Available online at www.sciencedirect.com
SCIENCE DIRECT
Journal of Mathematical Economics 40 (2004) 953–962
JOURNAL OF Mathematical ECONOMICS
www.elsevier.com/locate/jmateco

Consistency in the probabilistic assignment model

Christopher P. Chambers*

Division of the Humanities and Social Sciences, 228-77, California Institute of Technology, Pasadena, CA 91125, USA

Received 15 May 2003; received in revised form 28 August 2003; accepted 7 October 2003
Available online 20 January 2004

Abstract

This paper introduces a notion of *consistency* for the probabilistic assignment model, which we call *probabilistic consistency*. We show that the axioms *equal treatment of equals* and *probabilistic consistency* characterize the uniform rule, which is the rule which randomizes uniformly over all possible assignments.

JEL classification: D63; D70; D71

Keywords: Assignment model; Consistency; Probabilistic assignment model; Uniform rule

1. Introduction

This paper is devoted to the study of the *consistency* concept (prominent in the theory of resource allocation) in the probabilistic assignment model. To motivate this study, we first discuss the deterministic assignment model. Think of allocating indivisible objects to agents who have strict preferences over these objects. Each agent can consume one and only one object. The idea is to discuss general methods (henceforth rules) of assigning objects to agents, as a function of agents' preferences. As each object may be given to at most one agent, there may be conflicts of interests. The literature on the assignment model is devoted to the study of well-behaved rules for resolving these conflicts.

In this deterministic setting, *consistency* is a stability concept which deals with variable populations of agents. Suppose a rule recommends an assignment for a particular group of agents, objects, and preferences (henceforth referred to as an "economy"). Imagine that a subgroup of the agents "reapplies" the rule to the economy consisting of themselves, the

* Tel.: +1-626-395-3559.
E-mail address: chambers@hss.caltech.edu (C.P. Chambers).

0304-4068/$ – see front matter
doi:10.1016/j.jmateco.2003.10.004

objects they were assigned, and their induced preferences over these objects. *Consistency* requires that the assignment recommended by the rule for this "reduced" economy assigns each agent in the subgroup to the object he was initially assigned. *Consistency* (and a large class of rules satisfying this condition) is studied by Ergin (2000) (see also Ehlers and Klaus, 2003).

The probabilistic assignment model is a simple generalization of this deterministic environment, motivated by fairness considerations. For example, a common requirement in models of fair allocation is that rules should not discriminate between agents with similar characteristics. This is the well-known property of *equal treatment of equals*. Clearly, as objects are indivisible, this property cannot generally be satisfied. By allowing rules to randomize over assignments; however, we can ensure that agents with similar characteristics are treated symmetrically (simply by flipping a fair coin) in an ex ante stage.

As in most of the literature, we assume that only ordinal preferences over objects are observable (Abdulkadiroğlu and Sönmez, 1998, 2003; Bogomolnaia and Moulin, 2001, 2002; Hylland and Zeckhauser, 1979; Zhou, 1990).[1] However, our main result holds more generally. Suppose all agents strictly rank all objects and rank lotteries in a way that is "monotonic" with respect to stochastic dominance (for example Bogomolnaia and Moulin, 2001). If we allow a rule to consider these risk preferences, the main result holds.

Our main contribution is to discuss a natural analogue of the *consistency* concept for the probabilistic environment and study its compatibility with basic fairness requirements. We introduce one such notion, *probabilistic consistency*. We suggest two requirements that a natural analogue of *consistency* must satisfy. The first is that when restricted to rules which select only deterministic outcomes, it should coincide with the standard *consistency* concept. The second requirement is that it should be intuitive and should be motivated by a similar story as the standard *consistency* concept.

