is assigned to them. Agents could refuse an assignment for different reasons. If the assignment rule depends on the current truncated assignment, agents may prefer to wait and obtain a better object after the assignment of other agents have changed . If the assignment rule is such that agents with better objects have a lower probability of obtaining objects of higher value, agents may prefer to wait in order to guarantee that they will eventually obtain the best objects.

Interestingly, the four assignment rules we have focused on (rank, seniority, uniform and replacement rules) are all immune to this strategic behavior. Because the rank, uniform and replacement rules are independent, agents have no incentive to wait for the assignment of other agents to change. For the seniority rule, waiting can never be beneficial either: if an agent is the most senior eligible agent to receive an object, he will remain the most senior eligible agent independently of the assignment of objects to other agents. Furthermore, in the rank and seniority rules, an agents' probability of receiving object is increasing in the quality of the object held by the agent. In the uniform and replacement rules, the agents' probability of receiving an object is independent of the object currently held. Hence, in all four cases, agents have no incentive to wait in order to "jump" to an object of higher quality.[17]

5.2 Conclusions

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 experience the same history) and analyze the dynamical properties of the Markov chains generated by assignment rules. When agents are homogeneous, we characterize independent convergent assignment rules. 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.

While our analysis represents a first step in the understanding of dynamic assignment processes, it is based on a number of simplifying assumptions. As we discussed above, assumptions on the number of objects and the sizes of cohorts could be relaxed without changing the analysis. Other assumptions cannot be so easily dispensed with. 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 in a deterministic way. Studying dynamic assignment mechanisms with stochastic entry and exit is a challenging task that deserves further study.

[17] In fact, all independent rules such that $\alpha_j(i)$ is weakly increasing in $\nu(i)$ will be immune to strategic behavior on the part of agents.

6 References

1. Abdulkadiroglu A. and S. Loertscher (2007)“Dynamic house allocations,”, mimeo., Duke University and University of Melbourne.

2. Abdulkadiroglu, A. and T. Sonmez (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. Välimäki (2010) “The dynamic pivot mechanism,”*Econometrica* **78**, 771-789.

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. Gershkov, A. and B. Moldovanu (2009a) “Learning about the future and dynamic efficiency,” *American Economic Review* **99**, 1576-1587.

10. Gershkov, A. and B. Moldovanu (2008b) “Dynamic revenue maximization with heterogeneous objects: A mechanism design approach,” *American Economic Journal: Microeconomics* **1**, 168-198.

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

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

13. Kurino, M. (2008) “House allocation with overlapping agents: A mechanism design approach,” mimeo., University of Pittsburgh.

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

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

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

17. 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)*

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

19. Sonmez, T. and U. Ünver (2008) “House allocation with existing tenants: A charcetrization, ”forthcoming, *Games and Economic Behavior*.

20. Talluri, K. and G. van Ryzin (2004) *The theory and practice of revenue management*, Springer: New York.

21. Thomson, W. (2007) “Fair allocation rules,” mimeo., University of Rochester.

22. Ünver, U. (2010) “Dynamic kidney exchange,” *Review of Economic Studies* **77**, 372-414.

7 Proofs

Proof of Lemma 1: in the text.

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:

$$\prod_j \alpha_j(\tilde{\nu}^j, j) = 1, \tag{2}$$

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. At period 2, object 2 is reassigned to agent 2 (who currently holds object 1) and object 1 is reassigned to the entering agent. Following this argument, is it easy to see that the Markov chain will be absorbed into the identity assignment in at most n periods.

In the seniority rule, notice that the entering agent receives object 1 with probability 1. As for the rank rule, this implies that, starting from any assignment, in period 1, agent 1 holds object 1 ; in period 2, agent 2 holds object 2 and agent 1 object 1, etc. The identity assignment is reached in at most n steps.

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 Proposition 2: Consider any assignment μ. We will show the existence of a path to the identity assignment ι. Because object 1 is reassigned at least every n periods, the Markov chain will eventually reach an assignment μ^1 where $\mu^1(1) = 1$. will be reassigned and, by the condition in the Proposition, there is a positive probability that all objects $j = 1, 2, ..., \mu^1(n)$ will be reassigned to the agents holding object $j-1$ for $j \geq 1$ and object 1 to the entering agent. There is thus a new assignment μ^2 with $\mu^2(1) = 1$ and $\mu^2(2) = 2$ which can be reached with positive probability. Repeating the argument, we reach the identity assignment ι in at most n steps.

