An Equivalence result in School Choice

Jay Sethuraman*

May 2009

Abstract

The main result of the paper is a proof of the equivalence of single and multiple lottery mechanisms for the problem of allocating students to schools in which students have strict preferences and the schools are indifferent. This solves a recent open problem proposed by Pathak, who was motivated by the practical problem of assigning students to high schools in New York City. In proving this result, a new approach is introduced, that simplifies and unifies all the known equivalence results in the house allocation literature. Along the way, two new mechanisms—Partitioned Random Priority and Partitioned Random Endowment—are introduced for the house allocation problem. These mechanisms generalize several known and well-studied mechanisms for the house allocation problem and are particularly appropriate for many-to-one versions of the problem, the school choice problem being the most prominent.

*IEOR Department, Columbia University, New York, NY; jay@ieor.columbia.edu.

1 Introduction

This paper is motivated by recent work of Pathak [6] on competing mechanisms for the problem of assigning students to schools in New York City. The public school system in NYC has been using a new mechanism for assigning students to high schools [1]. This mechanism consists of a main round, followed by a supplementary round that is meant to assign the applicants who are unassigned after the main round. The assignment process in the main round takes into account several factors including the student priorities at each school based on standardized test scores, and is therefore a two-sided matching problem. In contrast, the assignment process in the supplementary round is quite simple: all unassigned applicants are invited to rank order high schools with vacant capacity; all students at this stage have the *same* priority to attend any school. Thus we are led naturally to a setting in which agents rank heterogeneous "objects," for which they have equal claims. This raises the following question: if many students rank the same school as their top choice, which of these students are assigned to that school? To achieve fairness in the assignment process, it thus becomes necessary to use a lottery mechanism[1]. Motivated by these considerations, Pathak [6] recently studied two "lottery" mechanisms: first is the *single* lottery mechanism, in which a single random ordering of the agents is drawn; any ties (at *any* school) are broken in favor of the student whose lottery number is lower. A natural alternative—which *seems* fairer to the students at first glance—is the *multiple* lottery mechanism, which allows each school to conduct its own lottery. The actual assignment is made by the core mechanism applied to the preference profiles of the students and the priority profiles of the schools [3]. Pathak [6, pp. 3] notes that during the course of the design of the new assignment mechanism, policymakers from the Department of Education believed that the single lottery mechanism is less equitable than the multiple lottery mechanism. Remarkably, Pathak shows that, for the special case of the problem in which each school has exactly one vacant spot, the distribution of assignments is exactly the same under both mechanisms! To prove this result, Pathak uses the following approach: he associates with any given ordering of the students (in the single lottery mechanism) a *set* of $(n!)^{n-1}$ priority profiles of schools (in the multiple lottery mechanism) such that the resulting final assignment is identical, and such that the sets of priority profiles associated with distinct student orderings are disjoint.

[1]We assume, as is standard in this literature, that money cannot be used.

The construction is insightful, requires a great deal of ingenuity and care to ensure that priority profiles associated with distinct orderings do not overlap; this last requirement distinguishes Pathak's proof from earlier equivalence proofs in the literature (more details in §2).

Our main result, in §3, is a proof of the equivalence of the single and multiple lottery mechanisms in full generality. In particular, we do away with the restriction that each school can admit only one student. To prove this result, we retain the restriction that "schools" can admit only one student, but instead introduce a "new" mechanism that we call *Partitioned Random Priority* (PRP). Under the PRP mechanism, we are given an arbitrary partition $S_1, S_2, \ldots, S_k$ of the "schools;" the schools within each S_i use a *common* lottery, and distinct S_i's use an independent lottery[2]. We show that the distribution of assignments under the PRP mechanism is the same, *regardless of the partition of the schools.* The analog of Pathak's result when schools have multiple seats follows: if there are k schools and school i can admit q_i students, then make q_i copies of school i, and let S_i consists of these q_i copies.

We prove a similar equivalence result for another new mechanism that we introduce called *Partitioned Random Endowment* (PRE), which can be thought of as a generalization of both the single lottery mechanism and the Random Endowment (RE) mechanism. In the RE mechanism, each agent is endowed with a distinct house, with each of the $n!$ endowment vectors equally likely; the resulting allocation is computed by find the core of the associated trading game, first analyzed by Shapley and Scarf [7]. In PRE, we are given a partition of $S_1, S_2, \ldots, S_k$ of the schools; the students are partitioned into sets $N_1, N_2, \ldots, N_k$ such that N_i has exactly $|S_i|$ students, and every school in S_i has a common ranking of the students in N_i. Thus, if each school is in a partition by itself, we recover the RE mechanism; if all the schools form a single partition, we recover the single lottery mechanism. Our equivalence result thus subsumes the earlier equivalence result, due independently to Knuth [5] and Abdulkadiroglu & Sönmez [2], of the RE and single lottery mechanisms.