To understand *probabilistic consistency*, imagine the following scenario. Suppose that a rule recommends some probabilistic assignment for a particular economy. Choose a particular agent. We may easily calculate the induced probabilities that this agent should receive each particular object. Suppose we randomize and determine which object the agent receives. Whichever object the agent receives, the question arises of how to allocate the remaining objects among the remaining agents. A natural proposal is to assign the remaining objects to the remaining agents according to the Bayesian update (when it exists) of the initial probabilistic assignment, conditional on the event that the agent had received that particular object. *Probabilistic consistency* is the requirement that this Bayesian update is indeed recommended for the reduced economy in which all of the remaining objects and agents are present, but the chosen agent has left with his object. Analogous concepts are discussed by Moulin (2002) and Moulin and Stong (2002) for the probabilistic discrete rationing model and by Moulin and Stong (2003) for the probabilistic "urn-filling" model.

We show that there is exactly one rule satisfying *equal treatment of equals* and *probabilistic consistency*. This is the rule which randomizes uniformly across *all* possible assignments. This rule is independent of preference–hence it violates such basic properties as *efficiency* (of any kind, be it ex ante, ordinal, or ex post). This is unfortunate, as it states that *proba-*

[1] In fact, this assumption is much more pervasive in the theory of fair allocation, dating back at least to Gibbard (1977, 1978), who employs a similar assumption in a probabilistic social choice model.

bilistic consistency is incompatible with the most basic notions of fairness and *efficiency*. However, this result should not greatly surprise us. Intuitively, *probabilistic consistency* is much stronger than its deterministic counterpart. This is because for any given n-agent economy, *consistency* in the deterministic model imposes restrictions on exactly n associated reduced economies. But *probabilistic consistency* imposes restrictions on at least n reduced economies, and can impose restrictions on as many as $n!$ (this occurs when all assignments occur with positive probability, which will be the case under the seemingly trivial *equal treatment of equals* property).

Finally, we mention the independent and closely related work of Haluk Ergin. In an unpublished master's thesis (Ergin, 1999, pp. 54–58), Ergin proposes the same notion of consistency for the probabilistic assignment model. His main result is that there is no *single-valued* rule that satisfies *ex post efficiency*, *anonymity*, and *probabilistic consistency*.[2] *Anonymity* implies *equal treatment of equals*; however, the proof of Ergin's result does not use the full force of *anonymity*. Our result is slightly more general than Ergin's as we do not require *ex post efficiency* in our characterization.

Section 2 introduces the formal model. Section 3 presents the main results. Finally, Section 4 concludes.

2. The model

2.1. Preliminaries

Let $\mathbb{N}$ be a set of potential agents such that $|\mathbb{N}| \geq 3$. The set of nonempty, finite subsets of $\mathbb{N}$ is denoted $\mathcal{N}$, with typical element N. Let $\mathbb{X}$ be a set of potential objects such that $|\mathbb{X}| \geq 3$. The set of nonempty finite subsets of $\mathbb{X}$ is denoted $\mathcal{X}$, with typical element X.

Let $X \in \mathcal{X}$. A **preference** R over X is a binary relation over X which is complete, transitive, and anti-symmetric. Thus, preferences are strict. Let the set of preferences over X be denoted $\mathcal{R}(X)$. An **economy** is a tuple consisting of a pair $(N, X) \in \mathcal{N} \times \mathcal{X}$ such that $|N| = |X|$, and a preference profile $R \in \mathcal{R}(X)^N$. The set of economies is denoted by $\mathcal{E}$.

An **assignment for** $(\mathbf{N}, \mathbf{X})$ is a bijection between N and X. A typical assignment will be written μ. The set of all assignments for (N, X) is denoted $\mathcal{A}(N, X)$. A **probabilistic assignment for** $(\mathbf{N}, \mathbf{X})$ is a probability distribution over $\mathcal{A}(N, X)$. The set of probabilistic assignments for (N, X) is denoted $\Delta(\mathcal{A}(N, X))$. Note that risk preferences of agents are not observed; so that all assignment decisions must be made using only ordinal information.

A **rule** r is a correspondence which associates with each economy (N, X, R) a nonempty subset of $\Delta(\mathcal{A}(N, X))$. We will be particularly interested in rules which enjoy certain properties.