Proof of Theorem 2: We prove the existence of a path from any assignment to any assignment. First notice that, in at most n steps, the Markov chain will reach an assignment μ where $\mu(n) = n$. Furthermore, as $\alpha_j(1) > 0$, any assignment can be reached from some assignment μ' where $\mu'(n) = n$ in at most n steps by successively assigning the good held by the retiring agent to the entering agent. Hence, in order to prove the theorem, it suffices to prove that there exists a path from any assignment μ such that $\mu(n) = n$ to any assignment μ' such that $\mu'(n) = n$.

Consider two such assignments, and let $\pi = \mu' \circ \mu^{-1}$ be the permutation over the set of objects J such that $\mu'(i) = \pi[\mu(i)]$ for all i. Notice that π leaves the last object invariant, $\pi(n) = n$. Recall that a *transposition* τ_{ij} is a bijective mapping on some index set such that $\tau_{ij}(i) = j, \tau_{ij}(j) = i$ and $\tau_{ij}(k) = k$ for all $k \neq i, j$. The permutation π can be decomposed as a product of transpositions, and we let $i^1, i^2, .., i^Q$ denote the sequence of transpositions such that

$$\pi = \tau_{i^Q i^{Q-1}} \circ ... \circ \tau_{i^2 i^1}.$$

We now exhibit a path from μ to $\tau_{i^2 i^1} \circ \mu$. As transpositions are symmetric, we can assume without loss of generality that $i^1 > i^2$. Because the graph $G(\alpha)$ is connected, there exists a sequence $j^1 = i^1, ..., j^t, ..., j^T = i^2$ such that $\alpha_{j^t}(j^{t-1}) > 0$ for all $t = 2, ..., T$. We can decompose the transposition τ_{i^2, i^1} as:

$$\begin{aligned} \tau_{i^1, i^2} &= \tau_{j^1, j^2} \circ ... \circ \tau_{j^{T-2} j^{T-1}} \circ \tau_{j^{T-1} j^T} \circ \tau_{j^{T-2} j^{T-1}} \circ .... \tau_{j^1 j^2} \\ &\equiv \chi \end{aligned}$$

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

$$\begin{aligned} \tau_{j^{T-1} j^T} \circ \tau_{j^{T-2} j^{T-1}} \circ .... \tau_{j^1 j^2}(i^1) &= i^2, \\ \tau_{j^1, j^2} \circ ... \circ \tau_{j^{T-2} j^{T-1}}(i^2) &= i^2, \end{aligned}$$

so that $\chi(i^1) = i^2$. Similarly,

$$\begin{aligned} \tau_{j^{T-2} j^{T-1}} \circ .... \tau_{j^1 j^2}(i^2) &= i^2, \\ \tau_{j^1, j^2} \circ ... \circ \tau_{j^{T-2} j^{T-1}} \circ \tau_{j^{T-1} j^T}(i^2) &= i^1, \end{aligned}$$

so that $\chi(i^2) = i^1$. Finally, for any $j^t \neq i^1, i^2$ in the sequence,

$$\begin{aligned} \tau_{j^{T-2} j^{T-1}} \circ .... \tau_{j^1 j^2}(j^t) &= j^{t-1}, \\ \tau_{j^1, j^2} \circ ... \circ \tau_{j^{T-2} j^{T-1}} \circ \tau_{j^{T-1} j^T}(j^{t-1}) &= j^t, \end{aligned}$$

so that $\chi(j^t) = j^t$. We now construct a path from μ to $\tau_{j^1 j^2} \circ \mu$. First, consider giving the object of the retiring agent to the entering agent until you reach an assignment where agent n holds object j^1. Consider then the reassignment where j^1 is assigned to the agent holding j^2 and j^2 is assigned to the entering agent. Continue assigning the object of the retiring agent to the entering agent until agent n holds object n. In the new assignment, the agent who held object j^1 now holds object j^2 ; the agent who held object j^2 now holds object j^1, and all other agents i still hold the same object

$\mu(i) \neq j^1, j^2$. Hence, the new assignment is $\tau_{j^1 j^2} \circ \mu$. This construction can be repeated to generate a path from μ to $\tau_{i^2 i^1} \circ \mu$, and then to $\pi \circ \mu = \mu'$, concluding the proof of the Theorem.

Proof of Lemma 2: in the text.

Proof of Theorem 3: in the text.

Proof of Theorem 4: 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. Consider a state $s = (\mu, \theta)$ in S'. We 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, 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.