Finally, in contrast to the proofs of earlier equivalence results, our proof technique is elementary, inductive, and relies on basic properties of the cycle structure of permutations. We illustrate the usefulness of this approach by providing, in §4, a simple alternative proof

[2] If each school is in a partition by itself we recover the multiple lottery mechanism; if all the schools belong to a single partition, we recover the single lottery mechanism.

of another equivalence result in the context of house allocation with existing tenants.

2 The Model

Let N denote the set of students and S the set of schools, and suppose $|N| = |S| = n$. Each student has a strict preference ordering of the schools, and wishes to attend exactly one school. Each school wishes to admit exactly one student, but does not care which one it admits. Because the school preferences exhibit extreme indifference, we can view this problem as a *house allocation* problem in which the agents are students, and the houses are schools—this analogy is possible only because the schools are treated as objects that do not care who they are allocated to. Not surprisingly, the mechanisms proposed for the school choice problem have been inspired by this literature [3, 1, 4, 6].

2.1 Existing Mechanisms

An important ingredient in house allocation mechanisms is the precise manner in which *conflicts* between agents are resolved: if many agents rank the same house as their most-preferred house, which one of them gets it? If houses had strict preferences (as would be the case in a *two-sided* matching problem), such conflicts can be resolved based on the preference ordering of the house in question. As houses do *not* have strict preferences, a natural idea is to *generate* a *priority list* of agents for each house that will be used to resolve conflicts. The various mechanisms proposed for the house allocation problem differ only in how these priority orderings are generated (equivalently, how such conflicts are resolved). We review the prominent mechanisms before describing the two new mechanisms we propose.

Random Endowment (RE). Each agent is given top priority at a distinct house, with each of the $n!$ possibilities being equally likely. For a given priority list, the resulting allocation is the (unique) core of the associated trading game in which the top-priority agent at any house is said to "own" that house. This allocation can be found using the TTC algorithm: we first construct a graph with one node for each agent, and there is an arc from agent i to agent j if j owns i's most-preferred house (note that j could be i, in which case there will be a self-loop). This graph must have a cycle as every node has out-degree 1 and the graph is finite; each agent in the cycle is assigned to the house owned by the agent he points to, all these nodes are deleted, and the computation is repeated on the residual problem.

Random Priority (RP). The agents are ordered randomly, with each of the $n!$ orderings being equally likely. Every house uses this ordering as its priority list. For any given priority list, the resulting allocation is found as before by the TTC algorithm, with the modification that each house is owned by the agent who appears first in its priority list. Note that, with this new definition, an agent may own multiple houses at a time; moreover, a house owned by some agent may be owned by a different agent later on. Neither one of these aspects is present in the RE mechanism. Finally, observe that this mechanism is equivalent to the more familiar description of the RP mechanism: the agents are ordered uniformly at random, make their choices according to this ordering, and pick their most-preferred house that is still available when it is their turn.

Independent Random Priority (IRP). *Each* house draws a random ordering of the agents, independently of the other houses, with each of the $n!$ orderings being equally likely. Thus the priority profile of the houses is equally likely to be any one of the $(n!)^n$ possibilities. For any given priority profile, the resulting allocation is found as before by the TTC algorithm.

The three mechanisms introduced so far—RE, RP and IRP—all have very different descriptions. It is therefore remarkable that all of these mechanisms are equivalent! The RP and RE assignments are obtained by running the TTC algorithm on $n!$ profiles. That they find the same assignment was shown independently by Abdulkadiroglu & Sönmez [2] and Knuth [5]—remarkably, both papers construct a bijective proof, i.e., they associate any given ordering of the agents (RP) with a distinct endowment (RE) such that the resulting final assignment is identical. In contrast, the IRP assignment is found by running the TTC algorithm on $(n!)^n$ profiles. Pathak [6] proves the equivalence of RP and IRP using the following approach: he associates with any given ordering of the students a *set* of $(n!)^{n-1}$ priority profiles of schools such that the resulting final assignment is identical, and such that the sets of priority profiles associated with distinct student orderings are disjoint. The construction is insightful, and requires a great deal of ingenuity. His result, coupled with the earlier equivalence result of Abdulkadiroglu and Sönmez [2], establishes the equivalence of RE and IRP mechanisms.