Any assignment can be identified with a permutation matrix. Thus, by the theorem of Birkhoff and von Neumann, any probabilistic assignment induces a bistochastic matrix, and conversely (for example, see Birkhoff, 1946). Therefore, previous works assume a reduced model in which assignments are specified as the set of bistochastic matrices. For

[2] What we call *probabilistic consistency*, Ergin calls simply *consistency*. Furthermore, the condition of *anonymity*, which states that a rule is independent of agents' names, is referred to as ex ante *anonymity* in Ergin's work.

our purposes, we are not justified in working directly with the set of bistochastic matrices. The reason is that the mapping which takes the set of probabilistic assignments into the set of bistochastic matrices is not one-to-one. Thus, many probabilistic assignments may induce the same bistochastic matrix. This is without loss of generality when discussing concepts which relate only to the welfare levels of agents; but when discussing conditions which relate *directly to the structure* of assignments, it is not without loss of generality. We will present an example that shows that *probabilistic consistency* is such a condition.

2.2. Properties

Given $(N, X, R) \in \mathcal{E}$ and $\mu \in \mathcal{A}(N, X)$, say that $\boldsymbol{\mu}$ **is efficient for** $(\mathbf{N}, \mathbf{X}, \mathbf{R})$ if for all assignments $\mu' \in \mathcal{A}(N, X)$ such that $\mu' \neq \mu$, if there exists $i \in N$ such that $\mu'(i) P_i \mu(i)$, then there exists $j \in N$ such that $\mu(j) P_j \mu'(j)$. Given $p \in \Delta(\mathcal{A}(N, X))$, say **p is ex post efficient for** $(\mathbf{N}, \mathbf{X}, \mathbf{R})$ if its support consists only of assignments efficient for (N, X, R). *Ex post efficiency* is the weakest possible form of *efficiency* in environments with objective uncertainty. It amounts to saying that the outcome of a randomization is always *efficient*. However, if all agents behave "monotonically" with respect to risk, then generally we can do better than *ex post efficiency*, see Abdulkadiroğlu and Sönmez (2003), Bogomolnaia and Moulin (2001, 2002), and McLennan (2002).

Ex post efficiency

For all $(N, X, R) \in \mathcal{E}$, for all $p \in r(N, X, R)$, p is ex post efficient for (N, X, R).

The introduction of lotteries into the assignment model is based on "fairness" considerations. Many common notions of fairness are not compatible with environments in which goods come in discrete units. One such notion is that two agents with the same characteristics should be treated equally. *Equal treatment of equals* is the requirement that in a given economy, two agents with identical observable characteristics should be treated equally in terms of their induced lottery over objects. Normally, this requirement states that agents should be treated equally in terms of welfare. However, since agents' preferences over lotteries are unobservable, a natural way of capturing the idea that they are treated equally is to require that they have the same induced probabilistic consumption.[3]

Equal treatment of equals

For all $(N, X, R) \in \mathcal{E}$ for which there exists $i, j \in N$ such that $R_i = R_j$, for all $p \in r(N, X, R)$, $p_i = p_j$, where p_i and p_j are the induced lotteries over the objects consumed by i and j, respectively.

We informally discuss our consistency notion. Let r be a rule. Let (N, X, R) be an economy. Suppose that $p \in r(N, X, R)$. We want to imagine a scenario in which some agent receives an object and then "leaves." We can imagine the following "two-stage" procedure. For any agent, a probabilistic assignment induces a unique lottery over objects

[3] When preferences over lotteries are unobservable, *equal treatment of equals* clearly has more bite than if they are observable. This is because two agents may have different preferences over lotteries, but have the same preferences over degenerate objects, so that two agents with different preferences over lotteries are treated equally. Moreover, if agents' preferences over lotteries *are* observable, then *equal treatment of equals* is stated in terms of welfare. Our results remain unchanged when preferences over lotteries are observable.

this agent may receive. For example, the probability that agent i receives object x is $p(\{\mu : \mu(i) = x\})$. Suppose we determine the object that i receives according to these probabilities in the "first-stage." Suppose that the result of this randomization is that agent i receives object x. We then obtain a new economy, in which agent i and object x are no longer present. In the "second stage," the probabilistic assignment is determined according to the rule's recommendation for the "remaining" agents and objects, where agents' preferences are now their induced preferences over the remaining objects. This two-stage lottery over assignments is naturally identified with a one-stage lottery over assignments; that is, an element $p' \in \Delta(\mathcal{A}(N, X))$. We will say a rule is *probabilistically consistent* if $p' = p$. Alternatively, *probabilistic consistency* requires that for all $i \in N$ and $x \in X$, the Bayesian update of a recommended probabilistic assignment conditional on the event $\{\mu : \mu(i) = x\}$ when it exists is exactly the recommended probabilistic assignment when the rule is applied to the economy $(N \setminus \{i\}, X \setminus \{x\}, R|_{N\setminus\{i\}, X\setminus\{x\}})$.[4]

As each $\mu \in \mathcal{A}(N, X)$ can be viewed as a function $\mu : N \to X$, for all $N' \subset N$, the notation $\mu|_{N'}$ refers to the restriction of μ to the domain N'. Thus, $\mu|_{N'}$ is an element of $\mathcal{A}(N', X \setminus \mu(N \setminus N'))$. The formal consistency notion follows.

Probabilistic consistency

For all $(N, X, R) \in \mathcal{E}$, for all $p \in r(N, X, R)$, for all $(i, x) \in N \times X$ such that $p(\{\mu : \mu(i) = x\}) > 0$, $p^* \in r(N \setminus \{i\}, X \setminus \{x\}, R|_{N\setminus\{i\}, X\setminus\{x\}})$, where $p^* \in \Delta(\mathcal{A}(N \setminus \{i\}, X \setminus \{x\}))$ satisfies for all $\mu' \in \mathcal{A}(N, X)$ such that $\mu'(i) = x$,

$$p^*(\mu'|_{N\setminus\{i\}}) = \frac{p(\mu')}{p(\{\mu : \mu(i) = x\})}.$$

Consistency concepts in the assignment model carry with them an implicit assumption of "*independence of irrelevant alternatives*." As agents are only presumed to have preferences over present objects, when discussing reduced economies, all information about preferences over infeasible objects is discarded. In a generalized model in which preferences over all objects (feasible and infeasible) are observable, *consistency* has much less bite. Thus, the intuitive principle embodied in *consistency* has stronger implications when working within the confines of a specific model. For more on this point, see Thomson (2003), Section 2.

Clearly, a deterministic rule which is *consistent* in a deterministic environment is *probabilistically consistent* when viewed as a probabilistic rule.

3. Results

3.1. On bistochastic matrices

We demonstrate why it is not without loss of generality to identify two probabilistic assignments which induce the same bistochastic matrix. Consider a three-agent environment;

[4] Implicit in this motivation is an informal "temporal" element. For such a story to make sense, we implicitly assume that all agents identify two-stage lotteries with their reduced one-stage counterpart. For example, see Segal (1990).

$N \equiv \{1, 2, 3\}$, $X \equiv \{1, 2, 3\}$. An element $\mu \in \mathcal{A}(N, X)$ can be represented as a permutation matrix P, in which

$$P_{ij} = \begin{cases} 0 & \text{if } \mu(i) \neq j \\ 1 & \text{if } \mu(i) = j \end{cases}.$$

The theorem of Birkhoff and von Neumann states that the set of bistochastic matrices is the convex hull of the set of permutation matrices. Hence, any element of $\Delta(\mathcal{A}(N, X))$ can be represented as a bistochastic matrix P, in which P_{ij} represents the induced probability that agent i receives object j. By the theorem of Birkhoff and von Neumann, the converse is also true; any bistochastic matrix can be identified with an element of $\Delta(\mathcal{A}(N, X))$. However, there may be more than one element of $\Delta(\mathcal{A}(N, X))$ identified with the same bistochastic matrix. For example, a uniform randomization over the three permutation matrices

$$\left\{ \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} \right\}$$

is identified with the bistochastic matrix

$$\begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{pmatrix}.$$

However, a uniform randomization over the following three permutation matrices

$$\left\{ \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix} \right\}$$

is identified with the same bistochastic matrix. By imagining the economy in which agent 3 has left with object 3, the updated assignment for the first probabilistic assignment is

$$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},$$

and the updated assignment for the second probabilistic assignment is

$$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

Thus, depending on *which probabilistic assignment* we take as primitive, the updated assignment may differ.