2.2 New Mechanisms

The mechanisms discussed so far were all on one-to-one allocation problems: each agent wishes to be assigned a house and each house can be assigned to exactly one agent. The assignment problems arising in school choice, however, are typically many-to-one: a school can admit many students, whereas each student still wishes to be assigned to one school. We can, of course, formulate an equivalent one-to-one version of the problem: if a school can admit 13 students, we can make 13 copies of that school, and have each student rank all the copies of this school in an arbitrary (but fixed) way. Note also that there is no difficulty in extending the RP mechanism. But the IRP and the RE mechanisms are somewhat unnatural in the equivalent one-to-one formulation. Ideally, we would like the IRP mechanism to allow the priority list of each *school* to be generated independently; in the equivalent one-to-one formulation, IRP allows each "seat" to generate its own priority list, so that different copies of the same school may have different priority lists. A similar problem plagues the RE mechanism. To remedy this difficulty, we introduce the following two new mechanisms for the school choice problem. (To emphasize the application, we switch to students and schools instead of agents and houses.)

Partitioned Random Priority (PRP). Fix an arbitrary partition of the schools S into sets $S_1, S_2, \ldots, S_k$. School priority lists are determined in the following way: For each i, the schools in S_i have identical priority lists, which are generated independently and uniformly. Thus, each of the $(n!)^k$ priority profiles is equally likely. The final assignment matrix is found by running the TTC algorithm on each of the $(n!)^k$ possible priority profiles, and averaging the assignments obtained. Note that when each school is in a partition by itself, we recover the IRP mechanism; and when all the schools are in one partition, we recover the RP mechanism.

Partitioned Random Endowment (PRE). Fix an arbitrary partition of the schools S into sets $S_1, S_2, \ldots, S_k$. School priority lists are determined in the following way: the students are arranged uniformly at random, all $n!$ orderings being equally likely. From this random ordering, the first $|S_1|$ students form the top $|S_1|$ choices of each of the schools in S_1; the next $|S_2|$ students form the top $|S_2|$ choices of each of the schools in S_2, etc. The final assignment matrix is found by running the TTC algorithm on each of the $n!$ possible priority profiles, and averaging the assignments obtained. Note that when each school is in a partition by itself, we recover the RE mechanism; and when all the schools are in one partition, we recover the RP mechanism.

Our main results (in §3) are that, for both mechanisms, the assignment matrix for *any partition* of the schools is identical to the assignment matrix for the RP mechanism. Thus our main results generalize the earlier results of Pathak [6], Abdulkadiroglu & Sönmez [2], and Knuth [5]. In contrast to the earlier equivalence proofs, our proof is not bijective in nature. Instead the ingredients for the equivalence proofs presented here include standard results about the cycle structure of permutations and mathematical induction. The proofs are simpler, and the same framework is used in §4 to prove an equivalence result for the house allocation problem with existing tenants.

3 The Equivalence Results

Let $S_1, S_2, \ldots, S_k$ be a partition of the schools. What does it mean to generate a random preference profile with respect to this partition? We start with the PRP mechanism. Imagine n cards—one for each student— arranged in a sequence next to each school. The cards are all face-down, and the sequences are identical for all the schools in each S_i. Consider now the PRE mechanism. As in the PRP mechanism, all the schools within a given S_i have an identical priority list. The difference is that, in the PRE mechanism, the schools in S_i have an identical ranking of $|S_i|$ students, and these students are not in the "priority" list of any of the other schools. That is, the students are partitioned so that exactly $|S_i|$ students appear in the priority list of the schools in S_i, and they are ranked identically by the schools in S_i, with each of the $|S_i|!$ rankings being equally likely.

When can the TTC algorithm find a given cycle $C = (i_1, i_2, \dots, i_m)$? The TTC algorithm can find C only if, for every $1 \leq \ell \leq m$, student $i_{\ell+1}$ has the top priority at student i_ℓ's most-preferred school (where i_{m+1} is to be read as i_1). In particular, the m students in C should all have distinct most-preferred schools. Suppose the TTC algorithm finds cycle C in the given priority profile. We can delete the students in C along with their most-preferred schools; we can also remove the cards bearing the name of any student in C from the remaining $(n - m)$ schools, as these students have no further role to play in the TTC algorithm. We call the instance thus obtained as the *residual instance* after eliminating cycle C. Let $S'_1, S'_2, \dots, S'_k$ be the partition of the remaining $(n-m)$ schools where S'_i contains those schools in S_i that are still in the residual instance. (Note that each S_i can lose at most one school. Also, if some S'_j is empty, it can be deleted, so the number of non-empty classes in the partition for the residual instance may be strictly smaller than k.) The critical observation is that—for both PRP and PRE mechanisms—this residual instance is *equally likely to be any of the possible profiles* with the partition $S'_1, S'_2, \dots, S'_{k'}$. This follows from the observation that eliminating cycle C does not convey any information about the priority lists of any of the remaining schools with respect to the remaining students.

It is clear that "cycles" play an important role in both mechanisms. Observe that any set of students whose most-preferred schools are all distinct can be in a first cycle found by the mechanism, and vice-versa. Moreover, the cycles that *could* be found as a first cycle are the same for both PRP and PRE, so the discussion that follows applies to both mechanisms.

To make these notions precise, we introduce some notation. Recall that $S_1, S_2, \dots, S_k$ is the given partition of the schools. For $i = 1, 2, \dots, k$, let N_i be the set of students whose most preferred school is in S_i. From the earlier discussion, any cycle that can be found as a first cycle can have at most one student from each N_i. Let $\mathcal{C}$ be the collection of all such cycles. We use C to denote a generic member of $\mathcal{C}$, and let $\overline{C}$ be the *set* of students involved in the cycle C. Note that it is possible for $C, C' \in \mathcal{C}$, $C \neq C'$, and $\overline{C} = \overline{C'}$. Given two cycles $C, C' \in \mathcal{C}$, we say that C and C' *conflict* if for some i, $N_i \cap \overline{C}$ and $N_i \cap \overline{C'}$ are both nonempty. For any $\ell \geq 1$, $\mathcal{C}^\ell$ consists of products of ℓ cycles of $\mathcal{C}$ such that no pair of cycles conflict. Note that $\mathcal{C}^1$ is just $\mathcal{C}$, and $\mathcal{C}^\ell$ consists of products of ℓ distinct cycles of C, all of which can be simultaneously present in some priority profile (and each of which can potentially be uncovered as a first cycle). We illustrate these definitions on the following example.

Example 1. Consider the following instance with 5 students and 5 schools, each with one seat. The strict student preferences are:

1	2	3	4	5
a	a	b	d	c
b	d	a	b	e
c	c	d	a	a
d	b	c	c	b
e	e	e	e	d

Suppose the given partition of the schools is $S_1 = \{a, d\}$, $S_2 = \{b, e\}$, and $S_3 = \{c\}$. Then, $N_1 = \{1, 2, 4\}$, $N_2 = \{3\}$ and $N_3 = \{5\}$. In this case

$$\begin{aligned}\mathcal{C} = \Big\{&(1), (2), (3), (4), (5), (13), (23), (43),\\ &(15), (25), (45), (35), (135), (153), (235), (253), (435), (453)\Big\}.\end{aligned}$$

Note that cycle (3) can be uncovered only if b has 3 as its top priority student; cycle (25) can be uncovered only if c has 2 as its top priority student and a has 5 as its top priority student; finally (435) can be uncovered only if schools d, b, and c have students 3, 5, and 4 as their top priority students respectively. We distinguish the cycle (453) from (435) because (453) can be uncovered only if schools d, b, and c have students 5, 3, and 4 as their top priority choices respectively. (Thus the top priority students of d and b are switched in these two cycles.) We turn now to $\mathcal{C}^2$. By definition, $\mathcal{C}^2$ consists of products of two distinct cycles of $\mathcal{C}$ that can be simultaneously present in a priority profile. Thus

$$\begin{aligned}\mathcal{C}^2 = \Big\{&(1)(3), (1)(5), (2)(3), (2)(5), (4)(3), (4)(5), (3)(5),\\ &(13)(5), (23)(5), (43)(5), (15)(3), (25)(3), (45)(3), (35)(1), (35)(2), (35)(4)\Big\}.\end{aligned}$$

Note that the product of cycles (1)(4) is *not* in $\mathcal{C}^2$ even though the individual cycles (1) and (4) are both in $\mathcal{C}$. This is because it is not possible for both agents 1 and 4 to be part of a cycle in the first iteration (as their most preferred houses both belong to S_1). Similarly, it is easy to check that

$$\mathcal{C}^3 \;=\; \Big\{(1)(3)(5), (2)(3)(5), (4)(3)(5)\Big\},$$

and $\mathcal{C}^\ell = \emptyset$ for any $\ell > 3$.

Before stating the main result, we prove the following lemma on the cycle-structure of permutations. This is a well-known result, but we present a proof for the sake of completeness.

Lemma 1 *Let $c(i, \ell)$ be the number of permutations of m elements that can be written as a product of ℓ cycles. Then, for any $m > 1$,*

$$\sum_{\ell=1}^{m} (-1)^{\ell-1} c(m, \ell) \;=\; 0.$$

Proof. For any $m > 1$,

$$c(m, \ell) \;=\; c(m-1, \ell-1) + (m-1)c(m-1, \ell). \tag{1}$$

The LHS counts the number of permutations of m elements that can be written as a product of ℓ cycles. The RHS counts the same set in a different way: either the element m is in a cycle by itself, in which case the rest of the elements have to be written as a product of $\ell - 1$ cycles; or, the element m is in a cycle with some other element(s), in which case the remaining elements have to be written as a product of ℓ cycles, and element m can be inserted into this product in $(m-1)$ ways. We use the recurrence relation (1) to prove the Lemma by induction.