3.2. The uniform priority rule

To illustrate the intuition behind our main result, we first provide a simple example. Let $(N, X, R) \in \mathcal{E}$. Let Σ be the set of linear orders of N (called "priorities"). A typical element of Σ will be denoted σ. Given $\sigma \in \Sigma$, we define the **σ-priority assignment** as follows. The agent ranked highest according to σ chooses her most preferred object among X, and then leaves. The next highest ranked agent according to σ chooses her most preferred object from the remaining objects, and then leaves. This process continues so that the highest ranked remaining agent chooses her most preferred object from the remaining objects. The assignment which results is the σ-priority assignment. Clearly, the σ-priority assignments are all *efficient*, but highly "unfair." To construct a rule satisfying *equal treatment of equals*, for any economy, we simply randomize uniformly over all priority assignments—the resulting probabilistic assignment is the assignment recommended by the **uniform priority rule**.[5] The uniform priority rule is *ex post efficient* and satisfies *equal treatment of equals*.

The uniform priority rule is not *probabilistically consistent*. Let $(N, X, R) \in \mathcal{E}$ so that $N = \{i, j, k\}$, $X = \{x, y, z\}$, $xR_i yR_i z$, $xR_j zR_j y$, and $xR_k zR_k y$.[6] In this economy, all agents prefer x to the other objects, but two agents prefer z to y whereas only one prefers y to z. Thus, z has a "higher demand" than y, where x has the highest demand. This is reflected in the probabilistic assignment recommended by the uniform priority rule. This probabilistic assignment features agent j receiving z half of the time, and agent k receiving z the other half of the time. Agent k receives z under the priorities which let him choose second, and under the priorities which let agent j choose first (exactly three priorities in total). Whenever agent j chooses first (under two priorities), she chooses x. Thus, conditional on agent k receiving z, agent j receives x exactly two-thirds of the time. In a sense, agent j is "compensated" for the fact that she had a higher demand for z than did agent i, but did not receive z.

The economy $(\{i, j\}, \{x, y\}, R|_{\{i,j\},\{x,y\}})$ is the economy induced by (N, X, R) after agent k receives object z. If the uniform priority rule were *probabilistically consistent*, agent j would receive x exactly two-thirds of the time, and agent i would receive x exactly one-third of the time. However, the uniform priority rule clearly recommends that agent j receive x half of the time, and agent i receive x the other half of the time. This is because the uniform priority rule disregards all information about the relative preference of z in the original economy. In fact, *all* rules suffer from this drawback. This is not so much an implication of the intuitive principle which *probabilistic consistency* attempts to capture as it is an artifact of the implicit "*independence of irrelevant alternatives*" idea forced upon us by the underlying structure of the model.

3.3. The main result

We show that, unfortunately, the axioms *ex post efficiency*, *equal treatment of equals*, and *probabilistic consistency* are incompatible. This result is somewhat surprising as in many other models without discrete goods, *efficiency*, *equal treatment of equals*, and *consistency* are compatible.

To show this, we actually prove a stronger statement. For a given economy, define the **uniform assignment** as the probabilistic assignment which places equal probability on all

---

[5] This rule is usually called the "random priority rule;" however, we feel that the term "random" is vague, and does not indicate that the orderings are equiprobable.

[6] This example was suggested by an anonymous referee.

deterministic assignments. Define the **uniform rule** as that rule which, for all economies, recommends the uniform assignment. The uniform rule clearly violates *ex post efficiency*, as it is preference independent. However, it is the only *single-valued* rule which satisfies both *equal treatment of equals* and *probabilistic consistency*. This proposition holds when the set of potential agents is as low as three.