For $m = 2$, it is easy to check that $c(m, 1) = c(m, 2) = 1$, so the Lemma is true. For any $m > 2$, we have:

$$\begin{aligned}
\sum_{\ell=1}^{m}(-1)^{\ell-1}c(m,\ell) &= \sum_{\ell=1}^{m}(-1)^{\ell-1}c(m-1,\ell-1)+\sum_{\ell=1}^{m}(-1)^{\ell-1}(m-1)c(m-1,\ell) \\
&= \sum_{r=0}^{m-1}(-1)^{r}c(m-1,r)+(m-1)\sum_{\ell=1}^{m-1}(-1)^{\ell-1}c(m-1,\ell) \\
&= (-1)\Big(-c(m-1,0)+\sum_{r=1}^{m-1}(-1)^{r-1}c(m-1,r)\Big) \\
&= 0.
\end{aligned}$$

The first equality follows from the recurrence relation (1) and the second from a simple change of variable. The next two equalities follow from the induction hypothesis, and the observation that $c(m,0)=0$ for any $m \geq 1$. ■

We are now ready to state the main result.

Theorem 2 *The random assignment matrix determined by the RP mechanism is identical to the random assignment matrix of the PRP mechanism, for any partition of the schools.*

Proof. The result is clearly true when there is one student and one school. Consider an instance with n students and n schools, and suppose the result holds for all smaller problems. Fix a partition $S_1, S_2, \ldots, S_k$ of the schools.

Let $RP(X)$ denote the outcome of the RP mechanism when the agents in $N \setminus X$ are removed along with their most-preferred schools. Since each agent is equally likely to be the first agent in the uniform random ordering chosen by the RP mechanism, we can conclude that

$$\mathrm{RP}(N) = \sum_{i=1}^{n}\frac{1}{n}\Big\{Y_i + \mathrm{RP}(N\setminus\{i\})\Big\}, \tag{2}$$

where Y_i indicates the assignment of student i to his most-preferred school.

Now we turn to the PRP mechanism. As discussed earlier, the PRP mechanism works by finding a cycle (that is necessarily in $\mathcal{C}$ by definition). Let $P(C)$ be the probability that cycle C *can* be uncovered (i.e., cycle C is present in a priority profile that is drawn uniformly at random from all possible priority profiles), and let $Y_{\overline{C}}$ indicates the assignment in which each student in $\overline{C}$ is allocated his most-preferred school. Thus, analogous to Eq. (2), we may be tempted to write

$$\mathrm{PRP}(N) \;=\; \sum_{C\in\ \mathcal{C}} P(C)\Big\{Y_{\overline{C}} \;+\; \mathrm{PRP}(N\setminus\overline{C})\Big\},$$

This, however, is incorrect: it is possible for cycles C and C' to both be present in a given profile, in which case that profile is counted once for C and once for C'. The profiles that are over-counted are precisely those in which a cycle in $\mathcal{C}^2$ is present; but subtracting these would under-count profiles that contain three distinct cycles from $\mathcal{C}$: these are counted thrice, once for each cycle, but subtracted thrice as well, once for each pair of cycles, so we need to add back the profiles that contain a cycle in $\mathcal{C}^3$. Taking this argument to its logical conclusion, we find:

$$\mathrm{PRP}(N) \;=\; \sum_{\ell=1}^{k}\sum_{C\in\ \mathcal{C}^{\ell}} (-1)^{\ell-1}P(C)\Big\{Y_{\overline{C}} \;+\; \mathrm{PRP}(N\setminus\overline{C})\Big\}, \tag{3}$$

where $Y_{\overline{C}}$ and $\mathcal{C}^{\ell}$ are as defined before.

Consider any subset $\overline{X}$ of students, and suppose $|\overline{X}| = m$. Suppose there is a cycle C in $\mathcal{C}^{\ell}$ such that $\overline{C} = \overline{X}$. Observe that $P(C) = 1/n^m$—for any agent in $\overline{X}$, the cycle C assigns his most-preferred school to a particular agent (also in $\overline{X}$) and this occurs with probability $1/n$, independently, for each of the m agents in $\overline{X}$. If, on the other hand, there is no cycle in $\mathcal{C}^{\ell}$ involving the agents in $\overline{X}$, we can take $P(\overline{X}) = 0$. In either case, we can write $P(\overline{X})$ instead of $P(C)$. Furthermore, the term inside the braces in the RHS of Eq. (3) depends only on $\overline{X}$, not on the cycle X itself. Using these observations, we can rewrite Eq. (3) as

$$\mathrm{PRP}(N) \;=\; \sum_{\overline{X}:\overline{X}\subseteq N}\Big(Y_{\overline{X}} \;+\; \mathrm{PRP}(N\setminus\overline{X})\Big)P(\overline{X})\Bigg\{\sum_{\ell=1}^{k}\sum_{\substack{C\in\ \mathcal{C}^{\ell}\\ \overline{C}=\overline{X}}}(-1)^{\ell-1}\Bigg\}. \tag{4}$$

We now focus on the term inside the braces in Eq. (4). As $|\overline{X}|$ has only m students, it is clear that for any $\ell > m$, no cycle $C\in\mathcal{C}^{\ell}$ can have $\overline{C} = \overline{X}$. This means that the upper limit of the first summation in that term can be changed from k to m. Therefore, we have

$$\sum_{\ell=1}^{k}\sum_{\substack{C\in\ \mathcal{C}^{\ell}\\ \overline{C}=\overline{X}}}(-1)^{\ell-1} \;=\; \sum_{\ell=1}^{m}\sum_{\substack{C\in\ \mathcal{C}^{\ell}\\ \overline{C}=\overline{X}}}(-1)^{\ell-1} \;=\; \sum_{\ell=1}^{m}(-1)^{\ell-1}c(m,\ell) \;=\; \begin{cases}1, & \text{if } m=1,\\ 0, & \text{if } m>1,\end{cases}$$

where the last expression follows from the definition of $c(m, \ell)$ and Lemma 1. Thus the only terms that survive in the RHS of Eq. (4) are those for which $|\overline{X}| = 1$. In this case, however, $P(\overline{X}) = 1/n$, and Eq. (4) simplifies to

$$\text{PRP}(N) \;=\; \sum_{i=1}^{n} \frac{1}{n} \Big\{ Y_i \;+\; \text{PRP}(N \setminus \{i\}) \Big\} \tag{5}$$

By the induction hypothesis, $PRP(N \setminus \{i\}) \;=\; RP(N \setminus \{i\})$ regardless of the partition of the schools (in the problem without student i). We see that the expressions determining $PRP(N)$ and $RP(N)$ are identical, establishing the result. ■

An important consequence of Theorem 2 is the analog of Pathak's result for school choice in which schools can admit more than one student.

Theorem 3 *Consider a school choice problem with n students and k schools. Suppose school i can admit q_i students, and suppose $\sum_i q_i = n$. The expected assignment matrix obtained by the TTC algorithm is the same, whether a common lottery is used for all the schools, or whether each school uses its own independent lottery.*

Proof. We make q_i copies of school i and arrive at an equivalent one-to-one matching problem. Place all the copies of a given school into one partition, so that the partition structure of the "schools" for the PRP mechanism is precisely $S_1, S_2, \ldots, S_k$, where S_i consists of all copies of school i. The RP mechanism finds the outcome when a common ordering used by all the schools; the PRP mechanism finds the outcome in which each school uses an independent ordering of the students. The result now follows from Theorem 2. ■

We state and prove a parallel result for the PRE mechanism.

Theorem 4 *The random assignment matrix determined by the RP mechanism is identical to the random assignment matrix of the PRE mechanism, for any partition of the schools.*

Proof. The proof is identical to the proof of Theorem 2 with the only change being the expression for $P(C)$. For the PRP mechanism we saw that $P(C) = 1/n^m$ whenever C is a cycle involving m students; for the PRE mechanism, however,

$$P(C) = \frac{1}{n}\frac{1}{n-1}\cdots\frac{1}{n-m+1}.$$

This is so because the most-preferred schools of the students involved in the cycle C all are required to have distinct top-priority students. The exact expression for $P(C)$ does not matter whenever C has more than one student because of the "cancellation" of terms (by Lemma 1). When C has exactly one student, or equivalently when $m = 1$, $P(C) = 1/n$ as in the proof of Theorem 2, and the result follows. ■

Remarks.

1. Theorem 4 provides an alternative proof of the equivalence of RP and RE mechanisms: this alternative proof relies on induction and counting, instead of the original bijective argument, due independently to Knuth [5] and Abdulkadiroglu & Sönmez [2]. A similar comment applies to the alternative equivalence proof of RP and IRP, implicit in Theorem 2.

2. The discussion preceding Lemma 1 suggests a simple equivalence proof of the RE and IRP mechanisms. Recall that in the IRP mechanism, each school is in a partition by itself and generates its priority list uniformly at random, independently of the other schools. As the sets $\mathcal{C}^\ell$ are identical for both RE and IRP for any $\ell \geq 1$, we may conclude that the expressions— like Eq. (3)—for the RE and IRP mechanisms will be identical, making Lemma 1 unnecessary. Note, however, that Lemma 1 is still needed: although RE and IRP are very similar, $P(C)$ is different for these two mechanisms, so the recursive expressions are indeed different.