Proposition. *A single-valued rule satisfies equal treatment of equals and probabilistic consistency if and only if it is the uniform rule*.

Proof. It is straightforward to establish that the uniform rule satisfies the two axioms. The converse statement is divided into two steps.

Step 1 (For all two-agent economies, the rule coincides with the uniform rule). Let $(N, X, R) \in \mathcal{E}$ such that $|N| = 2$. Without loss of generality, label $N \equiv \{i, j\}$ and $X \equiv \{x, y\}$. Suppose that $xR_i y$ and $xR_j y$. Let $p' \in r(N, X, R)$. By *equal treatment of equals*, for all $\mu \in \mathcal{A}(N, X)$, $p'(\mu) = 1/2$. Thus, for all two-agent economies, if agents' preferences coincide, the probabilistic assignment recommended is the uniform assignment.

We now establish the same result for two-agent economies for which preferences differ.

Let $(N, X, R) \in \mathcal{E}$ such that $|N| = 2$. Without loss of generality, label $N \equiv \{i, j\}$ and $X \equiv \{x, y\}$. Suppose that $xR_i y$ and $yR_j x$. We claim that if $p' \in r(N, X, R)$, then for all $\mu \in \mathcal{A}(N, X)$, $p'(\mu) = 1/2$. Let $(N', X', R') \in \mathcal{E}$ satisfy $N' \equiv \{i, j, k\}$, $X' \equiv \{x, y, z\}$, and $xR'_i yR'_i z$, $yR'_j xR'_j z$, and $xR'_k yR'_k z$. We write (a, b, c) to denote the assignment μ such that $\mu(i) = a$, $\mu(j) = b$, and $\mu(k) = c$. Let $p \in r(N', X', R')$. We claim that p puts equal weight on all elements of $\mathcal{A}(N', X')$.

Suppose that two agents' preferences coincide over a pair of objects in X'. Then the two elements of $\mathcal{A}(N', X')$ in which the remaining agent consumes the remaining object are equiprobable under p. We verify one case; the remaining cases are similar. Thus, $R'_j|_{\{y,z\}} = R'_k|_{\{y,z\}}$. Suppose that $p'' \in r(\{j, k\}, \{y, z\}, (R'_j|_{\{y,z\}}, R'_k|_{\{y,z\}}))$. By the previous paragraph, $p''(\{(x, y, z)|_{\{j,k\}}\}) = p''(\{(x, z, y)|_{\{j,k\}}\}) = 1/2$. Suppose that $p(\{(x, y, z), (x, z, y)\}) > 0$. By *probabilistic consistency*,

$$\frac{p(\{(x, y, z)\})}{p(\{(x, y, z), (x, z, y)\})} = p''(\{(x, y, z)|_{\{j,k\}}\}) = p''(\{(x, z, y)|_{\{j,k\}}\})$$

$$= \frac{p(\{(x, z, y)\})}{p(\{(x, y, z), (x, z, y)\})}.$$

Thus, $p(\{(x, y, z)\}) = p(\{(x, z, y)\})$. If $p(\{(x, y, z), (x, z, y)\}) = 0$, then by the definition of probability, $p(\{(x, y, z)\}) = p(\{(x, z, y)\}) = 0$.

So, $p(\{(x, z, y)\}) = p(\{(x, y, z)\})$. Continuing in a parallel fashion, as $R'_i|_{\{x,z\}} = R'_k|_{\{x,z\}}$, $p(\{(x, y, z)\}) = p(\{(z, y, x)\})$. As $R'_i|_{\{y,z\}} = R'_j|_{\{y,z\}}$, $p(\{(z, y, x)\}) = p(\{(y, z, x)\})$. As $R'_j|_{\{x,z\}} = R'_k|_{\{x,z\}}$, $p(\{(y, z, x)\}) = p(\{(y, x, z)\})$. As $R'_i|_{\{y,z\}} = R'_k|_{\{y,z\}}$, $p(\{(y, x, z)\}) = p(\{(z, x, y)\})$. Thus, p randomizes over all elements of $\mathcal{A}(N', X')$ with equal probabilities. In particular, $p(\{(x, y, z)\}) = p(\{(y, x, z)\})$ and $p(\{(x, y, z), (y, x, z)\}) > 0$.

As $R_i = R'_i|_{\{x,y\}}$ and $R_j = R'_j|_{\{x,y\}}$, by *probabilistic consistency*,

$$p'(\{(x, y, z)|_{\{i,j\}}\}) = \frac{p(\{(x, y, z)\})}{p(\{(x, y, z)(y, x, z)\})} = \frac{1}{2}.$$

Thus, for all $\mu \in \mathcal{A}(N, X)$, $p'(\mu) = 1/2$.

Step 2 (For all economies, the rule coincides with the uniform rule). The proof of Step 2 is by induction on the cardinality of the set of agents.

Let $(N', X', R') \in \mathcal{E}$. If $|N'| = 1$, for all $p' \in r(N', X', R')$, p' is the uniform assignment.

Let $(N', X', R') \in \mathcal{E}$. If $|N'| = 2$, then by Step 1, for all $p' \in r(N', X', R')$, p' is the uniform assignment.

Let $K \in \mathbb{N}$ such that $K \geq 3$. Suppose that for all $(N, X, R) \in \mathcal{E}$ such that $|N| < K$ and for all $p \in r(N, X, R)$, p is the uniform assignment. Let $(N', X', R') \in \mathcal{E}$ such that $|N'| = K$. Let $p' \in r(N', X', R')$. We claim that for all $\mu', \mu'' \in \mathcal{A}(N', X')$, $p'(\mu') = p'(\mu'')$.

For all $\boldsymbol{\mu}', \boldsymbol{\mu}'' \in \mathcal{A}(N', X')$, say that **$\boldsymbol{\mu}'$ and $\boldsymbol{\mu}''$ are adjacent** if there exists $i \in N'$ such that $\mu'(i) = \mu''(i)$. We claim that if μ' and μ'' are adjacent, then $p'(\mu') = p'(\mu'')$. Let $i \in N'$ and $x \in X'$ such that $\mu'(i) = x$ and $\mu''(i) = x$. Suppose that $p'(\{\mu : \mu(i) = x\}) > 0$. Let $p^* \in r(N' \setminus \{i\}, X' \setminus \{x\}, R'|_{N\setminus'\{i\}, X'\setminus\{x\}})$. By the induction hypothesis, $p^*(\mu'|_{N'\setminus\{i\}}) = p^*(\mu''|_{N'\setminus\{i\}})$. By *probabilistic consistency*, $p'(\mu') = p^*(\mu'|_{N'\setminus\{i\}})p'(\{\mu : \mu(i) = x\}) = p^*(\mu''|_{N'\setminus\{i\}})p'(\{\mu : \mu(i) = x\}) = p'(\mu'')$. Thus, $p'(\mu') = p'(\mu'')$. Suppose that $p'(\{\mu : \mu(i) = x\}) = 0$. By the definition of probability, $p'(\mu') = 0 = p'(\mu'')$.

Let $\mu', \mu'' \in \mathcal{A}(N', X')$ be arbitrary. Let $i \in N'$. Let $j \in N'$ such that $\mu'(j) = \mu''(i)$. Let $\mu''' \in \mathcal{A}(N', X')$ be such that $\mu'''(i) = \mu''(i)$, $\mu'''(j) = \mu'(i)$, and for all $k \in N' \setminus \{i, j\}$, $\mu'''(k) = \mu'(k)$. Thus, μ''' is the assignment in which, starting from μ', agent i "trades" his object with the agent who possesses the object that i receives under μ''. Since $|N'| \geq 3$, there exists k such that $\mu'(k) = \mu'''(k)$; hence μ' is adjacent to μ'''. Moreover, μ''' is adjacent to μ'' as $\mu'''(i) = \mu''(i)$. Thus, by the preceding paragraph, $p'(\mu') = p'(\mu''') = p'(\mu'')$, so that all assignments are equiprobable under p'. □

Probabilistic consistency is intuitively a much "stronger" requirement than *consistency* in the deterministic case. For example, in a three person economy, *probabilistic consistency* has implications for at least three other economies, and can have implications for up to six. The deterministic version has implications for exactly three (in the single-valued case).