4 House Allocation with Existing Tenants

We can use these ideas to prove alternative and perhaps simpler proof of another equivalence result, due to Sönmez and Ünver [8], on house allocation with existing tenants. In an instance of this problem there are n agents and n houses; $(n - r)$ of the agents are existing tenants, each of whom already occupying a distinct house; the remaining r agents and the r unoccupied houses constitute the "new" arrivals to the problem. There is a natural generalization of the RP and RE mechanisms to this setting. First we describe the RP mechanism: order the agents uniformly at random, and let agents make their choices in this order, except that whenever an agent j demands a house already occupied by an agent j', then j' gets promoted to just before j in the ordering. In this case, j' is invited to make his choice, and j picks only after j' is assigned a house (that may possibly be the one that he already occupies). For the RE mechanism: endow the *new* houses uniformly at random to the *new* agents, each existing tenant is endowed the house he occupies, and apply the TTC algorithm to this trading game. It is not hard to see that, as defined, the RE mechanism is *not* equivalent to the RP mechanism. This is because the new agents are treated preferentially in the RE mechanism: if j' is an existing tenant who happens to occupy a house that is the last choice for all the agents (including himself!), j' cannot be assigned any other house under the RE mechanism, but he could get other houses under RP. Sönmez and Ünver [8] modify the RP mechanism in the following way: the new agents are ordered uniformly at random, and the existing tenants appear *after* the new agents in an *arbitrary* order [3]. They show that this modified RP mechanism is equivalent to the RE mechanism. We prove this result next. (To emphasize the distinction in the context, we use the terms "agents" and "houses" in this section instead of students and schools.)

Theorem 5 *For the house allocation mechanism with existing tenants, the RE mechanism (which assigns the new houses randomly to the new agents) is equivalent to the RP mechanism in which all the new agents are ordered uniformly at random before any of the existing tenants.*

Proof. We prove the result by induction on the number of existing tenants. If there are no existing tenants, the result follows from Theorem 4.

Suppose there are ℓ existing tenants. We may assume that initially there is no cycle involving the existing tenants alone: that is, there is no set of existing tenants $i_1, i_2, \ldots, i_m$ such that the most preferred house of agent i_j is the one occupied by i_{j+1} (where i_{m+1} is read as i_1). If there is one, observe that the RE and RP mechanisms will *always* assign these houses to these agents such that each agent in this set gets their most-preferred house; we can then delete these agents and their respective houses, and the result would follow by induction (as we have an instance with fewer existing tenants).

[3]We could of course look at uniform random orderings of the existing tenants as well, but as it turns out their exact ordering does not matter as long as they appear *after* the new applicants. This will be apparent from the proof.

Suppose there are n agents and n houses in all, and suppose there are r new applicants, which, without loss of generality, we assume to be agents $1, 2, \ldots, r$. Then, in the RP mechanism, each of the r new agents is equally likely to be the first. Suppose (new) agent i is first. As we assume that there is no cycle involving the existing tenants alone, observe that agent i will necessarily receive his first choice—either i's first choice is an unoccupied house, or i's first choice is an occupied house, in which case the owner of that house will ask for his best choice; this latter house is either unoccupied (in which case the chain stops), or is occupied by another existing tenant, who will now ask for *his* best choice, etc. By our assumption, this chain will have to end with an existing tenant asking for an unoccupied house. The agents involved in this chain, including agent i, will all get their most-preferred house. Let C_i be the set of agents involved in the chain started by agent i. Note that C_i contains i and does not contain any other new applicant. The relevant recursion for the RP mechanism is

$$\mathrm{RP}(N) \;=\; \sum_{i=1}^{r} \frac{1}{r}\Big\{Y_{C_i} \;+\; \mathrm{RP}(N \setminus \{C_i\})\Big\},$$

where Y_{C_i} indicates the assignment in which the agents in C_i get their most-preferred houses.

Let $\mathcal{C}$ be the set of cycles that can be uncovered as a first cycle by the RE mechanism. As before, we think of cycles as involving just the agents as it is understood that the houses involved are all the most-preferred houses of the agents in the cycle. Let $\mathcal{C}^\ell$ be defined analogously as well, for $\ell = 1, 2, \ldots$. Finally, note that the existing tenants are always endowed with the houses they occupy. Clearly,

$$\mathrm{RE}(N) \;=\; \sum_{\ell=1}^{r} \sum_{C \in \mathcal{C}^\ell} (-1)^{\ell-1} P(C)\Big\{Y_{\overline{C}} \;+\; \mathrm{RE}(N \setminus \overline{C})\Big\}, \qquad (6)$$

where $Y_{\overline{C}}$ and $\mathcal{C}^\ell$ are as defined before.