A corollary of the main result is that there is no *single-valued* rule which satisfies *ex post efficiency*, *equal treatment of equals*, and *probabilistic consistency*.

4. Conclusion

This paper provides a notion of *consistency* in models where randomization of assignments is permitted. The main reason for introducing lotteries into the assignment model is to accommodate the inherent indivisibilities associated with it. However, *probabilistic consistency* is not limited to the assignment model. Any model which currently has a *consistency* concept admits a *probabilistic consistency* concept when the model is extended to incorporate lotteries, as has been done by the previously cited papers by Moulin, and Moulin and Stong.

Although we have shown that *probabilistic consistency* is incompatible with the most basic fairness concepts, work still remains to be done with the concept. For example, in the deterministic assignment model, Ergin (2000) characterizes the class of rules satisfying *neutrality* and *consistency*. An analogous result for the probabilistic model is desirable. It is not clear how to generalize Ergin's family of rules to the probabilistic environment; in fact, the uniform rule turns out to be both *neutral* and *probabilistically consistent* when *neutrality* is defined appropriately for the probabilistic assignment model.

Acknowledgements

I would like to thank Onur Kesten and William Thomson for comments and suggestions. I would especially like to thank Haluk Ergin for pointing out his closely related work. I am also very grateful to an anonymous referee, whose comments proved to be very helpful. All errors are my own.

References

Abdulkadiroğlu, A., Sönmez, T., 1998. Random serial dictatorship and the core from random endowments in house assignment problems. Econometrica 66, 689–701.

Abdulkadiroğlu, A., Sönmez, T., 2003. Ordinal efficiency and dominated sets of assignments. Journal of Economic Theory 112, 152–172.

Birkhoff, G., 1946. Three observations on linear algebra. Universidad Nacional de Tucuman Revista, Serie A 5, 147–151.

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

Bogomolnaia, A., Moulin, H., 2002. A simple random assignment problem with a unique solution. Economic Theory 19, 623–635.

Ehlers, E., Klaus, B., 2003. Consistent house allocation, manuscript in preparation.

Ergin, H., 1999. Consistency in House Allocation Problems. Master's Thesis. Bilkent University.

Ergin, H., 2000. Consistency in house allocation problems. Journal of Mathematical Economics 34, 77–97.

Gibbard, A., 1977. Manipulation of voting schemes that mix voting with chance. Econometrica 45, 665–681.

Gibbard, A., 1978. Straightforwardness of game forms with lotteries as outcomes. Econometrica 46, 595–614.

Hylland, A., Zeckhauser, R., 1979. The efficient allocation of individuals to positions. Journal of Political Economy 87, 293–314.

McLennan, A., 2002. Ordinal efficiency and the polyhedral separating hyperplane theorem. Journal of Economic Theory 105, 435–449.

Moulin, H., 2002. The proportional random allocation of indivisible units. Social Choice and Welfare 19, 381–413.

Moulin, H., Stong, R., 2002. Fair queuing and other probabilistic allocation methods. Mathematics of Operations Research 27, 1–30.

Moulin, H., Stong, R., 2003. Filling a multicolor urn: an axiomatic analysis. Games and Economic Behavior 45, 242–269.

Segal, U., 1990. Two-stage lotteries without the reduction axiom. Econometrica 58, 349–377.

Thomson, W., 2003. Consistent Allocation Rules. Manuscript. University of Rochester.

Zhou, L., 1990. On a conjecture by Gale about one-sided matching problems. Journal of Economic Theory 52, 123–135.