Consider an arbitrary set $\overline{X} \subseteq N$ of agents. Suppose there is a cycle C in $\mathcal{C}^\ell$ such that $\overline{C} = \overline{X}$. As the endowment of the occupied houses is fixed, it is clear that for any agent $i \in C$, if i's most-preferred house is an occupied house, then i *must* be followed by the existing tenant occupying that house. So the only agents in C for whom a successor is not already determined are those whose most-preferred house is *new*. Suppose there are exactly m of these agents in C. Then, observe that

$$P(C) = \frac{1}{r}\frac{1}{r-1}\cdots\frac{1}{r-m+1}.$$

This is because the new houses that each of these agents desire must each be endowed to a distinct new applicant—there are r choices for the first new applicant, $r-1$ for the second, and so on. As before, if there is no cycle in $\mathcal{C}^\ell$ involving the agents in $\overline{X}$, we can take $P(\overline{X}) = 0$. In either case, we can write $P(\overline{X})$ instead of $P(C)$, as the probability does not depend on the cycle, but only on the set of agents involved in the cycle (in fact, it depends only on m and r). Furthermore, the term inside the braces in the RHS of Eq. (6) depends only on $\overline{X}$, not on the cycle X itself. Using these observations, we can rewrite Eq. (6) as

$$\mathrm{RE}(N) \;=\; \sum_{\overline{X}:\overline{X}\subseteq N} \Big(Y_{\overline{X}} \;+\; \mathrm{RE}(N\setminus\overline{X})\Big)P(\overline{X})\Big\{\sum_{\ell=1}^{r}\sum_{\substack{C\in\,\mathcal{C}^\ell\\ \overline{C}=\overline{X}}}(-1)^{\ell-1}\Big\}. \tag{7}$$

We now focus on the term inside the braces in Eq. (7). As $|\overline{X}|$ has only m agents whose most-preferred house is new, and as each cycle in $\mathcal{C}$ has at least one of these agents, it is clear that for any $\ell > m$, no cycle $C \in \mathcal{C}^\ell$ can have $\overline{C} = \overline{X}$. This means that the upper limit of the second summation can be changed from r to m. Therefore, we have

$$\sum_{\ell=1}^{r}\sum_{\substack{C\in\,\mathcal{C}^\ell\\ \overline{C}=\overline{X}}}(-1)^{\ell-1} \;=\; \sum_{\ell=1}^{m}\sum_{\substack{C\in\,\mathcal{C}^\ell\\ \overline{C}=\overline{X}}}(-1)^{\ell-1} \;=\; \sum_{\ell=1}^{m}(-1)^{\ell-1}c(m,\ell) \;=\; \begin{cases}1, & \text{if } m=1,\\ 0, & \text{if } m>1,\end{cases}$$

where the last expression follows from the definition of $c(m,\ell)$ and Lemma 1. Thus the only terms that survive in the RHS of Eq. (7) are those for which there is exactly one agent whose most-preferred house is new. This also implies that there is exactly one new applicant in the cycle—as the new houses are all endowed to the new applicants, the number of new houses in any cycle should equal the number of new applicants in the cycle. In this case, however, $P(C) = 1/r$, and Eq. (6) simplifies to

$$\mathrm{RE}(N) \;=\; \sum_{i=1}^{r}\frac{1}{r}\Big\{Y_{C_i} \;+\; \mathrm{RE}(N\setminus\{C_i\})\Big\} \tag{8}$$

By the induction hypothesis, $RE(N\setminus\{i\}) \;=\; RP(N\setminus\{i\})$. We see that the expressions determining $RE(N)$ and $RP(N)$ are identical, establishing the result. ■

Acknowledgments

I am grateful to Thiam Lee and Hervé Moulin for many insightful discussions during the course of this research. It was an early conversation with Thiam Lee that led to a formulation of the PRP mechanism; and Hervé Moulin suggested that I think about the PRE mechanism. I thank Stergios Athanassoglou for detailed comments on an earlier draft of the paper.

References

[1] A.Abdulkadiroglu, P. A. Pathak, and A. E. Roth, "The New York City high school match," *American Economic Review, Papers and Proceedings*, 95, 364–367, 2005.

[2] A.Abdulkadiroglu and T.Sönmez, "Random serial dictatorship and the core from random endowments in house allocation problems", *Econometrica*, 66(3):689–701, 1998.

[3] A.Abdulkadiroglu and T.Sönmez, "School Choice: A Mechanism Design Approach", *American Economic Review*, 93(3):729–747, 2003.

[4] A. Erdil and H. Ergin, "What's the Matter with Tie-Breaking? Improving Efficiency in School Choice," *American Economic Review*, 98(3):669–689, 2008.

[5] D. E. Knuth, "An exact analysis of stable allocation," *Journal of Algorithms*, 20(2):431–442, 1996.

[6] P. A. Pathak, "Lotteries in Student Assignment: The Equivalence of Queueing and a Market-Based Approach," *Working Paper*, 2008.

[7] L. Shapley and H. Scarf (1974), "On cores and indivisibility," *Journal of Mathematical Economics*, 1, 23–28.

[8] T. Sönmez and M. Utku Ünver, "House Allocation with Existing Tenants: An Equivalence," *Games and Economic Behavior*, 52(1):153–185, 2005.